

Reprint Series  
5 June 1992, Volume 256, pp. 1443-1445

**SCIENCE**

## **Exhaustive Matching of the Entire Protein Sequence Database**

Gaston H. Gonnet, Mark A. Cohen, and Steven A. Benner\*

# Exhaustive Matching of the Entire Protein Sequence Database

Gaston H. Gonnet, Mark A. Cohen, Steven A. Benner\*

The entire protein sequence database has been exhaustively matched. Definitive mutation matrices and models for scoring gaps were obtained from the matching and used to organize the sequence database as sets of evolutionarily connected components. The methods developed are general and can be used to manage sequence data generated by major genome sequencing projects. The alignments made possible by the exhaustive matching are the starting point for successful *de novo* prediction of the folded structures of proteins, for reconstructing sequences of ancient proteins and metabolisms in ancient organisms, and for obtaining new perspectives in structural biochemistry.

A decade has passed since questions were raised (1) about the general validity of conclusions drawn from alignments of protein sequences (2). Today, virtually every biochemical analysis routinely begins with, contains, or concludes with an alignment of sequences of proteins that are presumed to be homologous (3). Alignments are also the starting point for methods of predicting *de novo* the secondary structure of proteins (4-6), for all knowledge-based structure predictions (7), for estimating the number of different types of protein folds (8), for interpreting data from the human genome project (9), and for resolving phylogenetic issues (3, 10).

Despite the varied applications of sequence alignments, it has proved difficult to construct sequence alignments correctly. This is not because of inadequate theory; an algorithm that achieves the optimal alignment of two homologous protein sequences was provided over 20 years ago by Needleman and Wunsch (11). Rather, the problem arises because there are simply too many sequence data to analyze and because the parameters needed to correctly score mutations, deletions, and insertions are unavailable.

Today, mutations (mismatches) in an alignment are usually scored with a mutation matrix developed by Dayhoff and her co-workers in the 1970s (12). However, this matrix was derived from alignments of an extremely small set of proteins that are very similar in sequence, and is therefore unsuitable for alignments between two proteins whose sequences are sufficiently similar to suggest that they might be homologous, but not similar enough to make homology obvious.

The difficulties in constructing alignment routines are further complicated by the requirement that they handle deletions and insertions. Even random sequences can be aligned if gaps are introduced at no

penalty. Most alignment programs therefore assign penalties to gaps of the form ( $ak + b$ ), where  $k$  is the length of the gap and  $a$  and  $b$  are arbitrarily chosen constants. There is no justification, either theoretical or empirical, for this treatment. Indeed, many of the questionable conclusions drawn from alignments arise because of inappropriately placed gaps. Conversely, correctly placed gaps provide information that is critical to the *de novo* prediction of the folded structure of proteins from sequence data alone (4-6). Such information led to the remarkably accurate predictions of the folded structures of tryptophan synthase and protein kinase before crystallographic information was available (4-6).


The amount of sequence data presently available should make it possible to do an exhaustive matching of the entire sequence database (defined here as the result of an attempted Needleman-Wunsch alignment of every subsequence in the database with every other subsequence), obtain empirical probabilities of mutations between amino acids, determine empirical gap scoring penalties, and use these to obtain high-quality alignments. This has not been done previously because the Needleman-Wunsch algorithm is slow (about one pairwise comparison per second). Because a typical contemporary database (such as MIPS Version 64) contains 8,344,353 ( $n$ ) amino acids, exhaustive matching of all subsequences could involve some  $35 \times 10^9$  pairwise

comparisons (on the order of  $n^2$ ) and more than  $10^6$  years of computer time. Not surprisingly, an exhaustive matching of a modern sequence database with the Needleman-Wunsch algorithm has been thought to be essentially impossible (13).

We report the exhaustive matching of an entire protein sequence database. Neither the Needleman-Wunsch algorithm nor any of the rigor that it implies was sacrificed. The key to matching in a reasonable time lies in the step preceding the application of the Needleman-Wunsch algorithm: a reorganization of the sequence data by indexing on a patricia tree (14) (Fig. 1). In an indexed database, pairs of identical sequences are found instantaneously because they lie together on the tree. Similar sequences lie near each other in the tree. Thus, all pairs of sequences that might be significantly similar can be found in an indexed database by far fewer than  $n^2$  matching operations and aligned with the Needleman-Wunsch algorithm. Thus, our exhaustive matching required only 405 days of CPU time and was obtained in the background (otherwise idle CPU capability) from up to six workstations running in parallel for only 19 weeks.

Classical mutation matrices and gap penalties were used in the first phase of the exhaustive matching. A liberal target score ensured that every match with potentially significant sequence similarity was examined. The  $6.5 \times 10^6$  matched pairs of subsequences that were found in the first phase were then refined by running the Needleman-Wunsch algorithm from the point where each match began in one direction along the sequence alignment to the point where the alignment was optimized (or the sequences exhausted), running the algorithm in the reverse direction to achieve the same goal, and repeating the process until the alignment score was no longer improved. After refining,  $1.7 \times 10^6$  matches remained, each optimally aligned, which were then used to calculate new mutation matrices and a model for scoring gaps. These new scoring parameters were then used to further refine the matches to self-consistency. The parameters provide

**Fig. 1.** Reorganization of sequences to form semi-infinite strings placed in "alphabetical order" (left to right in this diagram) on a patricia tree (13), idealized here for sequences built from just two letters. Reorganization time is almost linear with database size and requires negligible computation. Exhaustive matching is achieved by comparing patricia subtrees from the top. Time is saved because the matching of patricia subtrees is aborted when the score falls below a liberally chosen similarity limit. Because all subsequences in the database are indexed, the fact that two similar protein sequences do not begin identically does not diminish the generality of the search.



Institute for Scientific Computation and Institute for Organic Chemistry, Swiss Federal Institute of Technology, Universitatstr. 16, 8092 Zurich, Switzerland.

\*To whom correspondence should be addressed.