þÿ P A N N Z E R   a p r a c t i c a l t o o l f o r p r o t e i n f u n c t i o n

Törönen, Petri

2022-01

"PANNZER – a practical tool for protein function prediction"

Petri Törönen[1] and Liisa Holm[1,2]

[1]Institute of Biotechnology, Helsinki Institute of Life Sciences, and [2]Organismal and Evolutionary Biology Research Program, Faculty of Biosciences, University of Helsinki, Finland

ABSTRACT

The facility of next-generation sequencing has led to an explosion of gene catalogs for novel genomes, transcriptomes and metagenomes, which are functionally uncharacterized. Computational inference has emerged as a necessary substitute for first-hand experimental evidence. PANNZER (Protein ANNotation with Z-scoRE) is a high-throughput functional annotation web server that stands out among similar publically accessible web servers in supporting submission of up to 100,000 protein sequences at once and providing both Gene Ontology (GO) annotations and free text description predictions. Here, we demonstrate the use of PANNZER and discuss future plans and challenges. We present two case studies to illustrate problems related to data quality and method evaluation. Some commonly used evaluation metrics and used evaluation datasets promote methods that that favor unspecific and broad classes over more informative and specific classes. We argue that this can bias the development of automated function prediction methods. The PANNZER web server and source code are available at http://ekhidna2.biocenter.helsinki.fi/sanspanz/.

INTRODUCTION

Climate change, environmental problems, new emerging pathogens and growing human population have created new critical research questions. Lowering sequencing costs have made it reasonable to study relevant species, linked to these research questions, with next-generation sequencing projects. Genome and transcriptome projects have sequenced various agricultural plants, microbes used in biochemical processing or biofuel generation, animal and human pathogens, or microbes used in food processing.[1-4] These genome projects generate massive amounts of sequences. Often the most relevant set of sequences are the predicted genes, extracted from the sequenced genomes. These extracted genes lack any information on what they

actually do. Therefore, in order to make this new genome data useful, these genes must be associated with the relevant biological features. Annotations can be used to generate hypotheses how the studied agricultural plant reacts to pathogens or what differentiates a multidrug-resistant bacterium from its relative strains. Annotations could be done manually, but even with one genome, often with 20 000 - 30 000 genes, it is too overwhelming a task. Therefore, we need Automated Function Prediction methods (AFP methods) that can automatically predict these features.

The goal of the AFP methods is to predict the functions for the studied sequences and communicate these generated predictions to end-users. This function can be presented in a few ways: i) sequence can have a short descriptive text, ii) sequence can have a longer more detailed description, iii) sequence can be classified to a sequence family with conserved function,[5] and/or iv) sequence can be classified to known signaling or reaction pathways. Short descriptions give a quick summary on the sequence, especially when one is working with FASTA sequence files. They are also required for submitting the sequences to a database. Longer detailed descriptions, presented by databases like SwissProt[6] and RefSeq,[7] are comprehensive descriptions on various sequence functions including literature citations. Sequences can be also classified to various functional classification hierarchies, with Gene Ontology (GO)[8] being the most popular. The classes (also referred to as terms) in GO represent varying levels of information on the gene's function. Still, even the most specific classes often lose some details on the protein function, when compared to a free text description. Classes still have benefits. One sequence can have multiple functions and these can be represented with multiple classes. Furthermore, the classes also allow studies on groups of sequences, where scientists can look for over-represented functions from the set of sequences. Finally, the function can be presented with various signaling and reaction pathways. These can be used similarly to the aforementioned classes. Here KEGG (Kyoto

Encyclopedia for Genes and Genomes[9]) pathway database is the most commonly predicted

pathway collection. Most of these data types can be predicted with AFP tools.

None of the existing AFP methods invents the function for novel sequence out of nowhere. They

rather use mainly classifiers on various sequence and interaction network derived features,[10,11]

pre-annotated protein families[12-14] and/or K-nearest sequence principles.[15,16] Altogether, AFP

methods often look for various similarities, between the sequences with known function and the

query sequence, and next transfer the function(s) to the query sequence if the similarities are

strong enough (Figure 1). This means that even the most advanced AFP methods are only as good

as the available training data they use. Most of the existing annotations cumulate to model

organisms. Therefore, the more exotic species, with no closely related model organism, is likely to

get weakly annotated by all methods. The annotation efficiency is skewed across the different

functions and different species. One solution to the lack of data is to use predicted functions also

as an input in our predictions. Unreviewed GO annotations are tagged as such, allowing their

inclusion or exclusion from the training data. Inclusion has its pros and cons. 99% of the current

GO annotations for Uniprot are predictions, making our training data significantly larger. The

drawback is the inclusion of wrong predictions into the training data. Thus, the repetitive use of

predictions could generate a cascade of decisions, where a prediction is used as an input for

prediction, which in turn is used as an input for a prediction, leading to error propagation.[17]

The hierarchical structure of the GO means that any sequence, annotated with GO term X, will

automatically also get annotated with parent terms of X. For a sequence, predicted to belong to

ribosome, the GO hierarchy implies, among others, the following classes: ribosome, translation,

protein metabolic process, organic substance metabolic process, cellular metabolic process

(http://amigo.geneontology.org/amigo/term/GO:0044391#display-lineage-tab). Notice how

information, related to function, decreases as we go forward on the list. These broader classes represent parent terms. This hierarchical GO structure causes larger, less informative GO terms to be more frequent in any training or testing datasets. This in turn causes AFP methods to easily report only unspecific, vague GO terms. Therefore, many AFP methods have been based on enrichment statistics.[15,16] Rather than looking what GO classes are frequent in the neighborhood of the query sequence, they look for annotations that are more frequent in the sequence neighborhood than in the whole database. The hierarchical structure of GO also complicates the evaluation of AFP methods. There is ongoing debate how to properly evaluate AFP models.[18] One of the biggest issues with the evaluation metrics is that one can get very good results, with some evaluation metrics, by simply reporting the GO classes in decreasing order of their frequency in the database, for every tested gene.[19,20] This baseline model, often referred as naïve model,[21] constitutes no information to end user.

We have developed PANNZER, a weighted K-nearest neighbor classifier for protein function prediction. This paper demonstrates the use of PANNZER and recently added features (taxonomic filtering and gene name prediction), future plans and challenges. We present two case studies to illustrate problems related to data quality and AFP method evaluations. We point that some used evaluation metrics and used evaluation datasets promote methods that that favor unspecific and broad classes over more informative and specific classes. We argue that this can bias the development of AFP methods.

PANZZER Resource

PANNZER is a fast and fully automated tool and web server for the AFP task. PANNZER is one of the few annotation tools that allow genome sized queries and estimate the reliability of the given

prediction (Table 1). PANNZER can predict short sequence description to any species (plants, bacteria, animals). PANNZER outperformed other tools in our comparisons across all datasets (8% - 45% higher Fmax scores, for example[22]). In addition, PANNZER has performed well in the international competitions, organized for the AFP methods (ranked overall 3rd in 2013 and frequently between 5th and 7th in 2019[21,23,24]) and also in comparisons done by other groups.[25-27]

Inputs

The inputs are FASTA formatted protein sequences. These sequences can be translated coding sequences from a genome project. Bacterial protein coding genes can be predicted from the genomic sequence using e.g. Prodigal.[28] Eukaryotic gene prediction is more complicated. Gene models can be cross-mapped from a related genome, guided by homologous proteins or RNA-seq data, or predicted *ab initio*.[29] RNA-seq data should be either mapped to a reference genome, or assembled *de novo* and translated to peptides using e.g. the Trinity and TransDecoder.[30]

Parameters

PANNZER2 uses a fast suffix array neighborhood search (SANSparallel[31]) to find homologous sequences in the UniProt database. We refer to homology search results as the *sequence neighborhood*. By default, PANNZER2 uses a maximum of 100 database hits. As we are transferring annotations based on sequence similarity, it is necessary for sequence matches to meet several criteria for inclusion in the sequence neighborhood. Search results must have at least 40% sequence identity, 60% alignment coverage of both the query and target sequences, and a minimum of 100 aligned residues. We refer to this step as *sequence filtering*. The criteria for sequence filtering can be changed from the Advanced parameters on the web form. Shorter alignment lengths and relaxed coverage requirements can be useful in situations where the query sequences are short fragments or come from an exotic organism.

Output is usually reduced to only one predicted description and to non-redundant GO terms (see Figure 2). PANNZER ranks GO predictions according to an enrichment statistic, which compares the frequency of a GO class in the sequence neighborhood to its frequency in the database.[15,22] The non-redundant subset removes all GO classes that have a higher ranked parent or descendant.

Sometimes homology transfer results in biologically implausible annotations such as predicting an organ in plants that only exists in animals. Per a user request, we have implemented a set of taxonomic branch specific GO subsets. These limit the GO predictions only to GO terms, occurring in manually curated annotations in fungi, plants, vertebrates, arthropods or bacteria, respectively. This branch specific GO filter is off by default.

Outputs

Pannzer predicts short descriptions (DE), GO-terms, gene names, EC classes and KEGG pathways. The web server produces nested tables for browsing (Figure 2) and a tabular format for downloading. PANNZER returns Positive Predictive Value, PPV, as a reliability estimate. PPV estimates the probability of the annotation being correct. Here, correct descriptions are a short TF-IDF distance (see[15]) from the most precise terms in the ground truth as given by the Uniprot database. The Uniprot database only reports the most specific terms of a protein, leaving out the implicit parent terms. Pannzer predictions for a query protein are targeted to these biologically most informative GO terms. Note that our definition of correct annotations differs from the probability of class membership of any and all GO classes in the GO hierarchy, which is evaluated e.g. in CAFA.[21,23,24]

Web server

PANNZER2 provides a fast, publically accessible web server for functional annotation.[22] PANNZER2 uses SANSparallel for high-performance homology searches, making bulk annotation based on

sequence similarity practical. It accepts query sets up to 100,000 protein sequences. Its throughput exceeds 1 M query sequences per day, roughly corresponding to 200 bacterial genomes or 10 eukaryotic transcriptomes. The results can be either downloaded for local post-processing or browsed via a web application. PANNZER2 outputs both GO and DE predictions. The web server front-end provides links to homology search results for each query sequence, enabling users to see how predictions were derived.

Software

PANNZER2 is implemented in a Python framework called SANSPANZ. The SANSPANZ framework frees method developers from data management, as homology searches and fetching metadata can be delegated to remote servers, though it is also possible to install the servers locally. The framework includes "functors" performing computations on these data. For example, SANSPANZ includes implementations of the scoring functions from ARGOT,[16] BLAST2GO[32] and PANNZER,[15] as well as hypergeometric enrichment and best informative hit. Novel application workflows are implemented by chaining "functors" one after another. PANNZER2[15] and AAI-profiler for taxonomic profiling[33] are fully integrated into SANSPANZ, while LazyPipe for virus identification[34] is partly built on it. SANSPANZ is open source (http://ekhidna2.biocenter.helsinki.fi/sanspanz/#tab-2).

Data servers

Both the SANSparallel server and DictServer are maintained by our group. The SANSparallel server takes a protein sequence as input and returns a list of similar sequences from the Uniprot database. Both GO annotations and free text descriptions are gathered for each search result by calling the DictServer. The DictServer is a store of key/value pairs. The servers work locally as a client-server or they can be accessed remotely via CGI. We keep a local copy of the databases used

by PANNZER2 (Appendix 2). The databases are updated on a regular schedule, ensuring that predictions benefit from new data. Our database update cycle is synchronized with the bimonthly new releases of the Uniprot database. The Uniprot database is mirrored from EBI, indexed and imported to the SANSparallel server. GOA annotations, the GO structure and taxonomy lists are downloaded and imported to the DictServer with recomputed background statistics for DE and GO enrichment analysis.

Case study 1: data conflicts

Database annotations have varying quality. Here, we have deliberately chosen an example, which does not run smooth in Pannzer analysis. The example is a predicted L-galactonate dehydratase-like protein from the cork oak (*Quercus suber*) genome assembly CorkOak1.0.[35] The protein sequence (Appendix 1) was submitted to the Pannzer web server (Figure 2). Clicking the "search" link in the web server output calls the SANSparallel server and displays the sequence neighborhood on which the predictions are based. Surprisingly, the sequence neighborhood consists of proteins from unexpected taxa: they are all fungi. This is not an isolated case, either. Taxonomic profiling by PANNZER2's companion tool, AAI-profiler,[33] shows that 21% of the proteins predicted from the cork oak genome have a closest match in fungi. The proteome is 97% complete with respect to BUSCO's single-copy ortholog set for *Ascomycota* and 96% complete with respect to that for eudicots.[36] We conclude that the cork oak genome assembly includes a fungal co-isolate.

Gene names and descriptions are free text. Pannzer applies spam filters and clustering to identify informative entries with the most support in the data. The gene name is identified as *lgd1* (Figure 2, column 2). This gene was first characterized in *Trichoderma reesei* and shown to convert L-galactonate to 2-dehydro-3-deoxy-L-galactonate and to be required for growth on D-

galacturonate.[37] *lgd1* occurs eight times in the sequence neighborhood; all other gene names are unique genomic locus identifiers typified by the use of underscores. Descriptions fall into three clusters (Figure 2B, column 3). The alternative descriptions represent different levels of evolutionary classification. Galactonate dehydratase is a subgroup of the mandelate racemase/muconate lactonizing (MR/MLE) family, which in turn belong to the broad, functionally diverse enolase superfamily.[38] L-galactonate dehydratase is the preferred description based on sequence similarity weighting. The MR/MLE family and enolase superfamily annotations are based on InterPro profile matches.[5] Phylogenetic analysis confirms that our text example is firmly embedded in the L-galactonate dehydratase subgroup (Figure 3).

The second surprise is that, at first glance, the predicted description and predicted GO classification contradict each other (Figure 2, column 3 vs. column 5). L-galactonate dehydratase activity is a sister class of the predicted L-fuconate dehydratase activity in GO's molecular function ontology. The GO prediction has high confidence, because 97 of 100 entries in the sequence neighborhood are annotated with L-fuconate dehydratase activity. Perusal of the literature shows that neither of these predictions is wrong.

L-galactonate dehydratase (gaaB) and L-fuconate dehydratase (fucD) are promiscuous, accepting both L-fuconate and L-galactonate as substrate.[39,40] L-Galactonate and L-fuconate have a similar chemical structure; the only difference is that L-fuconate has no hydroxyl group at the C6. The source material for Pannzer's GO predictions consisted solely of annotations automatically inferred using Interpro2GO. The experimentally characterized proteins fall outside the sequence neighborhood of our test example, which was restricted to 100 hits. The GO annotations for the Lgd1 protein are inferred from direct assay (IDA) as carbonate dehydratase activity and D-galacturonate catabolic process, a descendant of carboxylic acid catabolic process (cf. Figure 2B,

column 4). These annotations have not propagated as far in the database as the gene name or description, which were picked up by Pannzer.

L-fuconate catabolism has not been demonstrated in fungi.[39] Accordingly, if the taxonomic branch specific GO filter for fungi is activated in Pannzer, L-fuconate dehydratase activity is not reported and hydro-lyase activity becomes the top prediction and also the KEGG pathway prediction (Figure 2, column 7) drops away. Here, the taxonomic filter is set to fungi based on the previous taxonomic observations.

PANNZER results presented on the web page (Figure 2) are hyperlinked to sequences, metadata, and sequence databases with further links to literature. This facilitates tracking down the origins of unexpected results as illustrated here.

Case Study 2: comparison of results by AFP methods

Our second example represents results from various web servers for a query sequence KAG7012684.1. This is a ribosomal protein from a recently sequenced pumpkin genome (GenBank date 16-JUN-2021).[41] This should be clearly an easy case for AFP methods, as a) there are several reasonably well annotated plant genomes available and b) ribosome is well annotated across the species. So here we look what different methods report in quite a trivial case. All the compared AFP methods of Table 1 have been among the top 10 in CAFA competitions. In addition, we tested also eggNOG,[14] a popular annotation tool. The results (Table 2) also include the naïve model. The naïve model does not use any information on the sequence. It simply reports the frequency of each GO term in the whole database.

The results show different behavior of methods: PANNZER,[22] ARGOT[16] and PFP[42] report a group of informative GO term as their top predictions. These are linked to ribosomal RNA binding and to ribosome as expected. By default, PANNZER shows a non-redundant list of GO terms; without the

GO filtering, it also shows parent or child terms of the strongest predictions. FunFam[12] maps the query sequence to superfamilies, which include functionally diverse members. FunFam predictions included mRNA 3'/5'-UTR binding and ubiquitin ligase binding activities, as well as the rather ubiquitous Protein Binding. EggNOG[14] surprisingly did not predict any GO classes. INGA[43] reports a bag of GO classes with the same score. NetGO[10,44,45] shows different behavior from other methods. NetGO is most confident in predicting the root class and classes close to root resembling, in some ways, the naïve model. Indeed, NetGO uses the naïve model as one of its information sources. This kind of result ranking will improve performance in CAFA competitions, when the evaluation is done with traditional evaluation metrics. These near root classes convey, however, little information to end users. NetGO was actually the clear winner, with some evaluation metrics, in the recent CAFA3 competition.[23] Our discussion covers this topic more in detail. Altogether, these results show how different methods find and emphasize different GO terms.

DISCUSSION

Managing a comprehensive annotation pipeline involves keeping databases up-to-date and ensuring that growing disk space and memory requirements are met. This has made public web servers convenient and popular for annotation. PANNZER can be used via a web server or installed locally with programmatic access to our group's data servers. The benefits of PANNZER include its high throughput, which has allowed users to run large collection of genomes in pangenome projects and sequence collections from metagenome or transcriptome projects. PANNZER also generates independently DE and GO predictions. Time-consuming manual curation can focus on contradictory or borderline predictions and checking if they can be confirmed. Current weaknesses

include the lack of alternative information sources, like protein-protein interaction data or gene expression data. Notice that these are difficult to apply, when we are studying a novel genome. Large enough collections of interaction data and gene expression data are mostly available for model organisms and other well-studied organisms. So the analysis should first map the query sequence to a similar sequence in the interaction database, for example, and after it process the actual interaction data. Fortunately, there is currently research on this topic.[46-48] Furthermore, various profile alignment based features, like protein domains and sequence motifs, are expected to be useful for function prediction. These two sources are currently missing from PANNZER. We are, however, currently working on these data sources. In addition, the recent surge of structural models[49,50] should rekindle research into specificity determining residues.[51,52]

AFP method development and publishing require that the results from various AFP methods are compared and evaluated. PANNZER selects more informative level from the GO tree than many competing methods (see Table 2 for example). Comparison of AFP methods also requires that we have established a ground truth, i.e., suitable evaluation datasets that contain sequences with known function. The case study illustrated that defining the known function can be a precarious business. The current datasets have practically only knowledge that a protein has a certain function. We can rarely state that a protein does not have a certain function. This is often referred as positive – unlabeled learning task. This problem is further heightened by some very vague annotations that sequences can have in the databases, like Protein Binding. Any more precise annotation would be assessed as wrong prediction.[53] Notice that also the biologically correct function predictions will be evaluated as bad predictions, if this function is absent from the evaluation dataset. Other researchers have looked at these research questions in detail.[54] When the evaluation dataset has mostly annotations with very broad unspecific classes, it is better for AFP methods not to predict specific functions but rather broad and vague classes. This is why we

have used our own evaluation datasets, with detailed annotations, parallel with the popular CAFA evaluation datasets.[22] Our gold standard has been a set of GO-annotated sequences, selected from the UniProt database. We especially required that a) each sequence has, at least one, manually curated GO annotation, b) these GO annotations must occur in small informative GO classes, c) sequences should not have strong sequence similarities (see[22]). We have been able to show that PANNZER shows good performance especially when the evaluation dataset contains detailed information on functions.

The evaluation metric compares the generated predictions to known correct functions. The selection of this metric for a hierarchical classification like GO is not trivial and it has a strong effect on the results: ranking of the methods varies drastically between different metrics.[20,55] This has been shown also outside the bioinformatics field.[56,57] Our recent work used simulated data to show that some evaluation metrics fail to separate different amounts of error in the predictions.[19] Furthermore, some used evaluation metrics ranked earlier discussed naïve model, among the best predictions. This will lead to situation where the field is promoting methods that predict quite uninformative results. This can be corrected by using more advanced metrics, like weighted Jaccard correlation (SimGIC), term-centric AUCPR etc.; these metrics and further alternatives are explained in [19]. We propose that AFP researchers would demonstrate the results on several evaluation datasets and would use mainly evaluation metrics that are insensitive to biases in class sizes.

References

1. Wang B, Lin Z, Li X et al (2020) Genome-wide selection and genetic improvement during modern maize breeding. Nat Genet 52:565–571.

2. Morabito C, Aiese Cigliano R, Maréchal E, Rébeillé F, Amato A (2020) Illumina and PacBio DNA sequencing data, de novo assembly and annotation of the genome of *Aurantiochytrium limacinum* strain CCAP_4062/1. Data Brief 31:105729.

3. Simonet C, McNally L (2021) Kin selection explains the evolution of cooperation in the gut microbiota. Proc Natl Acad Sci USA 118:e2016046118.

4. Duru IC, Ylinen A, Belanov S et al (2021) Transcriptomic time-series analysis of cold- and heat-shock response in psychrotrophic lactic acid bacteria. BMC Genom 22:28.

5. Blum M, Chang H, Chuguransky S, Grego T, Kandasaamy S, Mitchell A, Nuka G, Paysan-Lafosse T, Qureshi M, Raj S, RichardsonL, Salazar GA, Williams L, Bork P, Bridge A, Gough J, Haft DH, Letunic I, Marchler-Bauer A, Mi H, Natale DA, Necci M, Orengo CA, Pandurangan AP, Rivoire C, Sigrist CJA, Sillitoe I, Thanki N, Thomas PD, Tosatto SCE, Wu CH, Bateman A, Finn RD (2021) The InterPro protein families and domains database: 20 years on. Nucleic Acids Res 49:D344-D354.

6. The UniProt Consortium (2021) UniProt: the universal protein knowledgebase in 2021. Nucleic Acids Res 49:D480-D489.

7. O'Leary NA, Wright MW, Brister JR, Ciufo S, Haddad D, McVeigh R, Rajput B, Robbertse B, Smith-White B, Ako-Adjei D, Astashyn A, Badretdin A, Bao Y, Blinkova O, Brover V, Chetvernin V, Choi J, Cox E, Ermolaeva O, Farrell CM, Goldfarb T, Gupta T, Haft D, Hatcher E, Hlavina W, Joardar VS, Kodali VK, Li W, Maglott D, Masterson P, McGarvey KM, Murphy MR, O'Neill K, Pujar S, Rangwala SH, Rausch D, Riddick LD, Schoch C, Shkeda A, Storz SS, Sun H, Thibaud-Nissen F, Tolstoy I, Tully RE, Vatsan AR, Wallin C, Webb D, Wu W, Landrum MJ, Kimchi A, Tatusova T, DiCuccio M, Kitts P, Murphy TD, Pruitt KD. Reference sequence (RefSeq) database at NCBI: current status, taxonomic expansion, and functional annotation. Nucleic Acids Res 44:D733-745.

8. Ashburner M, Ball C, Blake J et al (2000) Gene Ontology: tool for the unification of biology. Nat Genet 25:25–29.

9. Kanehisa M, Furumichi M, Sato Y, Ishiguro-Watanabe M, Tanabe M (2021) KEGG: integrating viruses and cellular organisms. Nucleic Acids Res 49:D545-D551.

10. Yao S, You R, Wang S et al (2021) NetGO 2.0: improving large-scale protein function prediction with massive sequence, text, domain, family and network information. Nucleic Acids Res 49:W469-W475.

11. Cozzetto D, Minneci F, Currant H et al (2016) FFPred 3: feature-based function prediction for all Gene Ontology domains. Sci Rep 6:31865.

12. Scheibenreif L, Littmann M, Orengo C et al (2019) FunFam protein families improve residue level molecular function prediction. BMC Bioinformatics 20:400.

13. Profiti G, Martelli PL, Casadio R (2017) The Bologna Annotation Resource (BAR 3.0): improving protein functional annotation. Nucleic Acids Res 45:W285-W290.

14. Huerta-Cepas J, Szklarczyk D, Heller D, Hernández-Plaza A, Forslund SK, Cook H, Mende DR, Letunic I, Rattei T, Jensen LJ, von Mering C, Bork P (2019) eggNOG 5.0: a hierarchical, functionally and phylogenetically annotated orthology resource based on 5090 organisms and 2502 viruses. Nucleic Acids Res 47:D309–D314.

15. Koskinen P, Törönen P, Nokso-Koivisto J, Holm L (2015) PANNZER: high-throughput functional annotation of uncharacterized proteins in an error-prone environment. Bioinformatics 31:1544-1552.

16. Lavezzo E, Falda M, Fontana P, Bianco L, Toppo S (2016) Enhancing protein function prediction with taxonomic constraints – The Argot2.5 web server. Methods 93:15-23.

17. Schnoes AM, Brown SD, Dodevski I, Babbitt PC (2009) Annotation error in public databases: misannotation of molecular function in enzyme superfamilies. PLoS Comput Biol 5:e1000605.

18. Makrodimitris S, van Ham R, Reinders MJT (2020) Automatic gene function prediction in the 2020's. Genes 11:1264.

19. Plyusnin I, Holm L, Törönen P (2019) Novel comparison of evaluation metrics for gene ontology classifiers reveals drastic performance differences. PLoS Comput Biol 15:e1007419.

20. Gillis J, Pavlidis P (2013) Characterizing the state of the art in the computational assignment of gene function: lessons from the first critical assessment of functional annotation (CAFA). BMC Bioinformatics 14:S15.

21. Radivojac P, Clark WT, Oron TR, Schnoes AM, Wittkop T, Sokolov A et al (2013) A large-scale evaluation of computational protein function prediction. Nat Methods 10:221-227.

22. Törönen P, Medlar AJ, Holm L (2018) PANNZER2: a rapid functional annotation web server. Nucleic Acids Research 46:W84-W88.

23. Zhou N, Jiang Y, Bergquist TR, Lee AJ, Kacsoh BZ, Crocker AW et al (2019) The CAFA challenge reports improved protein function prediction and new functional annotations for hundreds of genes through experimental screens. Genome Biol 20:1-23.

24. Jiang Y, Oron TR, Clark WT, Bankapur AR, D'Andrea D, Lepore R et al (2016) An expanded evaluation of protein function prediction methods shows an improvement in accuracy. Genome Biol 17:1-19.

25. Hippe K, Gbenro S, Cao R (2020) ProLanGO2: Protein Function Prediction with Ensemble of Encoder-Decoder Networks. In Proceedings of the 11th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics (BCB '20). Association for Computing Machinery, New York, NY, USA, Article 103, 1–6. DOI:https://doi.org/10.1145/3388440.3414701

26. Cao R, Freitas C, Chan L, Sun M, Jiang H, Chen Z (2017) ProLanGO: Protein Function Prediction Using Neural Machine Translation Based on a Recurrent Neural Network. Molecules 22:1732.

27. Wimalanathan K, Friedberg I, Andorf CM, Lawrence-Dill CJ (2018) Maize GO Annotation- Methods, Evaluation, and Review (maize-GAMER). Plant Direct 2:e00052.

28. Hyatt D, Chen GL, Locascio PF, Land ML, Larimer FW, Hauser LJ (2010) Prodigal: prokaryotic gene recognition and translation initiation site identification. BMC Bioinformatics 11:119.

29. Yandell M, Ence D (2012) A beginner's guide to eukaryotic genome annotation. Nat Rev Genet 13:329–342.

30. Grabherr MG, Haas BJ, Yassour M, Levin JZ, Thompson DA, Amit I, Adiconis X, Fan L, Raychowdhury R, Zeng Q, Chen Z, Mauceli E, Hacohen N, Gnirke A, Rhind N, di Palma F, Birren BW, Nusbaum C, Lindblad-Toh K, Friedman N, Regev A (2011) Full-length transcriptome assembly from RNA-seq data without a reference genome. Nat Biotechnol 29:644-652.

31. Somervuo P, Holm L (2015) SANSparallel: interactive homology search against Uniprot. Nucleic Acids Res 43:W24-W29.

32. Götz S, Garcia-Gomez JM, Terol J, Williams TD, Nagaraj SH, Nueda MJ, Robles M, Talon M, Dopazo J, Conesa A (2008) High-throughput functional annotation and data mining with the Blast2GO suite. Nucleic Acids Res 36:3420-3435.

33. Medlar AJ, Törönen P, Holm L (2018) AAI-profiler: fast proteome-wide exploratory analysis reveals taxonomic identity, misclassification and contamination. Nucleic Acids Res 46:W479-W485.

34. Plyusnin I, Kant R, Jääskeläinen AJ, Sironen T, Holm L, Vapalahti O, et al. (2020) Novel NGS pipeline for virus discovery from a wide spectrum of hosts and sample types. Virus Evol 6:veaa091

35. Ramos AM et al (2017) The draft genome sequence of cork oak. Sci Data 5:180069.

36. Manni M, Berkeley MR, Seppey M, Simão FA, Zdobnov EM. 2021. BUSCO update: novel and streamlined workflows along with broader and deeper phylogenetic coverage for scoring of eukaryotic, prokaryotic, and viral genomes. Mol Biol Evol, in press. Available from: https://doi.org/10.1093/molbev/msab199

37. Kuorelahti S, Jouhten P, Maaheimo H, Penttila M, Richard P (2006) L-galactonate dehydratase is part of the fungal path for D-galacturonic acid catabolism. Mol Microbiol 61:1060-1068.

38. Babbitt PC, Hasson MS, Wedekind JE, Palmer DR, Barrett WC, Reed GH, Rayment I, Ringe D, Kenyon GL, Gerlt JA (1996) The enolase superfamily: a general strategy for enzyme-catalyzed abstraction of the alpha-protons of carboxylic acids. Biochemistry 35:16489-16501.

39. Motter FA, Kuivanen J, Keränen H, Hilditch S, Penttilä M, Richard P (2014) Categorisation of sugar acid dehydratases in Aspergillus niger. Fungal Genet Biol 64:67-72.

40. Yew WS, Fedorov AA, Fedorov EV, Rakus JF, Pierce RW, Almo SC, Gerlt JA (2006) Evolution of enzymatic activities in the enolase superfamily: L-fuconate dehydratase from Xanthomonas campestris. Biochemistry 45:14582–14597.

41. Barrera-Redondo J, Ibarra-Laclette E, Vázquez-Lobo A, Gutiérrez-Guerrero YT, Sánchez de la Vega G, Piñero D, Montes-Hernández S, Lira-Saade R, Eguiarte LE (2019) The fenome of Cucurbita argyrosperma (Silver-Seed Gourd) reveals faster rates of protein-coding gene and long noncoding RNA turnover and neofunctionalization within cucurbita. Mol Plant 12:506-520.

42. Hawkins T, Chitale M, Luban S, Kihara D (2009) PFP: Automated prediction of gene ontology functional annotations with confidence scores using protein sequence data. Proteins 74:566-582

43. Piovesan D, Tosatto S (2019) INGA 2.0: improving protein function prediction for the dark proteome. Nucleic Acids Res 47:W373-W378.

44. You R, Yao S, Xiong Y, Huang X, Sun F, Mamitsuka H, Zhu S (2019) NetGO: improving large-scale protein function prediction with massive network information. Nucleic Acids Res 47:W379-W387.

45. You R, Zhang Z, Xiong Y, Sun F, Mamitsuka H, Zhu S (2018) GOLabeler: improving sequence-based large-scale protein function prediction by learning to rank. Bioinformatics 34:2465-2473.

46. You R, Yao S, Mamitsuka H, Zhu S (2021) DeepGraphGO: graph neural network for large-scale, multispecies protein function prediction. Bioinformatics 37:i262-i271.

47. Fan K, Guan Y, Zhang Y (2020) Graph2GO: a multi-modal attributed network embedding method for inferring protein functions. GigaScience 9:giaa081.

48. Barot M, Gligorijević V, Cho K, Bonneau R (2021) NetQuilt: deep multispecies network-based protein function prediction using homology-informed network similarity. Bioinformatics 37:2414–2422.

49. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, Bates R, Žídek A, Potapenko A, Bridgland A, Meyer C, Kohl SAA, Ballard AJ, Cowie A, Romera-Paredes B, Nikolov S, Jain R, Adler J, Back T, Petersen S, Reiman D, Clancy E, Zielinski M, Steinegger M, Pacholska M, Berghammer T, Bodenstein S, Silver D, Vinyals O, Senior AW, Kavukcuoglu K,

Kohli P, Hassabis D (2021) Highly accurate protein structure prediction with AlphaFold. Nature 596:583-589.

50. Baek M, DiMaio F, Anishchenko I, Dauparas J, Ovchinnikov S, Lee GR, Wang J, Cong Q, Kinch LN, Schaeffer RD, Millán C, Park H, Adams C, Glassman CR, DeGiovanni A, Pereira JH, Rodrigues AV, van Dijk AA, Ebrecht AC, Opperman DJ, Sagmeister T, Buhlheller C, Pavkov-Keller T, Rathinaswamy MK, Dalwadi U, Yip CK, Burke JE, Garcia KC, Grishin NV, Adams PD, Read RJ, Baker D (2021) Accurate prediction of protein structures and interactions using a three-track neural network. Science 373:871-876.

51. Marttinen P, Corander J, Törönen P, Holm L (2006) Bayesian search of functionally divergent protein subgroups and their function specific residues. Bioinformatics 22:2466-2474.

52. Ward RM, Venner E, Daines B, Murray S, Erdin S, Kristensen DM, Lichtarge O (2009) Evolutionary Trace Annotation Server: automated enzyme function prediction in protein structures using 3D templates. Bioinformatics 25:1426-1427.

53. Clark WT, Radivojac P (2013) Information-theoretic evaluation of predicted ontological annotations. Bioinformatics 29:i53-61.

54. Warwick Vesztrocy A, Dessimoz C. (2020) Benchmarking gene ontology function predictions using negative annotations. Bioinformatics 36:i210-i218.

55. Kahanda I, Funk CS, Ullah F, Verspoor KM, Ben-Hur A (2015) A close look at protein function prediction evaluation protocols. GigaScience 4:41.

56. Valverde-Albacete FJ, Peláez-Moreno C (2014) 100% classification accuracy considered harmful: The normalized information transfer factor explains the accuracy paradox. PloS one 9:e84217.

57. Ferri C, Hernández-Orallo J, Modroiu R (2009) An experimental comparison of performance measures for classification. Pattern Recognition Lett 30:27–38.

58. Akiva E, Brown S, Almonacid DE, Barber AE 2nd, Custer AF, Hicks MA, Huang CC, Lauck F, Mashiyama ST, Meng EC, Mischel D, Morris JH, Ojha S, Schnoes AM, Stryke D, Yunes JM, Ferrin TE, Holliday GL, Babbitt PC (2014) The structure-function linkage database. Nucleic Acids Res 42:D521-530.

59. Katoh K, Rozewicki J, Yamada KD (2019) MAFFT online service: multiple sequence alignment, interactive sequence choice and visualization. Brief Bioinformatics 20:1160-1166.

60. Letunic I, Bork P (2016) Interactive tree of life (iTOL) v3: an online tool for the display and annotation of phylogenetic and other trees. Nucleic Acids Res 44:W242-245.

Table 1: Feature comparison between selected annotation servers

Server schedule	GO prediction	DE prediction	>1000 query sequences	Open source	Last database update/update schedule
ARGOT	Yes	No	Yes	No	Nov-2016
eggNOG	Yes	Yes	Yes	Yes	Jan-2019
FFpred	Yes	No	No	Yes	Unknown
FunFam	Yes	No	Yes (API)	Yes	Daily
INGA	Yes	No	No	No	Feb-2019
NetGO	Yes	No	By request	No	Unknown
PANNZER2	Yes	Yes	Yes	Yes	Jun-2021 / bimonthly
PFP	Yes	No	No	No	Sep-2020

DE prediction stands for free text protein descriptions. Last database update is taken from explicit statements on annotation servers (at time of writing 13/07/21). The web servers were accessed at URLs http://www.medcomp.medicina.unipd.it/Argot2-5/, http://eggnog-mapper.embl.de/, http://bioinf.cs.ucl.ac.uk/psipred/, http://cathdb.info/search/by_sequence, https://inga.bio.unipd.it/, https://issubmission.sjtu.edu.cn/netgo/, http://ekhidna2.biocenter.helsinki.fi/sanspanz/, https://bar.biocomp.unibo.it/bar3, https://kiharalab.org/pfp.php.

Table 2: Rank of MF predictions by web servers for novel 60S ribosomal protein L5

GO class	naive	ARGOT	FunFam	INGA	NetGO	PANNZER filtered	PANNZER unfiltered	PFP	Description
GO:0003674	**1**			1	**1**				molecular_function
GO:0005488	3			1	2		9		+ binding
GO:0097159	4			1	5		8		++ organic cyclic compound binding
GO:1901363	5			1	5		7	5	++ heterocyclic compound binding
GO:0003676	7			1	8		6	4	+++ nucleic acid binding
GO:0003723	37	4	4	1	8		5		++++ RNA binding
GO:0005198	40			1	3		4		+ structural molecule activity
GO:0003735	49	2	**3**	1	3	2	**3**	**1**	++ structural constituent of ribosome
GO:0005515	96		**1**		7				++ protein binding
GO:0019843	137	3		1	8		**2**	**2**	+++++ rRNA binding
GO:0008097	1229	1	**2**	1	8	1	**1**	3	++++++ 5S rRNA binding

The NCBI identifier of the query sequence is KAG7012684.1. Subclasses in GO hierarchy are shown by indentation, padded by '+'. Non-redundant subsets of predictions are bold. Top 10 classes above reliability threshold by any AFP method are included, with the exception of FunFam which had an aberrant profile. Ranks 1-3 are shaded. We show two versions of PANNZER, with and without the filtering of redundant GO terms. FunFam GO terms were ranked by the number of annotations in the top family.

FIGURE LEGENDS

Figure 1: Weighted K-nearest neighbor approach. A query sequence (star in center) is associated with sequence neighbors. Sequence proximity and database background are taken into account by enrichment statistics. Random mislabels can be rejected – systematic mislabels cannot.

Figure 2: Output of Pannzer web server for test case. This view is for browsing; users can also download the results in a parseable format. (A) Non-redundant result obtained using default parameters presents a filtered view. Bottom: sequence neighborhood opens on clicking the "Search" link. (B) Result with parameter settings "show only one DE" off and "show only non-redundant GO terms" off. This view shows alternative predictions ordered by reliability. Inset: GO hierarchy plotted with QuickGO (https://www.ebi.ac.uk/QuickGO/slimming).

Figure 3: Phylogenetic tree of example test sequence (red) with seed sequences of the L-galactonate dehydratase (green) and L-fuconate dehydratase (yellow) subgroups of SFLD.[58] Tree generated by MAFFT,[59] displayed in iToL.[60]

Appendix 1: Test sequences used in this work

60S ribosomal protein L5, partial [Cucurbita argyrosperma subsp. argyrosperma]
> KAG7012684.1
MNSPFLELTASMFAFAKAQKTKAYFKRYQVKFKRRREGKTDYRARIRLINQDKNKYNTPKYRIVVRFSNKDITAQI
ISASIAGDLVLASAYSHELPRYGLEVGLTNYAAAYCTGLLLARRVLKQLEMDDEYEGNVEATGEDYSVEPADTRRP
FRALLDVGLLKTTTGNRVFGALKGALDGGLDIPHSDKRFAGFSKDSKQLDADVHRKYIYGGHVAAYMRTLMEDE
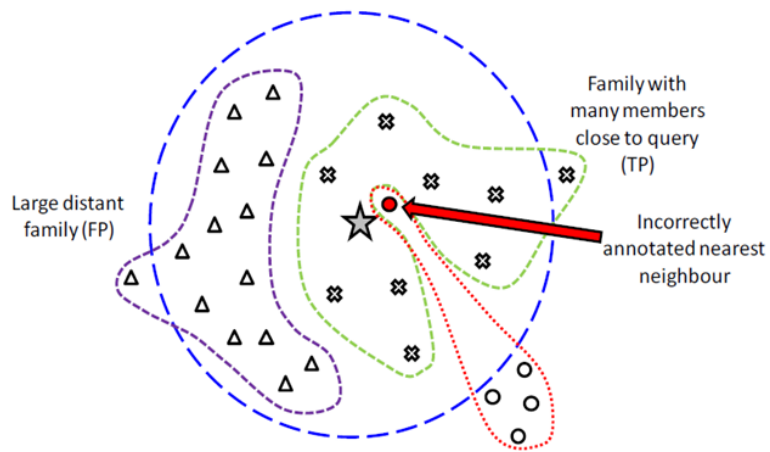PEKYQTHFSEYIKKGIEADDIEGLYKKVHAAIRADPSVKKSDKPQPKAHKRYNLKKLTYDERKARLVERLNALNSAA
NADDDDDDEDDE

L-galactonate dehydratase-like [Quercus suber]
>XP_023870721.1
MVLITHATTRDVRFPTSLDKTGSDAMNAAGDYSAAYVMLHSDTSHTGHGMTFTIGRGNEIVCKAISVLAQRVE
GKQLEDLVADWGKTWRYLVSDSQLRWIGPEKGVIHLALGAVVNAIWDLWAKVLGKPVWRIVAEMSPQEFVR
CIDFRYITDAITPEEAISMLEKEEAGKAQRIKEAEQNRAVPAYTTSAGWLGYGEAKMKGLLEETLAKGYKHFKLKV
GTSLEADKQRLAIARDVIGYDNGNVLMVDANQVWSVPEAITYMKELARFKPWFIEEPTSPDDVFGHKAIREALK
PYNIGVATGEMCQNRVMFKQLIVQGAIDVCQIDACRIGGVNEVMAVMLIAKKYGVPIVPHSGGVGLPEYTQHL
STIDYVVVSGKLSVLEYVDHLHEHFLHPSIIESGYYVTPTMPGYSVEMKAESMEQYEFPGTEGVSWWRSAQAKG
ILEGEKI

Appendix 2: Data sources used by PANNZER2

Uniprot	ftp://ftp.ebi.ac.uk/pub/databases/uniprot/current_release/knowledgebase/complete/uniprot_{sprot,trembl}.fasta.gz
Taxonomy	https://www.uniprot.org:443/taxonomy/?query=*&compress=yes&format=tab
GO assignments	ftp://ftp.ebi.ac.uk/pub/databases/GO/goa/UNIPROT/goa_uniprot_all.gaf.gz
GO cross-mappings	http://geneontology.org/external2go/{ec,kegg}2go
GO structure	http://geneontology.org/ontology/go-basic.obo

Large distant family (FP)

Family with many members close to query (TP)

Incorrectly annotated nearest neighbour
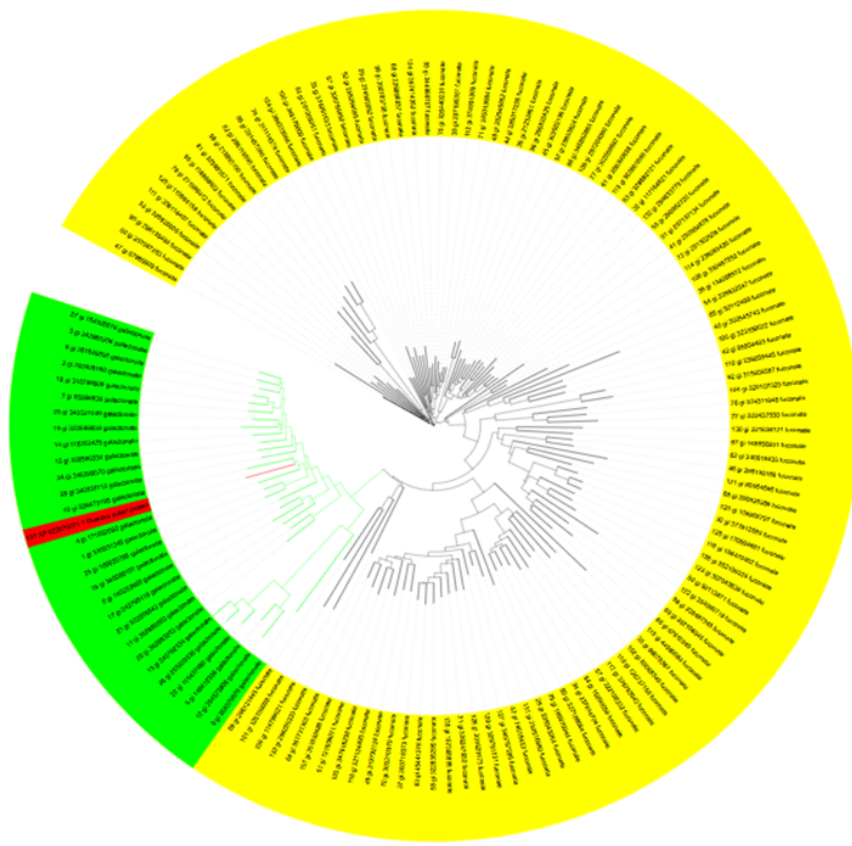
PRO_4193_FIgure1.tif

Query header	gene name	Description Estimated PPV, description		Biological process Estimated PPV, GO-id, description			Molecular function Estimated PPV, GO-i d, description			Cellular component Estimated PPV, GO-id, description	Inverse ec2go, kegg2go		
XP_023870721.1 XP_023870721.1 Search	LGD1	0.54	L-galactonate dehydratase	0.69 0.67	GO:0009063 GO:0016052	cellular amino acid catabolic process carbohydrate catabolic process	0.84	GO:0050023	L-fuconate dehydratase activity		0.84 0.84	EC:4.2.1.68 KEGG:R03688	GO:0050023 GO:0050023

Rank	Vote	Identity	Ranges	Ali length	Bitscore	E-value	Identifier	Description	Species	Gene name
1	817	0.88	1-447:1-447	447	898.4	0.00000000E+00	tr|A0A4U0UAD4|A0A4U0UAD4_9PEZI	L-galactonate dehydratase	Hortaea thailandica	B0A50_02252
2	7789	0.87	1-447:1-447	447	888.7	0.00000000E+00	tr|A0A3M7J293|A0A3M7J293_HORWE	MR_MLE domain-containing protein	Hortaea werneckii	D0859_04002
3	7775	0.87	1-447:1-447	447	887.9	0.00000000E+00	tr|A0A3M7FFW6|A0A3M7FFW6_HORWE	MR_MLE domain-containing protein	Hortaea werneckii	D0861_05205
4	7775	0.87	1-447:1-447	447	887.9	0.00000000E+00	tr|A0A1Z5TNB1|A0A1Z5TNB1_HORWE	L-galactonate dehydratase	Hortaea werneckii EXF-2000	BTJ68_02248
5	7743	0.87	1-447:1-447	447	885.8	0.00000000E+00	tr|A0A3M7G386|A0A3M7G386_HORWE	MR_MLE domain-containing protein	Hortaea werneckii	D0862_08742

PRO_4193_Figure2A.tif

Query header	gene name	Description Estimated PPV, description	Biological process Estimated PPV, GO-id, description	Molecular function Estimated PPV, GO-i d, description	Cellular component Estimated PPV, GO-id, description	Inverse ec2go, kegg2go
L-galactonate dehydratase-lik [Quer] Search	LGD1	0.54 L-galactonate dehydratase 0.49 Mitochondrial enolase superfamily member 1 0.42 Mandelate racemase/muconate lactonizing enzyme domain-containing protein	0.57 GO:0009063 cellular amino acid catabolic process 0.56 GO:0046395 carboxylic acid catabolic process 0.56 GO:0016054 organic acid catabolic process 0.55 GO:0016052 carbohydrate catabolic process 0.55 GO:0044282 small molecule catabolic process 0.54 GO:1901565 organonitrogen compound catabolic process 0.52 GO:0044248 cellular catabolic process 0.51 GO:1901575 organic substance catabolic process 0.51 GO:0009056 catabolic process 0.50 GO:0006520 cellular amino acid metabolic process	0.68 GO:0050023 L-fuconate dehydratase activity 0.56 GO:0016836 hydro-lyase activity 0.55 GO:0016835 carbon-oxygen lyase activity 0.52 GO:0016829 lyase activity 0.39 GO:0003824 catalytic activity		0.68 EC:4.2.1.68 GO:0050023 0.68 KEGG:R03688 GO:0050023

PRO_4193_FIgure2B.tif

PRO_4193_Figure3.tif