

The LabelHash Algorithm for Substructure Matching

Mark Moll*¹, Drew H. Bryant¹ and Lydia E. Kavraki^{1,2,3}

¹Department of Computer Science, Rice University, Houston, TX 77005, USA

²Department of Bioengineering, Rice University, Houston, TX 77005, USA

³Structural and Comp. Biology and Molec. Biophysics, Baylor College of Medicine, Houston, TX 77005, USA

Email: Mark Moll* - mmoll@rice.edu; Drew H. Bryant - dbryant@rice.edu; Lydia E. Kavraki - kavraki@rice.edu;

*Corresponding author

Abstract

Background: There is an increasing number of proteins with known structure but unknown function. Determining their function would have a significant impact on understanding diseases and designing new therapeutics. However, experimental protein function determination is expensive and very time-consuming. Computational methods can facilitate function determination by identifying proteins that have high structural and chemical similarity.

Results: We present LabelHash, a novel algorithm for matching substructural motifs to large collections of protein structures. The algorithm consists of two phases. In the first phase the proteins are preprocessed in a fashion that allows for instant lookup of partial matches to any motif. In the second phase, partial matches for a given motif are expanded to complete matches. The general applicability of the algorithm is demonstrated with three different case studies. First, we show that we can accurately identify members of the enolase superfamily with a single motif. Next, we demonstrate how LabelHash can complement SOIPPA, an algorithm for motif identification and pairwise substructure alignment. Finally, a large collection of Catalytic Site Atlas motifs is used to benchmark the performance of the algorithm. LabelHash runs very efficiently in parallel; matching a motif against all proteins in the 95% sequence identity filtered non-redundant Protein Data Bank typically takes no more than a few minutes. The LabelHash algorithm is available through a web server and as a suite of standalone programs at <http://labelhash.kavrakilab.org>. The output of the LabelHash algorithm can be further analyzed with Chimera through a plugin that we developed for this purpose.

Conclusions: LabelHash is an efficient, versatile algorithm for large-scale substructure matching. When LabelHash is running in parallel, motifs can typically be matched against the *entire* PDB on the order of minutes. The algorithm is able to identify functional homologs beyond the twilight zone of sequence identity and even beyond fold similarity. The three case studies presented in this paper illustrate the versatility of the algorithm.

Background

High-throughput methods for structure determination have greatly increased the number of proteins with known structure in the Protein Data Bank (PDB) [1]. Structural genomics initiatives [2–4] have contributed not only to this increase in *number*, but have also increased the *diversity* of known protein structures. The function of most proteins is still poorly understood or even completely unknown. Automated functional annotation methods make it possible to fill in some of the gaps of missing information. Such methods can be a critical component of computational drug design and protein engineering. Their applicability, however, goes beyond traditional applications. For example, a nuanced and detailed understanding of protein function can also provide insight into the roles hub proteins play in protein interaction networks. Sequence-based methods are an established way for detecting functional similarity [5–8], but sequence similarity does not always imply functional similarity and vice versa. Structural analysis using the entire PDB allows for the discovery of similar function in proteins with very different sequences and even different folds [9]. For an overview of current approaches in sequence- and structure-based methods see [2, 10, 11].

Although it is possible to compare structures at the fold level [12, 13], or by comparing pockets and clefts [14–16], this work focuses on substructure matching methods. Substructure matching methods aim to find a common substructure *motif* (sometimes also called a *template*) within one or more protein structures. Substructure matching methods can be used to identify functional similarity in cases where sequence similarity or fold similarity between homologous proteins is low (such as in cases of convergent evolution). The *identification* of an important substructure that forms a motif can be separated from the process of *matching* the motif against a number of structures. Several methods have been proposed to identify residues that are of functional importance [17–19]. Typically, such methods require as input a family of proteins that are known to be functionally similar. Once a structural motif has been identified that characterizes a given function or family, it is still a challenging problem to screen all the structures in the PDB for occurrences of substructures similar to this motif, and determine functional similarity. A wide variety of substructure matching methods have been proposed, such as: TESS [20], SPASM [21], CavBase [22], eF-site [23], ASSAM [24], PINTS [25], Jess [26], SuMo [27], SiteEngine [28], Query3D [29], ProFunc [30], ProKnow [31], SitesBase [32], GIRAF [33], MASH [34], SOIPPA [35, 36], FEATURE [37], and pevoSOAR [38]. These methods mainly differ in (1) the

representation of structural motifs, (2) the motif matching algorithm, and (3) the statistics used to determine significance of match. Representations for substructural motifs include: C_α coordinates with residue labels (with side-chain centroid [21] or without [25, 34]), physicochemical pseudo-centers [28], graphs [22, 24, 35], general sets of constraints on atom positions and residue types [26], binding surfaces annotated with evolutionary conservation [38], or learnt vectors of features (such as the presence of a residue type of metal ion) occurring at certain distance ranges [37]. Such representations can then be matched using a variety of techniques such as: geometric hashing [20, 28, 32, 39–41], depth-first search [21, 25, 29, 34], graph algorithms (clique detection [22, 23, 27], subgraph isomorphism [24]), and constraint solvers [26]. To assess the statistical significance of matches the use of Extreme Value Distributions [36, 42], mixtures of Gaussians [26], and a non-parametric model [35, 43] have been proposed.

This paper describes a novel method for rapidly matching a motif against all known structures in the PDB (or any arbitrary subset thereof). It addresses several algorithmic and system design issues that allow it to be run in parallel and obtain near real-time performance. The method makes very few assumptions about the motif. For instance, a motif does not necessarily have to represent a cavity or binding site. The method was designed to be easy to use by both novice and expert users: the default parameters work in variety of scenarios, but can be easily changed to control the desired output. Through three case studies we demonstrate the versatility of the method and the ability to obtain highly sensitive and specific results.

Results and Discussion

The LabelHash Algorithm

We are interested in matching a structural motif against a large set of target structures. The structural motif is defined by the backbone C_α coordinates of a number of residues and (optionally) allowed residue substitutions for each motif residue which are encoded as labels. Previous work has established that this is a feasible representation because it can find biologically relevant results [17, 34, 42, 44].

The method presented below is called LabelHash. We will first give a high-level description. The method builds a hash table for n -tuples of residues that occur in a set of targets. In spirit LabelHash is reminiscent of the geometric hashing technique [40], but the particulars of the approach are very different. The n -tuples are hashed based

on the residues' labels. Each n -tuple has to satisfy certain geometric constraints. The data in the hash table is indexed in a way that allows fast parallel access. Using this table we can look up partial matches of size n in constant time. These partial matches are augmented in parallel to full matches with an algorithm similar to MASH [34]. Compared to geometric hashing [40], our method significantly reduces storage requirements. Relative to MASH, we further improve the specificity. Furthermore, in the LabelHash algorithm it is no longer required to use importance ranking of residues to guide the matching (as was done in MASH). In our previous work, this ranking was obtained using Evolutionary Trace (ET) information [45]. The LabelHash algorithm was designed to improve the (already high) accuracy of MASH and push the envelope of matching with only very few geometric constraints. We want motifs to be as general as possible to allow for future extensions and to facilitate motif design through a variety of methods. Our simple-to-use and extensible LabelHash algorithm is extremely fast and can be a critical component of an exploratory process of iterative and near-interactive design and refinement of substructure templates. The algorithm consists of two stages: a pre-processing stage and a stage where a motif is matched against the preprocessed data.

Preprocessing Stage

The preprocessing stage has to be performed only once for a given set of targets; any motif can then be matched against the same preprocessed data. The targets can consist of single chains, but it is also possible to treat entire domains as a single target. This is useful for motifs that span multiple chains. During the preprocessing stage we aim to find possible candidate partial matches. This is done by finding all n -tuples of residues that satisfy certain geometric constraints. We will call these n -tuples *reference sets*. Typically, n is small; in our experiments we use 3-tuples. All valid reference sets for all targets are stored in a hash map, a data structure for *key/value* pairs that allows for constant time insertions and lookups (on average). In our case, each *key* is a sorted n -tuple of residue labels, and the *value* consists of a list of reference sets that contain residues with these labels in any order. So for any reference set in a motif we can instantly find all occurrences in all targets. Notice that in contrast to geometric hashing [40] we do not store transformed copies of the targets for each reference set, which allows us to store many more reference sets in the same amount of memory.

In our current implementation the geometric con-

straints apply to the C_α coordinates of each residue, but there is no fundamental reason why other points such as C_β 's or physicochemical pseudo-centers [22, 46] cannot be used instead. We have defined the following four constraints on *valid* reference sets:

- Each C_α in a reference set has to be within a distance $d_{\max\text{mindist}}$ from its nearest neighboring C_α .
- The maximum distance between any two C_α 's within a reference set is restricted to be less than d_{diameter} .
- Each residue has to be within distance $d_{\max\text{depth}}$ from the molecular surface. The distance is measured from the atom closest to the surface.
- At least one residue has to be within distance $d_{\max\text{mindepth}}$ from the surface.

The first pair of constraints requires points in valid reference sets to be within close proximity of each other, and the second pair requires them to be within close proximity of the surface. The distance parameters that define these constraints should be picked large enough to allow for at least one valid reference set for each motif that one is willing to consider, but small enough to restrict the number of seed matches in the targets. One would perhaps expect that the storage requirements would be prohibitively expensive, but—as described in the Implementation and Methods sections—the required storage is still very reasonable. The values for the four distance parameters described above were chosen empirically and kept fixed for all experiments (see *Methods*).

Matching Stage

For a given motif of size m ($m \geq n$), the LabelHash algorithm can look up all matches to a submotif of fixed size n , and expand each partial match to a complete match using a depth-first search. The partial match expansion is a variant of the match augmentation algorithm [34] that consists of the following steps. First, it solves the residue label correspondence between a motif reference set and the matching reference sets stored in the LabelHash table. (If more than one correspondence exists, all of them are considered.) Next, the match is augmented one residue at a time, each time updating the optimal alignment that minimizes the RMSD. If a partial match has an RMSD greater than some threshold ϵ , it is rejected. For a given motif point, we find all residues in a target that are within some threshold distance (after alignment). This threshold is for simplicity usually also set to ϵ . The threshold ϵ is

set to be sufficiently large (7 Å in our experiments) so that no interesting matches are missed.

The algorithm recursively augments each partial match with the addition of each candidate target residue. The residues added to a match during match augmentation are *not* subject to the geometric constraints of reference sets. In other words, residues that are not part of a reference set are allowed to be further from each other and more deeply buried in the core. For small-size reference sets, the requirement that a motif contains at least one reference set is therefore only a very mild constraint. As we will see in the next section, our approach is still highly sensitive and specific.

For a given motif, we generate all the valid reference sets for that motif. Any of these reference sets can be used as a starting point for matching. However, those reference sets that have the smallest number of matching reference sets in the LabelHash table may be more indicative of a unique function. Reference sets with a large number of matches are more likely to be common structural elements or due to chance. We could exhaustively try all possible reference sets, but for efficiency reasons we only process a fixed number of least common reference sets. Note that the selection of reference sets as seed matches is based only on frequency. In contrast, in our previous work, only one seed match was selected based on importance ordering frequently based on evolutionary importance [34]. Because of the preprocessing stage it now becomes feasible to expand matches from many different reference sets. The information stored inside a LabelHash file is stored so that only the relevant parts of the file need to be read from disk during matching.

The matching algorithm is flexible enough to give users full control over the kind of matches that are returned. It is possible to keep multiple matches per target or partial matches that match at least a certain minimum number of residues. The latter option can be useful for larger motifs where the functional significance of each motif point is not certain. In such a case, a 0.5 Å RMSD partial match of, say, 9 residues, might be preferable over a 2 Å complete match of 10 residues. With partial matches, the matches can be ranked by a scoring function that balances the importance of RMSD and the number of residues matched. One can also choose between keeping only the lowest RMSD match per target or all matches for a target, which may be desirable if the top-ranked matches for targets have very similar RMSD's. Finally, the number of motif reference sets that the algorithm uses for match augmentation can also be varied. Usually most matches are found with the first couple reference sets, but occasionally a larger number of reference sets need to be tried

before the smallest RMSD match for each target is found. With the default settings (used for all results in this paper), the number is set to a conservative threshold of 15.

Statistical Significance of Matches

There is no universal RMSD cut-off that can be used to decide whether a match is significant, and picking a cut-off for a given motif and corresponding protein family is non-trivial. The cut-off depends on the structural variation within a protein family and the likelihood that a match to a non-homologous protein occurs due to chance. In our work we use a nonparametric model to compute the statistical significance of each match. This model is briefly summarized below and described in more detail in [43,47]. The model assumes that the matching algorithm returns for each target only the lowest RMSD, complete match to a motif. Keeping partial matches or multiple matches per target complicates the determination of the statistical significance of each match. This is an issue we plan to investigate in future work.

For a set of matches we can compute a probability density function over RMSD by smoothing the RMSD distribution using the Sheather-Jones optimal bandwidth [48]. An example distribution is shown in figure 1. In an ideal case, functionally homologous targets would have a low RMSD and would be well-separated from the non-homologous targets. From this RMSD distribution, which we will call a *motif profile*, we can assign a p -value to a match by dividing the area under the curve to the left of the match's RMSD by the total area under the curve. For instance, for the match at the cut-off in figure 1, the p -value would be $\frac{A}{A+B}$, where A and B are the areas under the curve to the left and right of the dotted line, respectively. However, the value of the distance cut-off parameter ϵ introduces an algorithmic bias that affects the computation of the statistical significance of a match. Other matches could be found if the value of ϵ were increased. To correct for this bias, we model the existence of these 'missed' matches by placing a point-weight proportional to their relative frequency at infinity. The *corrected* p -value is then $\frac{A}{A+B+C}$, where C is the weight of the missed matches. Finding missed matches is straightforward: for each target where no match was found, we simply check whether there are enough residues of the right types so that a match is possible. The residue frequencies are pre-computed for each target and stored inside the LabelHash tables.

It can be shown that for a motif of n residues our statistical model computes the *exact* p -value of matches with RMSD less than ϵ/\sqrt{n} , i.e., their p -value would not change if no ϵ threshold was used [43]. For example, for

a 6-residue motif and $\epsilon = 7 \text{ \AA}$, the p -values of all matches within 2.9 \AA of the motif are exact.

Implementation

Data Layout

We aimed for LabelHash to be scalable to all available structures in the Protein Data Bank, even as it continues to grow. The structural information, reference sets and indexing information require significant storage. The data layout is determined not only by the content, but also by the expected access patterns. Space-efficient storage of all data that supports computationally efficient access can contribute significantly to the overall speed of the algorithm. Typically, only a very small fraction of the data is accessed in matching a motif for two reasons. First, only a small number of reference sets in a motif is used for match expansion. Second, often we are only interested in matching a motif against a subset of all known structures (such as the non-redundant PDB or a given class/family of proteins). The data format should also be extensible (so that we can store additional attributes for each target in the future) and allow concurrent access to facilitate parallelized matching (see next subsection). In our current implementation we have chosen to use the Hierarchical Data Format (HDF5) [49], a standard file format used for large data sets. Conceptually, the HDF5 software library creates a file system within a file. With HDF5 one can easily create and change hierarchical groups of data sets, without having to worry about keeping all the indexing information up-to-date. It also allows for sophisticated compression schemes (which can be enabled per data set) and includes support for parallel I/O (by providing a software layer on top of MPI-IO).

The data is laid out as follows. For each target there is a group that contains data sets with the target's structural information, residue types, and any other information specific to the target that we may add in the future. For every possible set of n residue labels, we store all the reference sets with those labels in a large matrix with n columns, where n is the reference set size. Associated with each matrix is some additional indexing information that keeps track of which block of rows in each matrix contains indices into which target structure. The n -tuples take up the bulk of the data that needs to be stored, but, luckily, they are also very compressible. We chose to use SZIP [50] compression because of its high compression ratio and fast decompression speed. The LabelHash HDF5 file for the *entire* PDB (based on a snapshot from March 25, 2009), including reference sets and all metadata takes up 65 GB. Although this is a very large file, it still fits eas-

ily on commodity hard drives (and even some solid state drives), and a significant portion of the file could be kept in memory (e.g., in page cache) on a dedicated server.

Large-Scale Matching

Multi-core processors and distributed computing clusters are increasingly commonplace, and naturally we would like to take advantage of that. Both the preprocessing stage and the matching stage are parallelized, and a near-linear speed-up with the number of CPU cores can be achieved. In the preprocessing phase a master node asynchronously sends PDB id's to slave nodes, which read the corresponding PDB file, compute the reference sets, and send all data back to the master. The master node writes all data to disk. Although this seems suboptimal in terms of communication, it avoids the difficulties associated with parallel write access (if all nodes write to a single file) or a sequential merge of several files (if nodes write to separate files).

Matching in LabelHash is also easily parallelized. The targets to be matched are evenly divided over the nodes and each node matches a given motif against its targets independently. Once matching is finished, the match results are aggregated into one output file by one of the cores. This parallelization scheme could lead to load imbalance if matching against some targets takes significantly longer than others. In our experiments the number of targets was usually large enough and arbitrarily distributed over the nodes such that there was no imbalance. If load imbalance were to become an issue, it would be relatively easy to implement schemes that dynamically assign batches of targets to the nodes. Another potential performance bottleneck is the simultaneous disk access by the nodes. If all nodes independently try to read data from the LabelHash table, they can spend a significant amount time in system calls, waiting for disk access or seeking the right disk sectors. The bulk of the data that needs to be read consists of the n -tuples for the selected targets. This data is read synchronously using the HDF5-provided software layer on top of MPI-IO. The amount of data to be read by each node is roughly equal, so that idle time is minimized. Once the n -tuples have been read, the match augmentation phase starts. For each of its targets, a node will load the structural information from the LabelHash table and compute the best match(es). The computation time tends to be significantly larger than the disk access time, so that the nodes can read the data asynchronously with only minimal disk contention.

Algorithmic Issues

In the original implementation of match augmentation [34], the runtime was dominated by computing (a) RMSD alignments and (b) nearest neighbors for a motif point in a target structure. In addition to the performance improvements obtained through the pre-computation of reference sets, the LabelHash algorithm uses a fast, new algorithm and pre-computation, respectively, to speed up these two components.

The match augmentation algorithm iteratively updates the RMSD alignment with each residue added to a partial match. Since match augmentation is started from many matching reference sets and for each partial match there can be many possibilities for match augmentation, the computation of RMSD alignments can potentially be a performance bottleneck. Computing such alignments typically involves computing a covariance matrix and the largest eigenvalue/eigenvector pair of this matrix [51, 52]. For small substructures the computational cost is dominated by the eigenvalue/eigenvector calculation. We use a recently proposed method to compute the optimal alignment by finding the largest root of the characteristic polynomial, instead of computing a matrix decomposition of the covariance matrix [53, 54]. Using this method in our implementation, the time needed for RMSD calculations is drastically reduced to less than one tenth of the traditional approach.

To expand a partial match by one residue, the match augmentation algorithm looks for matching residues for the next motif point under the current optimal alignment. The matching residues are found using a proximity data structure. Several such data structures exist; in our implementation we use Geometric Near-neighbor Access Trees [55]. The key observation to speeding up proximity queries is that the nearest neighbors of a motif point are very similar to the nearest neighbors of the point in the target closest to this motif point. Suppose we are interested in all nearest neighbors within 7 \AA of a motif point and the closest point in the target is $x \text{ \AA}$ away. Then the set of nearest neighbors of this point that are within $(7+x) \text{ \AA}$ includes all the nearest neighbors of the motif point. Those within $(7-x) \text{ \AA}$ are guaranteed to be also nearest neighbors of the motif point, and only those within the $(7-x) \text{ \AA}$ to $(7+x) \text{ \AA}$ range need to be checked. So the nearest neighbors of an *arbitrary* point can simply be found by computing *one* nearest neighbor, and looking up the nearest neighbors of this neighbor. To be able to look up all nearest neighbors for an arbitrary point within a radius of $r \text{ \AA}$, we need to precompute nearest neighbors for each target within a radius of $2r \text{ \AA}$, since x can be at most $r \text{ \AA}$. This is exactly what is currently implemented;

the indices of the nearest neighbors and corresponding distances are stored (in compressed format) for each target point. This adds only marginally to the total table size, while providing significant speedups.

The LabelHash Server and Chimera Interface

The LabelHash algorithm has been made accessible through a web interface at <http://labelhash.kavrakilab.org>. A user can specify a motif by a PDB ID and a number of residue ID's. For each motif residue the user can optionally specify a number of alternate labels. The user can match against either the full PDB or the 95% sequence identity filtered non-redundant PDB (nrPDB₉₅). Once the matches for a motif have been computed, an email is sent to the user which includes a URL for a web page with the match results. This page lists information for the top matches, including: the matched residues, the RMSD, the p -value, Enzyme Commission (EC) numbers and Gene Ontology (GO) annotations (if applicable), and a graphical rendering of the match aligned with the motif.

From the results page, one can also download an XML file that contains *all* the matches found. This match file can then be loaded in Chimera, a popular molecular visualization and analysis program [56]. For this, we have developed a plugin called ViewMatch. (It is derived from the Chimera ViewDock extension for visualizing results from the DOCK program.) The ViewMatch plugin allows the user to scroll through the list of matches. Figure 2 shows the user interface. In the main window a selected match is shown superimposed with the motif. Recall that, although all atoms in the matched residues are shown, only the C_α atoms were used to compute the alignment. The C_α atoms are shown as spheres. In the controller window on the right, all matches are listed in the top half with their RMSD to the motif, p -value and other attributes. By specifying constraints on the match attributes, the user can restrict the matches that are shown. The bottom half of the window shows additional information for the selected match, such as the EC classification and GO terms. By clicking the PDBsum button, the PDBsum web pages [57] are shown for the selected matches. This gives the user an enormous amount of information about a match.

We expect that advanced users may want to have more control over the matching parameters and the creation of LabelHash tables. For that reason, we have made available for download on the LabelHash web site a suite of command line tools (for Linux and OS X) to do just that. It includes programs to run LabelHash in parallel on a cluster of machines, as well as a Python interface.

Testing

The LabelHash algorithm has been evaluated for speed, scalability, and functional annotation precision through a carefully selected set of case studies that demonstrate orthogonal features of the algorithm. First, the ability of LabelHash to identify common substructures at the superfamily level of SCOP [58] classification is validated within the Enolase Superfamily (ES). Second, the ability of LabelHash to construct motifs from the output of state-of-the-art methods and then quickly match them against the entire PDB is demonstrated with the output of SOIPPA [35], a motif discovery and structural alignment algorithm. Third, as a comprehensive functional annotation benchmark, LabelHash motifs are created from catalytic residues documented by the Catalytic Site Atlas (CSA) [59] and then matched against the nrPDB₉₅ to assess per-motif functional annotation sensitivity and specificity. Finally, we evaluate the speed and scalability of the LabelHash algorithm by matching all motifs used in this paper against the nrPDB₉₅. Together, these tests demonstrate the versatility and generality of LabelHash for a variety of functional annotation problems.

Identifying Members of the Enolase Superfamily

To demonstrate the ability of LabelHash to successfully identify motifs at the superfamily scale, spanning multiple EC classes, LabelHash was used to identify shared catalytic substructures among Enolase Superfamily (ES) [60] proteins. Previous work by Meng et al. [61] used the SPASM [62] substructure comparison method to investigate a conserved substructure within ES proteins. As a challenging validation experiment, LabelHash was used to identify the conserved substructure (ESmotif) among the same benchmark set of ES structures (ESdb) that were both defined previously by Meng et al. [61]. LabelHash was able to identify those structures included in the ESdb with high sensitivity and specificity. In addition to those structures included in the ESdb, matches were identified by LabelHash to additional ES proteins which have more recently been identified and to several proteins with only putative ES-like function.

The ES currently includes 7 major sub-groups spanning 20 families and 14 well-defined EC classes as defined by the Structure-Function Linkage Database (SFLD) [63]. As shown in Fig. 3, a core of five residues directly mediates the conserved partial reaction among ES members. Modeling this substructure as a five-residue motif as done by Meng et al. [61] results in the following superfamily-specific motif (ES-motif): 164^{KH}, 195^D, 221^E, 247^{EDN}, 297^{HK}; residues are numbered with respect to mandelate

racemase structure [PDB:2MNR]. A benchmark set of ES structures defined by Meng et al. [61] (ESdb) was used as a positive test set; the nrPDB₉₅ was used as a background structure set which also includes several overlapping structures belonging to the ES. Residues making up the conserved substructure among ES members are known to vary widely in C_α RMSD (up to 2.3 Å C_α RMSD) although the side-chains of the residues superimpose closely (as shown in Fig. 3). By using LabelHash to enumerate all reasonable substructure matches based on C_α RMSD and then filtering the resulting potential matches by side-chain RMSD, LabelHash successfully identifies ES members with high sensitivity and specificity using the approach outlined below.

To identify matches based on a different distance metric rather than C_α deviation involves adding only a simple post-processing step to the initial substructure matches identified by LabelHash. For each target structure in the ESdb, many possible matches are first identified based on amino acid label compatibility and C_α distance cutoffs to the ES motif. From these possible matches, a single “best” match is selected per target based upon minimum side-chain centroid RMSD to the motif rather than C_α RMSD as done in previous experiments. The distribution of matches yielded by this process differs from the minimum C_α distance distribution. For example, enumerating all possible matches in enolase structure [PDB:1EBH] results in 84 possible matches and from these possible matches, a “best” match can be selected based on *any* user-defined criteria. For ES we used side-chain centroid RMSD, but in general C_α RMSD, surface accessibility, distance to a known ligand, or even colocation with a pocket/cavity could be used instead. For the ES, side-chain centroid RMSD is preferred over full-atom side-chain RMSD due to the fact that several different types of amino acids are possible at each position and defining an appropriate one-to-one atom mapping between side-chains of different amino acids is difficult to define. However, C_β alignment is a viable alternative to the side-chain centroid if a pseudo atom is defined for glycine.

Comparing the prediction performance of C_α versus side-chain centroid RMSD for ES, as shown in Fig. 4, demonstrates that alternative match selection measures, such as side-chain deviation, outperforms C_α deviation for the ES. Using minimum C_α RMSD as a match selection criteria per target, only 31% of structures in the ESdb are matched with statistical significance (*p*-value threshold of $\alpha = 0.01$, 29 false negative matches). However, with minimum side-chain centroid RMSD as a match selection criterion per target, LabelHash achieves 95% sensitivity. The match RMSD distributions in Fig. 4 reveal the higher

discriminating power of side-chain centroid RMSD for the ES. Using side-chain centroid RMSD, the distribution of matches for the ESdb structures is easily separable from the much larger number of matches to nrPDB₉₅ structures while C_α RMSD alone results in the majority of ESdb matches to overlap with unrelated matches to structures in the nrPDB₉₅.

Many matches to nrPDB₉₅ structures still fall below the statistical significance threshold (as shown in Additional File 1) and these possibly false positive (FP) matches were further investigated. Out of the 47 total FP matches, 35 matches actually corresponded to structures that have now been identified to belong to ES as documented by the SFLD [63]. An additional 10 FP matches corresponded to structures that have only putatively defined functions, but may be related to the ES: [PDB:2QGY], [PDB:3CK5], [PDB:2OO6], [PDB:2PPG], [PDB:2OQH], [PDB:3CYJ], [PDB:1WUF], [PDB:1WUE], [PDB:2OZ8], and [PDB:2POZ]. The final 2 matches correspond to an aminoacylase from the M20 family [PDB:1YSJ] and human arginase I [PDB:2AEB]; both matches correspond to dimetal binding sites within the two different enzymes which bind Mn²⁺ and Ni²⁺, respectively. Altogether, these results highlight the potential of LabelHash to identify similarities between remote homologs.

Combining LabelHash and SOIPPA

Recent work by Xie and Bourne on Sequence Order-Independent Profile-Profile Alignment (SOIPPA) [35] and SMAP [36] allows the alignment and identification of structurally similar, but sequentially distinct, sequence motifs. To further validate the LabelHash method, the alignment output of SOIPPA was used to create LabelHash motifs that are demonstrated to maintain the high-specificity of SOIPPA. LabelHash provides the means to search the entire nrPDB₉₅ for matches to SOIPPA-derived motifs in a matter of minutes. These experiments illustrate how LabelHash can be used to further enhance available structural analysis methods by converting identified residues of interest directly to LabelHash motifs that can be efficiently scanned against large structural databases.

Several of the aligned sequence motifs identified by SOIPPA were used to construct the LabelHash motifs listed in Table 1. For each template structure in this table, a LabelHash motif was defined using the residues and alternate amino acid labels identified by SOIPPA alignment [35]. Specifically, to construct a LabelHash motif from the SOIPPA alignment output data, the coordinates of the SOIPPA-identified motif residues from *only* the

template structure are combined with the set of alternate amino acid labels identified by SOIPPA-alignment; the coordinates of the structures aligned by SOIPPA (i.e., structures other than the template structure) are not used. Each LabelHash motif was then matched against the corresponding SOIPPA alignment structures in the table, in addition to the nrPDB₉₅ as a background structure set, in order to assess the statistical significance of matches in each corresponding alignment structure. In all cases the correct match was identified with extremely low *p*-values. Both the LabelHash motif substructure and corresponding match substructures can be seen in Figs. 5 and 6.

Combining multiple pairwise SOIPPA alignments allows for the construction of a single LabelHash motif that comprises the common SOIPPA motif residues shared between each pairwise alignment. For example, based upon SOIPPA pairwise alignments between SAM-dependent methyltransferase [PDB:1ZQ9] and each of urocanase [PDB:1UWK], carbonyl reductase [PDB:1CYD], and flavocytochrome c3 fumarate [PDB:1D4D], a single LabelHash motif was derived using the subset of residues from [PDB:1ZQ9] identified by SOIPPA [35] within each of the three aforementioned pairwise alignments and the alternative amino acid labels for residues in each pairwise SOIPPA alignment. The resulting LabelHash motif correctly identifies the corresponding residues in each of [PDB:1UWK], [PDB:1CYD], and [PDB:1D4D] without LabelHash taking into account any knowledge of each of these three target structures beyond the alternate labels required. The [PDB:1ZQ9]-based LabelHash motif was then matched against the full PDB to search for other structures with statistically significant similarity to the shared substructure modeled by the motif.

Using LabelHash to match the [PDB:1ZQ9]-based motif versus the PDB reveals many additional matches with low *p*-values to binding sites for adenine-containing molecules including FAD, NAD, and SAM. Although LabelHash does not take into account the presence, position, or absence of a bound ligand at any point in the matching process, many of the identified matches could be confirmed as binding sites because of a bound adenine-containing ligand. Matches to lipoamide dehydrogenase [PDB:2YQU], urocanase [PDB:1W1U], and MnmC2 [PDB:2E58] were identified with *p*-values less than 0.0001, 0.0005, and 0.0006, respectively. Each match identified is shown in Fig. 7.

Catalytic Site Atlas Motifs

In our final experiment we matched a large number of motifs from the Catalytic Site Atlas (CSA) [59] against

the nrPDB₉₅. The CSA is a manually-curated database of literature-documented functional sites. All experiments were conducted using CSA version 2.2.11 which contains 968 literature-documented sites. The subset of CSA sites used in our experiments includes only literature-documented sites that adhere to the following criteria:

1. The PDB structure corresponding to the site must have a fully qualified EC classification (i.e., 3.4.1.1 is fully qualified but 3.4.1.– or 3.2.– or 3.– are all not); each site will be compared against proteins within the same EC classification to assess sensitivity.
2. The EC class corresponding to the documented site must contain more than 10 structures so that meaningful per-site statistics can be computed for all sites included in the benchmark.
3. The LabelHash motif based upon the CSA site must contain at least one valid reference set (prerequisite for matching; see *Methods* for details on reference set constraints).
4. The site must contain a minimum of 4 non-hetero residues (i.e., residues referencing prosthetic groups, coordinated metal ions, or other ligands do not count towards the 4 residue minimum) because the current LabelHash implementation does not support matching heteroresidues/heteroatoms.

From the available 968 CSA sites, 147 adhered to our criteria. From each of these 147 CSA sites, a LabelHash motif was constructed using the PDB structure that corresponded to each CSA record in every case. Because the CSA does not include allowable amino acid substitutions or mutations for site records, only the single amino acid label that corresponds to the amino acids present in the PDB structure were used in all cases. The motifs analyzed here range in size from 4 to 8 residues, while previous work by Torrance et al. [64] only investigated the performance of CSA motifs ranging from 3–5 residues in size.

To assess the ability of each CSA-based motif to identify functionally equivalent catalytic sites in protein structure, each CSA-based motif was matched against the EC class corresponding to the PDB structure of each site. Structures sharing EC classification with the CSA-based motif are considered positive matches while matches to structures outside of the motif EC class within the nrPDB₉₅ were considered negative matches. For the CSA benchmark, the individual chains in both positive and negative protein structures are matched individually. For the

147 CSA-based motifs, 118 unique EC classes are represented spanning all 6 top-level EC classifications (oxidoreductases, transferases, hydrolases, lyases, isomerases, ligases).

Overall, the ability of CSA motifs to identify members of the same EC class (4th-level EC classification) varied wildly. This result is not unexpected due to the fact that per-EC class coverage was not a consideration in the design of the CSA [64]; the aforementioned study by Torrance et al. [64] considered the set of positive structures for each motif to be the “CSA family” of PSI-BLAST identified relatives rather than the full, 4th-level EC class. The analysis performed here widens the study by considering not only those structures with sequence similarity identifiable by PSI-BLAST, but the entire set of structures sharing an EC class with each motif. Highly successful motifs, such as the pyruvate kinase and xylose isomerase motifs based upon structures [PDB:1PKN] and [PDB:2XIS], respectively, correctly identify more than 100 matches to their corresponding EC classes, achieving > 90% EC-class sensitivity with at least 99.8% specificity against the entire nrPDB₉₅.

The CSA contains multiple catalytic site definitions for a single EC class in several cases (as can be seen in Additional File 2) and the motifs based on each definition sometimes had large differences in function prediction performance. Consider, for example, two motifs are defined by the CSA for EC:2.1.1.45 (thymidylate synthases): the motif based upon structure [PDB:1LCB] from *Lactobacillus casei* was defined as {198^C, 219^S, 221^D, 257^D, 259^H, 60^E} and matched 192/218 chains with significance resulting in 88.1% sensitivity at a *p*-value threshold of 0.002, while the motif based upon structure 1TY5 from *Escherichia coli* was defined as {146^S, 166^R, 169^D, 58^E, 94^Y} matched only 4/218 chains with significance resulting in a drastically lower 1.8% sensitivity at the same *p*-value threshold. Examining the EC:2.1.1.45 matches produced by the poorer performing [PDB:1TY5] motif revealed that a single motif residue (146^S) was matching a non-cognate serine residue in the majority of EC:2.1.1.45 structures that caused the RMSD of these matches to increase drastically and fall outside of the significance threshold. However, pruning 146^S from the [PDB:1TY5] motif increased the sensitivity to 80.3%, matching 175/218 chains in EC:2.1.1.45.

In other cases where multiple motifs were defined by the CSA for a single EC class, each motif was found to match a different subset of functionally related structures. For example, the CSA defines a separate motif for both cellobiohydrolase (CBH) I and II from *Trichoderma reesei* which both belong to EC:3.2.1.91 (cellulose

1,4- β -cellobiosidases). Individually, the CBH I and II motifs have a sensitivity of 43.7% and 21.1%, respectively, at $p = 0.002$, but because both motifs match a mutually exclusive set of structures within EC:3.2.1.91, the combination of both motifs identifies over 60% of structures within EC:3.2.1.91. However, combining the results from multiple motifs by taking the union of all positive matches will result in a decrease in overall specificity due to lack of multiple testing correction. In recent work [65], we have developed a technique called Motif Ensemble Statistical Hypothesis testing that allows multiple (potentially overlapping) motifs to be statistically combined into a single function prediction test.

Many of the CSA motifs included in the benchmark could be further optimized to increase function prediction sensitivity and specificity. The previously developed Geometric Sieving (GS) method for motif refinement [34] identifies subsets of motif residues within a larger motif that preserve the overall specificity of the larger motif while increasing sensitivity by removing motif residues that do not contribute to specificity. Also in previous work [34], evolutionary methods based upon Multiple Sequence Alignment (MSA) and phylogenetic tree construction, such as Evolutionary Trace [45] and ConSurf [66], have been used to automatically define motifs and identify alternate residue labels in order to represent residue substitutions present in a MSA. Therefore, while the CSA motifs examined in this benchmark serve as an objective set of expert-defined catalytic site definitions, many motifs could be improved further by applying our previously investigated optimization techniques.

Performance and Scalability

Figure 8 shows the average performance of matching all motifs mentioned above against *all* 21,745 targets in the nrPDB₉₅. The main observation in the left graph is that most motifs can be matched against the entire nrPDB₉₅ in a few seconds up to a few minutes. To the best of our knowledge no other general purpose substructure matching algorithm implementation comes even close to such performance. The right graph shows the speedup. The speedup for n cores is simply $\frac{T_1}{T_n}$, where T_i is the wallclock time for i cores. The ideal, linear speedup is shown with a black dotted line. Benchmark results for up to 128 cores are shown in Additional File 3.

There were 4 motifs used in the comparison with SOIPPA, 14 motifs in the Enolase Superfamily study, and 147 Catalytic Site Atlas motifs. In general, wallclock time can be reduced significantly by using more cores. The motifs that take more time to match tend to be larger and

have more alternate residue labels. For small, simple motifs (3–4 residues without alternate labels) the wallclock time with a single core is already very small and there is not much room for improvement by adding more cores. In this case the runtime is dominated by file access and communication between the cores, with only minimal time needed for computation. The difference between the solid and dashed lines show the performance increase when switching from a shared Panasas filesystem to a local solid-state drive. This difference confirms that much of the time is spent simply reading data. In other words, for simple motifs there is limited room for performance improvement by modifying the match augmentation algorithm. For motifs that take longer to match (the SOIPPA and Enolase superfamily motifs) we see that an almost linear speedup can be expected.

Conclusions

We have presented LabelHash, a novel algorithm for matching structural motifs. It quickly matches a motif consisting of residue positions, and possible residue types to all structures in the PDB (or any subset thereof). We have shown that LabelHash often achieves very high sensitivity and specificity. The statistical significance of matches is computed using a non-parametric model. Typically, the number of false positive matches is much smaller than the number of true positive matches, despite the large number of targets in our background database. This greatly speeds up the analysis of match results. Our algorithm uses only a small number of parameters whose meaning is easy to understand and which usually do not need to be changed from their default values. Matching results are easily visualized through a plugin for Chimera, a molecular modeling program.

Extensibility was an important factor in the design of the LabelHash implementation. Our program is easily extended to incorporate additional constraints or use even conceptually different types of motifs. For instance, matching based on physicochemical pseudocenters [22, 46] could easily be incorporated, and we plan to offer this functionality in the future. Input and output are all in XML format (except for the LabelHash table itself), which enables easy integration with other tools or web services. As we demonstrated in the study of the enolase superfamily, LabelHash can be part of a multi-stage matching pipeline, where matches produced by LabelHash can be given to the next program, which can apply additional constraints to eliminate more false positives. As long as the set of matches produced by LabelHash

include all functional homologs, this is a viable strategy. Of course, the output of LabelHash can also easily be passed on to any clustering algorithm or a visualization front-end.

Our matching algorithm has the capability to keep partial matches and multiple matches per target. This makes the statistical analysis significantly more complicated. Currently, we just disable the p -value computation when either option is selected, but we plan to investigate the modeling of the statistical distribution of these matches.

Methods

Data sets

For most of our tests we used one LabelHash table for the entire Protein Database (PDB), based on a snapshot taken at March 25, 2009. Each chain was inserted separately in the LabelHash file. This resulted in roughly 138,000 targets. Molecular surfaces (used to compute each residue's depth) were computed with the MSMS software [67]. We chose to use reference sets of size 3. The following parameter values were used for the reference sets: $d_{\text{maxmindist}} = 16 \text{ \AA}$, $d_{\text{diameter}} = 25 \text{ \AA}$, $d_{\text{maxmindepth}} = 1.6 \text{ \AA}$, and $d_{\text{maxdepth}} = 3.1 \text{ \AA}$. The same parameters were used in all experiments. These values were chosen such that motifs generally contained at least one reference set of size 3. They are very generous in the sense that most motifs we have used for evaluating the LabelHash algorithm contain many reference sets. If reference sets of more than 3 residues are used, the values of the distance parameters need to be increased to guarantee that each motif contains at least one reference set. The advantage of larger reference sets is that we instantly match a larger part of a motif. However, increasing these values also quickly increases the number of reference sets in the targets. So the number of reference sets to perform match augmentation on will also quickly increase. Finally, the storage required for the hash tables grows rapidly with reference set size. After the preprocessing phase the total size of the LabelHash table for the entire PDB was 65 GB, while with reference sets of size 4 (and the same parameters) the table would grow to approximately 5 TB (estimated based on a table constructed for a subset of the PDB).

Benchmarking

The computational experiments were carried out on a Linux cluster where each node has two quad core 2.4 GHz Intel Xeon (Nehalem) CPUs with 8 MB cache. Each node

has 12 GB of memory per node shared by all cores on the node. So for the results in figure 8 only 2 nodes were used, while for the results in Additional File 3 up to $128/8 = 16$ nodes were used. The 65 GB table did not fit on our solid-state drives, so a separate table for just the nrPDB₉₅ (21,745 targets, 9.7 GB) was used in the benchmarking tests.

Authors contributions

MM and LEK designed the LabelHash algorithm. All authors collectively conceived and designed the experiments and analyzed the resulting data. DHB and MM contributed computational analysis tools/software and performed experiments. All authors wrote the paper, and read and approved the final manuscript.

Acknowledgements

The authors are indebted to V. Fofanov for many useful discussions on the use of statistical analysis and for his comments on LabelHash. They are also deeply grateful for the help of B. Chen with MASH. This work has benefited from earlier contributions by O. Lichtarge, M. Kimmel, D. Kristensen and M. Lisewski within the context of an earlier NSF funded project. This work was supported in part by NSF grant DBI-0960612 and by the John & Ann Doerr Fund for Computational Biomedicine, at Rice University. This material is based in part upon work supported by the Texas Norman Hackerman Advanced Research Program under Grant No. 01907. LK has also been supported by a Sloan Fellowship. DHB was additionally supported by NSF GRFP grant DGE-0237081. The computers used to carry out experiments of this project were funded by NSF CNS-0454333 and NSF CNS-0421109 in partnership with Rice University, AMD and Cray. Molecular graphics images were produced using the UCSF Chimera package from the Resource for Biocomputing, Visualization, and Informatics at the University of California, San Francisco (supported by NIH P41 RR-01081) [56].

References

1. Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE: **The Protein Data Bank**. *Nucleic Acids Research* 2000, **28**:235–242.
2. Zhang C, Kim SH: **Overview of structural genomics: from structure to function**. *Current Opinion in Chemical Biology* 2003, **7**:28–32.

3. Chandonia JM, Brenner SE: **The impact of structural genomics: expectations and outcomes.** *Science* 2006, **311**(5759):347–351.
4. Kouranov A, Xie L, de la Cruz J, Chen L, Westbrook J, Bourne PE, Berman HM: **The RCSB PDB information portal for structural genomics.** *Nucleic Acids Res* 2006, **34**(Database issue):D302–5.
5. Altschul SF, Gish W, Miller W, Myers EW, Lipman DJ: **Basic local alignment search tool.** *J. Mol. Biol* 1990, **215**(3):403–410.
6. Thompson JD, Higgins DG, Gibson TJ: **CLUSTAL W: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice.** *Nucleic Acids Research* 1994, **22**(22):4673–4680.
7. Eddy SR: **Hidden Markov models.** *Curr Opin Struct Biol* 1996, **6**(3):361–365.
8. Finn RD, Tate J, Mistry J, Coghill PC, Sammut SJ, Hotz HR, Ceric G, Forslund K, Eddy SR, Sonnhammer ELL, Bateman A: **The Pfam protein families database.** *Nucleic Acids Res* 2008, **36**(Database issue):D281–8.
9. Hermann JC, Marti-Arbona R, Fedorov AA, Fedorov E, Almo SC, Shoichet BK, Rauschel FM: **Structure-based activity prediction for an enzyme of unknown function.** *Nature* 2007, **448**(7155):775–779.
10. Gherardini PF, Helmer-Citterich M: **Structure-based function prediction: approaches and applications.** *Briefings in functional genomics & proteomics* 2008, **7**(4):291–302.
11. Watson J, Laskowski R, Thornton J: **Predicting protein function from sequence and structural data.** *Current Opinion in Structural Biology* 2005, **15**(3):275–284.
12. Holm L, Sander C: **Protein structure comparison by alignment of distance matrices.** *J Mol Biol* 1993, **233**:123–138.
13. Zhu J, Weng Z: **FAST: a novel protein structure alignment algorithm.** *Proteins* 2005, **58**(3):618–627.
14. Binkowski TA, Freeman P, Liang J: **pvSOAR: detecting similar surface patterns of pocket and void surfaces of amino acid residues on proteins.** *Nucleic Acids Res* 2004, **32**:W555–W558.
15. Dundas J, Ouyang Z, Tseng J, Binkowski A, Turpaz Y, Liang J: **CASTp: computed atlas of surface topography of proteins with structural and topographical mapping of functionally annotated residues.** *Nucleic Acids Research* 2006, **34**(Web Server issue):W116–W118.
16. Laskowski RA: **SURFNET: a program for visualizing molecular surfaces, cavities, and intermolecular interactions.** *J Mol Graph* 1995, **13**(5):323–330.
17. Kristensen DM, Ward RM, Lisewski AM, Chen BY, Fofanov VY, Kimmel M, Kavradi LE, Lichtarge O: **Prediction of enzyme function based on 3D templates of evolutionary important amino acids.** *BMC Bioinformatics* 2008, **9**(17).
18. Glaser F, Rosenberg Y, Kessel A, Pupko T, Ben-Tal N: **The ConSurf-HSSP database: the mapping of evolutionary conservation among homologs onto PDB structures.** *Proteins* 2005, **58**(3):610–617.
19. Chakrabarti S, Lanczycki C: **Analysis and prediction of functionally important sites in proteins.** *Protein Science* 2007, **16**:4.
20. Wallace AC, Borkakoti N, Thornton JM: **TESS: A geometric hashing algorithm for deriving 3D coordinate templates for searching structural databases. Application to enzyme active sites.** *Protein Science* 1997, **6**(11):2308.
21. Kleywegt GJ: **Recognition of spatial motifs in protein structures.** *J Mol Biol* 1999, **285**(4):1887–1897.
22. Schmitt S, Kuhn D, Klebe G: **A new method to detect related function among proteins independent of sequence and fold homology.** *J Mol Biol* 2002, **323**(2):387–406.
23. Kinoshita K, Furui J, Nakamura H: **Identification of protein functions from a molecular surface database, eF-site.** *J Struct Funct Genomics* 2002, **2**:9–22.
24. Spriggs RV, Artymiuk PJ, Willett P: **Searching for patterns of amino acids in 3D protein structures.** *J Chem Inf Comput Sci* 2003, **43**(2):412–21.
25. Stark A, Russell RB: **Annotation in three dimensions. PINTS: Patterns in Non-homologous Tertiary Structures.** *Nucleic Acids Research* 2003, **31**(13):3341–3344.
26. Barker JA, Thornton JM: **An algorithm for constraint-based structural template matching: application to 3D templates with statistical analysis.** *Bioinformatics* 2003, **19**(13):1644–1649.
27. Jambon M, Imberty A, Deléage G, Geourjon C: **A new bioinformatic approach to detect common 3D sites in protein structures.** *Proteins* 2003, **52**(2):137–45.
28. Shulman-Peleg A, Nussinov R, Wolfson HJ: **Recognition of functional sites in protein structures.** *J Mol Biol* 2004, **339**(3):607–633.
29. Ausiello G, Via A, Helmer-Citterich M: **Query3d: a new method for high-throughput analysis of functional residues in protein structures.** *BMC Bioinformatics* 2005, **6** Suppl 4:S5.
30. Laskowski RA, Watson JD, Thornton JM: **ProFunc: a server for predicting protein function from 3D structure.** *Nucleic Acids Research* 2005, **33**:W89–W93.
31. Pal D, Eisenberg D: **Inference of protein function from protein structure.** *Structure* 2005, **13**:121–130.
32. Gold ND, Jackson RM: **Fold independent structural comparisons of protein-ligand binding sites for exploring functional relationships.** *J Mol Biol* 2006, **355**(5):1112–1124.
33. Kinjo AR, Nakamura H: **Similarity search for local protein structures at atomic resolution by exploiting a database management system.** *Biophysics* 2007, **3**:75–84.
34. Chen BY, Fofanov VY, Bryant DH, Dodson BD, Kristensen DM, Lisewski AM, Kimmel M, Lichtarge O, Kavradi LE: **The MASH pipeline for protein function prediction and an algorithm for the geometric refinement of 3D motifs.** *J. Comp. Bio.* 2007, **14**(6):791–816.
35. Xie L, Bourne PE: **Detecting evolutionary relationships across existing fold space, using sequence order-independent profile-profile alignments.** *Proc Natl Acad Sci U S A* 2008, **105**(14):5441–5446.
36. Xie L, Xie L, Bourne PE: **A unified statistical model to support local sequence order independent similarity searching for ligand-binding sites and its application to genome-based drug discovery.** *Bioinformatics* 2009, **25**(12):i305–312.
37. Halperin I, Glazer DS, Wu S, Altman RB: **The FEATURE framework for protein function annotation: modeling new functions, improving performance, and extending to novel applications.** *BMC Genomics* 2008, **9** Suppl 2:S2.
38. Tseng YY, Dundas J, Liang J: **Predicting protein function and binding profile via matching of local evolutionary and geometric surface patterns.** *J Mol Biol* 2009, **387**(2):451–464.
39. Wolfson HJ, Rigoutsos I: **Geometric hashing: an overview.** *IEEE Computational Science and Engineering* 1997, **4**(4):10–21.

40. Nussinov R, Wolfson HJ: **Efficient detection of three-dimensional structural motifs in biological macromolecules by computer vision techniques.** *Proc Natl Acad Sci U S A* 1991, **88**(23):10495–10499.
41. Chang DTH, Chen CY, Chung WC, Oyang YJ, Juan HF, Huang HC: **ProteMiner-SSM: a web server for efficient analysis of similar protein tertiary structures.** *Nucleic Acids Res* 2004, **32**(Web Server issue):W76–82.
42. Stark A, Sunyaev S, Russell RB: **A Model for Statistical Significance of Local Similarities in Structure.** *Journal of Molecular Biology* 2003, **326**(5):1307–1316.
43. Fofanov VY, Chen BY, Bryant DH, Moll M, Lichtarge O, Kavraki LE, Kimmel M: **A statistical model to correct systematic bias introduced by algorithmic thresholds in protein structural comparison algorithms.** In *IEEE Intl. Conf. on Bioinformatics and Biomedicine Workshops (BIBMW)* 2008:1–8.
44. Yao H, Kristensen DM, Mihalek I, Sowa ME, Shaw C, Kimmel M, Kavraki L, Lichtarge O: **An accurate, sensitive, and scalable method to identify functional sites in protein structures.** *J Mol Biol* 2003, **326**:255–261.
45. Lichtarge O, Bourne HR, Cohen FE: **An evolutionary trace method defines binding surfaces common to protein families.** *J Mol Biol* 1996, **257**(2):342–358.
46. Artymiuk PJ, Poirrette AR, Grindley HM, Rice DW, Willett P: **A Graph-theoretic Approach to the Identification of Three-dimensional Patterns of Amino Acid Side-chains in Protein Structures.** *Journal of Molecular Biology* 1994, **243**(2):327–344.
47. Fofanov VY: **Statistical Models in Protein Structural Alignments.** *PhD thesis*, Department of Statistics, Rice University, Houston, TX 2008.
48. Sheather SJ, Jones MC: **A reliable data-based bandwidth selection method for kernel density estimation.** *J Royal Statistical Society. Series B (Methodological)* 1991, **53**(3):683–690.
49. The HDF Group: *Hierarchical data format version 5.* <http://www.hdfgroup.org/HDF5> 2000–2010.
50. Yeh PS, Xia-Serafino W, Miles L, Kobler B, Menasce D: **Implementation of CCSDS Lossless Data Compression in HDF.** In *Earth Science Technology Conference* 2002.
51. Kabsch W: **A solution of the best rotation to relate two sets of vectors.** *Acta Crystallographica A* 1976, **32**:922–923.
52. Coutsiaris EA, Seok C, Dill KA: **Using quaternions to calculate RMSD.** *J. Comp. Chemistry* 2004, **25**(15):1849–1849.
53. Liu P, Agrafiotis DK, Theobald DL: **Fast determination of the optimal rotational matrix for macromolecular superpositions.** *J Comput Chem* 2010.
54. Theobald DL: **Rapid calculation of RMSDs using a quaternion-based characteristic polynomial.** *Acta Crystallogr A* 2005, **61**(Pt 4):478–480.
55. Brin S: **Near Neighbor Search in Large Metric Spaces.** In *Proc. 21st Conf. on Very Large Databases* 1995:574–584.
56. Pettersen EF, Goddard TD, Huang CC, Couch GS, Greenblatt DM, Meng EC, Ferrin TE: **UCSF Chimera—A visualization system for exploratory research and analysis.** *Journal of Computational Chemistry* 2004, **25**(13):1605–1612.
57. Laskowski RA: **PDBsum: summaries and analyses of PDB structures.** *Nucleic Acids Research* 2001, **29**:221–222.
58. Murzin AG, Brenner SE, Hubbard T, Chothia C: **SCOP: a structural classification of proteins database for the investigation of sequences and structures.** *J Mol Biol* 1995, **247**(4):536–540.
59. Porter CT, Bartlett GJ, Thornton JM: **The Catalytic Site Atlas: a resource of catalytic sites and residues identified in enzymes using structural data.** *Nucleic Acids Res* 2004, **32**(Database issue):D129–33.
60. Babbitt PC, Hasson MS, Wedekind JE, Palmer DR, Barrett WC, Reed GH, Rayment I, Ringe D, Kenyon GL, Gerlt JA: **The enolase superfamily: a general strategy for enzyme-catalyzed abstraction of the alpha-protons of carboxylic acids.** *Biochemistry* 1996, **35**(51):16489–16501.
61. Meng EC, Polacco BJ, Babbitt PC: **Superfamily active site templates.** *Proteins* 2004, **55**(4):962–976.
62. Kleywegt GJ: **Recognition of spatial motifs in protein structures.** *Journal of Molecular Biology* 1999, **285**(4):1887–1897.
63. Pegg SCH, Brown SD, Ojha S, Seffernick J, Meng EC, Morris JH, Chang PJ, Huang CC, Ferrin TE, Babbitt PC: **Leveraging enzyme structure-function relationships for functional inference and experimental design: the structure-function linkage database.** *Biochemistry* 2006, **45**(8):2545–2555.
64. Torrance JW, Bartlett GJ, Porter CT, Thornton JM: **Using a library of structural templates to recognise catalytic sites and explore their evolution in homologous families.** *J Mol Biol* 2005, **347**(3):565–81.
65. Bryant DH, Moll M, Chen BY, Fofanov VY, Kavraki LE: **Analysis of substructural variation in families of enzymatic proteins with applications to protein function prediction.** *BMC Bioinformatics* 2010, **11**(242).
66. Glaser F, Pupko T, Paz I, Bell RE, Bechor-Shental D, Martz E, Ben-Tal N: **ConSurf: identification of functional regions in proteins by surface-mapping of phylogenetic information.** *Bioinformatics* 2003, **19**:163–164.
67. Sanner MF, Olson AJ, Spehner JC: **Reduced surface: an efficient way to compute molecular surfaces.** *Biopolymers* 1996, **38**(3):305–320.
68. Zhang E, Hatada M, Brewer JM, Lebioda L: **Catalytic metal ion binding in enolase: the crystal structure of an enolase-Mn²⁺-phosphonoacetohydroxamate complex at 2.4-Å resolution.** *Biochemistry* 1994, **33**(20):6295–6300.

Figures

Figure 1 - Non-parametric model of statistical significance of matches

Accounting for the 'missed' matches due to the RMSD cut-off parameter with a point-weight at ∞ allows for accurate estimation of the p -value of matches.





Figure 2 - Match visualization using Chimera

Left: A match (in green) shown superimposed with a motif (in white), while the rest of the matching protein is shown in ribbon representation. **Right:** The ViewMatch controller window.



ViewMatch - /Users/mmolli/labelhash/motifs/lady-7A.xml

S	name	rmsd	pvalue	depth	score
	V1uzg8	1.11678	0.000678561	1.43279	6
	V1dbiA	1.12611	0.000730073	2.83062	6
	V2pifA	1.12612	0.000730074	1.46846	6
	V1qe0A	1.13561	0.000785013	1.50039	6
	V4pgaB	1.15275	0.000961279	1.58947	6
	V2v4wA	1.15585	0.00096148	1.92598	6
	V2jd4B	1.15912	0.00102634	1.58436	6

Chimera Model #16

EC class: 6.1.1.21
GO terms:
MOLECULAR FUNCTION: ligase activity, ATP binding, aminoacyl-tRNA ligase activity, nucleotide binding, histidine-tRNA ligase activity
BIOLOGICAL PROCESS: translation, histidyl-tRNA aminoacylation, tRNA aminoacylation for protein translation
CELLULAR COMPONENT: cytoplasm
HEADER LIGASE 12-JUL-99 IQE0
COMPND MOL_ID: 1;

Change Match State
 Viable Deleted Purged

Load PDB file PDBsum Hide Quit Help

Figure 3 - The enolase superfamily (ES) motif.

(a) The substructure responsible for the conserved partial reaction among members of the ES consists of the five residues shown above from which the ES motif used here is derived [61]. (b) Enolase from *Saccharomyces cerevisiae* [PDB:1ELS] with demonstrated divalent ion coordination to the conserved superfamily substructure [68]. A single Mn^{2+} is coordinated by a carboxylate triad (Asp, Glu, Asp in this enzyme) and a second Mn^{2+} binds in a substrate specific manner [68].




Figure 4 - LabelHash results for ES motif based on [PDB:2MNR].

Dark gray denotes matches to structures in the ESdb as defined by Meng et al. [61] while light gray denotes matches to structures in the nrPDB₉₅. The dashed line in each plot corresponds to the p -value threshold at $\alpha = 0.01$. (a) The clear separation of the distributions of ESdb and nrPDB₉₅ matches identified using minimum side-chain RMSD illustrates the high-specificity of the 2MNR-based motif. (b) Examining the distributions of matches based upon C_{α} RMSD demonstrates the inseparability of ESdb matches from the majority of nrPDB₉₅ structures if minimum C_{α} RMSD is used alone for match selection.




Figure 5 - SOIPPA motifs

In each image, the motif substructure is colored green while the matched substructure is colored purple. All ligands neighboring either the motif or matched substructure are shown colored by atom type for reference; the ligand atoms are not used at any point during the LabelHash matching process. **(a)** LabelHash motif based on [PDB:1HQC] in C_{α} RMSD alignment to the matched substructure from [PDB:1ZTF]; 8 matched residues. **(b)** LabelHash motif based on [PDB:1ECJ] in C_{α} RMSD alignment to the matched substructure from [PDB:1H3D]; 7 matched residues. **(c)** LabelHash motif based on [PDB:1AYL] in C_{α} RMSD alignment to the matched substructure from [PDB:1P9W]; 11 matched residues.




Figure 6 - SOIPPA-derived motif

LabelHash motif constructed by combining multiple pairwise SOIPPA alignments in order to identify a common set of residues shared amongst all SOIPPA-aligned structures. Each of the four substructures used to construct the LabelHash motif are shown in red, green, purple, and fuchsia; the identified common set of 5 residues is shown for each substructure. The ligands neighboring each substructure are shown for reference and colored by atom type.




Figure 7 - Matches to adenine-containing ligand binding sites

Several high-confidence matches identified by the [PDB:1ZQ9]-based motif were identified and subsequently confirmed by the presence of an adenine-containing ligand in the matched site. The matched residues are colored green with the bound ligand colored by atom type. From top to bottom the matches shown above correspond to [PDB:2YQU], [PDB:1W1U], and [PDB:2E58].





Figure 8 - Performance

Average wallclock time (**left**) and speedup (**right**) for matching the different categories of motifs against the entire nrPDB₉₅. The solid lines indicate performance using a shared Panasas filesystem, while the dashed lines indicate performance with solid-state drives.



Tables

Table 1 - SOIPPA-derived Motifs

Template Structure	SOIPPA-derived LabelHash Motif	SOIPPA-aligned Structure(s)	<i>p</i> -value of LabelHash-computed match
1HQC	9 ^E , 10 ^{YI} , 11 ^{IF} , 12 ^{GE} , 13 ^Q , 169 ^{LV} , 171 ^{QY} , 172 ^{GA}	1ZTF	< 0.0009
1ECJ	367 ^D , 369 ^{IV} , 371 ^{RT} , 372 ^G , 373 ^{TA} , 374 ^T , 375 ^{SL}	1H3D	< 0.0005
1AYL	232 ^H , 250 ST , 251 ^G , 252 ^{TS} , 253 ^{GA} , 254 ^K , 255 ^{TS} , 256 ^T , 257 ^{LT} , 268 ^{DG} , 269 ^{DE}	1P9W	< 0.0006
1ZQ9	29 ^G , 50 ^{ET} , 51 ^{LRKS} , 52 ^{DTEQ} , 79 ^{VLA}	1CYD, 1D4D, 1UWK	< 0.0005

Additional Files

Additional File 1 - Enolase Superfamily (ES) motifs based on different structures

The ES motif was defined using different PDB structures as templates and although the amino acid labels are identical across all motifs, the 3D coordinates of each motif point vary according to the structure on which a motif was based [61]. This causes significant variation in the ability of the motifs to accurately classify ESdb structures.

Additional File 2 - Matching results for CSA motifs

Several motifs obtained from the Catalytic Site Atlas (CSA) were matched against the corresponding protein family and the nrPDB₉₅. From the resulting matches the Area Under the Curve (AUC) for the Receiver Operating Characteristic (ROC) curve and Precision-Recall (PR) curve were computed. A ROC curve plots sensitivity ($\frac{TP}{TP+FN}$) as a function of the false positive rate ($\frac{FP}{FP+TN}$), while a PR curve plots precision ($\frac{TP}{TP+FP}$) as a function of recall (which is equivalent to sensitivity). The sensitivity and specificity at a p -value threshold of 0.001 were also computed.

Additional File 3 - LabelHash parallel performance

Average wallclock time (**left**) and speedup (**right**) for matching different motifs against the nrPDB₉₅ on a high-performance computing cluster. For the motifs used in the SOIPPA comparison and the Enolase Superfamily study substantial (but diminishing) improvements can be obtained by increasing the number of cores. These motifs are larger and have more alternate labels than the CSA.

Additional File 1 for:
M. Moll, D.H. Bryant, L.E. Kavraki, The LabelHash Algorithm for Substructure Matching, *BMC Bioinformatics*, 2010.

Enolase Superfamily (ES) motifs based on different structures

The ES motif was defined using different PDB structures as templates and although the amino acid labels are identical across all motifs, the 3D coordinates of each motif point vary according to the structure on which a motif was based [1]. This causes significant variation in the ability of the motifs to accurately classify ESdb structures. Below are the enlarged results for the 2MNR-based motif and the following page contains the corresponding results for the remaining structures.

[1] Meng EC, Polacco BJ, Babbitt PC: **Superfamily active site templates.** *Proteins* 2004, **55**(4):962–976.




Figure 1: Dark gray denotes matches to structures in the ESdb as defined by Meng et al. [1] while light gray denotes matches to structures in the nrPDB₉₅. The dashed line in each plot corresponds to the p -value threshold at $\alpha = 0.01$. The p -value threshold is relative to the distribution of side-chain RMSD in (a) and (b) while relative to C α RMSD in (c) and (d). (a): Selecting the “best” match for each individual structure as the match with minimum side-chain RMSD results in a set of ESdb matches that are largely distinguished from the bulk of nrPDB₉₅ matches. (b): The clear separation of the distributions of ESdb and nrPDB₉₅ matches identified using minimum side-chain RMSD illustrates the high-specificity of the 2MNR-based motif. (c): Selecting minimum C α RMSD matches to the ES motif as the “best” matches results in decreased sensitivity at the same $\alpha = 0.01$ threshold due to the high C α RMSD among ES member substructures. (d): Examining the distributions of matches based upon C α RMSD demonstrates the inseparability of enolase matches from the majority of nrPDB₉₅ structures if minimum C α RMSD is used alone for match selection.




Figure 2: The ES motif based on 1EBH has difficulty separating more distant members of the ES from the bulk of matches to unrelated nrPDB₉₅ structures, while the 1MUC more accurately distinguishes distant members of the ES (in terms of structural deviation), indicating that 1MUC serves as a better consensus structure for the ESdb.

Additional File 2 for:

M. Moll, D.H. Bryant, L.E. Kavraki, The LabelHash Algorithm for Substructure Matching, *BMC Bioinformatics*, 2010.

Matching results for CSA motifs

Several motifs obtained from the Catalytic Site Atlas (CSA) were matched against the corresponding protein family and the nrPDB₉₅. From the resulting matches the Area Under the Curve (AUC) for the Receiver Operating Characteristic (ROC) curve and Precision-Recall (PR) curve were computed. A ROC curve plots sensitivity ($\frac{TP}{TP+FN}$) as a function of the false positive rate ($\frac{FP}{FP+TN}$), while a PR curve plots precision ($\frac{TP}{TP+FP}$) as a function of recall (which is equivalent to sensitivity). The sensitivity and specificity at a p -value threshold of 0.001 were also computed.

Protein family size For both the positive examples (chains belonging to the given EC class) and negative examples (chains within the nrPDB₉₅ but not within the EC class), the total number of

examples in each case is given by the number of *possible* matches according to label existence criteria. This number is shown in parentheses in the first column. A possible match is one where, within the entirety of all residues constituting a single chain, one or more sets of atoms that share label compatibility with the motif exist. Chains from both the positive and negative sets which do not contain a possible match are excluded from the set of structures matched against *a priori*; this set of excluded structures will vary based upon the given motif and alternate amino acid labels present. For example, if a motif consisting of 5 tryptophan residues is matched against a chain that contains only 4 tryptophan residues, then no match, regardless of the LRMSD value, is possible and thus this chain is excluded. However, any chain that contains 5 or more tryptophan residues would contain a possible match in this case.

Domain of the ROC and PR curves LabelHash limits the maximum LRMSD of matches returned during matching to a user specified limit (the ϵ distance cutoff). The matches with an LRMSD beyond the distance cutoff are, however, accounted for statistically by using a pointweight at $+\infty$ as discussed in the text. Therefore, the domain of the ROC and PR curves is $[0, 1 - \textit{pointweight})$ and the curves are only plotted within this domain.

Calculation of AUC values In order to calculate AUC_{ROC} and AUC_{PR} values for each motif over the full interval $[0, 1]$, the integral is interpolated linearly between $[1 - \textit{pointweight}, 1]$. Therefore, for a motif with $AUC_{ROC} = 1.0$ the corresponding AUC_{PR} may be calculated to be less than 1.0 due to the linear interpolation of the AUC values within the undefined point weight region $[1 - \textit{pointweight}, 1]$.

Matching results for CSA motifs

Protein family (#)	Motif	AUC_{ROC}	AUC_{PR}	Sensitivity (#TP)	Specificity (#FP)
1.1.1.35 (50)	[PDB:2HDH] 137 ^S , 158 ^H , 170 ^E , 208 ^N	0.75	0.31	38.0% (19)	99.9% (28)
1.1.1.44 (18)	[PDB:2PGD] 130 ^G , 183 ^K , 187 ^N , 190 ^E	1.00	0.33	100.0% (18)	99.9% (27)
1.1.2.3 (22)	[PDB:1FCB] 143 ^Y , 254 ^Y , 282 ^D , 373 ^H , 376 ^R	0.99	0.25	90.9% (20)	99.8% (26)
1.1.3.9 (12)	[PDB:1GOG] 228 ^C , 272 ^Y , 290 ^W , 495 ^Y	0.85	0.14	66.7% (8)	99.8% (23)

Matching results for CSA motifs

Protein family (#)	Motif	AUC _{ROC}	AUC _{PR}	Sensitivity (#TP)	Specificity (#FP)
1.11.1.10 (41)	[PDB:1A8S] 95 ^W , 224 ^S , 28 ^T , 94 ^D , 253 ^H	0.10	0.00	0.0% (0)	99.8% (27)
1.11.1.10 (41)	[PDB:1QI9] 486 ^K , 341 ^H , 418 ^H , 411 ^H	0.68	0.06	19.5% (8)	99.9% (18)
1.14.13.39 (284)	[PDB:3NOS] 356 ^C , 187 ^R , 184 ^W , 361 ^E	0.99	0.96	96.1% (273)	99.9% (11)
1.14.99.3 (113)	[PDB:1DVE] 25 ^H , 143 ^Y , 140 ^T , 136 ^R , 135 ^G , 139 ^D , 58 ^G	0.84	0.55	54.9% (62)	99.9% (26)
1.17.4.1 (112)	[PDB:3R1R] 437 ^C , 441 ^N , 439 ^C , 225 ^E , 462 ^C	0.74	0.26	42.9% (48)	99.9% (13)
1.17.4.1 (133)	[PDB:1QFN] 8 ^R , 13 ^G , 72 ^Y , 18 ^K , 10 ^Y	0.66	0.02	0.8% (1)	99.8% (32)
1.18.1.2 (102)	[PDB:1E6E] 55 ^H , 159 ^D , 377 ^I , 376 ^T	0.73	0.04	3.9% (4)	99.9% (24)
1.2.1.2 (27)	[PDB:1KQF] 800 ^W , 803 ^S , 805 ^W , 806 ^T , 807 ^G , 808 ^N , 809 ^D , 810 ^I	0.50	0.06	7.4% (2)	99.8% (21)
1.2.1.2 (27)	[PDB:2NAC] 313 ^N , 146 ^R , 284 ^Q , 332 ^H	0.06	0.00	0.0% (0)	99.8% (28)
1.2.1.3 (159)	[PDB:1O04] 192 ^K , 268 ^E , 302 ^S , 399 ^E	0.66	0.09	10.1% (16)	99.8% (31)
1.3.3.1 (50)	[PDB:1D3G] 149 ^F , 215 ^S , 218 ^T , 255 ^K	0.87	0.19	24.0% (12)	99.9% (17)
1.3.99.1 (97)	[PDB:1D4C] 364 ^H , 401 ^R , 503 ^H , 544 ^R	0.66	0.15	24.7% (24)	99.9% (24)
1.4.1.1 (40)	[PDB:1PJB] 117 ^K , 269 ^H , 74 ^E , 95 ^D	0.85	0.02	5.0% (2)	99.8% (35)
1.5.1.3 (205)	[PDB:1VIE] 32 ^K , 67 ^Q , 68 ^I , 69 ^Y	0.50	0.02	2.9% (6)	99.8% (35)
1.5.1.3 (207)	[PDB:1RA2] 20 ^I , 27 ^M , 28 ^D , 31 ^L , 5 ^F , 54 ^L , 94 ^I	0.70	0.25	25.6% (53)	99.8% (27)
1.6.4.2 (25)	[PDB:1GET] 177 ^C , 181 ^C , 42 ^K , 47 ^Y , 50 ^E	0.95	0.30	88.0% (22)	99.8% (25)
1.8.1.2 (16)	[PDB:1AOP] 153 ^R , 215 ^R , 217 ^K , 483 ^K , 83 ^C	0.90	0.19	75.0% (12)	99.8% (25)
1.9.3.1 (229)	[PDB:1EHK] 233 ^H , 237 ^Y , 384 ^H , 385 ^F , 386 ^H , 449 ^R , 450 ^R	0.59	0.22	20.1% (46)	99.9% (19)
2.1.1.45 (218)	[PDB:1LCB] 198 ^E , 219 ^C , 221 ^S , 257 ^D , 259 ^D , 60 ^H	0.98	0.90	88.1% (192)	99.8% (24)
2.1.1.45 (218)	[PDB:1TYS] 146 ^E , 166 ^Y , 169 ^S , 58 ^R , 94 ^D	0.85	0.05	1.8% (4)	99.8% (36)
2.1.1.63 (14)	[PDB:1EH6] 137 ^N , 145 ^C , 146 ^H , 172 ^E	0.82	0.57	71.4% (10)	99.9% (22)
2.1.2.2 (48)	[PDB:1C2T] 106 ^N , 108 ^H , 135 ^S , 144 ^D	0.88	0.17	20.8% (10)	99.8% (37)
2.1.2.2 (48)	[PDB:1GRC] 106 ^N , 108 ^H , 135 ^S , 144 ^D	0.79	0.03	14.6% (7)	99.8% (32)
2.1.3.2 (235)	[PDB:1AT1] 105 ^R , 134 ^T , 54 ^R , 55 ^H	0.82	0.52	53.2% (125)	99.9% (25)
2.1.3.3 (49)	[PDB:1AKM] 106 ^R , 133 ^H , 136 ^Q , 231 ^D , 273 ^C , 319 ^R	1.00	0.89	93.9% (46)	99.9% (19)
2.3.1.41 (173)	[PDB:1DD8] 163 ^C , 298 ^H , 328 ^K , 333 ^H , 390 ^F , 392 ^F	0.71	0.42	42.8% (74)	99.9% (18)
2.3.1.41 (173)	[PDB:1KAS] 163 ^C , 303 ^H , 340 ^H , 400 ^F	0.78	0.52	52.6% (91)	99.9% (17)
2.3.1.74 (41)	[PDB:1CGK] 164 ^C , 215 ^F , 303 ^H , 336 ^N	0.65	0.08	29.3% (12)	99.8% (26)
2.3.2.13 (43)	[PDB:1IU4] 255 ^C , 269 ^D , 272 ^K , 274 ^W , 64 ^H	0.62	0.03	9.3% (4)	99.8% (22)
2.4.1.1 (189)	[PDB:1GPA] 568 ^K , 569 ^R , 574 ^K , 676 ^T	0.81	0.18	33.9% (64)	99.8% (31)
2.4.1.19 (61)	[PDB:1CDG] 227 ^R , 229 ^D , 257 ^E , 327 ^H , 328 ^D	0.97	0.71	80.3% (49)	99.8% (34)

Matching results for CSA motifs

Protein family (#)	Motif	AUC _{ROC}	AUC _{PR}	Sensitivity (#TP)	Specificity (#FP)
2.4.2.6 (29)	[PDB:1F8X] 207 ^Y , 272 ^D , 292 ^D , 298 ^E	0.82	0.42	55.2% (16)	99.9% (27)
2.4.2.8 (76)	[PDB:1BZY] 133 ^E , 134 ^D , 137 ^D , 165 ^K , 169 ^R	0.91	0.45	55.3% (42)	99.8% (33)
2.5.1.19 (45)	[PDB:1G6T] 22 ^K , 313 ^D , 341 ^E , 385 ^H , 411 ^K	0.69	0.09	35.6% (16)	99.8% (32)
2.5.1.31 (38)	[PDB:1F75] 197 ^R , 203 ^R , 33 ^R , 42 ^R	0.12	0.00	0.0% (0)	99.8% (30)
2.5.1.6 (28)	[PDB:1FUG] 14 ^H , 165 ^K , 244 ^R , 245 ^K	1.00	0.88	100.0% (28)	99.9% (26)
2.5.1.7 (45)	[PDB:1UAE] 115 ^N , 23 ^C , 305 ^D , 397 ^R	0.05	0.00	0.0% (0)	99.8% (25)
2.5.1.9 (130)	[PDB:1I8D] 102 ^F , 2 ^S , 41 ^C , 48 ^M , 64 ^H	0.46	0.03	2.3% (3)	99.9% (16)
2.5.1.9 (185)	[PDB:1KZL] 146 ^S , 185 ^C , 41 ^H , 48 ^S , 97 ^D	0.29	0.01	2.2% (4)	99.8% (25)
2.6.1.16 (37)	[PDB:1JXA] 481 ^E , 485 ^K , 488 ^E , 603 ^K	0.72	0.28	35.1% (13)	99.9% (27)
2.6.1.16 (37)	[PDB:1MOQ] 481 ^E , 485 ^K , 488 ^E , 603 ^K	0.73	0.11	35.1% (13)	99.9% (23)
2.7.1.21 (101)	[PDB:1KIM] 163 ^G , 222 ^E , 59 ^R , 83 ^R	0.68	0.06	16.8% (17)	99.8% (37)
2.7.1.40 (119)	[PDB:1PKN] 119 ^R , 269 ^R , 327 ^K , 361 ^T , 363 ^S , 72 ^E	1.00	0.92	91.6% (109)	99.8% (31)
2.7.1.69 (55)	[PDB:1GPR] 66 ^T , 68 ^H , 83 ^H , 85 ^G	0.56	0.04	16.4% (9)	99.9% (17)
2.7.4.3 (73)	[PDB:1ZIO] 127 ^K , 13 ^R , 160 ^R , 162 ^D , 163 ^D , 171 ^R	0.15	0.00	0.0% (0)	99.9% (24)
2.8.1.1 (16)	[PDB:1E0C] 231 ^Q , 232 ^T , 233 ^H , 234 ^H , 235 ^R	0.60	0.14	18.8% (3)	99.8% (25)
2.8.1.1 (16)	[PDB:1RHS] 186 ^R , 248 ^R , 249 ^K , 250 ^G , 251 ^V , 252 ^T	0.75	0.31	43.8% (7)	99.8% (38)
3.1.1.1 (115)	[PDB:1AUO] 114 ^L , 115 ^S , 168 ^Q , 199 ^D , 23 ^H	0.64	0.07	8.7% (10)	99.8% (31)
3.1.1.1 (115)	[PDB:1CI8] 181 ^S , 348 ^K , 351 ^Y , 75 ^W , 78 ^V	0.55	0.03	3.5% (4)	99.8% (27)
3.1.1.1 (115)	[PDB:1L7Q] 117 ^Y , 118 ^A , 166 ^Y , 259 ^W , 287 ^D , 44 ^H	0.73	0.02	0.9% (1)	99.8% (26)
3.1.1.3 (164)	[PDB:1R4Z] 12 ^I , 133 ^S , 156 ^M , 77 ^D , 78 ^H	0.61	0.10	20.7% (34)	99.8% (24)
3.1.1.3 (192)	[PDB:2LIP] 17 ^L , 264 ^S , 286 ^Q , 87 ^D , 88 ^H	0.79	0.19	16.1% (31)	99.8% (32)
3.1.1.4 (347)	[PDB:1CJY] 197 ^G , 198 ^G , 228 ^S , 549 ^D	0.45	0.01	0.6% (2)	99.8% (32)
3.1.1.47 (31)	[PDB:1BWP] 104 ^S , 192 ^G , 195 ^N , 47 ^D , 74 ^H	0.07	0.00	0.0% (0)	99.8% (33)
3.1.11.2 (23)	[PDB:1AKO] 151 ^N , 153 ^D , 229 ^N , 259 ^D , 7 ^H	0.29	0.01	13.0% (3)	99.9% (26)
3.1.2.15 (51)	[PDB:1CMX] 484 ^Q , 490 ^C , 566 ^H , 581 ^D	0.77	0.07	7.8% (4)	99.8% (25)
3.1.2.15 (51)	[PDB:1NBF] 218 ^N , 223 ^C , 464 ^H , 481 ^D	0.74	0.03	9.8% (5)	99.8% (23)
3.1.2.15 (51)	[PDB:1UCH] 169 ^Q , 184 ^C , 89 ^H , 95 ^D	0.77	0.04	5.9% (3)	99.8% (23)
3.1.21.1 (52)	[PDB:1FR2] 100 ^R , 102 ^E , 103 ^H , 5 ^H	0.69	0.07	26.9% (14)	99.9% (21)
3.1.27.5 (315)	[PDB:1RBN] 119 ^H , 12 ^K , 120 ^H , 41 ^F	0.90	0.74	72.1% (227)	99.9% (17)
3.1.3.16 (107)	[PDB:1S95] 274 ^D , 275 ^R , 303 ^N , 304 ^H , 400 ^R , 427 ^H	0.69	0.16	35.5% (38)	99.9% (21)
3.1.3.2 (131)	[PDB:1D2T] 150 ^H , 183 ^R , 189 ^H , 193 ^D	0.62	0.13	13.7% (18)	99.8% (25)

Matching results for CSA motifs

Protein family (#)	Motif	AUC _{ROC}	AUC _{PR}	Sensitivity (#TP)	Specificity (#FP)
3.1.3.2 (131)	[PDB:1RPT] 11 ^R , 12 ^H , 15 ^R , 257 ^R , 258 ^H , 79 ^D	0.25	0.00	0.0% (0)	99.8% (24)
3.1.3.2 (65)	[PDB:1PNT] 12 ^C , 129 ^C , 17 ^R , 18 ^D	0.73	0.10	12.3% (8)	99.9% (17)
3.1.3.3 (25)	[PDB:1L7N] 100 ^D , 11 ^F , 12 ^D , 13 ^G , 144 ^K , 171 ^D	0.79	0.16	20.0% (5)	99.8% (33)
3.1.3.48 (328)	[PDB:1BZC] 181 ^D , 215 ^C , 221 ^R , 222 ^S	0.77	0.18	27.4% (90)	99.9% (20)
3.1.3.48 (328)	[PDB:1PA9] 194 ^D , 241 ^C , 247 ^R , 248 ^T	0.76	0.08	8.5% (28)	99.9% (20)
3.1.3.48 (328)	[PDB:1YTW] 290 ^E , 356 ^D , 402 ^H , 403 ^C , 409 ^R , 410 ^T	0.85	0.46	16.8% (55)	99.9% (13)
3.1.3.48 (346)	[PDB:1D1Q] 13 ^A , 132 ^R , 19 ^S , 20 ^D	0.66	0.04	2.0% (7)	99.8% (31)
3.1.3.5 (60)	[PDB:1USH] 116 ^N , 117 ^H , 120 ^D , 375 ^R , 379 ^R , 410 ^R	0.50	0.01	5.0% (3)	99.8% (27)
3.1.4.10 (11)	[PDB:1PTD] 274 ^H , 32 ^R , 69 ^H , 82 ^D	0.79	0.40	54.5% (6)	99.9% (23)
3.1.4.10 (11)	[PDB:2PLC] 278 ^H , 45 ^D , 46 ^R , 84 ^H , 93 ^D	0.88	0.45	63.6% (7)	99.8% (26)
3.2.1.113 (19)	[PDB:1DL2] 132 ^E , 136 ^R , 275 ^D , 435 ^E	1.00	0.80	100.0% (19)	99.9% (27)
3.2.1.135 (38)	[PDB:1BVZ] 325 ^D , 354 ^E , 356 ^W , 421 ^D	0.86	0.28	36.8% (14)	99.8% (28)
3.2.1.17 (1276)	[PDB:1OBA] 10 ^D , 182 ^D , 92 ^E , 94 ^D	0.30	0.04	0.5% (6)	99.8% (33)
3.2.1.18 (270)	[PDB:7NN9] 151 ^D , 220 ^R , 277 ^E , 371 ^R , 412 ^D	0.85	0.14	8.1% (22)	99.8% (29)
3.2.1.35 (14)	[PDB:1FCQ] 111 ^D , 113 ^E , 184 ^Y , 227 ^Y , 301 ^W	0.87	0.08	42.9% (6)	99.8% (28)
3.2.1.73 (24)	[PDB:1AQ0] 232 ^E , 280 ^E , 283 ^K , 288 ^E	0.92	0.02	12.5% (3)	99.8% (37)
3.2.1.78 (24)	[PDB:1BQC] 127 ^N , 128 ^E , 196 ^H , 198 ^Y , 225 ^E	0.10	0.00	0.0% (0)	99.8% (29)
3.2.1.8 (226)	[PDB:1XYZ] 645 ^E , 723 ^H , 754 ^E , 756 ^D	0.57	0.19	38.5% (87)	99.9% (17)
3.2.1.91 (71)	[PDB:1CEL] 212 ^E , 214 ^D , 217 ^E , 228 ^H	0.74	0.38	43.7% (31)	99.9% (28)
3.2.1.91 (71)	[PDB:1QK2] 169 ^Y , 174 ^R , 175 ^D , 221 ^D	0.85	0.07	21.1% (15)	99.8% (30)
3.2.1.8 (226)	[PDB:1EXP] 127 ^E , 205 ^H , 233 ^E , 235 ^D	0.57	0.19	38.5% (87)	99.9% (19)
3.2.2.1 (32)	[PDB:1R4F] 10 ^D , 186 ^W , 260 ^N , 83 ^A	0.93	0.02	3.1% (1)	99.8% (25)
3.2.2.23 (17)	[PDB:1K82] 1 ^P , 2 ^E , 258 ^K , 56 ^R	0.95	0.09	41.2% (7)	99.8% (34)
3.2.3.1 (19)	[PDB:1MYR] 187 ^R , 190 ^Q , 328 ^S , 330 ^N , 409 ^Y , 95 ^E	0.99	0.22	89.5% (17)	99.8% (36)
3.3.1.1 (71)	[PDB:1B3R] 130 ^H , 185 ^D , 189 ^K , 190 ^D , 194 ^N , 300 ^C , 54 ^H	1.00	0.91	84.5% (60)	99.8% (20)
3.3.2.3 (24)	[PDB:1EHY] 107 ^W , 108 ^D , 152 ^F , 215 ^Y , 246 ^Y , 275 ^D , 38 ^H	0.99	0.23	25.0% (6)	99.8% (24)
3.4.15.1 (16)	[PDB:1O8A] 353 ^H , 354 ^A , 384 ^E , 513 ^H , 523 ^Y	0.97	0.75	93.8% (15)	99.8% (32)
3.4.19.3 (22)	[PDB:1AUG] 291 ^E , 301 ^R , 354 ^C , 378 ^H	0.94	0.77	86.4% (19)	99.9% (17)
3.4.21.12 (39)	[PDB:1SSX] 102 ^H , 193 ^D , 195 ^G , 214 ^S , 57 ^S	1.00	0.48	100.0% (39)	99.8% (43)
3.4.21.7 (29)	[PDB:1DDJ] 603 ^H , 646 ^D , 737 ^C , 738 ^Q , 739 ^G , 740 ^D , 741 ^A	0.57	0.19	27.6% (8)	99.8% (27)
3.4.21.81 (52)	[PDB:1DS2] 102 ^H , 193 ^D , 195 ^G , 57 ^S	0.55	0.20	50.0% (26)	99.8% (39)

Matching results for CSA motifs

Protein family (#)	Motif	AUC _{ROC}	AUC _{PR}	Sensitivity (#TP)	Specificity (#FP)
3.4.21.92 (180)	[PDB:1TYF] 122 ^G , 171 ^S , 68 ^M , 97 ^H , 98 ^D	0.92	0.80	81.7% (147)	99.9% (15)
3.4.21.97 (29)	[PDB:1IEC] 134 ^H , 157 ^S , 165 ^A , 166 ^R , 63 ^R	0.95	0.08	6.9% (2)	99.8% (35)
3.4.21.97 (29)	[PDB:1NKK] 132 ^H , 157 ^S , 165 ^H , 63 ^R	0.91	0.57	69.0% (20)	99.9% (23)
3.4.22.17 (24)	[PDB:1KFU] 105 ^Q , 262 ^C , 286 ^H , 99 ^N	0.31	0.01	8.3% (2)	99.9% (21)
3.4.22.17 (24)	[PDB:1KFX] 105 ^Q , 262 ^C , 286 ^H , 288 ^N , 99 ^W	0.43	0.01	8.3% (2)	99.8% (22)
3.4.22.27 (37)	[PDB:1GLO] 164 ^Q , 184 ^S , 19 ^H , 25 ^N	0.67	0.03	2.7% (1)	99.8% (31)
3.4.22.28 (20)	[PDB:2BHG] 163 ^H , 182 ^D , 46 ^A , 84 ^S	0.45	0.02	20.0% (4)	99.9% (28)
3.4.22.36 (55)	[PDB:2FQQ] 237 ^H , 238 ^G , 285 ^A , 286 ^R	0.63	0.07	9.1% (5)	99.8% (32)
3.4.23.20 (14)	[PDB:1APT] 213 ^D , 216 ^S , 33 ^D , 36 ^T	1.00	0.73	100.0% (14)	99.8% (36)
3.4.23.22 (22)	[PDB:1EED] 215 ^D , 218 ^S , 32 ^D , 35 ^T	1.00	0.75	100.0% (22)	99.8% (38)
3.4.23.25 (12)	[PDB:2JXR] 215 ^D , 218 ^T , 32 ^D , 33 ^T	1.00	0.17	100.0% (12)	99.8% (31)
3.4.23.6 (20)	[PDB:2APR] 218 ^D , 221 ^S , 35 ^D , 38 ^T	1.00	0.76	100.0% (20)	99.8% (35)
3.5.1.26 (34)	[PDB:1APY] 183 ^T , 201 ^T , 234 ^T , 235 ^G	0.65	0.15	20.6% (7)	99.9% (26)
3.5.1.31 (42)	[PDB:1BS4] 133 ^G , 45 ^Q , 50 ^L , 91 ^E	0.99	0.79	95.2% (40)	99.8% (36)
3.5.1.5 (106)	[PDB:1KRA] 219 ^H , 221 ^D , 320 ^H , 336 ^R	0.57	0.08	9.4% (10)	99.9% (17)
3.5.2.6 (506)	[PDB:1BTL] 130 ^S , 166 ^K , 70 ^S , 73 ^E	0.78	0.35	28.1% (142)	99.9% (22)
3.6.1.34 (90)	[PDB:1EF0] 1 ^A , 433 ^N , 434 ^T , 454 ^H , 455 ^G , 76 ^I , 78 ^A , 79 ^C	0.56	0.01	2.2% (2)	99.9% (20)
4.1.1.1 (38)	[PDB:1PVD] 114 ^D , 115 ^H , 28 ^H , 477 ^E	0.95	0.42	50.0% (19)	99.8% (27)
4.1.1.23 (137)	[PDB:1EIX] 44 ^K , 71 ^D , 73 ^K , 76 ^D	0.90	0.77	71.5% (98)	99.9% (22)
4.1.1.39 (318)	[PDB:1RBA] 166 ^K , 191 ^K , 192 ^N , 287 ^H , 321 ^H	0.78	0.15	3.1% (10)	99.8% (29)
4.1.3.7 (41)	[PDB:1AL6] 244 ^S , 274 ^H , 320 ^H , 375 ^D	0.94	0.21	56.1% (23)	99.8% (28)
4.2.1.10 (174)	[PDB:1UQR] 100 ^N , 107 ^R , 11 ^Y , 18 ^E , 23 ^H , 98 ^R	0.47	0.01	1.1% (2)	99.8% (29)
4.2.1.11 (64)	[PDB:1ELS] 168 ^E , 211 ^E , 373 ^H , 396 ^K	0.92	0.47	51.6% (33)	99.9% (27)
4.2.1.11 (64)	[PDB:5ENL] 168 ^E , 211 ^E , 345 ^K , 373 ^H	0.83	0.19	51.6% (33)	99.9% (26)
4.2.1.17 (57)	[PDB:1DUB] 141 ^A , 144 ^G , 164 ^E , 98 ^E	0.91	0.60	63.2% (36)	99.9% (27)
4.2.1.20 (152)	[PDB:1A50] 167 ^H , 305 ^K , 86 ^K , 87 ^D	0.74	0.09	9.9% (15)	99.8% (32)
4.2.1.92 (19)	[PDB:1U5U] 137 ^T , 193 ^H , 66 ^N , 67 ^Y	0.80	0.01	10.5% (2)	99.8% (30)
4.2.3.4 (34)	[PDB:1UJN] 131 ^K , 210 ^K , 220 ^E , 224 ^R , 228 ^N , 235 ^H	0.98	0.54	64.7% (22)	99.8% (33)
4.2.99.18 (22)	[PDB:1BIX] 171 ^Y , 210 ^D , 283 ^D , 309 ^H	0.87	0.67	81.8% (18)	99.9% (26)
5.1.1.3 (38)	[PDB:1B73] 178 ^D , 7 ^S , 70 ^C , 8 ^C	0.62	0.16	18.4% (7)	99.9% (15)
5.2.1.8 (285)	[PDB:1M9C] 102 ^R , 113 ^F , 122 ^Q , 55 ^N , 60 ^F , 63 ^L	0.67	0.28	49.1% (140)	99.9% (13)

Matching results for CSA motifs

Protein family (#)	Motif	AUC _{ROC}	AUC _{PR}	Sensitivity (#TP)	Specificity (#FP)
5.3.1.1 (271)	[PDB:1HTI] 11 ^N , 13 ^K , 165 ^H , 171 ^E , 95 ^G	0.08	0.01	0.0% (0)	99.8% (28)
5.3.1.1 (271)	[PDB:1TPH] 11 ^N , 13 ^K , 165 ^H , 171 ^E , 95 ^G	0.93	0.56	41.7% (113)	99.9% (22)
5.3.1.5 (155)	[PDB:2XIS] 181 ^H , 183 ^D , 54 ^E , 57 ^K	0.99	0.46	94.8% (147)	99.9% (20)
5.3.1.6 (65)	[PDB:1NN4] 10 ^D , 137 ^H , 66 ^C , 9 ^R	0.48	0.12	15.4% (10)	99.8% (24)
5.3.1.9 (97)	[PDB:1DQR] 210 ^K , 216 ^E , 271 ^G , 272 ^R , 357 ^E , 518 ^K	0.96	0.72	74.2% (72)	99.9% (29)
5.3.3.2 (50)	[PDB:1I9A] 1067 ^C , 1087 ^E , 1116 ^E , 1161 ^W	0.85	0.54	58.0% (29)	99.9% (17)
5.3.4.1 (26)	[PDB:1EEJ] 100 ^C , 101 ^Y , 125 ^C , 98 ^R	0.70	0.04	15.4% (4)	99.9% (16)
5.3.4.1 (26)	[PDB:1MEK] 36 ^C , 37 ^G , 38 ^H , 39 ^C	0.83	0.26	57.7% (15)	99.9% (13)
5.4.2.1 (47)	[PDB:1QHF] 181 ^H , 59 ^R , 8 ^E , 86 ^H	0.96	0.76	85.1% (40)	99.9% (21)
5.4.2.8 (23)	[PDB:1P5D] 109 ^R , 118 ^H , 20 ^K , 247 ^R , 329 ^H	0.70	0.22	30.4% (7)	99.8% (28)
5.4.99.5 (72)	[PDB:3CSM] 157 ^R , 16 ^R , 168 ^K , 246 ^E	0.78	0.07	6.9% (5)	99.8% (34)
5.99.1.2 (51)	[PDB:1A41] 167 ^K , 223 ^R , 265 ^H , 274 ^Y	0.81	0.03	9.8% (5)	99.9% (24)
5.99.1.2 (51)	[PDB:1ECL] 111 ^E , 319 ^D , 365 ^Y , 9 ^H	0.73	0.23	25.5% (13)	99.8% (31)
5.99.1.2 (53)	[PDB:1D6M] 328 ^E , 330 ^K , 7 ^Y , 8 ^R	0.87	0.06	15.1% (8)	99.8% (33)
6.1.1.1 (66)	[PDB:2TS1] 230 ^K , 233 ^R , 82 ^K , 86 ^K	0.79	0.12	12.1% (8)	99.8% (31)
6.1.1.11 (24)	[PDB:1SES] 256 ^R , 258 ^E , 261 ^S , 265 ^D , 271 ^R	0.94	0.08	12.5% (3)	99.9% (29)
6.3.2.1 (34)	[PDB:2A86] 160 ^M , 196 ^H , 197 ^H , 198 ^K , 40 ^S , 44 ^S , 47 ^R	1.00	0.87	94.1% (32)	99.9% (19)
6.3.2.4 (33)	[PDB:2DLN] 15 ^E , 150 ^S , 216 ^Y , 255 ^R , 276 ^G	0.80	0.15	21.2% (7)	99.8% (34)
6.3.3.3 (14)	[PDB:1DAE] 11 ^T , 15 ^K , 37 ^K , 41 ^S	0.98	0.62	85.7% (12)	99.9% (28)

**Additional File 3 for:
M. Moll, D.H. Bryant, L.E. Kavraki, The LabelHash Algorithm for Substructure
Matching, *BMC Bioinformatics*, 2010.**

LabelHash Parallel Performance

Average wallclock time (**left**) and speedup (**right**) for matching different motifs against the nrPDB₉₅ on a high-performance computing cluster. For the motifs used in the SOIPPA comparison and the Enolase Superfamily study substantial (but diminishing) improvements can be obtained by increasing the number of cores. These motifs are larger and have more alternate labels than the CSA motifs.

