

Identification of protein coding regions by database similarity search

Warren Gish & David J. States

Sequence similarity between a translated nucleotide sequence and a known biological protein can provide strong evidence for the presence of a homologous coding region, even between distantly related genes. The computer program BLASTX performed conceptual translation of a nucleotide query sequence followed by a protein database search in one programmatic step. We characterized the sensitivity of BLASTX recognition to the presence of substitution, insertion and deletion errors in the query sequence and to sequence divergence. Reading frames were reliably identified in the presence of 1% query errors, a rate that is typical for primary sequence data. BLASTX is appropriate for use in moderate and large scale sequencing projects at the earliest opportunity, when the data are most prone to containing errors.

National Center for Biotechnology Information, National Library of Medicine, Building 38A, Room 8N806, 8600 Rockville Pike, Bethesda, Maryland 20894-0001, USA

D.J.S. present address: Institute for Biomedical Computing, Washington University, 700 S. Euclid Avenue, St. Louis, Missouri 63110, USA

Correspondence should be addressed to W.G.

Most primary sequence data is obtained as nucleic acid, while much of the biological interest lies in the encoded protein. Inference of likely protein coding regions is often based on statistical features, such as codon usage¹⁻⁴ and the locations of putative splice site signals⁵, but significant false positive rates are common. In contrast, similarity between a conceptually translated nucleotide sequence and a known protein sequence may be highly significant statistically, which suggested a more discriminating approach to inferring coding potential. We present a software tool, BLASTX, that can be used to probe a nucleotide sequence directly for the presence of protein coding regions by identifying segments that encode significant similarity to members of a protein sequence database, a technique that may also be used to assign putative function.

Molecular sequence determination is a complex process, the course of which may be significantly altered when homologues or related sequences are identified by database search. Search tools should therefore be amenable to use in early stages of a sequencing project, even though the data may be more prone to containing errors. Protein-protein comparison methods are important in nucleotide sequencing projects because distant evolutionary relationships which appear to be merely coincidental at the nucleotide sequence level can be meaningfully discerned at the protein sequence level⁶. The amino acid alphabet is expected to be superior to the nucleic acid when sequences are diverged by as few as 50 amino acid PAMs (point accepted mutations per 100 residues), due to degeneracy in the genetic code and functional constraints on protein structure⁶. For example, the coding sequences for cystic fibrosis transmembrane regulatory (CFTR)

protein⁷ and the human multiple drug resistance protein⁸, two members of the family of ATP-binding proteins, share at best a region of 70% identity in 60 nucleotides (only marginally significant statistically), while the two sequences are highly similar at the protein level ($P < 10^{-13}$).

Early sequence data may not only be more prone to errors, but may also be more abundant. These factors cause increased dependence to be placed on the speed, sensitivity and convenience of the software tools used. The selectivity of a search method, or its ability to rank matches by precise statistical criteria, takes on greater importance as well. When mixed with biologically relevant matches, purely chance matches add to the tedium of data analysis and may obscure the presence of the former. These issues are addressed by the software tool presented here.

The BLASTX program has been successfully employed to identify likely protein coding sequences in thousands of partial cDNA sequences from human brain tissue (expressed sequence tags; ESTs)⁹ and in genomic cosmid clones from *Caenorhabditis elegans*¹⁰. BLASTX allows protein-protein comparisons to be considered when only uncharacterized nucleotide query sequence is available. The program conceptually translates query sequences in all six reading frames (three on each strand) and compares each of these full-length translation products with a comprehensive protein sequence database in a single pass. Non-coding sequences tend to yield alignments of marginal significance at best that are selectively excluded from BLASTX output on the basis of a statistically determined cutoff score^{11,13}. In Monte Carlo simulations, BLASTX is effective at recognizing statistically significant sequence similarities in the presence of 1% data errors, an

```

BLASTX 1.2.7MP [31-Mar-92]
Query= YSCMIS1A S.cerevisiae mitochondrial C-1-Tetrahydrofolate synthase gene
(4359 residues)
Translating both strands of query sequence in all 6 reading frames
Database: PIR 32.0 (complete), March 31, 1992
40,298 sequences; 11,831,134 total residues.

Sequences producing High-scoring Segment Pairs:
      Reading High   Smallest
      Frame Score Poisson
                        P(N)
A28174 C1-tetrahydrofolate synthase precursor, mitocho... +1 5008 0.0 1
A29550 C1-tetrahydrofolate synthase - Yeast (Saccharom... +1 1775 0.0 2
A31903 Methylene-tetrahydrofolate dehydrogenase (NADP+)... +1 800 1.4e-218 2
A35367 *C1-tetrahydrofolate synthase - Rat +1 797 1.0e-214 2
A35942 *Formate--tetrahydrofolate ligase - Clostridium... +1 382 9.2e-108 4
A28185 *Formate--tetrahydrofolate ligase - Clostridium... +1 344 2.2e-105 4
A35667 *TEC1 protein - Yeast (Saccharomyces cerevisiae) +3 581 1.2e-86 1
R5RT19 Ribosomal protein L19 - Rat -1 429 9.0e-59 1
A36554 Ribosomal protein L19 - Mouse -1 428 1.3e-58 1
R5D09E Ribosomal protein L19e - Slime mold (Dictyostel... -1 419 3.4e-57 1
J50662 *Methylene-tetrahydrofolate dehydrogenase (NADP+... +1 187 1.2e-42 3
DEHUMT Methylene-tetrahydrofolate dehydrogenase (NAD+)... +1 203 4.5e-39 2
A33267 *Methylene-tetrahydrofolate dehydrogenase/ methy... +1 199 4.6e-39 2
R5MXE Ribosomal protein E - Methanococcus vannielii [... -1 157 6.4e-16 1
S16540 Ribosomal protein L19eR - Haloarcula marismortui... -1 112 7.9e-09 1
R5HS4 Ribosomal protein HL24 - Haloarcula marismortui... -1 104 1.4e-07 1

( Several alignments deleted )

>R5RT19 Ribosomal protein L19 - Rat Length = 196
Score = 429 (224.4 bits), Expect = 9.0e-59, P = 9.0e-59
Identities = 77/144 (53%), Positives = 109/144 (75%), Frame = -1
Query: 4341 KLIVNGTIVKKSIVTVHSKSRTRAHAQSKREGRHSGYGRKRGKGTREARLPSQVWVIRRLRVL 4162
      KL+K+G I++K+VTVHS++R R ++ +R GRH G GKRRGT +AR+P V W+RR+R+L
Sbjct: 43 KLIKDGLIIRKPVTVHSRARCRCNTLARRKGRHMGIGIKRKGKTANARMPKVTWMRRL 102

Query: 4161 RRLLAKYRDAGKIKDKLHYHVLYKESKNAPFKHKKRALVEHI IQAKADAQREKALNEEAEAR 3982
      RRLI +YR++ KID+H+YH LY KGN FK+KR L+EHI KAD R K L ++AEAR
Sbjct: 103 RRLLRRYRESKKIDRHMYSLSLYKVGKGNVFNKRI LMEHILKADKARKRLLDQAEAR 162

Query: 3981 RLKNAARDRRAQRVAEKRDALLK 3910
      R K + AR RR +R+ K++ +K
Sbjct: 163 RSKTKEARKRREERLQAKKEITK 186

Parameters:
E = 0.100, S = see table
W = 3, T = see table, X = see table
M = PAM120
C = 0 (Standard genetic code)
H = n/a, V = 500, B = 250

Frame Length E S T X
+3 1452 0.076 65 13 19
+2 1452 0.094 67 13 20
+1 1453 0.082 71 13 21
-1 1453 0.097 67 13 20
-2 1452 0.094 68 13 20
-3 1452 0.081 65 13 19

Statistics:
Frame Lambda Lambda/ln2 K H High Score High Score
+3 0.378 0.546 0.212 1.45 58 (31.6 bits) 581 (317.0 bits)
+2 0.363 0.524 0.205 1.21 60 (31.4 bits) 59 (30.9 bits)
NOTE: the cutoff score is greater than the expected high score
+1 0.343 0.495 0.187 0.879 63 (31.2 bits) 508 (2480.2 bits)
-1 0.363 0.523 0.203 1.24 60 (31.4 bits) 429 (224.4 bits)
-2 0.357 0.515 0.195 1.07 61 (31.4 bits) 63 (32.5 bits)
NOTE: the cutoff score is greater than the expected high score
-3 0.378 0.545 0.218 1.41 58 (31.6 bits) 55 (30.0 bits)
NOTE: the cutoff score is greater than the expected high score

# of neighborhood words in query = 114,650
# of exact words scoring below T = 1461
# of word hits against database = 82,603,593
# of failed hit extensions = 70,774,341
# of successful extensions = 42
# of overlapping HSPs discarded = 6
# of HSPs reportable = 36
# of database sequences with at least one HSP = 16
No. of states in DFA: 598 (59 KB) Total size of DFA: 1010 KB (1024 KB)
Time to generate neighborhoods: 1.15u 0.09s 1.24t
No. of processors used: 8
Time to search database: 141.78u 0.24s 142.02t
Total cpu time: 143.61u 2.63s 146.24t

```

Fig. 1 Sample BLASTX output. A yeast MIS1 tetrahydrofolate synthase gene sequence¹⁷ (GenBank accession number J03724) was used as a BLASTX query sequence against release 32 of the PIR database¹⁸. The complete list of one-line descriptions of matching database sequences is shown for a probability cutoff of 0.1. The reading frame of the single highest-scoring segment is indicated alongside its score, followed by the Poisson probability for the most significant cluster of segments (which can conceivably involve segments in other reading frames, but in this example never did). As well as the expected synthase protein alignments, highly significant alignments between several L19 ribosomal proteins²⁰ (data not shown) and an unannotated region of the query sequence were identified. Only a short fragment of the yeast L19 sequence has been published²¹. The presence of TEC1 coding sequence located 5' to the MIS1 gene has been noted¹⁹. Residue coordinate numbers for the query sequence are given with respect to the original nucleotide query sequence. Cpu times were reported in seconds for user (u), system (s), and the total (t) of user + system for one of the 8 processors.

error rate that is typical for raw molecular sequence data. Homologues were identifiable by the program in approximately 17% of the raw, human EST sequences examined⁹.

BLASTX offers distinct advantages over related tools such as BLASTP¹³, by combining the conceptual translation and database search procedures into a single step and by collating the results from searches with all six reading frames into a single report; furthermore, with its knowledge of the matches appearing in all reading frames, BLASTX applies Poisson statistics to combine marginally significant matches in different frames from the same strand to yield higher estimates of statistical significance than would be obtained in separate BLASTP searches.

The converse problem, that of searching with protein query sequence against a translated nucleotide sequence database, is addressed by other programs such as TFASTA¹⁴ and PATMAT^{15,16}. Coercing TFASTA to perform the converse of its designed function would be impractical for even modest sized protein databases, as each protein sequence would require a separate invocation of the program. PATMAT does however provide the flexibility of using nucleic acid sequence as the query in a search, translating the nucleic acid in all six reading frames, and comparing the full-length products against either a standard protein sequence database or a concise database of protein sequence "blocks"¹⁶. Searches against a blocks

Table 1 Comparison of BLASTX and BLASTP database searches to identify ESTs

Query sequence	Maximum HSP score	Number of database sequence matches BLASTX and BLASTP	BLASTX only	BLASTP only
CDC2	40-49	835	1850	37
	50-59	307	75	0
	60-69	282	1	0
	70-79	258	0	0
	80-89	210	0	0
	90-99	124	0	0
	100-109	65	0	0
	110-119	53	0	0
	≥120	17	0	0
	CFTR	40-49	295	132
50-59		231	196	6
60-69		92	4	0
70-79		76	0	1
80-89		56	0	0
90-99		34	0	0
100-109		23	0	0
110-119		19	0	0
≥120		17	0	0

The sensitivity of searches performed using BLASTP and BLASTX was compared in searching release 22.0 of the SWISS-PROT database²⁹. Two sample query sequences were used, human cdc2 kinase (CDC2)²², and the cystic fibrosis transmembrane conductance regulator (CFTR)⁷, a large protein with many distant homologues but few closely related ones. For both genes, BLASTP searches were performed using the protein sequence from SWISS-PROT and BLASTX searches were performed using the corresponding cDNA (mRNA) sequence from the GenBank database. Low compositional complexity and repetitive sequences were eliminated from the query sequences by pre-filtering with XNU (ref. 34). Alignments were scored with the PAM120 matrix^{12,36}. The number of matches for each score interval are reported.

database can be practically performed with lower thresholds of significance than against a comprehensive protein database, thus permitting increased search sensitivity; the smaller size of a blocks database may yield shorter search times, as well. In contrast, the use of Karlin-Altschul statistics¹¹ by BLASTX permits precise significance estimates and thresholds to be established; and for comprehensive database searches, BLASTX is severalfold faster.

Acceptable computational efficiency is required for practical use in modest to large scale sequencing applications. Without the aid of an automated sequencer, an individual researcher may generate ten to twenty runs of raw sequence daily, each 300 or more nucleotides in length and totalling several kilobases; automated sequencing laboratories may exceed this throughput by severalfold. Practical database search tools must be able to maintain pace with the laboratory, using available computational resources and without overburdening the individual with random, uninformative matches that must each be evaluated in subsequent steps. If the sense strand or reading frame has not been clearly established, it may be necessary to compare all six reading frames in these data against a comprehensive database of published sequence to find significant similarities. BLASTX is effective in this capacity due to its speed, sensitivity, selectivity and convenience; additional speed is achieved by employing parallel processing on a common multi-purpose, multi-processor computer platform.

Results and discussion

To search a protein sequence database with a nucleotide query sequence, BLASTX translated the query in all six reading frames, built one neighbourhood table¹³ containing pointers into each of these reading frames, and searched the database in a single pass. An example of the effectiveness of this strategy is shown in Fig. 1, the partial BLASTX output generated using the gene sequence for *Saccharomyces cerevisiae* tetrahydrofolate synthase¹⁷ as a query to search the PIR protein sequence database¹⁸. The enzyme itself and several closely related homologues were seen at the top of the list of matches, along with the previously noted TEC1 protein encoded in flanking sequence¹⁹, which were translated in positive- or plus-strand reading frames. These matches were followed by a set of highly significant ($P < 10^{-56}$) alignments with several L19 ribosomal protein sequences²⁰, but the reading frame involved in these alignments resided on the opposite strand in an unannotated segment of the query sequence. It appeared very likely that the yeast homologue of the L19 protein was cloned and sequenced incidentally. Two fragments of the L19 protein have been identified by direct peptide sequencing²¹, but the segment identified here was more extensive and provided the encoding nucleotide sequence.

The reliability of simultaneous translation and alignment is described in Table 1, for which comparative searches were performed with BLASTP and BLASTX using sequences for human cdc2 kinase (CDC2)²² and the human CFTR⁷ as queries. All of the alignments scoring over 80 with direct protein searches of the protein database (BLASTP) were identified by BLASTX using mRNA queries, and no erroneous alignments were introduced (Table 1, column 3). All 258 of the distantly related CDC2 homologs found by BLASTP (maximum HSP scores

between 70–79, see Methodology) were correctly identified by BLASTX without introducing false positive matches. All but one of the 77 distantly related CFTR homologues found by BLASTP were correctly identified by BLASTX, again without introducing false positive matches. For alignments with maximum HSP scores of less than 70, many discrepancies were seen between the BLASTP and BLASTX searches. These alignments were not statistically significant and the discrepancies can therefore be attributed to random events. The majority of low-scoring matches found by BLASTX but not BLASTP were in reading frames different from the true one; the matches found by BLASTP but not BLASTX were not found by BLASTX either because of its use of a larger neighbourhood threshold T or due to EST fragmentation.

The cost of missing data

Karlin and Altschul derived an analytic expression for the probability of finding an alignment above any set score starting at given positions in two sequences¹¹. The effective number of such starting points in a database search is proportional to the product of the query sequence length and the size of the database. In BLASTX searches, the reading frame and orientation of the query sequence are unknown prior to the search. In querying with all six reading frames, BLASTX searches a space roughly six times the size of a single-frame BLASTP search, thus decreasing the significance of any alignments found by a factor of about 6, or by a score of ≈ 2.6 bits. In practice, the impact of this ambiguity is minimal (Table 1). To ensure a manageable size for the neighbourhood word list (see Methodology) with nucleotide query sequences of potentially great length, BLASTX uses a marginally higher neighbourhood word score threshold, T , as compared to the value used by BLASTP. This results in a slight relative reduction in sensitivity of BLASTX in its detection of lower-scoring alignments and accounts for the BLASTP-BLASTX difference seen with CFTR (Table 1).

Sensitivity to sequence data errors

Surveys suggest that the GenBank sequence database²³ contains errors at an overall rate of approximately 1 in 300 nucleotides, of which about half are insertions or deletions (indels)²⁴. For maximal-scoring segment pair (MSP)-based search algorithms such as BLAST (see Methodology), sequencing errors may alter a statistically significant MSP and prevent its recognition in two ways: indels in one or both sequences may break the MSP into two or more shorter segments that individually score below the threshold of significance; and substitution errors may decrease the score of the MSP to a value below the cutoff and perhaps alter its end-points.

The degradation of an alignment score by individual substitution errors is often small relative to the cutoff score for reporting matches, so little or no significant loss of search sensitivity is expected even at error rates encountered in raw molecular sequence data (< 1 substitution per 100 amino acids)²⁵. In contrast, frameshift sequencing errors have a significant impact on search sensitivity, particularly when the query and target sequences are highly diverged²⁵. Even in the absence of sequencing errors, naturally occurring indel mutations can prevent homologue detection. With increasing divergence, the expected contribution to an alignment score by each pair of aligned residues falls while the

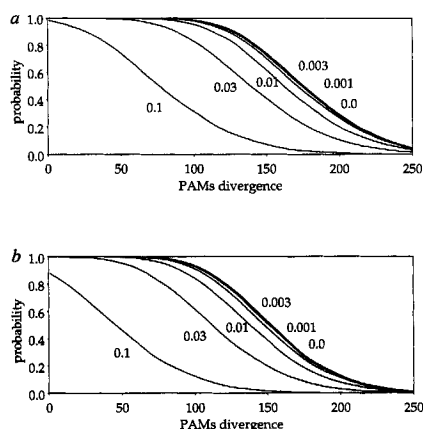


Fig. 2 Probability of finding an MSP in the presence of indel errors and mutations. The probability that a query sequence 300 nucleotides (100 codons) long would retain at least one ungapped alignment of sufficient length to permit its identification in a database search was plotted versus PAM distance between the encoded amino acid sequence and a homologue for alignments of (a) at least 35 bits desired significance and (b) at least 45 bits desired significance. Frameshift errors were introduced at rates of 0, 0.001, 0.003, 0.01, 0.03 and 0.1 per nucleotide, as indicated. Naturally occurring frameshift mutations were modeled at a rate of 5×10^{-5} per nucleotide per (amino acid) PAM; this rate was chosen by extrapolating from results of multiple protein sequence alignments (data not shown; ref. 26). Scoring was performed in all cases with a PAM120 substitution matrix. The value of K (Methodology, Equation 1) in the Karlin-Altschul analysis¹¹ is typically about 0.2, so a score of 45 bits corresponds to a less than 1 in 10,000 chance of finding an alignment of equal or greater significance, when comparing a query of length 750 nucleotides in all six reading frames against a protein sequence database of 10 million residues; and 35 bits corresponds to about a 1 in 10 chance under the same conditions, or a marginal threshold for statistical significance¹².

expected length of a statistically significant HSP rises¹². The increased length provides greater opportunity for indel mutations and frameshift errors to impart their deleterious effects on search sensitivity, a situation which is exacerbated by the tendency of indel mutations to appear more frequently with increasing PAM distance²⁶.

Frameshift sequencing errors are shown in Fig. 2 to have reduced the probability that an MSP was long enough to accumulate a score detectable above random. Under optimal conditions of error-free data at 120 PAMs divergence, an alignment significant to at least 35 bits was detected with a probability of 94%; in the presence of 1% indel errors, the probability was 86%. Overall, sequencing errors at rates as high as 3 indels per 1,000 nucleotides made little difference in the probability of detection; but even with perfect sequencing data, homologues could not be reliably identified at 250 PAMs divergence, due in part to the use of the sub-optimal PAM120 scoring scheme^{6,12}. Neglecting indel mutations and all sequencing errors, if the optimal, PAM250 matrix for scoring homologues diverged by 250 PAMs is used instead, the expected score between true homologues is about 0.36 bits per aligned residue pair¹² or 36 bits over the entire modeled length of 100 codons; when the same homologues are scored with the PAM120 matrix, the expected score drops to 0.17 bits

per aligned residue pair¹², or 17 bits over the entire length. This is 18 bits (or a factor of $2^{18} \approx 10^{5.4}$) below the threshold of 35 bits deemed necessary to achieve marginal significance in a typical database search (Fig. 2 legend), but due to a large variance in the distribution of segment scores (data not shown), the observed probability of detection at 250 PAMs divergence was >4% (Fig. 2a).

The empirical behaviour of BLASTX in the presence of sequencing errors was assessed by comparing the results of BLASTX searches performed with 275 error prone cDNA sequence fragments (ESTs)^{9,10,27,28} to the results obtained by BLASTP with the corresponding finished protein homologue sequences used as queries against the SWISS-PROT database²⁹. EST identification was originally based on searches of the NCBI non-redundant database³⁰; the extended coverage of this database permitted more ESTs to be positively characterized. Figure 3 exhibits the effects upon search performance by sequence errors and EST fragmentation. For approximately one quarter of the ESTs, all of the matches found by BLASTP were correctly identified by BLASTX. At the other end of the spectrum, in twelve of the 275 cases all of the matches found by BLASTP were missed by BLASTX. Manual inspection revealed that in most of these cases, failure to identify a homologue was a result of frameshift errors or truncation of the cDNA sequence in the EST, leaving little or no intact protein coding sequence. For the remaining three quarters of the ESTs, BLASTX correctly identified at least some of the matches found by BLASTP, thus permitting a homologue to be identified without resorting to the use of the latter program.

In addition to the effects imparted by sequence errors and EST fragmentation, some BLASTP matches may have gone unreported by BLASTX because they fell just below the cutoff threshold for BLASTX yet just above the threshold for BLASTP, where the two score thresholds differ by the approximately 2.6 bits explained earlier; but with the expectation cutoff of 0.01 that was used, this threshold effect should have been relatively small. As a caveat to this study, the 275 ESTs were chosen on the basis of previous, positive characterization by BLASTX, which constitutes some degree of selection for ESTs of higher quality — particularly with fewer frameshift errors and more complete coding regions. While BLASTX was responsible for the initial characterization, the results in Fig. 3 may tend to present a best-case distribution.

The results in Figs 2 and 3 show that inaccurate sequence data can degrade the sensitivity of database searches. As more error prone data enters the public databases, it becomes increasingly difficult and costly for other investigators to check their results. When BLASTX reports adjacent high-scoring segment pairs (HSPs) appearing in different reading frames on the same strand of the query, the possibility of frameshift errors — the errors which are of greatest concern — should be investigated and corrected as appropriate. While minimizing the error rate is one concern in a sequencing project, practical considerations often demand that error prone data be released to the public. This provides some of the impetus for maintaining a specialized, public repository of EST sequence data^{27,28}.

Caveats: some causes of misleading results

Several phenomena complicate the statistical analysis of similarity searches. These are independent of the algorithm used to perform the search, but must be considered in the




Fig. 3 Comparison of EST and homologous protein database searches. For each of the 275 non-*Alu* containing cDNA sequences in pre-release 1.0 of the dbEST database^{27,28} for which a homologue was identified by a BLASTX search of the NCBI non-redundant database³⁰, the results of a BLASTX search against release 22.0 of SWISS-PROT²⁹ were compared to a BLASTP search using the highest scoring protein homologue as the query. Both the protein queries and the protein database sequences were pre-filtered with the XNU program (J.-M. Claverie & D.J.S., manuscript submitted) to eliminate improbable residue clusters and short period repetitive sequences. Based on the sizes and compositions of both the database and the queries, only matches with less than a 1% chance of random occurrence¹¹ were counted. For each EST, the Fig. shows the fraction of the BLASTP sequence similarities which were identified by BLASTX, after sorting by this fraction. The abscissa is labelled with both the EST index number and the fraction of the total number of ESTs.

interpretation of BLASTX output. Genomes contain local regions of strongly biased residue composition and reduced information content. Such regions of low compositional complexity sequence are present in the public databases and may be present in a query sequence or its translations, such that a large number of high scoring but biologically uninformative alignments are observed. Inappropriately high statistical significance may be assigned to these alignments because local biases in amino acid composition are not encompassed by the random sequence model assumed in Karlin-Altschul statistics. The biological significance of matches against low complexity regions must, therefore, remain suspect.

At the next level of order above low-complexity sequences, complex but repetitive sequence elements such as the human *Alu* family³¹ are present in genomes and are particularly frequent in higher eukaryotic genomes. It is common for a query sequence to contain a segment derived from such a repetitive element, even if the query is a cDNA. Searches performed with such a query may produce alignments that appear statistically significant wherever members of the same repetitive sequence family are present in the database. Proteins may be partially encoded by repetitive sequence elements, as well. For example, the PIR entries for human complement factor 5 (C5HU), decay accelerating factors 1 and 2 (A26359, B26359), the HLA-DR β 1 precursor (S01441), and platelet glycoprotein IIb (A28411) all appear to contain translated *Alu* sequences³². BLASTX searches with a query sequence containing an *Alu* element may score highly against these database sequences, even if the *Alu* sequence in the query is not translated *in vivo* and no further homology exists outside of the repetitive element.

Repetitive sequence elements may be filtered from a query sequence by searching against a concise database of exceptional sequences that includes the six frame translations of *Alu*³². Segments of the query that match against the exceptions database can then be excluded from the full database search by masking with the IUB nucleic acid code N (for "any"); such masked segments were translated by BLASTX into the IUB amino acid code X (for "unknown") and would not appear in alignments when the default PAM120 matrix was used. Low complexity sequences can be filtered by analysing local residue composition³³ or by identifying stretches of short period, internal repeats in the amino acid translations³⁴ and then masking these segments.

Our results show that BLASTX is a computationally efficient tool for finding gene homologues without prior knowledge of the coding regions or reading frames in a nucleotide query sequence. Pseudogenes are true homologues and may achieve significant alignment scores, as well, but they can be distinguished from functional coding regions by other criteria, such as the presence of stop codons and the absence of promoter elements and splice sites. BLAST identifies many related sequences even if the query sequence is error prone, but at lower sensitivity. The greatest impact of query errors on search performance is expected in comparisons between distantly related proteins, and this effect was confirmed empirically.

BLAST identifies local regions of similarity which are ungapped. Multiple local regions of similarity may contribute to the overall score, but algorithms such as the dynamic programming approach of Smith and Waterman³⁵ may also provide increased sensitivity when insertion and deletion errors are present²⁵, albeit at increased cost. Even with that algorithm, however, gap initialization is often heavily penalized, such that a small number of insertion or deletion errors rapidly degrades the significance of an alignment.

Methodology

Algorithms and implementations. The BLAST algorithm approximates a well defined measure of local sequence similarity based on a matrix of similarity or substitution scores for all possible pairs of residues¹³. The algorithm identifies ungapped, aligned pairs of sequence segments with locally maximum scores which meet or exceed a parameterized cutoff score, *S*. These segments are referred to as "high-scoring segment pairs" (HSPs), and the highest scoring segment pair derivable from any two sequences is their maximal-scoring segment pair (MSP). BLASTX, based on this rapid, probabilistic algorithm, is used to find statistically significant HSPs between a translated nucleotide query sequence and a target protein sequence database. When an HSP is found, the analysis of Karlin and Altschul¹¹ is used to estimate the significance of its score (Equation 1 below).

No prior knowledge of the reading frame or direction is assumed by BLASTX; all possible reading frames in both orientations of the query sequence are translated into protein sequence using the standard genetic code, with eight other genetic codes available to choose from using a simple command line parameter. The PAM (point accepted mutation) amino acid substitution model is typically used for scoring similarity between peptide sequences³⁶, wherein identities and conservative replacements receive positive scores, and non-conservative replacements, for example leucine for aspartic acid, receive negative scores. By default, BLASTX uses a PAM120 matrix scaled to 0.5 bit per unit score¹².

Stop codons are not explicitly included in development of the PAM

theoretical framework³⁶. We score alignments between amino acids and stop codons as equivalent to the least-favorable (most negative) substitution score observed between any two amino acids in the PAM matrix. All substitution scores are readily accessible to the user as rows and columns in the score matrix file read by the program and, as such, were user-modifiable. Alignments incorporating a stop codon can be effectively forbidden by applying a large negative penalty to any substitution for a stop codon, or scores might be chosen to represent the log-odds that a stop codon resulted from experimental error.

The expected number of alignments scoring S or greater in a comparison between two random sequences of lengths m and n is

$$E = mnKe^{-\lambda S} \quad \text{Equation 1}$$

where K and λ are parameters dependent on the amino acid compositions of the sequences¹¹. For values less than about 0.1, E is often an acceptable approximation to P , the probability of occurrence of one or more matches scoring S or greater. In a true coding region, one reading frame may have a predicted amino acid composition typical for biologically occurring proteins, while the other reading frames exhibit anomalous compositions^{13,32}. For this reason, BLASTX calculated separate K and λ values for each reading frame. Alignment scores are often cited here in units of bits (binary digits), such that their significance can be evaluated independently of the scale of the scoring system employed¹². Each unitary increase in the bits of information corresponds to a factor of two increase in the statistical significance.

The BLAST algorithm operates in two successive stages, "neighbourhood" word generation followed by the actual search, with an implicit trade-off in speed versus sensitivity imparted in the first stage¹³. A list of neighbourhood words of length W is generated from consecutive, overlapping words of length W in the query sequence, using a specified scoring matrix. The neighbourhood list contains all words which satisfy a threshold scoring parameter, T , when aligned with words in the query sequence. Raising T decreases the size of the neighbourhood and, consequently, increases the search speed in the algorithm's second stage, but at the expense of decreased sensitivity¹³. In BLASTX, the neighbourhood word list was built from the conceptual translations of all six reading frames on both strands of the query sequence, and this word list was stored in a class of data structures known as a deterministic finite-state automaton (DFA)³⁷. Depending on the scoring matrix and value chosen for T , a given word of length W in the query sequence may yield no neighbourhood words. The fundamental BLAST algorithm was enhanced by including any query word in the DFA, as long as the score of the word when aligned with itself was positive.

During the second stage of the BLAST algorithm, the neighbourhood words from the first stage are searched for in the database or "target" sequence; the presence of a neighbourhood word match indicates the possible location of an HSP. Individual neighbourhood word matches (or word "hits") are extended in both directions along the matrix diagonal until the ends are reached or the cumulative alignment score falls from its maximum achieved value by a parameterized quantity X . In BLASTX and BLASTP, the initial word hits were found by streaming database sequences through the automaton; after hit extension, only those segment pairs whose scores met or exceeded a cutoff score, S , were reported to the user. Rather than choose a value for S explicitly, users often found it more natural to specify a maximum expected frequency of occurrence, E , for HSPs to be reported by the program. From the specified value of E and the relationship shown in equation 1, BLASTX calculated the value for S necessary to achieve the desired level of significance, as a function of the length and amino acid composition of the query sequence in each reading frame, the length of the database, and the particular scoring matrix employed; a fixed set of amino acid frequencies characteristic of general protein databases was also used in these calculations¹³.

The BLAST algorithm is heuristic but has the property that any desired sensitivity in the detection of MSPs may be obtained at the cost of increased computation time. For the particular parameter values used to obtain the BLASTX results presented here (word length $W=3$ with T determined by an *ad hoc* formula and a PAM120 matrix), the predicted sensitivity of BLAST in finding an MSP with a score of 32 bits or greater is very nearly 100%¹³. (In contrast, the program BLASTN for comparing nucleotide sequences uses a longer word length of 12 and builds its DFA from a very restricted set of neighbourhood words — just the query words themselves — thus

achieving high speed at the expense of sensitivity¹³.)

Simultaneous translation and alignment strategies multiply the effective size of the query sequence. Instead of searching a single reading frame translation of the query against the database, BLASTX searches six reading frames, so the DFA used to implement the neighbourhood search in BLASTX is roughly six times the size it would be for one frame only. The search speed of the DFA itself is unaffected by the query length. Even for modest queries (>50 nucleotides or 100 amino acids) the rate limiting step is extension of the word hits found by the DFA. The number of such hits was expected to be proportional to the product of the query and database lengths. In one example, BLASTX searched a 20 million residue protein sequence database in 120 s for a 2,115 nucleotide (4,226 amino acids translated) query sequence (GenBank® 71 locus HUMBARR), or only about 6-fold longer than the 20 s required for BLASTP to search the same database with the encoded 413 amino acid protein sequence, a sequence roughly one-tenth as long. This modest discrepancy in relative execution times is largely attributable to the use by BLASTP of a lower (more sensitive) neighbourhood word score threshold, T , resulting in an increased number of word hits and subsequent word hit extensions.

Parallel processing of BLAST calculations. Multiprocessing capabilities were exploited when either BLASTX or BLASTP executed on Silicon Graphics, Inc. (SGI) PowerIRIS platforms, including the models 4D/240 and 4D/480. These systems provided respectively 4 and 8 processors, for which real-time processing improvements of about 3- and 6-fold were observed. For example, in probing a 703 nucleotide mRNA sequence encoding bovine lactalbumin against release 32.0 of the PIR database with BLASTX, 0.2 s of CPU time in serial mode was required to build the DFA, followed by 21 s real time in 8-processor parallel search mode, and finishing with 3 s of single-processor time to report the results, yielding an overall efficiency of 86%.

Construction of the DFA to recognize neighbourhood words was performed serially on one processor. The search was then run in parallel with individual processors accessing the same DFA (as read only memory) to search dynamically assigned segments of the database. While searching, each processor compiled a local list of HSPs in common heap storage, with only intermittent interprocessor communication required; the efficiency of multiple processor use during this phase approached 100%. After all of the processors had searched their database segments, the program reverted to single processor mode and the individual HSP lists were merged into a single list that was then sorted by statistical significance. No interprocessor communication overhead was incurred for merging the lists because of their placement in common heap storage.

To facilitate repetitive or concurrent searches, BLASTX and BLASTP optionally searched database files loaded semi-permanently in shared memory, thus significantly reducing the overhead of disk I/O and contention for disk-based resources. Shared memory was managed using AT&T UNIX System V interprocess communication (IPC) facilities, including semaphores and message queues to arbitrate access and signal updates to the memory-resident database files.

Analysis of sensitivity to sequence errors. The probability that an MSP is recognizable by the BLAST algorithm using an error-prone query sequence was estimated by sampling using randomly generated peptide sequence data constrained to an amino acid composition typical for that observed in protein sequence databases³⁶. For each query sequence, a target sequence was generated using the transition probabilities underlying the PAM model^{36,12}. Insertion and deletion error sites were then generated randomly in the query sequence and alignments of the remaining continuous segments were scored with a PAM120 matrix. 2,000 trials were performed at each integral value of PAM from one through 250, and a least-squares curve was fit through the resulting data. Calculations were implemented in the S statistical analysis language (AT&T).

Software compatibility and availability. The programs described here were implemented in the C programming language and are compatible with many UNIX-based computing platforms. Complete, public domain source code is available over the Internet via anonymous FTP (file transfer protocol) on ncbi.nlm.nih.gov (130.14.20.1) or by contacting W.G. at gish@ncbi.nlm.nih.gov.

Acknowledgements

We thank Stephen Altschul for informative discussion and David Landsman, Mark Boguski, Jean-Michel Claverie and David Lipman for their critical reading of the manuscript.

Received 28 October; accepted 3 December 1992.

1. Staden, R. & McLachlan, A. D. Codon preference and its use in identifying protein coding regions in long DNA sequence. *Nucl. Acids Res.* **10**, 141–156 (1982).
2. Fickett, J. W. Recognition of protein coding regions in DNA sequences. *Nucl. Acids Res.* **10**, 5303–5318 (1982).
3. Staden, R. Finding protein coding regions in genomic sequences. *Methods Enzymol.* **183**, 163–180 (1990).
4. Uberbacher, E. C. & Mural, R. J. Locating protein-coding regions in human DNA sequences by a multiple sensor-neural network approach. *Proc. natn. Acad. Sci. U.S.A.* **88**, 11261–11265 (1991).
5. Fields, C. A. & Soderlund, C. A. gm: a practical tool for automating DNA sequence analysis. *Comput. Appl. Biosci.* **6**, 263–270 (1990).
6. States, D. J., Gish, W. & Altschul, S. F. Improved sensitivity of nucleic acid database searches using application-specific scoring matrices. *Methods: A Compan. Meth. Enzymol.* **3**, 66–70 (1991).
7. Riordan, J. R. *et al.* Identification of the cystic fibrosis gene: cloning and characterization of complementary DNA. *Science* **245**, 1066–1073 (1989).
8. Chen, C.-J. *et al.* Genomic organization of the human multidrug resistance (MDR1) gene and origin of P-glycoproteins. *J. biol. Chem.* **265**, 506–514 (1990).
9. Adams, M. D. *et al.* Sequence identification of 2,375 human brain genes. *Nature* **355**, 632–634 (1992).
10. Sulston, J. *et al.* The *C. elegans* genome sequencing project: a beginning. *Nature* **356**, 37–41 (1992).
11. Karlin, S. & Altschul, S. F. Methods for assessing the statistical significance of molecular sequence features by using general scoring schemes. *Proc. natn. Acad. Sci. U.S.A.* **87**, 2264–2268 (1990).
12. Altschul, S. F. Amino acid substitution matrices from an information theoretic perspective. *J. molec. Biol.* **219**, 555–565 (1991).
13. Altschul, S. F., Gish, W., Miller, W. Myers, E. W. & Lipman, D. J. Basic local alignment search tool. *J. molec. Biol.* **215**, 403–410 (1990).
14. Pearson, W. R. & Lipman, D. J. Improved tools for biological sequence comparison. *Proc. natn. Acad. Sci. U.S.A.* **85**, 2444–2448 (1988).
15. Henikoff, S., Wallace, J. C. & Brown, J.P. Finding protein similarities with nucleotide sequence databases. *Methods Enzymol.* **183**, 111–132 (1990).
16. Wallace, J. C. & Henikoff, S. PATMAT: a searching and extraction program for sequence, pattern and block queries and databases. *Comp. appl. bio. Sci.* **8**, 249–254 (1992).
17. Shannon, K. W. & Rabinowitz, J. C. Isolation and characterization of the *Saccharomyces cerevisiae* MIS1 gene encoding mitochondrial C1-tetrahydrofolate synthase. *J. biol. Chem.* **263**, 7717–7725 (1988).
18. Barker, W. C., George, D. G. & Hunt, L. T. Protein sequence database. *Meth. Enzymol.* **183**, 31–49 (1990).
19. Laloux, I., Dubois, E., Dewerchin, M. & Jacobs, E. TEC1, a gene involved in the activation of Ty1 and Ty1-mediated gene expression in *Saccharomyces cerevisiae*: cloning and molecular analysis. *Molec. cell. Biol.* **10**, 3541–3550 (1991).
20. Chan, Y. L. *et al.* The primary structure of rat ribosomal protein L7. The presence near the amino terminus of L7 of five tandem repeats of a sequence of 12 amino acids. *J. biol. Chem.* **262**, 1111–1115 (1987).
21. Otaka, E., Higo, K. I. & Itoh, T. Isolation of seventeen proteins and amino-terminal amino acid sequences of eight proteins from cytoplasmic ribosomes of yeast. *Molec. gen. Genet.* **191**, 519–524 (1983).
22. Lee, M. G. & Nurse, P. Complementation used to clone a human homologue of the fission yeast cell cycle control gene *cdc2*. *Nature* **327**, 31–35 (1987).
23. Burks, C. *et al.* GenBank: current status and future directions. *Methods Enzymol.* **183**, 3–22 (1990).
24. Krawetz, S. A. Sequence errors described in GenBank: a means to determine the accuracy of DNA sequence interpretation. *Nucl. Acids Res.* **17**, 3951–3957 (1989).
25. States, D. J. & Botstein, D. Molecular sequence accuracy and the analysis of protein coding regions. *Proc. natn. Acad. Sci. U.S.A.* **88**, 5518–5522 (1991).
26. Gonnet, G. H., Cohen, M. A. & Benner, S. A. Exhaustive matching of the entire protein sequence database. *Science* **256**, 1443–1445 (1992).
27. Boguski, M. S. dbEST, a database of expressed sequence tagged sites. (National Center for Biotechnology Information, National Library of Medicine, Bethesda, MD 20894–0001, Internet electronic mail: boguski@ncbi.nlm.nih.gov, 1992).
28. Update on expressed sequence tag database. *NCBI News* 1(3), 6 (1992).
29. Balroch, A. & Boeckmann, B. The SWISS-PROT protein sequence data bank. *Nucl. Acids Res.* **20**, 2019–2022 (1992).
30. *Entrez* CD ROM Pre-release 6 (National Center for Biotechnology Information, National Library of Medicine, Bethesda, MD, 1992).
31. Rubin, C. M., Houck, C. M., Deininger, P. L., Friedmann, T. & Schmid, C. W. Partial nucleotide sequence of the 300-nucleotide interspersed repeated human DNA sequences. *Nature* **284**, 372–374 (1980).
32. Claverie, J.-M. Identifying coding exons by similarity search: Alu-derived and other potentially misleading protein sequences. *Genomics* **12**, 838–841 (1992).
33. Wootton, J.C. & Federhen, S. Statistics of local complexity in amino acid sequences and sequence databases. *Computers Chem.* (in the press).
34. Claverie, J.-M. & States, D.J. Information enhancement methods for large scale sequence analysis. *Computers Chem.* (in the press).
35. Smith, T. F. & Waterman, M. S. Identification of common molecular subsequences. *J. molec. Biol.* **147**, 195–197 (1981).
36. Dayhoff, M. O., Schwartz, R. M. & Orcutt, B. C. in *Atlas of Protein Sequence and Structure* (ed. Dayhoff, M. O.) 5(3), 345–352 (Natl. Biomed. Res. Found., Washington D.C., 1978).
37. Hopcroft, J. E. & Ullman, J. D. *Introduction to automata theory, languages, and computation*, 42–43 (Addison-Wesley Publishing, Reading, MA, 1979).