

Adapting the Knuth-Morris-Pratt Algorithm for Pattern Matching in Huffman Encoded Texts

Ajay Daptardar and Dana Shapira
{amax,shapird}@cs.brandeis.edu

Computer Science Department,
Brandeis University, Waltham MA 02454.

Abstract

In the present work we perform *compressed pattern matching* in binary Huffman encoded texts. A modified Knuth-Morris-Pratt (KMP) algorithm is used in order to overcome the problem of *false matches*, i.e., an occurrence of the encoded pattern in the encoded text that does not correspond to an occurrence of the pattern itself in the original text. We propose a bitwise KMP algorithm that can move one extra bit in the case of a mismatch since the alphabet is binary. To avoid processing any bit of the encoded text more than once, a preprocessed table is used to determine how far to back up when a mismatch is detected, and is defined so that we are always able to align the start of the encoded pattern with the start of a codeword in the encoded text. We combine our KMP algorithm with two practical Huffman decoding schemes which handle more than a single bit per machine operation; Skeleton trees defined by Klein [9], and numerical comparisons between special canonical values and portions of a sliding window presented in Moffat and Turpin [14]. Experiments show rapid search times of our algorithms compared to the “decompress then search” method, therefore, files can be kept in their compressed form, saving memory space. When compression gain is important, these algorithms are better than *cgrep* [6], which is only slightly faster than ours.

1 Introduction

The *Compressed Pattern Matching problem* which was first introduced by Amir and Benson [3], is of searching for a pattern directly in the compressed text without decompressing it. It is a variant of the classical pattern matching problem, in which one is given a pattern P and a text T , and the problem is to locate the first or all occurrences of P in T . In the compressed version of this problem, the text is supposed to be stored in some compressed form. More formally, given a pattern P and text T , and complementary encoding and decoding functions \mathcal{E} and \mathcal{D} , respectively. Our

goal is to search for the encoded pattern, $\mathcal{E}(P)$, in the encoded text $\mathcal{E}(T)$, rather than searching for the pattern, P , in the decompressed text, $\mathcal{D}(\mathcal{E}(T))$. Searching for the encoded pattern $\mathcal{E}(P)$ assumes that the encoded pattern is compressed in the same way through out the text. This assumption is not always true, especially with dynamic compressions, such as Lempel-Ziv variants, where the compression changes when one proceeds. In these cases the encoding of the pattern is also dynamic, and is computed at each point of the file [1, 10]. Here we deal with static Huffman files, and search for $\mathcal{E}(P)$ in it using a modified KMP algorithm. In order to speed-up the search we use the Skeleton trees introduced in the work of Klein [9] and the Huffman implementation of Moffat and Turpin [14].

Since the Huffman code is a variable length code, searching for encoded patterns in encoded texts raises the problem of *false matches*, i.e., finding an occurrence of the encoded pattern in the encoded text which does not correspond to an occurrence of the pattern in the original text, due to crossing codeword boundaries. Consider for example the Huffman code $\{00, 01, 100, 101, 110, 1110, 1111\}$ for the characters \mathfrak{t} , \mathfrak{a} , \mathfrak{g} , \mathfrak{d} , \mathfrak{b} , \mathfrak{o} and \mathfrak{c} , respectively. The binary string $101-1110-100$ is the encoding of the string \mathfrak{dog} . Suppose, however, that we are searching for the pattern \mathfrak{cat} : we could find $\mathcal{E}(\mathfrak{cat})$ starting at the third bit and extending to the end of $\mathcal{E}(\mathfrak{dog})$. The problem is thus one of verifying that the occurrence detected by the pattern matching algorithm is aligned on a codeword boundary. False matches can be avoided by using a code where no codeword is a prefix or suffix of any other codeword. This type of code is called *Affix* or a *Fixfree code*, and is extremely rare [7]. In this particular example, the code is not an affix code, since the codeword for \mathfrak{t} is a suffix of the codeword \mathfrak{g} .

One approach for performing direct pattern matching in encoded texts is to generate a compression method which is especially suitable for compressed pattern matching. Manber [12], for example, presents a static compression technique which is based on packing pairs of frequent characters in a single byte. Another work was done by Klein and Shapira [10], which modifies the LZSS algorithm by moving the pointers to the front. The dynamic nature of this compression requires a more powerful pattern matching algorithm.

Compressed pattern matching in Huffman encoded text has already been studied. Klein and Shapira [11] present a probabilistic algorithm for searching Huffman encoded texts. They use the tendency of the Huffman code to re-synchronize quickly after errors. After an occurrence of the compressed pattern in the compressed text has been detected, the search continues by jumping back in the compressed text by a fixed number of bits, and starts decompressing the text from there up to the point of the occurrence. If the decoding synchronizes with the occurrence, a match with high probability is announced. Although the probability of finding wrong matches is low, in this work we present a deterministic algorithm.

Turpin and Moffat [2] present an algorithm to directly search texts which were compressed using word-based Huffman codes, allowing only one word patterns. They construct an index of all words that occur in all files in a given directory, and apply a known pattern matching technique (such as *agrep* [17]) on the index, in order to search for a pattern in this set of files. The index, which lists all words and their

occurrences, is stored as part of the compressed file. We adapt the way the Huffman file is decoded and extend this work by allowing more than a single word pattern to be searched. Using their implementation, and our modified KMP algorithm, the file is searched directly, and more than one bit can be handled at a time.

Another word-based Huffman method is studied in the work of Moura et al. [5]. They present a compression and decompression technique where arbitrary portions of the compressed text can be decompressed, without the need for decompressing the entire file. Moreover, both exact and approximate pattern matching can be done directly on the compressed text. Their compression uses a word-based model and is based on byte oriented Huffman coding rather than bit oriented. Instead of using the original binary Huffman coding, their tree's degree is either 128 in what they call a *tagged Huffman code*, or 256 in their *plain Huffman code*. Thus the atomic unit of each codeword is a byte and traditional byte oriented algorithms can be employed for searching through the compressed text. However, the problem of false matches persists and some verification needs to be performed after a possible match has been reported. In order to not start the verification process at the start of the compressed text, the text is partitioned into appropriately sized blocks and verification can then proceed starting at the block boundary in which the match was reported rather than the start of the text. Although, they achieve fast searching times, the disadvantage of this approach is not only the need to re-compress Huffman encoded files in order to apply their pattern matching algorithm on the encoded file, but also the compression effectiveness using this method as opposed to Huffman.

The remainder of this paper is organized as follows. Section 2 presents our modified KMP algorithm. Section 3 shows how we can combine the modified KMP algorithm together with Moffat and Turpin's [14] decoding technique and Klein's skeleton trees [9], in order to perform compressed pattern matching. Section 4 presents experiments that compare both processing time and compression performance of our algorithm against the traditional "decompress then search", and *cgrep* [5].

2 Modifying the KMP algorithm

While searching for a given pattern in the compressed text, we first compress the pattern with the same code that was used for generating the compressed text. We then use a modified Knuth-Morris-Pratt algorithm [4] to search for the compressed pattern $\mathcal{E}(P)$ directly in the compressed text $\mathcal{E}(T)$. Note that the Boyer-Moore algorithm which searches the pattern from its right end, is not suitable for our purpose since one can not determine the codeword boundaries in the compressed text unless the text is decoded. Moreover, Boyer-Moore's algorithm, even with its sub-linear performance, is suitable for large alphabets rather than a binary alphabet as we have here.

The basic idea behind the original KMP algorithm is that each time a mismatch is detected, we know exactly how far to back-up the pointer in the pattern since this relies only on the characters in the pattern and not in the text. Consequently, the pointer in the text is never decremented. To accomplish this, the pattern is

P		b		a		c		e			
		⏟		⏟		⏟			⏟		
$\mathcal{E}(P)$		0	1	0	0	1	0	0	1	1	0
j		0	1	2	3	4	5	6	7	8	9
$next[j]$		-1	0	-1	1	0	-1	1	0	5	-1
after shifting				0	1	0	0	1	0	...	

Figure 1: *False matches*

preprocessed to obtain a table that gives an index in the pattern of the character to be used for the next comparison with the character that caused the mismatch in the text. We use this idea for searching for the encoded pattern in the encoded text by preprocessing the encoded pattern.

For any pattern $P = p[0..m - 1]$, during the preprocessing stage of the original KMP algorithm, a $next[0..m - 1]$ table is used to determine how far to back up when a mismatch is detected. The original preprocessing algorithm slides a copy of the first i characters of the pattern over itself, from left to right, starting with the first character of the copy over the second character of the pattern, and stopping when all overlapping characters match or when there are none left. These overlapping characters define the next possible place the pattern could match, if a mismatch is detected at $p[i]$. The distance to back up in the pattern ($next[i]$) is exactly the number of overlapping characters. Conventionally, $next[0] = -1$, which means that there is no overlap, and one must slide the pattern all the way to its beginning. For our purposes since codewords do not have a fixed length, we cannot apply this same algorithm to generate the KMP next table for individual bits of the compressed pattern, as the table might instruct us to perform comparisons in between codeword boundaries. These comparisons must be omitted since this would result a false match. This situation is illustrated in the following example and presented in Figure 1.

Let $a=00$, $b=01$, $c=100$, $d=111$ and $e=110$, and let $P = bace$, and assume that the encoded pattern is aligned on some codeword boundary. The original KMP table assigns $next[8] = 5$, i.e., if a mismatch at the 8th bit, the next bit to be processed is the fifth one. Even though there might be a match of $\mathcal{E}(P)$ at this position of $\mathcal{E}(T)$, this shifting will result in a mismatch, since the first bit of $\mathcal{E}(P)$ will be aligned on the second bit of the character a and thus not aligned on a codeword boundary of $\mathcal{E}(T)$. In addition to the problem of false matches, the following example illustrates the another drawback of using the $next$ table generated by the original KMP algorithm on characters of the uncompressed pattern rather than using the table obtained by our algorithm. Assume the same codewords as in the previous example, and let $P = bacbaebad$. Then $\mathcal{E}(P) = 01 - 00 - 100 - 01 - 00 - 110 - 01 - 00 - 111$. Suppose there is a mismatch when we are located on the last bit of $\mathcal{E}(P)$. The original KMP table generated from the characters, will direct us to the third character, c , while we already know that no occurrence occurs at the current position, since the last three bits read from $\mathcal{E}(T)$ are 110 and the codeword of c is 100. We can slide the pattern all the way so that the current position is aligned on the first bit of $\mathcal{E}(T)$. Note that shifting the pattern as instructed in the original KMP table, will force us to recheck




Figure 2: A mismatch between bits $\mathcal{E}(P)[j]$ and $\mathcal{E}(T)[i + j]$

bits.

The KMP next table for the encoded pattern, $next_bit[0..|\mathcal{E}(P)| - 1]$, then, must be such that each bit of the encoded text is processed exactly once while at the same time keeping codeword boundaries aligned. Preprocessing the compressed pattern can then be done in the following way. For each bit in the compressed pattern look for the longest prefix that matches the current suffix, so that they are both aligned on the first codeword boundary (and therefore on all codeword boundaries). Also, since we are dealing with the binary alphabet, we can improve on this algorithm by taking into account, the bit that caused the mismatch. If the pattern's bit that caused the mismatch is a 1, one can use this information and slide it to a point where the corresponding bit in the encoded pattern is a 0.

Given an encoded pattern $\mathcal{E}(P)$ of length m and an encoded text $\mathcal{E}(T)$ of length n , we consider an attempt to match the encoded pattern with the encoded text at position i , that is, the encoded pattern $\mathcal{E}(P)$ is aligned with the sub-pattern $\mathcal{E}(T)[i..i + m - 1]$ of the encoded text. Assume that the first mismatch occurs between bits $\mathcal{E}(P)[j]$ and $\mathcal{E}(T)[i + j]$ for some $0 \leq j < |\mathcal{E}(P)|$, as shown in Figure 2. Then $\mathcal{E}(P)[0..j - 1] = \mathcal{E}(T)[i..i + j - 1]$. The $next_bit[i]$, $0 \leq i < |\mathcal{E}(P)|$, gives the index of longest prefix v of $\mathcal{E}(P)$ that matches some suffix u of $\mathcal{E}(P)[k..j - 1]$, and they are both aligned on codeword boundaries. In addition, the bit following v in $\mathcal{E}(P)$ is equal to $\bar{a} = 1 - a$, where a is the bit following u in $\mathcal{E}(P)$.

We define a mapping I between indices in the pattern and indexes in the encoded pattern. Using a zero based index, the i^{th} character, $0 \leq i < |P|$, is mapped to the index of the last bit corresponding the i^{th} codeword in $\mathcal{E}(P)$. More formally, $I(i) = \sum_{j=0}^i |\mathcal{E}(P[j])| - 1$. For example, consider the codewords and pattern of the last example, and let us refer to the character c with index 2 (starting from 0), then $I(2) = 6$, since the index of the last bit of c is 6.

In order to compute the $next_bit$ table, we match the pattern against itself, as in the original KMP algorithm, but at the same time we also handle the codeword boundaries. The algorithm is presented in Figure 2. All entries of the first codeword are set to -1, since a mismatch in the first codeword enforces a mismatch of the entire pattern. In order to slide the pattern against itself, we use two indices i and j . If j is equal to -1 and i is not pointing to the last bit in a codeword, we advance i to the last bit of the current codeword, while setting each entry to -1. Otherwise, we are pointing on a codeword boundary, and we can use the original logic of the KMP

Algorithm: *bitwise-KMP-next*. Compute KMP *next_bit* table for $\mathcal{E}(P)$

```

begin
   $i \leftarrow 0$ 
  while ( $i \leq I(0)$ ) do
     $next\_bit[i] \leftarrow -1$  /* Set all entries of first codeword to -1. */
     $i \leftarrow i + 1$ 
  end while
   $i \leftarrow I(0)$  /* Set  $i$  to last bit in first codeword. */
   $j \leftarrow -1$ 
  while (true) do
    while( $j \geq 0$ ) and ( $\mathcal{E}(P)[i] \neq \mathcal{E}(P)[j]$ ) do
       $j \leftarrow next\_bit[j]$ 
    end while
    if ( $j = -1$ ) /* If  $j$  points to the last bit of the zeroth codeword, */
      while( $i \neq I(i)$ ) /* make sure  $i$  does too. */
         $i \leftarrow i + 1$ 
         $next\_bit[i] \leftarrow -1$ 
      end while
    end if
     $i \leftarrow i + 1$ 
     $j \leftarrow j + 1$ 
    if ( $i = |\mathcal{E}(P)|$ ) end
    if( $\mathcal{E}(P)[i] = \mathcal{E}(P)[j]$ )
       $next\_bit[i] \leftarrow next\_bit[j]$ 
    else
       $next\_bit[i] \leftarrow j$ 
    end if
  end while
end

```

Figure 3: *Computation of KMP next_bit table adapted for Huffman encoded patterns*

algorithm to compute the entries for the *next_bit* table.

3 Fast Pattern Matching on Huffman Texts

The number of Huffman trees for a given probability distribution is quite large. Many applications prefer to use a data structure defined by Schwartz and Kalick [16] known as a *canonical tree*. A tree is called canonical if when scanning its leaves from left to right (or right to left), they appear in non-decreasing order of their depth. An equivalent way for defining it is that when the codewords are sorted by the frequency of their corresponding symbols, they are ordered lexicographically. When using canonical trees, decoding can be done in a more efficient manner, requiring less memory space and fewer bitwise operations. Moffat and Turpin [14] and Klein [9] use the canonical

Algorithm: *KMP-search*. Search for $\mathcal{E}(P)$ in $\mathcal{E}(T)$

```

begin
  Preprocess  $\mathcal{E}(P)$  to obtain the next_bit table.
  while (not end of input) do
    read the next bit  $b$  from the input
    if ( $j \geq 0$ ) and ( $b \neq \mathcal{E}(P)[j]$ ) do
       $j \leftarrow \text{next\_bit}[j]$ 
    end if
    if ( $j = |\mathcal{E}(P)|$ )
      announce a match
       $j \leftarrow -1$ 
    end if
     $j \leftarrow j + 1$ 
  end while
end

```

Figure 4: *The modified KMP-search algorithm for Huffman encoded texts*

codes for fast decoding of Huffman texts, so that more than one bit can be processed in one machine operation. In this section we show how to combine the modified KMP algorithm with these methods.

The algorithm for searching Huffman encoded texts is given in Figure 4. Each bit in the encoded file is processed only once, using the modified KMP algorithm. Note that using the original KMP algorithm on the characters (symbols of the uncompressed text), effectively decompresses the file, which is something we wanted to avoid.

3.1 KMP with Skeleton trees

In the *KMP-search* algorithm, Figure 4, we are processing each bit of the encoded text in order to locate the encoded pattern. We present the *sk-KMP* algorithm for searching Huffman texts, which improves it by processing more than a single bit per machine operation, using *skeleton trees* [9]. This data structure represents canonical Huffman codes in a space efficient way and speeds up the decoding by handling more than a single bit at a time. A skeleton tree is a binary tree which is induced by the underlying Huffman tree, where all full subtrees of depth at least 1 have been pruned. That is, the nodes of the Huffman tree that remain in the skeleton tree are those up to the depth necessary to identify the length of the codeword with the prefix corresponding to the path from the root to that node. The leaves v of the skeleton tree then contain the length of the corresponding codewords $\ell(v)$. To search for a pattern in a Huffman encoded text using a skeleton tree, we start reading the encoded text one bit at a time while simultaneously traversing the skeleton tree and also keeping track of the position j in the *next_bit* table. If $\text{next_bit}[j] = -1$, then we must skip to the start of the next codeword in the encoded text and restart the process at the root of the skeleton tree. Skipping to the start of the next codeword

can be done as follows: read individual bits until a skeleton leaf has been reached and then read $\ell(v)$ at once. On the other hand if we reach a leaf v first, then the length of the current codeword can be identified, and the remaining $\ell(v)$ bits can be read at once in order to complete reading the current codeword.

Algorithm: *sk-KMP*. Search for $\mathcal{E}(P)$ in $\mathcal{E}(T)$ using skeleton tree *sk_tree*.

begin

Preprocess $\mathcal{E}(P)$ to obtain the *next_bit* table for the pattern and also initialize the *leaf_next* tables.

$v \leftarrow \text{root}(\text{sk_tree})$

$j \leftarrow 0$

while (not end of input) **do**

Let b be the next bit.

if ($b = 0$)

$v \leftarrow \text{left}(v)$

else

$v \leftarrow \text{right}(v)$

end if

if (v is a skeleton leaf)

Read next $\ell(v)$ bits from input into y .

$j \leftarrow \text{next_leaf}(v, j)[y]$

$v \leftarrow \text{root}(\text{sk_tree})$

else if ($j \geq 0$) **and** ($\mathcal{E}(P)(j) \neq b$)

$j \leftarrow \text{next_bit}[j]$

end if

if ($j = -1$)

Skip to the next codeword.

$v \leftarrow \text{root}(\text{sk_tree})$

end if

if ($j = |\mathcal{E}(P)|$)

Announce a match at the current position.

$j \leftarrow -1$

end if

$j \leftarrow j + 1$

end

Figure 5: *The sk-KMP search algorithm using skeleton trees for Huffman texts*

The algorithm in Figure 5 shows how to combine the skeleton tree with the modified KMP algorithm. We use v to point to the nodes in the skeleton tree, and j to point to bits in the encoded pattern. Each leaf, v , of the skeleton tree, that is also on a path of a codeword of $\mathcal{E}(P)$, has a *leaf_next* table, with 2^ℓ entries, where ℓ is the remaining bits from v to a Huffman leaf (codeword) rooted at v . A single lookup in $\text{next_leaf}(v, j)[y]$ will give us the position in $\mathcal{E}(P)$ after reading into y , the remaining ℓ bits of the codeword. The *sk-KMP* algorithm performs a single operation per node

in the skeleton tree, and not per node in the Huffman tree, saving in practice, about 50% of bit operations [9].

4 Experiments

The data files considered for the experiments were all natural language texts; these are as follows: *world192.txt* - CIA 1992 World Fact-book, *bible.txt* - King James Bible, *books.txt* - Random selection of texts from the Gutenberg Archives which is an archive of over 1000 documents in the English language in computer text format [8], *95-03-erp.txt* - US Economic Reports¹ from 1995 to 2003. All experiments were performed on an Intel PC with a 900 MHz AMD Athlon CPU with 256 KB cache memory and 256 MB of main memory running RedHat Linux-7.3.

The compression parsing model used on the input files uses a word-based alphabet together with the *spaceless words* method presented in the work of Moura et al. [5]. Every word is assumed to be followed with a space, in which case only the word is encoded, otherwise if the next symbol corresponds to a separator, in which case the separator must also be encoded. For compressing the files, we used Moffat and Katajainen’s algorithm [13] to first compute the lengths of the codewords and then proceed with the canonical code construction.

To speed-up the KMP algorithm we use Moffat and Turpin’s implementation of Huffman decoding [14, 2] who use the structure of canonical codes, that all codewords of a given length are consecutive binary integers. Only the first codeword of each length is stored, which also gives a sorted list of integers. A window, at least as long as the longest codeword, slides through the compressed stream, and its numerical value is compared against each one of these integers. The length of the next codeword is then determined, and the translation of the symbols number to the output string is done by a table look-up. We combine this with our modified KMP algorithm and call it *win-KMP*.

We compare our algorithms, *sk-KMP* and *win-KMP*, against *cgrep* of Moura et al. [5], and *agrep* [17] which searches the uncompressed text for patterns with or without errors. The alphabet of the *cgrep*, includes all words and separators of the text, and each element is assigned a sequence of bytes using only the 7 lower order bits of each byte. The most significant bit (MSB) in the first byte of each codeword is used as a separator between codewords by setting it to 1, while all other byte’s MSB is set to 0. This is done so that they can make sure that a reported match is not a false match. An implementation of *cgrep* can be downloaded from [6]. This particular implementation searches for single words with or without errors.

Table 1 compares the compression performance of the original Huffman compression (*Huff*) against the compression achieved by *cgrep* and *gzip*. The compression ratio presented in this table is the size of the compressed text as a percentage of the uncompressed text. Table 2 compares the processing time of pattern matching of these algorithms. The figures are given in seconds. The “decompress and search” methods,

¹The original files were in Adobe’s pdf file format from which they were converted to text using the UNIX utility *pdftotext* [15]

Files	Size (bytes)	Compression Ratio		
		<i>cgrep</i>	<i>Huff</i>	<i>gzip</i>
<i>world192.txt</i>	2,473,400	50.88	32.20	29.29
<i>bible.txt</i>	4,047,392	49.70	26.18	29.42
<i>books.txt</i>	12,582,090	52.10	30.30	37.04
<i>95-03-erp.txt</i>	23,976,547	34.49	25.14	22.53

Table 1: *Compression Performance*

first decode using skeleton trees (*sk-d*) or Moffat and Turpin’s sliding window (*win-d*) and then perform the search using *agrep*.

Files	Size (bytes)	Search Times (sec)				
		<i>cgrep</i>	<i>sk-KMP</i>	<i>win-KMP</i>	<i>agrep</i>	
					<i>sk-d</i>	<i>win-d</i>
<i>world192.txt</i>	2,473,400	0.07	0.13	0.08	0.21	0.13
<i>bible.txt</i>	4,047,392	0.05	0.22	0.13	0.36	0.22
<i>books.txt</i>	12,582,090	0.21	0.69	0.39	1.21	0.74
<i>95-03-erp.txt</i>	23,976,547	0.18	1.10	0.65	1.80	1.11

Table 2: *Search Performance*

As can be seen from these tables, using our modified KMP algorithm implemented with either Skeleton trees or Moffat and Turpin’s decoding, is faster than the “decompress and search” algorithms using the corresponding methods. When comparing the processing times of the KMP variants to *cgrep*, *cgrep* is faster, but its compression is less effective than Huffman. Not only should files be re-compressed in order to use *cgrep*, but it also harms the compression. Moreover, compressed files that can fit into the main memory, might exceed the memory space using the *cgrep* compression. Therefore, pattern matching that could have been performed in main memory would now have to include the time spent to transfer the file from secondary storage into main memory.

5 Conclusion

We have modified the Knuth-Morris-Pratt algorithm to perform *compressed pattern matching* in Huffman encoded texts. Our bitwise algorithm processes each bit of the encoded text exactly once. By combining it with the Skeleton trees defined by Klein, and the Huffman decoding implementation presented in Moffat and Turpin we are able to handle more than a single bit per machine operation. The processing times are better than the “decompress and search” method and slower than *cgrep*. However, when compression performance is important or when one does not want to re-compressed Huffman encoded files in order to use *cgrep*, the proposed algorithms are the better choice.

References

- [1] G. Benson A. Amir and M. Farach. Let sleeping files lie: Pattern matching in Z-compressed files. *Journal of Computer and System Sciences*, 52(2):299–307, April 1996.
- [2] A. Moffat A. Turpin. Fast file searching using text compression. In *Proc. 20th Australasian Computer Science Conference*, pages 1–8, Sydney, February 1997.
- [3] A. Amir and G. Benson. Efficient two-dimensional compressed matching. In *Proc. Data Compression Conference*, pages 279–288, Snowbird, Utah, March 1992. IEEE Computer Society.
- [4] J. H. Morris D. E. Knuth and V. R. Pratt. Fast pattern matching in strings. *SIAM Journal on Computing*, 6(2):323–350, 1977.
- [5] N. Ziviani E. S. de Moura, G. Navarro and R A. Baeza-Yates. Fast and flexible word searching on compressed text. *ACM Transactions on Information Systems*, 18(2):113–139, 2000.
- [6] P. Ferragina and A. Tommasi. Cgrep: C library to search over compressed texts. URL – <http://butirro.di.unipi.it/~ferrax/CompressedSearch>.
- [7] A. S. Fraenkel and S. T. Klein. Bidirectional huffman coding. *The Computer Journal*, 33(4):296–307, 1990.
- [8] Gutenberg archives. URL – <http://www.gutenberg.net/index.shtml>.
- [9] S. T. Klein. Skeleton trees for efficient decoding of huffman encoded texts. *Information Retrieval*, 3(1):7–23, July 2000.
- [10] S. T. Klein and D. Shapira. A new compression method for compressed matching. In *Proc. Data Compression Conference*, pages 400–409, 2000.
- [11] S. T. Klein and D. Shapira. Pattern matching in huffman encoded texts. In *Proc. Data Compression Conference*, pages 449–458, 2001.
- [12] U. Manber. A text compression scheme that allows fast searching directly in the compressed file. *ACM Transactions on Information Systems*, 15(2):124–136, 1997.
- [13] A. Moffat and J. Katajainen. In-place calculation of minimum-redundancy codes. In *Proc. Workshop on Algorithms and Data Structures*, pages 393–402, 1995.
- [14] A. Moffat and A. Turpin. On the implementation of minimum redundancy prefix codes. In *Proc. Data Compression Conference*, pages 170–179, 1996.
- [15] D. B. Noonburg. pdftotext - portable document format (pdf) to text converter (version 1.00). URL – <http://www.foolabs.com/xpdf/>.
- [16] E. S. Schwartz and B. Kallick. Generating a canonical prefix encoding. *Communications of the ACM*, 7(3):166–169, March 1964.

- [17] S. Wu and U. Manber. Fast text searching: Allowing errors. *Communications of the ACM*, 35(10):83–91, October 1992.