Efficient Selection of Unique and Popular Oligos for Large EST Databases

Jie Zheng[1], Timothy J. Close[2], Tao Jiang[1], and Stefano Lonardi[1]

[1] Dept. of Computer Science & Engineering
[2] Department of Botany & Plant Sciences
University of California
Riverside, CA 92521

Abstract. EST databases have grown exponentially in recent years and now represent the largest collection of genetic sequences. An important application of these databases is that they contain information useful for the design of gene-specific oligonucleotides (or simply, oligos) that can be used in PCR primer design, microarray experiments, and genomic library screening. In this paper, we study two complementary problems concerning the selection of short oligos, *e.g.*, 20–50 bases, from a large database of tens of thousands of EST sequences: (i) selection of oligos each of which appears (exactly) in one EST sequence but does not appear (exactly or approximately) in any other EST sequence and (ii) selection of oligos that appear (exactly or approximately) in many ESTs. The first problem is called the *unique oligo* problem and has applications in PCR primer and microarray probe designs. The second is called the *popular oligo* problem and is useful in screening genomic libraries (such as BAC libraries) for gene-rich regions. We present an efficient algorithm to identify all unique oligos in the ESTs and an efficient heuristic algorithm to enumerate the most popular oligos. By taking into account the distribution of the frequencies of the words in the EST database, the algorithms have been carefully engineered to achieve remarkable running times on regular PCs. Each of the algorithms takes only a couple of hours (on a 1.2 GHz CPU, 1 GB RAM machine) to run on a dataset 28 Mbases of barley ESTs from the HarvEST database. We present simulation results on synthetic data and a preliminary analysis of the barley EST database.

1 Introduction

Expressed sequence tags (ESTs) are partial sequences of expressed genes, usually 200–700 bases long, which are generated by sequencing from one or both ends of cDNAs. The information in an EST allows researchers to infer functions of the gene based on similarity to genes of known functions, source of tissue and timing of expression, and genetic map position. EST sequences have become widely accepted as a cost-effective method to gather information about the majority of expressed genes in a number of systems. They can be used to accelerate various research activities, including map-based cloning of genes that control

R. Baeza-Yates et al. (Eds.): CPM 2003, LNCS 2676, pp. 384–401, 2003.
© Springer-Verlag Berlin Heidelberg 2003

traits, comparative genome analysis, protein identification, and numerous methods that rely on gene-specific oligonucleotides (or *oligos*, for short) such as the DNA microarray technology.

Due to their utility, speed with which they may be obtained, and the low cost associated with this technology, many individual scientists and large genome sequencing centers have been generating hundreds of thousands of ESTs for public use. EST databases have been growing exponentially fast since the first few hundreds sequences obtained in the early nineties by Adams *et al.* [1], and now they represent the largest collection of genetic sequences. As of September 2002, the number of sequences deposited in NCBI's dbEST [2] has reached 13 million sequences out of the 18 million sequences which composes the entire GenBank.

With the advent of whole genome sequencing, it may appear that ESTs have lost some of its appeal. However, the genomes of many organisms that are important to society, including the majority of crop plants, have not yet been fully sequenced, and the prospects for large-scale funding to support the sequencing of any but a few in the immediate future is slim to none. In addition, several of our most important crop plants have genomes that are of daunting sizes and present special computational challenges because they are composed mostly of highly repetitive DNA. For example, the *Triticeae* (wheat, barley and rye) genomes, each with a size of about 5×10^9 base pairs per haploid genome (this is about twice the size of maize, 12 times the size of rice, and 35 times the size of the Arabidopsis genomes), are too large for us to seriously consider whole genome sequencing at the present time.

Of the many EST databases, we will be especially interested in the dataset of barley (*Hordeum vulgare*). Barley is premiere model for Triticeae plants due to its diploid genome and a rich legacy of mutant collections, germplasm diversity, mapping populations (see `http://www.css.orst.edu/barley/nabgmp/nabgmp.htm`), and the recent accumulation of other genomics resources such as BAC [3] and cDNA libraries [4,5]. Nearly 300,000 publicly available ESTs derived from barley cDNA libraries are currently present in dbEST. These sequences have been quality-trimmed, cleaned of vector and other contaminating sequences, pre-clustered using the software TGICL (`http://www.tigr.org/tdb/tgi/software/`) and clustered into final assemblies of "contigs" (*i.e.*, overlapping EST sequences) and "singletons" (*i.e.*, non-overlapping EST sequences) using CAP3 [6]. The collection of the singletons and consensus sequences of the contigs, called *unigenes*, form our main dataset. As of July 23, 2002, the collection has 46,145 unigene ESTs of a total of 28,475,017 bases. This dataset can be obtained from `http://harvest.ucr.edu/` using the HARVEST viewer.

In this paper, we study two computational problems arising in the selection of short oligos (*e.g.*, 20–50 bases) from a large EST database. One is to identify oligos that are *unique* to each EST in the database. The other is to identify oligos that are popular among the ESTs. More precisely, the *unique oligo* problem asks for the set of all oligos each of which appears (exactly) in one EST sequence

but does not appear (exactly or approximately) in any other EST sequence, whereas the *popular oligo* problem asks for a list of oligos that appear (exactly or approximately) in the largest number of ESTs. [1]

A unique oligo can be thought of as a "signature" that distinguishes an EST from all the others. Unique oligos are particularly valuable as locus-specific PCR primers for placement of ESTs at single positions on a genetic linkage map, on microarrays for studies of the expression of specific genes without signal interference from other genes, and to probe genomic libraries [7] in search of specific genes.

Popular oligos can be used to screen efficiently large genomic library. They could allow one to simultaneously identify a large number of genomic clones that carry expressed genes using a relatively small number of (popular) probes and thus save considerable amounts of money. In particular for the database under analysis, it has been shown previously by a number of independent methods that the expressed genes in Triticeae are concentrated in a small fraction of the total genome. In barley, this portion of the genome, often referred to as the *gene-space*, has been estimated to be only 12% of the total genome [8]. If this is indeed true, then at most 12% of the clones in a typical BAC library would carry expressed genes, and therefore also the vast majority of barley genes could be sequenced by focusing only on this 12% of the genome. An efficient method to reveal the small portion of BAC clones derived from the gene-space has the potential for tremendous cost savings in the context of obtaining the sequences of the vast majority of barley genes. The most commonly used barley BAC library has a 6.3 fold genome coverage, 17-filter set with a total of 313,344 clones [3]. This number of filters is inconvenient and costly to handle, and the total number of BAC clones is intractable for whole genome physical mapping or sequencing. However, a reduction of this library to a gene-space of only 12% of the total would make it fit onto two filters that would comprise only about 600 Mb. This is about the same size as the rice genome, which has been recently sequenced. A solution for the popular oligo problem should make it possible to develop an effective greedy approach to BAC library screening, enabling a very inexpensive method of identifying a large portion of the BAC clones from the gene-space. This would also likely accelerate progress in many crop plant and other systems that are not being considered for whole genome sequencing.

Our Contribution. In this paper, we present an efficient algorithm to identify all unique oligos in the ESTs and an efficient heuristic algorithm to enumerate the most popular oligos. Although the unique and popular oligos problems are complementary in some sense, the two algorithms are very different because unique oligos are required to appear in the ESTs while the popular oligos are not. In particular, the heuristic algorithm for popular oligos is much more involved than that for unique oligos, although their (average) running times are similar. The algorithms combine well-established algorithmic and data structuring techniques such as hashing, approximate string matching, and clustering, and take

[1] Note that, a popular oligo does not necessarily have to appear exactly in any EST.

advantage of the facts that (i) the number of mismatches allowed in these problems is usually small and (ii) we usually require a pair of approximately matched strings to share a long common substring (called a *common factor* in [9]). These algorithms have been carefully engineered to achieve satisfactory speeds on PCs, by taking into account the distribution of the frequencies of the words in the input EST dataset. For example, running each of the algorithms for the barley EST dataset from HarvEST takes only a couple of hours (on a 1.2 GHz AMD machine). This is a great improvement over other brute-force methods, like the ones based on BLAST. [2] Simulations results show that the number of missed positives by the heuristic algorithm for popular oligos is very limited and can be controlled very effectively by adjusting the parameters.

Previous related work. The problem of finding infrequent and frequent patterns in sequences is a common task in pattern discovery. A quite large family of pattern discovery algorithms has been proposed in the literature and implemented in software tools. Without pretending to be exhaustive, we mention Meme [10], Pratt [11,12], Teiresias [13], Consensus [14], Gibbs Sampler [15,16], Winnower [17,18], Projection [19,20], Verbumculus [21], Mitra [22], among others. Although these tools have been demonstrated to perform very well on small and medium-size datasets, they cannot handle large datasets such as the barley EST dataset that we are interested in. In particular, some of these tools were designed to attack the "challenge" posed by Pevzner and Sze [17], which is in the order of a few Kbases. Among the more general and efficient tools, we tried to run Teiresias on the 28 Mbases barley EST dataset on an 1.2GHz Athlon CPU with 1GB of RAM, without being able to obtain any result (probably due to lack of memory).

The unique oligo problem has been studied in the context of *probe design* [23,9,24]. The algorithms in [23,24] consider physical and structural properties of oligos and are very time consuming. (The algorithm in [24] also uses BLAST.) A very recent algorithm by Rahman [9] is, on the other hand, purely combinatorial. It uses suffix arrays instead of hash tables, and requires approximately 50 hours for a dataset of 40 Mb on a high-performance Compaq Alpha machine with 16 Gb of RAM. However, his definition of unique oligos is slightly different from ours (to be given in the next section).

The rest of the paper is organized as follows. Section 2 defines the unique and popular oligo problems formally. The algorithms are presented in Section 3. Experimental results on the barley EST dataset and simulation results can be found in Section 4. In Section 5, we draw some concluding remarks. The Appendix B explains the popular oligo algorithm with an example.

[2] For example, one can identify unique oligos by repeatedly running BLAST for each EST sequence against the entire dataset. This was the strategy previously employed by the HarvEST researchers.

2 The Unique and Popular Oligo Problems

We denote the input dataset as $X = \{x_1, x_2, \ldots x_k\}$, where the generic string x_i is an EST sequence over the alphabet $\Sigma = \{A, C, G, T\}$ and k is the cardinality of the set. Let n_i denote the length of the i-th sequence, $1 \leq i \leq k$. We set $n = \sum_{i=1}^{k} n_i$, which represents the total size of the input. A string (or oligo) from Σ is called an *l-mer* if its length is l.

Given a string x, we write $x_{[i]}$, $1 \leq i \leq |x|$, to indicate the i-th symbol in x. We use $x_{[i,j]}$ as a shorthand for the substring $x_{[i]}x_{[i+1]} \ldots x_{[j]}$ where $1 \leq i \leq j \leq n$, with the convention that $x_{[i,i]} = x_{[i]}$. Substrings in the form $x_{[1,j]}$ correspond to the *prefixes* of x, and substrings in the form $x_{[i,n]}$ to the *suffixes* of x. A string y *occurs* at position i of another string x if $y_{[1]} = x_{[i]}, \ldots, y_{[m]} = x_{[i+m-1]}$, where $m = |y|$. For any substring y of x, we denote by $f_x(y)$ the number of occurrences of y in x. $f_X(y)$ denotes the total number of occurrences of y in $x_1, \ldots, x_k$.

The *color-set* of y in the set $X = \{x_1, x_2, \ldots, x_k\}$ is a subset $col(y) = \{i_1, i_2, \ldots, i_l\}$ of $\{1, 2, \ldots, k\}$ such that if $i_j \in col(y)$ then y occurs at least once in x_{i_j}. We also say that y has *colors* $i_1, i_2, \ldots, i_l$. The number of colors, l, of y is denoted as $c_X(w)$. Clearly $f_X(y) \geq c_X(y)$.

Given two strings x and y of the same length, we denote by $H(x, y)$ the Hamming distance between x and y, that is, the number of mismatches between x and y. If $H(x, y) \leq d$, we say that x *d-matches* y and x is a *d-mutant* of y. The set of all the strings that d-match x is called the *d-neighborhood* of x. The notion of occurrences and colors can be extended to d-occurrence and d-colors by allowing up to d mismatches. If a string y has d-mutants at j distinct positions in a string x, we say that y has j d-occurrences in x. If a string y has at least one d-occurrence in each of j sequences in X, we say that y has j d-colors in X.

In the context of DNA hybridization, most papers define the specificity of an l-mer in terms of its mismatches to the length-l substrings of target sequences, although some also consider its physical and structural characteristics such as melting temperature, free-energy, GC-content, and secondary structure [23,24]. In [9], Rahman took a more optimistic approach and used the length of the longest common substring (called the longest common factor or LCF) as a measure of unspecificity. Given the nature of our target applications, we will take a conservative approach in the definitions of unique and popular oligos.

Definition 1. *Given the set $X = \{x_1, x_2, \ldots x_k\}$ of ESTs and integers l and d, a* unique oligo *is an l-mer y such that y occurs in at least one EST and the number of d-colors of y in X is exactly one. In other words, y appears exactly in some EST but does not appear approximately in any other EST.*

Suppose that strings x and y have the same length. For any given constants c, d, we say that string x (c, d)-matches string y if x and y can be partitioned into substrings as $x = x_1x_2x_3$ and $y = y_1y_2y_3$ with $|x_i| = |y_i|$ such that (i) $|x_2| = c$, (ii) $x_2 = y_2$, and (iii) the string x_1x_3 d-matches the string y_1y_3. In the above partition, we call x_2 a *core* in the (c, d)-match between x and y. (Note that a (c, d)-match may have many cores). The notion of d-occurrences and d-colors can be easily extended to (c, d)-occurrences and (c, d)-colors.

Definition 2. *Given the set $X = \{x_1, x_2, \ldots x_k\}$ of ESTs and integers l, d, c and T, a* popular oligo *is an l-mer y such that the number of (c, d)-colors of y in X is greater than or equal to T. In other words, the l-mer y appears approximately in at least T ESTs.*

The use of pooled oligo probes for BAC library screening [7] generally have lengths from 24 to 40 bases. Given this range and based on discussion with researchers from the Triticeae community, we consider $l = 33$ and $d = 5$ in the unique oligo problem. In the popular oligo problem, we consider $d = 1, 2, 3$ and $c = 20$.

3 The Algorithms

Our goal is to determine unique and popular oligos for a given set X of EST sequences. Although the objectives of the two problems seem complementary, our algorithms are quite different. The algorithm for the popular oligo problem turns out to be much more involved because popular oligos are not necessarily contained in the ESTs. Nevertheless, both algorithms share some common strategies such as the idea of separating dissimilar strings as early as possible to reduce the search space. To achieve this in the popular oligo problem, we first find cores (*i.e.*, c-mers) that appear exactly in at least two ESTs. Then we cluster all length-l substrings of the ESTs by their cores of length c using a hash table, and then, within each cluster, we cluster again the l-mers based on Hamming distance between regions flanking the cores. Candidate popular oligos are then enumerated from the small groups resulted from the two clusterings, and their number of colors are counted. For the unique oligo problem, the notion of cores, however, does not exist. On the other hand, due to the small number of mismatches allowed (relative to l), two l-mers that d-match each other must contain substrings that 1-match each other. Such substrings are called *seeds*. We can thus cluster the ESTs by their seeds using a dictionary. For each cluster, we compare the ESTs in the cluster by counting mismatches in the regions flanking the seeds. The details are given below.

3.1 Unique Oligos

Recall that the unique oligo problem is to identify length-l substrings of the ESTs that have exactly one d-color in the dataset X, for a given value of d. Our strategy is first to eliminate all the those l-mers that cannot be unique oligos. The algorithm is based on the following observation. Assume that x and y are two l-mers such that $H(x, y) \leq d$. Divide x and y into $t = \lfloor d/2 \rfloor + 1$ substrings. That is $x = x_1 x_2 \cdots x_t$ and $y = y_1 y_2 \cdots y_t$, where the length of each substring is $q = \lceil l/t \rceil$, except possibly for the last one. In practice, one can always choose l and d so that l is a multiple of t and hence x and y can be decomposed into t substrings of length q, which we call *seeds*. It is easy to see that since $H(x, y) \leq d$, at least one of the seeds of x has at most one mismatch with the corresponding seed of y.

```
UNIQUE-OLIGO-SELECTION(X, l, m)
Input:  EST sequences X = {x_1, x_2, ..., x_k}
        l: length of the oligos to be reported
        d: maximum number of mismatches for non-unique oligos
Output: Mark each unique l-mers in X
1  t, q ← ⌊d/2⌋ + 1, ⌈l/t⌉
2  table ← HIT(X, q)
3  EXTENSION(X, t, q, table, d)
```

```
COMPARE(table[i][j], table[i][k], d)
1  q_1 ← the q-mer located at table[i][j]
2  q_2 ← the q-mer located at table[i][k]
3  for each pair of l-mers that contain q_1 and q_2 as seeds respectively do
4        c ← H(q_1, q_2)
5        if c ≤ d then
6           mark the two l-mers as "non-unique"
```

```
HIT(X, q)
1  for i ← 1 to 4^q do
2      initialize table[i]
3      index[i] ← 0
4  for i ← 1 to k do
5      for j ← 1 to n_i − q + 1 do
6          key ← MAP(x_{i,[j,...,j+q−1]})
7          table[key][index[key]] ← ⟨i, j⟩
8          index[key] ← index[key] + 1
9  return table
```

Fig. 1. The algorithm for identifying unique oligos (continues in Figure 2).

Using this idea, we design an efficient two-phase algorithm. In the first phase, we cluster all the possible seeds from the ESTs into groups such that within each group, a seed has no more than one mismatch with the other seeds. In the second phase, we check whether extending the flanking regions of a seed would result in a d-match with the corresponding extension of any other seed in the same group. If so, the l-mer given by this extension is not a unique oligo.

Phase 1. (HIT) We file all q-mers (seeds) from the input ESTs into a dictionary with 4^q entries. (If 4^q cannot fit in the main memory, one could use a hash table of an appropriate size.) Each entry of the table points to a list of locations where the q-mer occurs in the EST sequences. Using the table we can immediately locate identical seeds in the ESTs.

```
EXTENSION(X, t, q, table, m)
1  for i ← 1 to 4^q do
2      List mut ← mutant list of table[i]
3      len ← # of records in table[i]
4      for j ← 1 to len do
5          for k ← j + 1 to len do
6              COMPARE(table[i][j], table[i][k], m)
7          for h ← 1 to # of records in mut do
8              mutlen ← # of records in table[mut[h]]
9              for k ← 1 to mutlen do
10                 COMPARE(table[i][j], table[mut[h]][k], m)
```

```
MAP(string S)
1  map S into an integer X with function f : {A,C,G,T} → {0, 1, 2, 3}
2  return X
```

Fig. 2. The algorithm for identifying unique oligos (continued).

Phase 2. (EXTENSION) We compare the corresponding flanking regions of each pair of matching seeds to determine whether they can be extended to a pair of l-mers that d-match each other. Here, we also collect seeds that have exactly one mismatch with each other as follows. For each table entry corresponding to a seed y, we record a list of other seeds that have exactly one mismatch with y, by looking up table entries that correspond to all the 1-mutants of y. This list is called a *mutant list* of y. We examine all the seeds in the mutant list, and compare the flanking regions of the q-mers and that of y in the same way as we did for identical seeds, except that now the cutoff for the number of mismatches in the flanking regions is $d - 1$.

The algorithm is summarized in Figures 1 and 2.

Time complexity. Suppose that the total number of bases in X is n. The time complexity of phase one is simply $\Theta(qn)$. The time complexity of phase two depends on the distribution of the number of seeds filed into each table entry. Simply speaking, if the distribution is more or less uniform (which is the case in our experiment) and each table entry contains $r \approx n/4^q$ identical seeds, the number of comparisons within the table entry is $O(r^2)$. The number for comparisons for each mutant lists of size $3q$ is $O(qr^2)$. Each comparison requires extension of the seeds and takes $2(l - q)$ time. Since there are 4^q entries in the table, the overall time complexity is $O((l - q)qr^2 4^q)$. Given the exponential dependency on q, one needs to make sure that q is not too large before using the algorithm. (Again, in our experiment on the barley dataset, $l = 33, d = 5$ and $q = 11$.)

In the practice of EST data analysis, we need also consider the reverse complementary strand of each EST, which implies more stringency in the choice of unique oligos. The above algorithm can be easily modified to take into account reverse complementary EST strands without a significant increase of complexity.

3.2 Popular Oligos

Recall that the objective is to find all l-mers that have sufficiently large number of (c, d)-colors in X. Since popular oligos are not required to appear exactly in the EST sequences, we cannot solve the problem by exhaustive enumerate all the substrings of length l in the EST sequences and count their (c, d)-colors. In fact, one can easily show that the problem is NP-hard in general.

A straightforward algorithm is to consider all l-mers occuring in the ESTs and for each l-mer, enumerate all its (c, d)-mutants and count their number of (c, d)-colors. However, the number of (c, d)-mutants of an l-mer over the DNA alphabet is more than $\binom{l-c}{d}3^d$. Hence, the "brute-force" method becomes computionally impractical due to its memory requirement as soon as the input size reaches the order of hundreds of thousands of bases (like the barley dataset). [3]

We can reduce the search space using the same idea as in the algorithm for unique oligos, except that here the role of seeds is played by cores. Observe that, if a (popular) oligo has (c, d)-occurrences in many ESTs, many of these ESTs must contain length-l substring that share common cores. Based on this observation, we propose a heuristic strategy that first clusters the l-mers in the EST sequences into groups by their cores, and then enumerates candidate l-mers by comparing the members of each cluster in a hierarchical way.

An outline of the algorithm is illustrated in Figure 3. Here, we determine the popularity of the cores (*i.e.*, length-c substrings) from the ESTs in the first step. For each popular core, we consider extension of the cores into l-mers by including flanking regions and cluster them using a well-known hierarchical clustering method, called *unweighted pair group method with arithmetic mean* (UPGMA) [25]. We recall that UPGMA builds the tree bottom-up in a greedy fashion by merging groups (or subtrees) of data points that have the smallest average distance. Based on the clustering tree, we compute the common oligos shared by the l-mers by performing set intersection. These common oligos shared by many l-mers become candidate popular oligos. Finally, we count the number of colors of these candidates, and output the oligos with at least T colors. A more detailed description is given below. A complete example of the algorithm on a toy dataset is also given in the appendix.

Phase 1. We compute the number of colors for all c-mers in the ESTs to determine whether they could be candidate cores for popular l-mers, using a hash table. According to our definition, a popular oligo should have a popular core. We therefore set a threshold T_c on the minimum number of colors of each

[3] When $d = 3$ and $c = 20$, $\binom{l-c}{d}3^d = \binom{13}{3}3^3 = 7,722$ for the barley dataset. Hence, the straightforward algorithm would have to count the number of colors for about $7,722 \cdot 28 \times 10^6 = 217 \times 10^9$ l-mers.

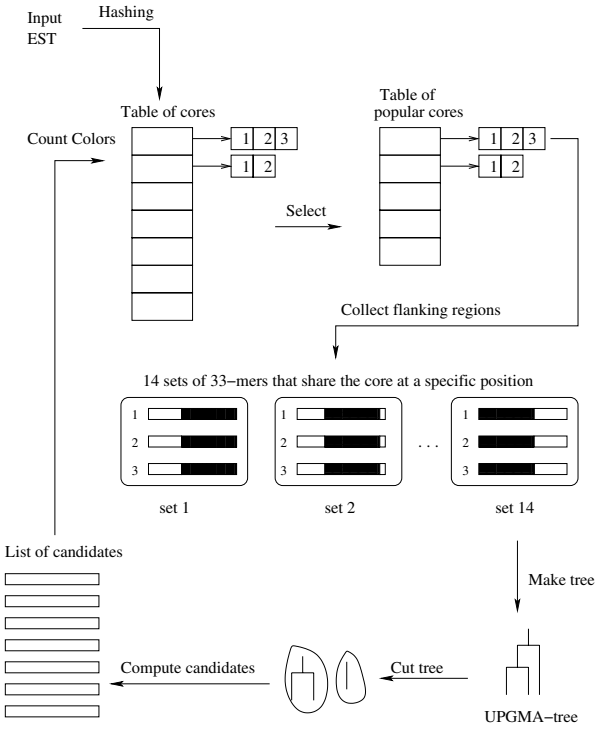

Fig. 3. An overview of the algorithm for selecting popular oligos. For convenience of illustration, the length of the oligos is assumed to be $l = 33$, and the length of the cores is assumed to be $c = 20$.

popular core, depending on T, c, l and X. All cores that have a number of colors below T_c are filtered out, and considered "unpopular". However, since an l-mer can (c, d)-match another l-mer with any of its $l - d + 1$ cores, it is possible that we might miss some popular oligos that critically depend on unpopular core. The parameter T_c represents a tradeoff between precision and efficiency. We will show in Section 4 the effect of changing T_c on the output. We will see that in practice we might miss only a negligible number of popular oligos.

Phase 2. Here we collect the substrings flanking the popular cores. For each popular core, we construct $l - c + 1$ sets of substrings, one for each possible extension of the core into an l-mer. Each set contains substrings of length $l - c$ constructed by concatenating the left and right flanking regions.

Phase 3. For each set of flanking substrings, we would like to identify all $(l - c)$-mers that have d-occurrences in many of these substrings. In order to achieve this efficiently, we first cluster the substrings according to their mutual Hamming distance using the well-known hierarchical clustering method UPGMA. In the process of building the clustering tree, whenever the Hamming distance between some leaves in the tree is zero we compress the distance matrix by combining the

identical strings into one entry. This significantly reduces the running time not only because the tree becomes smaller, but also because the number of common d-mutants of two different l-mers is much less than that of two identical ones. As we can see later, a significant proportion of the running time is spent on the intersection of the sets of d-mutants. Compressing the distance matrices avoids intersecting identical sets of d-mutants, which is expensive and also useless. We then create a set of d-mutants for each substring represented at the leaves and traverse the tree bottom-up. At each internal node u, we compute the intersection of the two sets attached to the children, using a hash table based on the hash function described in [26]. This intersection represents all the $(l-c)$-mers that have d-occurrences in all the leaves (substrings) under the node u. As soon as the intersetion becomes empty, we prune the tree. At the end of this process, we obtain a collection of sets of $(l-c)$-mers, each of which, together with the popular core, represents a candidate popular oligo.

Phase 4. Given the candidate popular oligos, we need to count their number of (c,d)-colors. Before counting, we radix-sort the candidates and remove duplicates. More precisely, due to the possibly very large number of candidates and duplicates (as in the barley case), we sort the candidates in several stages as follows. For each core, we radix-sort all the candidates derived from the core and remove duplicates. Then we merge the sorted list into the sorted list containing all the candidates from those cores that have already been processed.

Time complexity. Phase 1 costs time $O(cn)$. In phase 2, if the number of popular cores selected in the first step is p and the average number of occurrence of the cores is r, this phase costs $O(nr(l-c))$. For phase 3, the time for building a UPGMA tree, including the computation of the distance matrix, is $O((l-c)r^2)$, where r stands for the number of strings to be clustered. Since a (binary) UPGMA tree with r leaves has $2r-1$ nodes, the time for traversing (and pruning) the tree is $O(r\binom{l-c}{d}3^d)$, where $\binom{l-c}{d}3^d$ is the number of d-mutants at each leaf. Finally for phase 4, if the total number of candidates is m, counting the colors for the candidates, excluding the time for radix-sort, costs time $O(rm(l-c))$.

4 Implementation and Results

We have implemented both algorithms in C and tested the programs on a desktop PC with a 1.2GHz AMD Athlon CPU and 1GB RAM, under Linux. The main dataset is a collection barley ESTs from HarvEST containing $k = 46,145$ EST sequences with a total of $n = 28,475,017$ bases. Before doing the searches, we first cleaned the dataset by removing PolyT and PolyA repeats.

As mentioned above, our first task was to search for unique oligos of length $l = 33$ with a minimum number of mismatches $d = 5$. Based on these parameters, each oligo was divided into three seeds of length $q = 11$. Hence, our dictionary table had $4^{11} \approx 4$ million entries. The efficiency of our algorithm critically depends on the statistical distribution of the seeds in the dictionary. The statistics of the seeds in our experiment (before the extension phase) is shown in table in the left of Figure 4. Clearly, most seeds occur less than 20 times in the ESTs

# of occurrences	# of seeds
0	242399
1-9	3063288
10-19	708745
20-29	120698
30-39	31637
40-49	11908
50-5049	15629

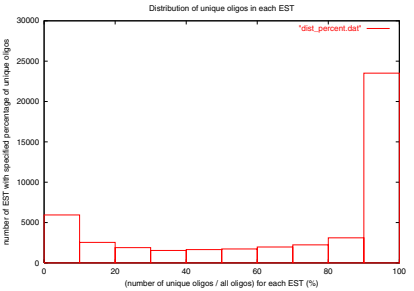

Fig. 4. LEFT: Distribution of frequencies of seeds in barley ESTs. RIGHT: Distribution of unique oligos. The horizontal axis stands for the percentage of unique oligos over all 33-mers in an EST, and the vertical axis stands for the number of ESTs whose unique oligos are at a certain percentage of all its 33-mers.

and this is the main reason why our algorithm was able to solve the dataset efficiently. The final distribution of unique oligos is shown in the right of Figure 4. Note that, there are many ESTs (slightly more than half of the entire dataset) whose length-33 substrings are almost all unique oligos. In particular, there are 13,430 ESTs whose length-33 substrings are all unique oligos and there are 2,159 ESTs that contain no unique oligos. The whole computation took 2 hours and 26 minutes and used about 200 MB of memory.

Table 1. Distribution of the number of colors of the cores. The left column is the range of the number of colors. The right column is the number of cores with a certain number of color.

colors	number of cores
1	22523412
2-10	2128677
11-20	5148
21-30	1131
31-40	492
41-50	346
51-60	242
61-70	77
71-80	34
81-90	29
91-100	43
101-176	19

Our second task was to search for popular oligos with length $l = 33$ and core length $c = 20$. We considered different choices for the maximum number of mismatches d ourside the core and the threshold T_c on the minimum number of

Table 2. Running time for enumerating popular oligos (in seconds).

	$T_c = 30$	$T_c = 40$	$T_c = 50$	$T_c = 60$
$d = 1$	184	166	156	153
$d = 2$	278	219	177	338
$d = 3$	3808	1788	730	359

colors for the popular cores. The distribution of the number of colors of the cores is shown in Figure 1. From the table we can see that the number of cores decreases almost exponentially as the number of colors increases. On the other hand, cores with low colors are unlikely to contribute to popular oligos. Therefore, it is important to filter them out to increase the efficiency.

The running time of this program varies with the parameters d and T_c, as shown in the Figure 2. The memory used in the program was mainly for storing the candidate popular oligos. In general, about 64 MB suffices since the program reuses the memory frequently.

4.1 Simulations

To evaluate the performance of our heuristics for selecting popular oligos we also have run a few simulations as follows. We generated a set of artificial ESTs by creating first a set of k random sequences and then injecting a controlled number of approximate occurrences of a given set of oligos. The initial set of oligos, denoted by $I_1, \ldots, I_s$, was also generated randomly over Σ. Each oligo I_i was assigned a predetermined number of colors C_i. We decided that the distribution of the C_i should be Gaussian, that is, we defined $C_i = C\, e^{-i^2/2}/\sqrt{2\pi}$ where C is a fixed constant which determines the maximum number of colors. As said, the positions in-between the oligos were filled with random symbols over the DNA alphabet.

We then ran our program for popular oligos on the artificial ESTs dataset and output a set of candidate oligos $O_1, \ldots, O_t$ with their respective colors $C'_1, \ldots, C'_t$. The output oligos were sorted by colors, that is $C'_i \geq C'_j$, if $i < j$. Since the output contained redudant candidates that came from the mutations of the original popular oligos, we removed those candidates that were a d-mutant of another oligo with an higher number of colors. More precisely, if O_i was a d-mutant of O_j, and $1 \leq i < j \leq t$, then O_j was discarded. This "compression step" did not eliminate good candidates for the following reason. Since the input oligos $I_1, \ldots, I_s$ were generated randomly they were very unlikely to be similar. As a consequence, the corresponding output oligos were also unlikely to be eliminated.

Finally, we compared the pair (I, C) with (O, C'). The more similar (O, C') is to (I, C), the better is the heuristic of our algorithm. Recall that I and O were sorted by decreasing number of of colors. We compared the entries in (I, C) with the ones in (O, C'), position by position. For each $1 \leq i \leq u$, where $u = \min(s, t)$, we computed the average difference between C and C'

Table 3. The average relative errors between the number of colors in the input and the number of colors in output for a simulated experiment ($n = 1,440,000$, $k = 720$, $c = 20$, $C = 100$, $s = 100$).

	$d = 2$	$d = 3$
$T_c = 10$	0.019	0.022
$T_c = 15$	0.000	0.002
$T_c = 20$	0.018	0.000
$T_c = 25$	0.000	0.001
$T_c = 30$	0.000	0.000

as $E = (1/u) \sum_{i=1}^{u} \frac{|C(i) - C'(i)|}{C'(i)}$. If we assume that I and O contain the same set of oligos, the smaller is E, the more similar is (I, C) to (O, C'). To validate this assumption, we also searched the list of oligos I in O, to determine whether we missed completely some oligos.

Figure 3 shows the value of E for one run of the program on a dataset of $n = 1,440,000$ bases composed by $k = 2,000$ sequences of size 720. We generated a set of $s = 100$ oligos with a maximum number of colors $C = 100$. In the analysis, we fixed the size of the core to be $c = 20$, whereas the maximum number of mismatches d outside the core and the threshold T_c were varied. The results show that the average relative error is below 2%. In the final version of the paper, we plan to repeat these simulations a few times to get more stable and reliable results. We also compared the list of input oligos with the list of output oligos and we found that sometimes the program misses one or two oligos out of 100. However, the number of colors of these missed oligos is always near the threshold T_c. We never miss an oligo whose number of color is above $T_c + 10$.

5 Conclusion

We have proposed two algorithms to find unique and popular oligos in large EST databases. The size of our dataset, in the order of tens of millions of bases, was the real challenge due to the limitation in the size of main memory in common 32-bits architectures. Our algorithms were able to produce a solution in a reasonable amount of time on a regular PC with a modest amount of memory. As far as we know, no other existing tools are capable of handling such a dataset with such limited resources. Simulations show that the number of missed oligos by the heuristic algorithm for popular oligos is negligible and can be controlled very effectively by adjusting the parameters. Although the algorithms were initially designed to address the challenges from the barley EST dataset, the methods can be easily adapted to solve similar problems concerning infrequent and frequent oligos on other large datasets. The software will be released in the public domain in the near future.

Acknowledgments. The authors would like to thank Steve Wanamaker for producing the sequence assemblies and unigene datasets that were used in this work.

References

1. Adams, M.D., Kelley, J.M., Gocayne, J.D., Dubnick, M., Polymeropoulos, M.H., Xiao, H., Merril, C.R., Wu, A., Olde, B., Moreno, R.F., Kerlavage, A.R., McCombie, W.R., Venter, J.C.: Complementary DNA sequencing: Expressed sequence tags and human genome project. Science **252** (1991) 1651–1656
2. Boguski, M., Lowe, T., Tolstoshev, C.: dbEST–database for "expressed sequence tags". Nat. Genet. **4** (1993) 332–3
3. Yu, Y., Tomkins, J.P., Waugh, R., Frisch, D.A., Kudrna, D., Kleinhofs, A., Brueggeman, R.S., Muehlbauer, G.J., Wise, R.P., Wing, R.A.: A bacterial artificial chromosome library for barley (*Hordeum vulgare L.*) and the identification of clones containing putative resistance genes. Theor Appl Genet **101** (2000) 1093–1099
4. Close, T., Wing, R., Kleinhofs, A., Wise, R.: Genetically and physically anchored EST resources for barley genomics. Barley Genetics Newsletter **31** (2001) 29–30
5. Michalek, W., Weschke, W., Pleissner, K., Graner, A.: Est analysis in barley defines a unigene set comprising 4,000 genes. Theor Appl Genet **104** (2002) 97–103
6. Huang, X., Madan, A.: CAP3: A DNA sequence assembly program. Genome Research **9** (1999) 868–877
7. Han, C., Sutherland, R., Jewett, P., Campbell, M., Meincke, L., Tesmer, J., Mundt, M., Fawcett, J., Kim, U., Deaven, L., Doggett, N.: Construction of a BAC contig map of chromosome 16q by two-dimensional overgo hybridization. Genome Research **104** (2000) 714–721
8. Barakat, A., Carels, N., Bernardi, G.: The distribution of genes in the genomes of gramineae. Proc. Natl. Acad. Sci. U.S.A. **94** (1997) 6857–6861
9. Rahmann, S.: Rapid large-scale oligonucleotide selection for microarrays. In: Proceedings of the First IEEE Computer Society Bioinformatics Conference (CSB'02), IEEE Press (2002)
10. Bailey, T.L., Elkan, C.: Unsupervised learning of multiple motifs in biopolymers using expectation maximization. Machine Learning **21** (1995) 51–80
11. Jonassen, I., Collins, J.F., Higgins, D.G.: Finding flexible patterns in unaligned protein sequences. Protein Science **4** (1995) 1587–1595
12. Jonassen, I.: Efficient discovery of conserved patterns using a pattern graph. Comput. Appl. Biosci. **13** (1997) 509–522
13. Rigoutsos, I., Floratos, A.: Combinatorial pattern discovery in biological sequences: The Teiresias algorithm. Bioinformatics **14** (1998) 55–67
14. Hertz, G., Stormo, G.: Identifying DNA and protein patterns with statistically significant alignments of multiple sequences. Bioinformatics **15** (1999) 563–577
15. Lawrence, C.E., Altschul, S.F., Boguski, M.S., Liu, J.S., Neuwald, A.F., Wootton, J.C.: Detecting subtle sequence signals: A Gibbs sampling strategy for multiple alignment. Science **262** (1993) 208–214
16. Neuwald, A., Liu, J., Lawrence, C.: Gibbs motif sampling: Detecting bacterial outer membrane protein repeats. Protein Science **4** (1995) 1618–1632
17. Pevzner, P.A., Sze, S.H.: Combinatorial approaches to finding subtle signals in DNA sequences. In: Proc. of the International Conference on Intelligent Systems for Molecular Biology, AAAI press, Menlo Park, CA (2000) 269–278

18. Keich, Pevzner: Finding motifs in the twilight zone. In: Annual International Conference on Computational Molecular Biology, Washington, DC (2002) 195–204
19. Tompa, M., Buhler, J.: Finding motifs using random projections. In: Annual International Conference on Computational Molecular Biology, Montreal, Canada (2001) 67–74
20. Buhler, J., Tompa, M.: Finding motifs using random projections. J. Comput. Bio. **9** (2002) 225–242
21. Apostolico, A., Bock, M.E., Lonardi, S.: Monotony of surprise and large-scale quest for unusual words (extended abstract). In Myers, G., Hannenhalli, S., Istrail, S., Pevzner, P., Waterman, M., eds.: Proc. of Research in Computational Molecular Biology (RECOMB), Washington, DC (2002)
22. Eskin, E., Pevzner, P.A.: Finding composite regulatory patterns in DNA sequences. In: Proc. of the International Conference on Intelligent Systems for Molecular Biology, AAAI press, Menlo Park, CA (2002) Bioinformatics S181–S188
23. Li, F., Stormo, G.D.: Selection of optimal DNA oligos for gene expression arrays. Bionformatics **17** (2001) 1067–1076
24. Rouillard, J.M., Herbert, C.J., Zuker, M.: Oligoarray: Genome-scale oligonucleotide design for microarrays. Bioinformatics **18** (2002) 486–487
25. Swofford, D. In: PAUP: Phylogenetic Analysis Using Parsimony version 4.0 beta 10. Sinauer Associates, Sunderland, Massachusetts (2002)
26. Wesselink, J.J., de la Iglesia, B., James, S.A., Dicks, J.L., Roberts, I.N., Rayward-Smith, V.J.: Determining a unique defining DNA sequence for yeast species using hashing techniques. Bionformatics **18** (2002) 1004–1010
27. Ito, M., Shimizu, K., Nakanishi, M., Hashimoto, A.: Polynomial-time algorithms for computing characteristic strings. In Crochemore, M., Gusfield, D., eds.: Proceedings of the 5th Annual Symposium on Combinatorial Pattern Matching. Number 807 in Lecture Notes in Computer Science, Asilomar, CA, Springer-Verlag, Berlin (1994) 274–288

Appendix: An Example of the Popular Oligo Algorithm

We show a small example to illustrate each step of the algorithm. Again, let l denote the length of oligos, c the length of cores, d the maximum number of mismatches between two oligos, and T_c the threshold on the minimum number of colors of popular cores. In this toy example, $l = 8, c = 5, d = 1, T_c = 3$. The input is composed of four artificial EST sequences as shown in Figures 5. `EST1` and `EST3` are highly similar, which represents the similarity between some of the EST sequences in real data.

Phase 1. Seven cores out of 148 possible cores are selected as popular cores. Each entry of the hash table points to a list of positions of the occurrences of the core in the input ESTs.

Phase 2. We collect the flanking regions for the seven cores. Figure 6 shows the four sets of flanking regions for `AAGGC`. Note that the fourth set has one fewer element than the other 3 sets. The reason is that the core `AAGGC` occurs at the right boundary of `EST0`, and therefore has a shorter flanking region.

Phase 3. We cluster the flanking regions using UPGMA. In Figure 8, we show the clusters for set 2 of the core `AAGGC`. Observe that the Hamming distance

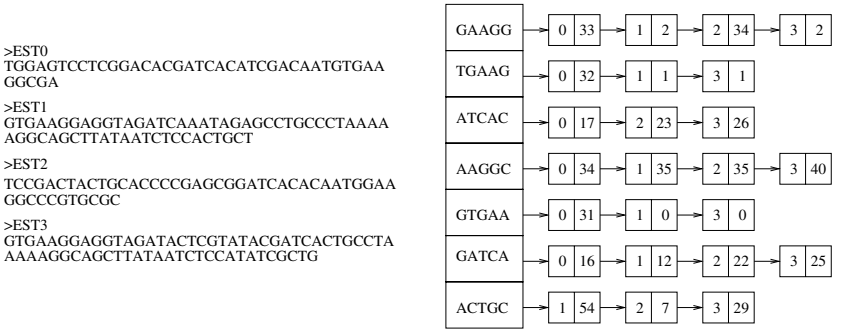

Fig. 5. The example dataset. The table of popular cores.

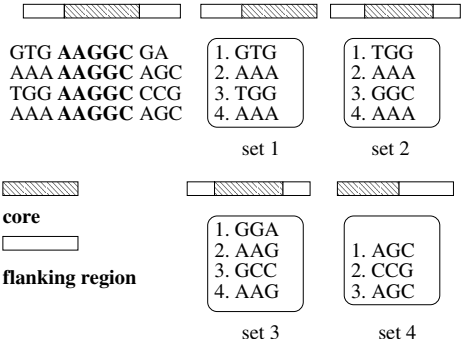

Fig. 6. Collecting flanking regions for the core `AAGGC`.

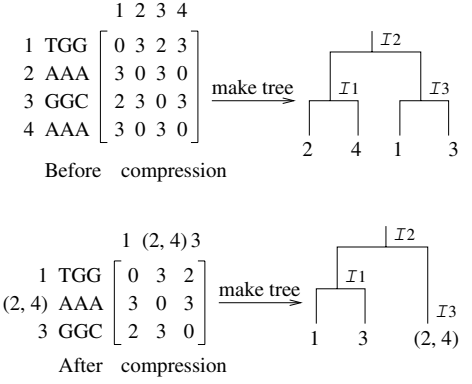

Fig. 7. UPGMA tree construction for set 2 of the core `AAGGC`.

between entry 2 and entry 4 is zero and therefore distance matrix is compressed by combining the identical strings into one entry. We then need to enumerate all the 1-mutants of the strings denoted by the leaves of the trees, that is `AAA`, `TGG`

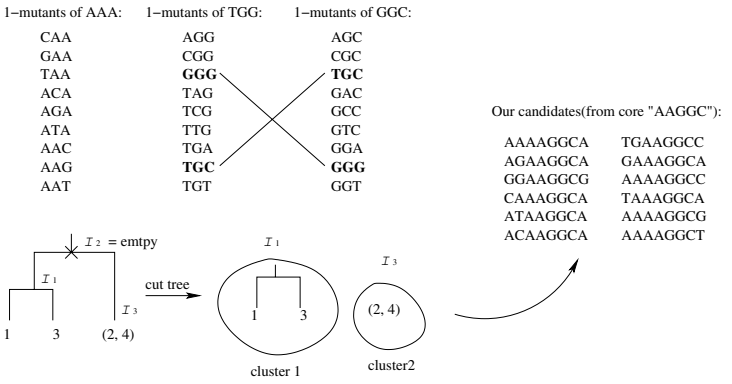

Fig. 8. Clustering set 2 of the core `AAGGC`

and `GGC`. Because leaf 1 and leaf 3 share the same parent, we apply intersection on their sets of 1-mutants and get the set I_1 = {`GGG, TGC`}. Then we apply intersection between I_1 and I_3, where I_3 represents the set of the 1-mutants of leaf `AAA`. Since the resulting intersection I_2 is empty, we prune the tree at I_2 and separately output the strings in I_1 and I_3 as flanking regions of candidate popular oligos. Note that, we output the 1-mutants of I_3 even if it is represented only by one node, because it is actually the intersection of two occurrences of `AAA`, and therefore all elements of I_3 have at least two d-matches in the EST sequences.