Published in final edited form as:

Similarity Search Appl. 2013 ; 8199 : 319–324. doi:10.1007/978-3-642-41062-8_32.

PFMFind: a system for discovery of peptide homology and function

Aleksandar Stojmirović^{1,*}, Peter Andreae², Mike Boland³, Thomas William Jordan⁴, and Vladimir G. Pestov⁵

¹ National Center for Biotechnology Information, National Library of Medicine, National Institutes of Health, Bethesda, MD 20894, United States ² School of Engineering and Computer Science, Victoria University of Wellington, PO Box 600, Wellington 6140, New Zealand ³ Riddet Institute, Massey University, PB 11 222, Palmerston North 4442, New Zealand ⁴ School of Biological Sciences, Victoria University of Wellington, PO Box 600, Wellington 6140, New Zealand ⁵ Department of Mathematics and Statistics, University of Ottawa, 585 King Edward Ave., Ottawa, ON K1N 6N5, Canada

Abstract

Protein Fragment Motif Finder (PFMFind) is a system that enables efficient discovery of relationships between short fragments of protein sequences using similarity search. It supports queries based on amino acid similarity matrices and position specific score matrices (PSSMs) obtained through an iterative procedure. PSSM construction is customisable through plugins written in Python. PFMFind consists of a GUI client, an index for fast similarity search and a relational database for storing search results and sequence annotations. It is written mostly in Python. The components of PFMFind communicate through TCP/IP sockets and can be located on different physical machines. PFMFind is freely available for download (under a GPL licence) from <http://pfmfind.stojmirovic.org>

Keywords

similarity search; indexing; protein fragments

1 Introduction

The biological functions of proteins are as much a function of particular motifs of peptide sequence as they are of the overall protein structure. It is of interest to the biologist to search for examples of convergent motifs as they are likely to indicate a functional role. While many approaches exist for finding longer sequence motifs (50 amino acids or more), finding relationships between short fragments (3–18 amino acids long) of full protein sequences also promises great rewards in understanding novel aspects of protein structure and function. These relationships might be evolutionary in origin or might arise by convergence, that is, by acquisition of the same biological function in evolutionarily distant species.

Finding short motifs presents significant challenges because many of the apparent relationships between short fragments could have arisen by chance and thus have no functional significance. Furthermore, most widely available tools for sequence database search and motif finding were designed with longer motifs in mind. For example, Watt and

* to whom correspondence should be addressed.

Doyle [11] observed that the NCBI BLAST [1] family of programs, the best known set of tools for searching biological sequence datasets, is not suitable for identifying shorter sequences with particular constraints and proposed a pattern search tool to find DNA or protein fragments that match a given sequence or a pattern exactly. This paper outlines the Protein Fragment Motif Finder (PFMFind), a tool that uses database search to identify the conserved short peptide motifs of a query sequence and associates them with the available functional annotations.

2 Overview

The PFMFind system consists of three major components: a search engine for fast similarity search of datasets of short peptide fragments called FSIndex, a relational database, and the PFMFind graphical user interface (GUI) client (Fig. 1). PFMFind client takes user input, and communicates with FSIndex and the database through its components. It passes search parameters in batches to FSIndex and receives the results of searches that are then stored in the database. It also retrieves the results from the database and displays them, together with available annotations, to the user. The annotations are stored in a separate (BioSQL) schema in the database.

Most of PFMFind was written in the Python programming language, and uses both the standard Python library and additional modules such as Biopython (<http://www.biopython.org>). The components communicate using the standard TCP/IP socket interface and can therefore be located on different machines. Since PFMFind is highly modular, the GUI client can be replaced by a Python script for non-interactive use.

2.1 Similarity Search

PFMFind supports searches of datasets of short peptide fragments of fixed length using an ungapped similarity score obtained by summing similarity scores at each position of the fragments being compared. The positional similarity scores can be defined by standard score matrices such as PAM [3] or BLOSUM [6], or by position specific score matrices (PSSMs) [4]. A fragment dataset consists of all fragments of a specified length from a given protein sequence dataset (where the fragments may overlap).

Iterative construction of PSSM, similar to that used by PSI-BLAST [1], is supported through plugins – Python routines that take the results of a previous search and construct a PSSM. The default plugin uses the weighting procedure of Heniko and Heniko [5] to assign weights to fragments and Dirichlet mixtures [9] for regularising the amino acid frequency counts at each position. Users with some knowledge of Python can create their own plugins and use them for searches by placing them in the appropriate directory.

Search criteria can be specified according to threshold raw similarity scores, distances, p-values and E-values (range search), as well as the number of closest data points to retrieve (kNN search). The probability model for calculation of p-values assumes that the score of each fragment is the sum of independent random variables corresponding to the score at each position and the score distribution is calculated using discrete Fast Fourier Transform.

2.2 FSIndex

The heart of PFMFind is FSIndex, an efficient indexing scheme for similarity search of very large datasets of short protein fragments of fixed length [8, 10]. FSIndex is based on two principles: reduction of the amino acid alphabet to clusters largely based on their biochemical properties (hydrophobic, polar, charged, aromatic ...) and combinatorial generation of neighbours. The design of FSIndex means that a typical search involves

scanning less than 1% of the fragment dataset, but ensures that no neighbours satisfying search criteria are ever missed.

FSIndex is implemented in the C programming language and embedded into Python, with the whole data structure as well as the indexed sequences stored in primary memory. For even greater efficiency, computation of searches can be distributed among several machines using a master/slave model: the master handles p-value computations, distributes queries to slaves, each of which is indexing a different part of the dataset, and communicates with the client.

2.3 Database

The second major component of PFMFind is a relational database, used both for storage of search results and the sequence annotation. We use PostgreSQL, a freely available modern database management system.

Each user of the system has their own schema for storing search results. The database also stores all search parameters, including PSSMs and the results of each iteration, facilitating reversion to a previous iteration without repeating the whole procedure.

The database stores sequence annotations in a standard BioSQL schema available to all users. PFMFind also contains scripts for loading four types of information beyond the protein sequence: Uniprot [2] keywords and features, Uniref clusters [2] and InterPro [7] domains. When retrieved for display, annotations are joined to search results through accession numbers.

2.4 GUI Client

The final PFMFind component is a GUI client that connects to both the FSIndex master and the database component. To perform fragment searches, the user specifies a query sequence, usually a long sequence that is broken into overlapping fragments of fixed length, and chooses the fragment lengths, threshold parameters and the actual fragments in the query sequences that will be used for the search.

The GUI client can display search results both as lists of hits associated with a particular location in the query sequence and as a feature vs location dot plot – each location matching a particular feature is marked by a coloured dot (Fig. 2). Dots are colour coded by the number of hits matching the feature to distinguish frequently represented features from those that appear only a few times in the hit list. The GUI client also performs all computations for constructing PSSMs.

3 Conclusion

PFMFind is an efficient, flexible, and extensible framework for similarity search of datasets of short peptide fragments. It supports fast similarity search with selectivity and sensitivity specified by PSSMs and associates search results with biological function by using sequence features and annotations.

Acknowledgments

We wish to thank Pavle Mogin and Danyl McLauchlan for their help with Post-greSQL and testing the software, respectively. A.S. was supported by a Bright Future PhD scholarship awarded by the NZ Tertiary Education Commission jointly with the Fonterra Research Centre, by a Fields Institute/University of Ottawa postdoctoral fellowship, and by the Intramural Research Program of the National Library of Medicine at the National Institutes of Health. V.G.P. and A.S. acknowledge support from NSERC discovery grant program and University of Ottawa internal grants.

References

1. Altschul SF, Madden TL, Schaffer AA, Zhang J, Zhang Z, Miller W, Lipman DJ. Gapped BLAST and PSI-BLAST: a new generation of protein database search programs. *Nucleic Acids Res.* 1997; 25:3389–3402. [PubMed: 9254694]
2. Bairoch A, Apweiler R, Wu CH, Barker WC, Boeckmann B, Ferro S, Gasteiger E, Huang H, Lopez R, Magrane M, Martin MJ, Natale DA, O'Donovan C, Redaschi N, Yeh LSL. The Universal Protein Resource (UniProt). *Nucleic Acids Res.* 2005; 33(Database Issue):154–159.
3. Dayhoff, MO.; Schwartz, RM.; Orcutt, BC. A model of evolutionary change in proteins.. In: Dayhoff, MO., editor. *Atlas of Protein Sequence and Structure*. Vol. 5. National Biomedical Research Foundation; 1978. p. 345-352.chap. 22
4. Gribskov M, McLachlan AD, Eisenberg D. Profile analysis: detection of distantly related proteins. *Proc Natl Acad Sci USA.* 1987; 84:4355–4358. [PubMed: 3474607]
5. Henikoff S, Henikoff JG. Position-based sequence weights. *J Mol Biol.* 1994; 243(4):574–578. [PubMed: 7966282]
6. Henikoff S, Henikoff J. Amino acid substitution matrices from protein blocks. *Proc Natl Acad Sci USA.* 1992; 89:10915–10919. [PubMed: 1438297]
7. Hunter S, Jones P, Mitchell A, Apweiler R, Attwood TK, Bateman A, Bernard T, Binns D, Bork P, Burge S, de Castro E, Coghill P, Corbett M, Das U, Daugherty L, Duquenne L, Finn RD, Fraser M, Gough J, Haft D, Hulo N, Kahn D, Kelly E, Letunic I, Lonsdale D, Lopez R, Madera M, Maslen J, McAnulla C, McDowall J, McMenamin C, Mi H, Mutowo-Muellenet P, Mulder N, Natale D, Orengo C, Pesseat S, Punta M, Quinn AF, Rivoire C, Sangrador-Vegas A, Selengut JD, Sigrist CJA, Scheremetjew M, Tate J, Thimmajananathan M, Thomas PD, Wu CH, Yeats C, Yong SY. Interpro in 2011: new developments in the family and domain prediction database. *Nucleic Acids Res.* Jan; 2012 40(Database issue):D306–12. [PubMed: 22096229]
8. Pestov V, Stojmirovi A. Indexing schemes for similarity search: an illustrated paradigm. *Fundam. Inform.* 2006; 70(4):367–385.
9. Sjlander K, Karplus K, Brown M, Hughey R, Krogh A, Mian I, Haussler D. Dirichlet mixtures: A method for improving detection of weak but significant protein sequence homology. *Comput. Appl. Biosci.* 1996; 12(4):327–345. [PubMed: 8902360]
10. Stojmirovi A, Pestov V. Indexing schemes for similarity search in datasets of short protein fragments. *Inf. Syst.* 2007; 32(8):1145–1165.
11. Watt TJ, Doyle DF. ESPSearch: a program for finding exact sequences and patterns in DNA, RNA, or protein. *Biotechniques.* 2005; 38(1):109–115. [PubMed: 15679092]

Fig. 1.
Structure of PFMFind system

Fig. 2.
A screenshot of PFMFind GUI showing search results associated with their annotations