

Performance and Scalability of a Large-Scale N-gram Based Information Retrieval System

Ethan Miller, Dan Shen, Junli Liu, and Charles Nicholas
University of Maryland Baltimore County
{elm,dshen,jliu,nicholas}@csee.umbc.edu

ABSTRACT

Information retrieval has become more and more important due to the rapid growth in the amount of all kinds of information. However, there are few suitable systems available. This paper presents a few approaches that enable large-scale information retrieval for the TELLTALE system. TELLTALE is an information retrieval system that provides full-text search for text corpora in which documents may be garbled by OCR (Optical Character Recognition) or transmission errors, and that may contain multiple languages. Given a kilobyte query document, Telltale can find similar documents from within a gigabyte of text data in 45 seconds on an ordinary PC-class machine. This remarkable performance is achieved by integrating new data structures and gamma compression into TELLTALE. This paper also compares several different types of query methods such as tf.idf and incremental similarity to the original technique of centroid subtraction. The new similarity techniques give better execution-time performance, but at some cost in retrieval effectiveness.

1 Introduction

Scientists, researchers, reporters and the rest of humanity all need to find documents relevant to their needs from a growing amount of textual information. For example, it was recently reported that the World Wide Web has over 320 million indexable pages containing over 15 billion words [1], and is growing at an astonishing rate. As a result, information retrieval (IR) systems have become more and more important. However, traditional IR systems for text suffer from several drawbacks, including the inability to deal well with different languages, susceptibility to optical character recognition errors and other minor mistakes common on the WWW. This paper describes our successful efforts to apply traditional techniques to an n-gram based IR system, showing which methods work, which don't, and describing new techniques we implemented for n-gram based retrieval.

The TELLTALE information retrieval system [2] was developed to address these concerns. TELLTALE uses n-grams (sequences of n consecutive Unicode characters) rather than words as the index terms across which retrieval is done. By using statistical IR techniques, the TELLTALE system can index text in any language; the current version has been used unmodified for documents in English, French, Spanish, and Chinese. Additionally, character n-grams provide resilience against minor errors in the text by allowing matches on portions of words rather than requiring the entire word to match.

Previously, however, the TELLTALE system was unable to index large volumes of text. While traditional word-based IR systems have a number of tricks and tools at their disposal, many of these methods must be modified or discarded when used in n-gram based systems. By using the techniques described in this paper, we were able to construct an n-gram based IR engine that per-

mitted full-document queries against a gigabyte of text. Both the size of the corpus and the speed with which our methods operate allow such a query to complete in a few seconds on inexpensive PC-class hardware. These improvements represent a hundred-fold increase in corpus size over earlier versions of TELLTALE. Moreover, the compression techniques we adapted from word-based IR systems reduced the size of the index file from seven times larger than the text corpus to approximately half the size of the original text, a fifteen-fold improvement.

2 Background

Our work builds on a large body of research in information retrieval covering both traditional word-based IR systems and systems based around n-grams. In this section, we discuss some of the most relevant previous work. Of course, a full treatment of prior work in information retrieval would require a full book (if not more), and such texts exist [3,4].

2.1 Word-based information retrieval systems

There are many information retrieval systems in existence, but space prevents us from mentioning more than a small selection. We will only discuss four of them: INQUERY, MG, SMART and TELLTALE; the reader is referred to [3] and [4] for a more detailed treatment of information retrieval systems.

2.1.1 INQUERY

The INQUERY system is a product of the Center for Intelligent Information (CIIR) at the University of Massachusetts at Amherst. INQUERY [5] uses a probabilistic model of information retrieval, based on Bayesian networks [6]. In the probabilistic model, the estimated relevance of a document to a query is a function of the (estimated) probabilities that each of the various terms in the document occur in at least one relevant document but in no irrelevant documents. Because a Bayesian network is a graphical model that encodes probabilistic relationships among variables of interest, it makes a good framework for this style of model.

INQUERY builds two sorts of network, namely a document net and a query net. The document net is static for a given collection. Nodes representing documents are connected to nodes representing terms. The query net is constructed by connecting terms in the query to nodes representing how those terms should be combined in relevant documents. To perform retrieval, the system connects these two networks together, and calculates the conditional probability that the information needed exists in a given document. The system then ranks the documents by this probability.

For more information on INQUERY, we refer the reader to the CIIR web site. In particular, the demo page lets the user try INQUERY on a number of different corpora. (<http://ciir.cs.umass.edu/ciirdemo/>).

2.1.2 SMART

SMART [7], developed by Gerard Salton and his students at Cornell University, uses a vector space model for representing documents. SMART performs automatic indexing by removing stopwords (words that are too common to be useful in distinguishing between documents) from a predetermined list, stemming (the process of removing prefixes and suffixes from words in a document or query), and term weighting. Given a new query, SMART converts it to a vector, and

then uses a similarity measure to compare it to the documents in the vector space. SMART ranks the documents, and returns the top n documents, where n is a number determined by the user. SMART can perform relevance feedback, a process of refining the results of retrieval using a given query, based on the results of the retrieval process.

The disk space requirements for the indexed collection are roughly 0.8 times the space of the text version. This space includes a dictionary, display information, and both an inverted file and a sequential representation of the indexed documents.

2.1.3 MG (Managing Gigabytes)

MG (Managing Gigabytes) [4] is a full-text retrieval system that gives fast access to a collection of documents while requiring far less storage space than the original documents. It is a word-based information retrieval system, using words as the basic terms on which matching and lookup are performed. MG uses a vector space model that represents documents and queries as vectors of term frequencies. Stemming, stopword removal, and *tf.idf* term weighting are used to keep the set of index terms small, and to improve retrieval effectiveness.

The MG system has two main parts: a program that turns a collection of documents into a compressed and indexed full-text database, and a program that services several types of interactive queries for words that appear in the documents. By default, queries are Boolean and are constructed using a collection of terms linked by the Boolean operators AND, OR, and NOT. MG also supports ranked queries, which take a list of terms and use frequency statistics from the documents to determine which of the terms should be given the most weight when they appear in a document.

2.2 N-gram based information retrieval using TELLTALE

TELLTALE [2,8] provides full-text information retrieval from text corpora using a hypertext-style user interface. The most important difference between TELLTALE and the systems described above is that TELLTALE is n-gram-based while the others are word-based. Because of its use of n-grams, TELLTALE has some unique features including language independence and garble tolerance. These features and others will be discussed in this section.

2.2.1 N-gram basics

An n-gram [9] is a character sequence of length n extracted from a document. Typically, n is fixed for a particular corpus of documents and the queries made against that corpus. To generate the n-gram vector for a document, a window n characters in length is moved through the text, sliding forward by a fixed number of characters (usually one) at a time. At each position of the window, the sequence of characters in the window is recorded. For example, the first three 5-grams in the phrase “character string” are “chara”, “harac” and “aract”.

The concept of n-grams was first discussed in 1951 by Shannon [10]. Since then, n-grams have been used in many areas, such as spelling-related applications, string searching, prediction and speech recognition. (In speech recognition, phoneme n-grams are used instead of character n-grams.)

For a given corpus, a word-based system may need much fewer index terms than a system based on n-grams. A word-based system may need fewer index terms for several reasons. First, for a

given corpus and $n > 3$, the number of unique words is smaller than the number of unique n-grams, as shown for example in Figure 1. Second, stemming techniques can be used in word-based sys-

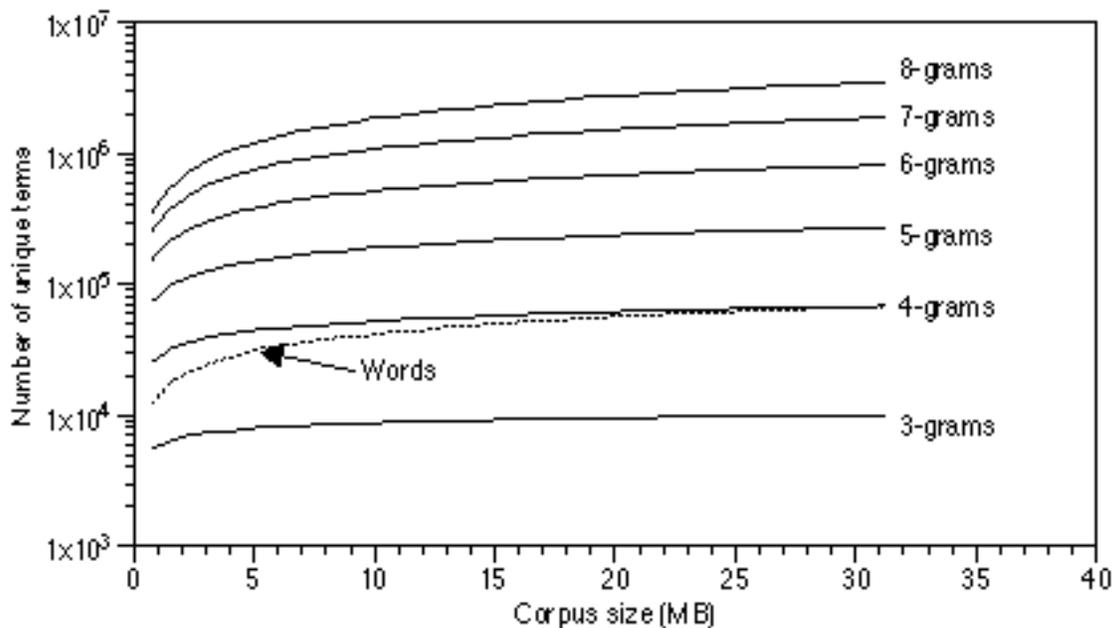


Figure 1. Number of unique terms (words and n-grams) in various subsets of the AP corpus.

tems. Stemming is a process that removes prefixes and suffixes from words in a document or query before placing them in the index. For example, a group of words with the same root, such as “walk”, “walked”, “walker” and “walking,” will be mapped to the same index term, namely “walk.” This reduces the size of the index set further, and also frees the user from needing to match the particular form of a word in a query and document. Finally, most word-based systems allow the removal of stopwords, which are words that occur often enough in every document so as to be of no use in distinguishing one document from another. Stopwords therefore are not helpful for retrieval, and are usually removed from queries as well as documents. Word based systems therefore have an advantage in terms of a smaller number of terms to be indexed.

On the other hand, there are several advantages to using n-grams. First, the system can become more garble tolerant by using n-grams as terms. If a document is scanned using OCR (Optical Character Recognition), there may be some missing or misread characters. For example, suppose “character” is scanned as “claracter”. A word-based system will not be able to match this word because it is misspelled, but an n-gram based system will still match the other n-grams such as “aract”, “racte”... and take their frequency into account. Pearce found that n-gram retrieval was still useful even when the garble rate was as high as 30% [3].

Second, by using n-grams the system can achieve language independence. In most word-based information retrieval systems, there is a level of language dependency. Stemming and stoplist processing are both language specific. Furthermore, one cannot assume that words are always separated by spaces. For example, in some Asian languages, different words are not separated by spaces, so a sentence is composed of many consecutive characters. Grammar knowledge is

needed to separate those characters into words, which is a very difficult task to perform. By using n-grams, the system does not need to separate characters into words.

An n-gram based system does not require a separate stopword removal process. Ekmekcioglu showed that stopword removal and stemming are give superior results in word-based systems, but do not improve results in n-gram based systems, because the use of n-grams provides similar benefits[11]. If a word or phrase occurs frequently in a collection, then so do its corresponding n-grams, and the n-gram term weighting and similarity calculations take this into account. Words derived from the same root word tend to generate many of the same n-grams, so a query using one form of a word will help cause documents containing different forms of that word to be retrieved. The sliding window approach allows us to capture n-grams corresponding to words, as well as pairs of words. The n-gram “of co”, for example, is the first n-gram in the phrase “of course.” An n-gram based system such as TELLTALE can therefore mimic stemming, stopword removal, and phrase indexing by using properties of n-grams.

In summary, most information retrieval systems are based on words, and employ stopword removal and/or stemming to reduce the number of index terms. N-gram based IR systems tend to have a larger number of index terms, but they provide a level of language independence and garble resistance that is important in many applications.

2.2.2 Document similarity computation

Similarity is the measure of how alike two documents are, or how alike a document and a query are. In the vector space model, this is usually interpreted as how close their corresponding vector representations are to each other. One way of determining this is to compute the cosine of the angle between the vectors.

In TELLTALE, each document is represented by a vector of n-grams. That is, a particular document i is identified by a collection of n-grams $ngram_1, ngram_2, \dots$. For each n-gram, a count c_{ik} records how many times $ngram_k$ occurred in document i . The frequency f_{ik} of $ngram_k$ is its count c_{ik} normalized by the total number of n-grams m_i in document i , or c_{ik}/m_i . The weight of each n-gram is the difference between f_{ik} and the average normalized frequency over all documents $a_k = \sum_i f_{ik}$ for $ngram_k$. This provides a weight for each n-gram in a document relative to the average for the collection. A document is then represented as a vector $\vec{d}_i = (d_{i1}, d_{i2}, \dots)$, where the individual elements $d_{ik} = (c_{ik}/m_i) - a_k$ have been normalized and the n-gram’s average value has been removed. The similarity between two document vectors \vec{d}_i and \vec{d}_j is then calculated as the cosine of the two representation vectors,

$$SIM_c(\vec{d}_i, \vec{d}_j) = \frac{\sum_k (d_{ik}d_{jk})}{\sqrt{\sum_k d_{ik}^2} \sqrt{\sum_k d_{jk}^2}} .$$

Since $\cos(\theta) = 1$ when the vectors are colinear and $\cos(\theta) = 0$ when the vectors are orthogonal, the similarity measure can be taken as the cosine of the angle between the document and query vector — the larger this cosine value, the greater the similarity.

2.2.3 Multilingual operation

The language independence of TELLTALE is achieved by its n-gram techniques, unique algorithms, Unicode [12] and display system based on Tcl/Tk [13]. As mentioned in Section 2.2.1, using n-grams can eliminate language-dependent features because the program need not know about grammar for stemming, stopwords, or even matters as simple as where to break individual words. In languages such as German, for example, words are often built from smaller words without including spaces between them. A German IR system would thus have to know how to break up a long word, or risk missing similarities between long compound words. For TELLTALE, however, this is not a problem because the text is broken into relatively short sequences of characters without knowledge of where individual words begin or end.

Second, the algorithms used in TELLTALE are independent of the language texts to be analyzed. We have found that our algorithms work well not only for English, but also for other languages such as Spanish and Chinese. We do not rely upon particulars of the English language to attain good retrieval accuracy.

Third, TELLTALE uses Unicode to represent a character. Unicode [12] is a 16-bit encoding standard in which each character is represented in two bytes. This encoding is necessary because many Asian languages require 16 bits for each character; one byte has no meaning in those languages. Thus, the use of Unicode in the algorithms is necessary to achieve language independence. While it may be necessary to convert a document from “native” format into Unicode, such conversion is mechanical. In this paper, storage requirements are expressed in units of bytes. Therefore, if a document consists of 500 Unicode characters, it requires 1000 bytes of storage. The user interface renders each (odd, even) pair of bytes as a single Unicode character.

TELLTALE’s fourth advantage is Tcl/Tk’s ability to use non-English fonts. This allows us to quickly build a system that has the ability to display a variety of fonts, such as Russian, Hebrew, and Chinese. Using this ability, the system can display documents in their native scripts.

2.2.4 TELLTALE interfaces

Because TELLTALE has a Tcl/Tk interface, implementing the user interface was relatively easy. A sample view of the interface is shown in Figure 2. The interface has several main areas. These include the main document window (A), the query text window (B), the document list (C), the status area (D), and various controls (E). Additionally, the interface supports the execution of Tcl/Tk commands via a small floating window.

When TELLTALE is first started, a user can initiate a search either by listing all documents and selecting the one he or she wants or by entering some text into the query window and finding documents “similar” to the entered text. In response, TELLTALE shows a ranked list of all documents that satisfy minimum similarity criteria specified by the controls. If the user clicks on one of the documents in the list, the full text of the document will be displayed in the document window. We refer to the interface as “hypertext style” since one can use mouse clicks to move from one document to another. Document similarity provides an implicit linking mechanism. The

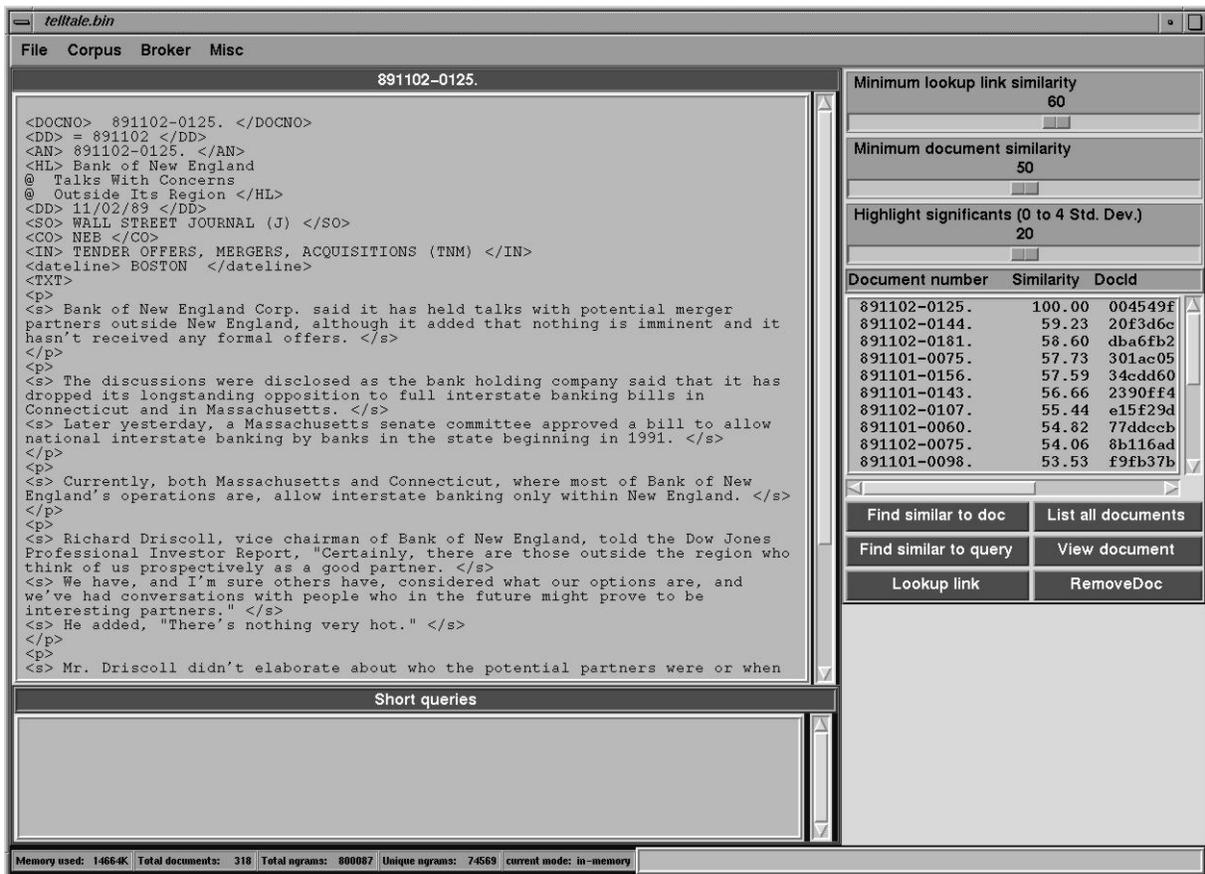


Figure 2. TELLTALE's graphical user interface.

searching process can be continued by either looking for documents similar to the one in the document window or by cutting some of the document text and pasting it into the query window. Of course, other text can be used in the query in addition to or instead of text from the displayed document. This method allows the user to incrementally approach the documents that match his or her needs. The user may also limit the documents selected by changing minimum similarity thresholds.

The display also includes status information about the TELLTALE system. This information includes corpus statistics such as the amount of text indexed, the number of unique n-grams, and the number of postings. (For each term, there is a list of documents in which that term occurs. This list is known as the *postings list*, and each entry in that list is a *posting*.) It also includes memory usage statistics, which are useful when trying to configure TELLTALE for best memory performance.

3 Approaches to large-scale retrieval in TELLTALE

Because TELLTALE is n-gram based and the number of n-grams in a document is in general much larger than the number of words in the same document, the index for TELLTALE is much larger than that of a word-based system. Thus, building a large-scale n-gram based retrieval sys-

tem is a technical challenge. In this section, we describe how we increased the capacity of TELLTALE to handle queries against gigabyte-sized corpora in a reasonably short time.

All of the performance measurements reported in this paper were made on a Silicon Graphics Origin200 with two 180 MHz MIPS R10000 processors, 256 MB of main memory, and 32 KB each of instruction and data cache. While this may seem an impressive machine, it is currently possible to purchase a more powerful machine relatively inexpensively from commodity PC vendors. Thus, we expect that our techniques will be applicable to those who can't afford large-scale computers as well as to those who can.

3.1 Textual data used in experiments

To allow practical comparison of various algorithms and techniques, we performed our experiments on real-world collections of data obtained from TIPSTER [14], a DARPA (Defense Advanced Research Projects Agency)-funded program of research and development in information retrieval and extraction. The TREC [15] (Text REtrieval Conference) is part of the TIPSTER Text Program, and provides a very large text data collection. Three types of text data from TREC are used in this paper: a selection of computer magazines and journals published by Ziff-Davis (ZIFF), the Associated Press newswire (AP), and the Wall Street Journal (WSJ). Here we use ZIFF1 to represent the collection from 1989's ZIFF, ZIFF2 to represent the text data from 1988, AP1 to represent the text data from 1989's AP, AP2 to represent the data from 1988 and WSJ to represent the data from 1989's Wall Street Journal.

	ZIFF1 (1989)	ZIFF2 (1988)	AP1 (1989)	AP2 (1988)	WSJ (1989)
Documents	75,029	56,903	83,719	78,789	12,046
Unique 5-grams	562,492	498,653	499,807	478,519	268,810
Total 5-grams	185,159,683	134,110,587	202,636,790	186,822,009	31,863,179
Size (MB)	257	180	260	240	40

Table 1. Statistics for the document collections.

Every corpus is composed of tens or hundreds of individual files, each of which averages one megabyte in length and contains one or more documents. Individual documents within a file are marked up with SGML (Standard Generalized Mark-Up Language) tags. The overall characteristics of the corpora on which we ran experiments are summarized in Table 1. Note that the figures for unique and total n-grams are calculated for $n = 5$; we used this value for n in all of our experiments. We chose $n = 5$ because this value of n gives reasonable retrieval performance (precision and recall) without requiring too much memory.

3.2 Data structures

The data structures used in TELLTALE are similar to those used in other information retrieval systems, but with modifications to make them efficient for managing tables of n-grams that are considerably larger than the word-based tables used elsewhere. Over the course of our work, structures evolved from relatively simple solutions to more advanced mechanisms that allow

TELLTALE to index gigabytes of textual data, a hundredfold increase over its original capacity as reported in [8].

3.2.1 In-memory data structures

The first set of data structures in TELLTALE are in-memory data structures similar to those used in traditional word-based IR systems. The three hash tables represent all of the information gathered from the raw text scanned into TELLTALE. There is one hash table for n-grams, one for document information, and a third detailing information about the files that contain the documents. The relationship between these three hash tables is shown in Figure 3. While all three data structures are crucial to TELLTALE, the n-gram hash table and associated postings lists consume by far the largest portion of memory.

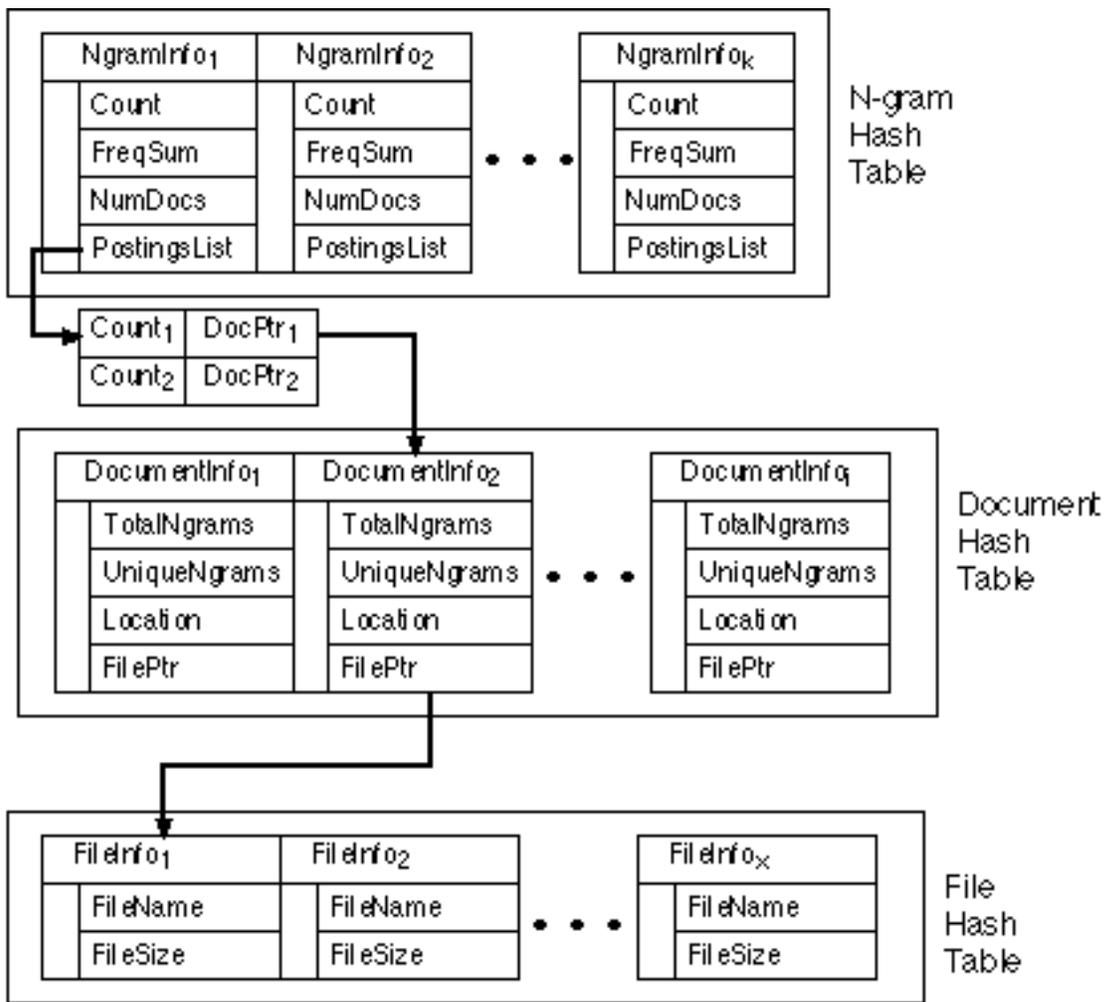


Figure 3. Relationships between hash tables in TELLTALE.

The file hash table provides a link between documents and the files that contain them. While it would be possible to fold this information into the document hash table, storing it separately results in a large memory savings at little cost because file names are long. For example, a corpus

with 500,000 documents at 2 KB apiece might pack an average of 500 documents into each 1 MB file. If file names average 60 bytes in length, the file table requires $60 + 4 = 64$ bytes of data and 4 bytes of overhead per file for a total of just 64 KB of storage. On the other hand, storing a file name with each document requires over 3 MB. Thus, large corpora benefit greatly from the savings provided by a separate file table. Additionally, this structure works well for systems that don't use traditional file systems. For example, a system might optimize performance for access via the WWW by consolidating references to a single URL together. If several documents are accessed from the same URL, then having a single pointer to that URL makes retrieval and caching simpler.

The document hash table contains information about each individual document. In addition to the usual information such as document length and location (file and offset), the document table contains a document serial number, which is allocated from a single integer incremented for each document scanned in. Thus, the document serial number is guaranteed unique in a given corpus. The document hash table also stores precomputed values for each document to assist in rapid similarity calculations. As will be described below, precomputed per-document values greatly reduce similarity calculation times for some similarity measures.

Additionally, each document in the hash table contains an identifier generated by cryptographic hash. In the current version of TELLTALE, this hash is generated by MD5 [16], although there is nothing to prevent switching to a different algorithm such as SHA [17]. This document ID is probabilistically unique across all corpora, with a chance of collision below 10^{-9} even for billions of documents. Thus, the document ID can be used to uniquely refer to a document in a massive environment that might contain tens or hundreds of TELLTALE engines. The document ID can also be used to remove duplicate documents from corpora; since the ID is based on the document's content, identical documents will have identical IDs. Memory usage for the document hash table is approximately 48 bytes per document, most of which is used to store the 128-bit (16 byte) document ID. In a system with 10^6 documents, this means that 48 MB will be required to store the document hash table.

The n-gram hash table is the central data structure in TELLTALE and, when the postings are included, the one that requires the most memory. This data structure is the one that is hardest to optimize for n-grams rather than words because of the far greater number of both unique n-grams in the corpus and unique n-grams in a document. For example, a typical 1 MB file from the WSJ corpus has around 300 documents with over 500,000 postings — one posting for each distinct n-gram in a document. When terms are words, however, the same file has around 60,000 word postings, a more than 8-fold reduction. It is this difference that makes it more difficult to build IR systems using n-grams rather than words as terms.

In TELLTALE, each posting contains a normalized frequency for n-gram k in document i , a pointer to document i , and a pointer to the next posting for n-gram k . Thus, each posting requires 12 bytes on a machine that supports 32-bit floating point numbers and 32-bit pointers. It is the space required for postings that consumes the lion's share of memory for corpora of almost any size. In typical documents, the number of unique 5-grams is about 65%-75% of the total number of 5-grams, so a 4 KB document will have 2500 - 3000 unique 5-grams, resulting in $12 \times 3000 = 36000$ bytes of storage. On the other hand, the word count for such a document will

total perhaps 800 words with perhaps 400 different words — a reduction of an order of magnitude.

To obtain good performance on the cosine similarity using these data structures, we broke the similarity formula down as follows:

$$\begin{aligned}
 SIM_c(\vec{d}_i, \vec{d}_j) &= \frac{\sum_k (d_{ik} d_{jk})}{\sqrt{\sum_k d_{ik}^2} \sqrt{\sum_k d_{jk}^2}} \\
 &= \frac{\sum_k ((f_{ik} - a_k)(f_{jk} - a_k))}{\sqrt{\sum_k d_{ik}^2} \sqrt{\sum_k d_{jk}^2}} \\
 &= \frac{\sum_k (f_{ik} f_{jk}) - \sum_k (f_{ik} a_k) - \sum_k (f_{jk} a_k) + \sum_k a_k^2}{\sqrt{\sum_k d_{ik}^2} \sqrt{\sum_k d_{jk}^2}}
 \end{aligned}$$

Note that, in the final equation, all of the terms with the exception of the first term in the numerator can be precalculated. The remaining term is non-zero only when a term appears in both the query and a document in the corpus. Thus, we can precompute all of the “constant” expressions in the formula for each document, and only need compute the sum of the term frequencies on the fly. Because there are relatively few n-grams in common between any pair of documents, this calculation can be done quickly once the precomputation time has been invested.

Based on our experiments with sample data, scanning in 10 MB of text data requires 88.9 MB for the resulting data structures. Using the data structures described above, TELLTALE can compute the similarity against a 1 KB document in two seconds. The performance is good, but too much memory is consumed.

The problem of memory consumption becomes worse as the corpus grows. In particular, the size of the posting list grows dramatically while the other data structures grow considerably more slowly. This is shown in Figure 4, which displays both the total memory consumed by each data structure and the percentage of memory consumed by each data structure for varying corpus sizes using 5-grams. As Figure 4 shows, by far the largest consumer of space is the postings list. Even for a small corpus of 1 MB, the postings list consumes 75% of the space. For larger corpora, the contribution of all other data structures shrinks further. By the time the corpus has reached 1 GB, the postings list consumes over 6 GB of memory, while all other data structures combined use less than 100 MB, or 1/60th the space. This is hardly unexpected — the number of unique n-grams in a corpus grows slowly after the corpus reaches a certain size because the number of unique n-grams in English (or any other language) grows rapidly for the first few megabytes of text but considerably more slowly for additional text. Moreover, some combinations (such as “zzyqv”) are unlikely to occur in any documents in a corpus, although this sentence itself shows that *any* n-gram is possible in any document in a given language. In essence, the first few documents define the “vocabulary,” and later documents add few new n-grams to it. However, the

number of postings grows linearly in the number of documents, and consumes far more space than document information or file information, both of which also grow linearly.

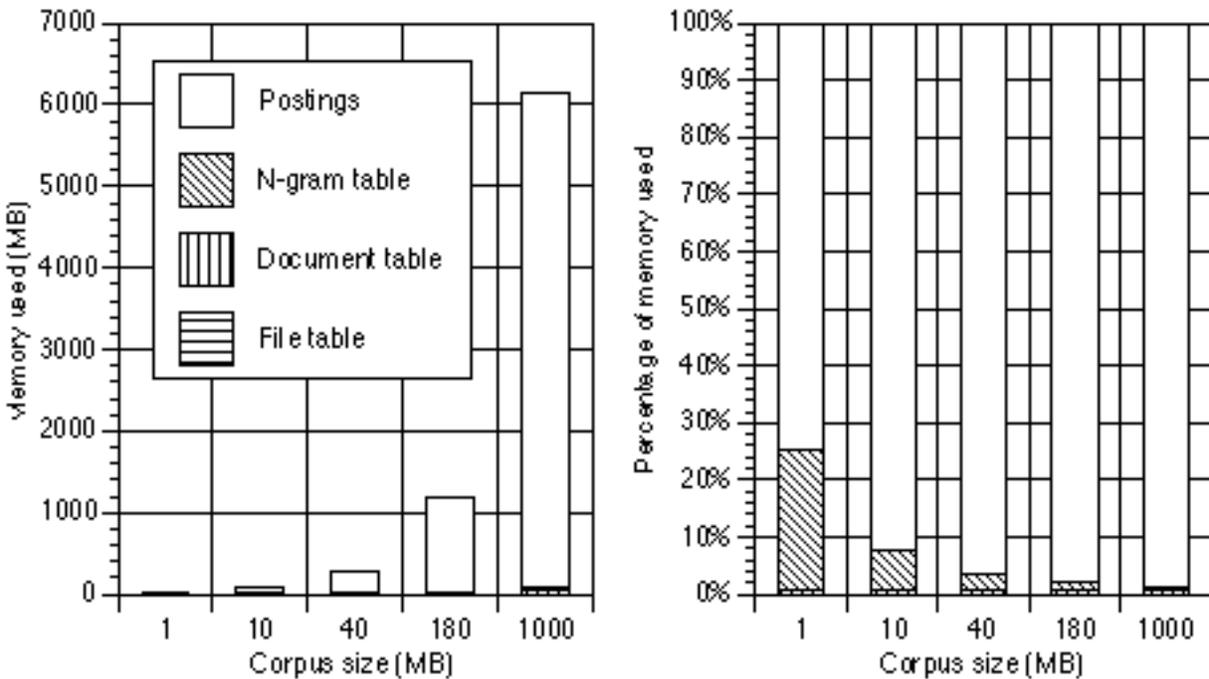


Figure 4. Space consumed by different data structures.

Because the postings list was clearly the largest impediment to scaling TELLTALE to handle gigabytes, we spent most of our effort optimizing its usage. The following sections describe our efforts.

3.2.2 On-disk data structures

The on-disk version of TELLTALE was developed to cope with the limitations of memory space by moving the postings list from memory to disk. Only the main data structures — n-gram hash table, document table, file table — are kept in memory. To accomplish this, TELLTALE generates an on-disk file to record all the information, including a posting list from the files scanned in.

We made several modifications to the original in-memory data structures to handle the on-disk data structures. First, we added variables to each n-gram's entry in the n-gram table to indicate where the postings list is stored in the on-disk file and how long it is. This is only a minor change, yet it increases the memory requirements by over 6 MB for a 1 GB corpus. However, it also allows TELLTALE to rapidly access postings lists on disk with the penalty of just a single disk seek.

The other change we made was to convert pointers in the in-memory data structures into integers, adding a table to translate from the integers into actual entries for files or documents. This change is necessary to allow the data structures to be stored to disk and retrieved, possibly with different memory addresses. We considered using the document ID, which was generated by hashing the document text, for this purpose. However, this ID requires 16 bytes of memory, which is too long

to include in every on-disk posting. Instead, we used the serial number assigned to each document. While this number is not unique across corpora, it is unique within a single on-disk file. While its use makes merging corpora together somewhat more complex, it saves a great deal of space; thus, we chose this option. A similar tactic was used to identify entries for particular files. The resulting structures are identical to those shown earlier in Figure 3, except that integers rather than pointers are used to link the tables together and the postings lists are stored on disk.

We chose to allow TELLTALE to operate in one of three modes: `memory`, `update` or `ondisk`. The `memory` mode is basically identical to the original version of TELLTALE; the only change is that the new version can load an on-disk corpus into memory and subsequently conduct similarity searches on it. The other two modes deal with building (or changing) the index for a corpus, and querying such a corpus, respectively. The `update` mode allows the system to create or update on-disk files. In this mode, TELLTALE can record information about documents that have been scanned in. The `ondisk` mode allows the user to issue queries against a corpus which has already been indexed in `update` mode. In `ondisk` mode the system loads the document hash table, file hash table and n-gram hash table header into memory. However, the posting list is kept on disk, and each bucket of the posting list has table must be loaded into memory before it is used. The format of the on-disk file is shown in Figure 5.

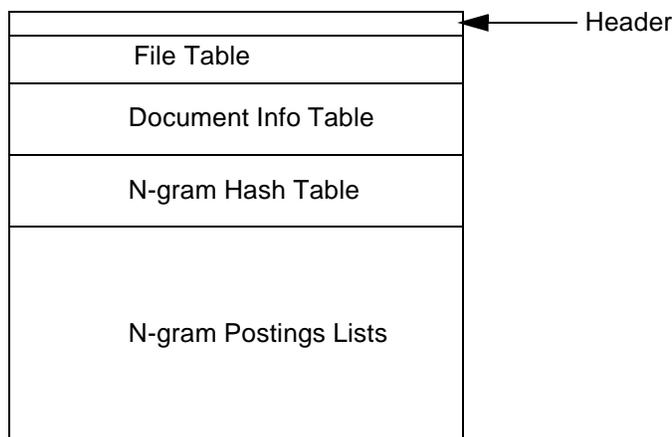


Figure 5. Format for an on-disk corpus index file.

Building an index is a memory-intensive process. On a machine with 128MB of memory, for example, we can only index 10 MB of raw text at a time. Indexing a larger corpus is therefore an iterative process, in which text is scanned in 10MB chunks, indexed, and then the index is written to disk and the next chunk of text read in. We use a Tcl script to automate the generation of a set of on-disk files. We then use the TELLTALE function `mergecorpus` to merge this set of on-disk files into one big on-disk file. Using this method, we can generate a 1 GB on-disk file containing all the information for about 300 MB of raw text data. After the big on-disk file is generated, the `usecorpus` command can load the file table, document table, and n-gram hashtable header into memory. We can then issue queries against this large corpus without having the memory space to read the corpus in its entirety into memory.

The file header records the general information for this file. It includes the number of corpus files, documents and n-grams in this text data set. It also contains the start offset for the other parts of the corpus index — file information, document information, n-gram index and n-gram postings lists. Additionally, several per-corpus summary values are stored here. The remaining sections of the on-disk file contain copies of the data stored in memory. However, the on-disk version must use integers rather than pointers for internal references such as a “pointer” from a document’s information to the structure describing the file that contains it. This conversion requires temporary tables on reading and writing the data, but does not require permanent data structures.

The n-gram hash table is split into two pieces in the on-disk file, one that contains the “header” for each n-gram bucket and the other listing all of the postings for the n-gram. This split is done so that the headers can be easily loaded in by a sequential read while the postings buckets are read from disk on demand. At 32 bytes per header, even a 1 GB text corpus indexed by 5-grams requires only 32 MB for the n-gram header hash table. This contrasts with the several gigabytes required for the postings list.

3.2.3 Performance

Using this naive strategy without compression, the on-disk file consumes 4.5-7 times the size of the raw text data. When the amount of raw text is small, this rate is even larger because the table of unique n-grams dominates the size of the on-disk index. As the corpus grows, however, the number of unique n-grams grows much more slowly, allowing the postings lists to dominate the space used for medium and large corpora.

We generated a 177 MB on-disk file representing the information from 40 MB of text data from the *Wall Street Journal*. Although the index requires a good amount of disk space, it increases the capability of TELLTALE from 10 MB to 40 MB of text data with the same amount of processor memory. The performance of similarity is also acceptable, though somewhat slow. However, we had to reduce the size of the on-disk index if we were to be able to perform retrieval on gigabyte corpora. The following section describes our compression techniques.

3.3 Compression

A large on-disk file not only takes up disk space, but also slows down the similarity calculation because of the time it takes to read posting lists from disk. Since I/O is very slow compared to CPU instructions, compressing the posting list results in two benefits: lower disk storage requirements, and faster similarity computations due to reductions in the time needed to read a bucket.

3.3.1 Strategy

The original on-disk file contained postings lists composed of a pair of numbers for each document in which the n-gram occurred: an integer identifying the document containing the n-gram and the normalized frequency for the n-gram in the document. This strategy required 8 bytes for each posting, 4 bytes each for the document number and a floating point number for the frequency.

First, we converted all of the n-gram frequencies into integers by storing the actual count of n-grams in a document rather than the normalized frequency. Since we already store the number of n-grams in a document, it is easy to regenerate the normalized frequency from the n-gram count

and the size of the document. This change allows us to use standard integer compression algorithms on our postings lists.

Second, we noted that if we sorted the postings in a list by document number, we could store the difference between an individual posting's document number and that of the previous posting. These gaps are smaller in magnitude than the document numbers themselves, and have the additional desirable property that they are smallest for large postings lists with many entries.

These two strategies can be used to greatly reduce the size of an on-disk index for n-gram based information retrieval. Moreover, compression can be more effective for n-grams than for words because the distribution of term frequencies is more skewed for n-grams than for words. For example, Figure 6 shows the distribution of term frequencies for 5-grams in the combined 1 GB corpus described in Table 1 alongside the distribution of integers describing the "gap" between document numbers in postings.

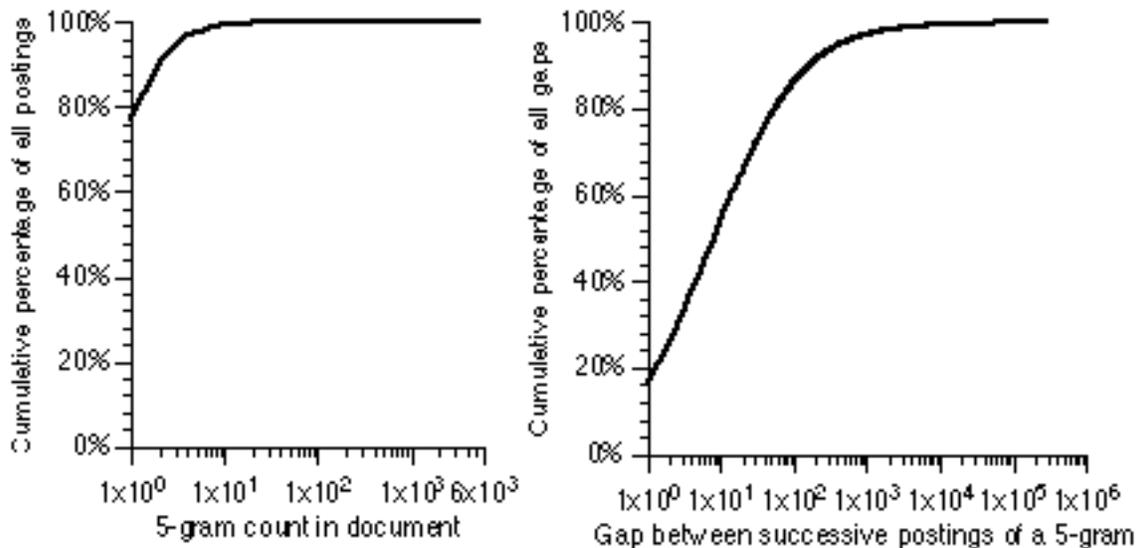


Figure 6. Distribution of 5-gram frequencies and document number gaps. There are a total of 4.62×10^8 postings in this corpus, which contains close to 1 GB of text.

The data in Figure 6 show that posting counts have the greatest potential for compression, though there is also some opportunity for compression of document number gaps as well. The default representation of integers in TELLTALE is 4 bytes, allowing values up to $2^{32}-1$. However, the count of a particular n-gram in a single document is usually very small; few n-grams have high frequency within a document. In the corpus we studied, 97.77% of all postings had counts of five or less, while over 77% had a count of exactly one. Based on this finding, we knew that the vast majority of counts could be stored in far less than the standard four bytes of an integer.

We also looked at the distribution of document serial number gaps, also shown in Figure 6. This graph shows that most serial number gaps are also relatively small. However, the curve falls off far more slowly than that for posting counts. Half of the gaps were or less, and 92.6% were 255 or less. This distribution means that over 92% of all gap values could each be stored in a single byte rather than the four bytes required by the default representation.

Based on these findings, we implemented two different compression schemes in TELLTALE. The first was simple to implement and provided reasonable compression for both n-gram counts and document serial number gaps. However, it was considerably below optimal; the `gzip` utility was able to compress the indices by a factor of two. We then switched to gamma compression, yielding files that were approximately the same size as the result of using `gzip` on the original files that used the first compression scheme.

3.3.2 Simple compression algorithm

Based on the statistics discussed in Section 3.3.1, we first considered a simple compression algorithm that saves space for small numbers. We used a single byte to represent numbers from 0 to 2^7-1 , two bytes for numbers 2^7 to $2^{14}-1$, and four bytes for numbers in the range 2^{14} to $2^{30}-1$. These numbers were stored in the format shown in Figure 7.

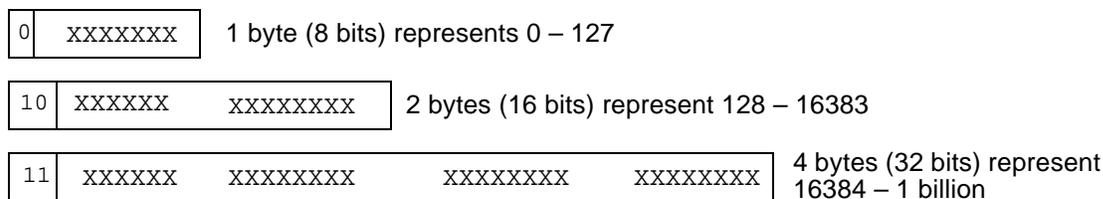


Figure 7. Number formats for the simple compression scheme.

We got good compression results from this scheme. We generated an on-disk file containing all of the documents in the ZIFF1, ZIFF2, AP1, and AP2 corpora using this compression. The combined corpus has 960 MB of raw text, including 294,440 documents and 889,125 unique n-grams, resulting in an on-disk file requiring 1.085 GB of storage. This provided better than a factor of four compression relative to the uncompressed on-disk file. Additionally, query response time improved greatly. After the system loads the in-memory tables for the large corpus, it can compute, sort, and display the similarity between the query and every document in the corpus at the rate of 5.5 minutes per 1000 characters in the query.

This simple compression algorithm showed that we can process n-gram queries against 1 GB of text data quite well. However, we did not achieve as much compression as we could. We noticed that `gzip` was able to compress our on-disk files by a factor of two, suggesting that we could devise a compression scheme that approached, or even surpassed, this level of compression. Doing so would reduce the size of on-disk indices and improve performance by reducing the amount of data that must be read for each query. Balanced against this reduction is the increased CPU time necessary to encode and decode a more complex compression scheme.

3.3.3 Gamma compression

Our initial experiments showed that compression was very effective at reducing resource requirements, but that we could achieve additional gains with more efficient compression schemes. We considered several standard schemes for our corpus, including unary code, gamma compression, and delta compression. All of these compression methods are relatively straightforward to implement and could provide significant improvement over our initial scheme.

We first considered unary code [4]. In this code, an integer $x \geq 1$ is coded as $x - 1$ one bits followed by a zero bit. For example, the unary code for 4 is 1110. This represents small integers such as n-gram counts within a document well, but is very inefficient at representing larger integers such as document serial number gaps.

We next examined gamma compression, which represents an integer x as two parts: a unary code for an integer m followed by a k -bit binary value y . The value for k is determined by taking the m th element of a vector K of integers that is constant across all compressed values (i.e., constant for a particular compression scheme). For a vector of integers $K = \langle k_0, k_1, \dots, k_n \rangle$, the value of a

representation my can be calculated as $\left(\sum_{i=0}^{m-1} 2^{k_i} \right) + y + 1$. Table 2 shows examples of representa-

tions using two different K vectors. Note that the largest number representable in each scheme is limited by the largest integer in the K vector. For this reason, gamma compression implementations often contain a large terminal value to handle the occasional integers larger than the range in which most values fall.

Prefix	$K = \langle 0, 1, 2, 3, 20 \rangle$		$K = \langle 1, 3, 5, 7, 15 \rangle$	
	k	Maximum value	k	Maximum value
0 (0)	0	$2^0 = 1$	1	$2^1 = 2$
1 (10)	1	$1 + 2^1 = 3$	3	$2 + 2^3 = 10$
2 (110)	2	$3 + 2^2 = 7$	5	$10 + 2^5 = 42$
3 (1110)	3	$7 + 2^3 = 15$	7	$42 + 2^7 = 160$
4 (11110)	20	$15 + 2^{20}$	15	$160 + 2^{15}$

Table 2. Sample gamma compression representations.

For compression, the algorithm first subtracts 1 from the value being encoded. It then searches through the table and finds the smallest “Maximum value” entry greater than the number being compressed. That Maximum value entry corresponds to component k_i of K , which indicates how many bits are needed for the binary portion. The unary prefix m is read directly from the table, and the binary portion y is obtained by subtracting the previous Maximum value entry (corresponding to component k_{i-1} of K) from the value to be compressed. For example, compressing the value 18 using the scheme $K = \langle 1, 3, 5, 7, 15 \rangle$ (on the right side of Table 2) would be done by first subtracting 1, yielding the value 17. Looking in the table, this requires a unary prefix of 2 (110), since 42 is the smallest Maximum value greater than 17. Since the previous Maximum value entry is 10, the binary portion is $17 - 10 = 7$ expressed in 5 bits (00111). Thus, the final representation is 11000111.

The decoding process is also relatively simple. The unary prefix m is extracted and used as the index i into the table. The previous Maximum value, corresponding to k_{i-1} , is then added to the k -bit binary portion y that follows the unary prefix using the k found in the table, and 1 is added to

the result. For example, decoding the compressed value 10110 using the scheme on the right side of Table 2 is done by looking up 2 (10) in the table, finding $k=3$ and the previous Maximum value is 2. The uncompressed number is thus $2 + 6(110) + 1 = 9$.

The delta code is a derivative of the gamma code that uses the gamma code, rather than unary, to encode the prefix. However, the delta code is more complex, and not as efficient for very small numbers such as those found in n-gram frequency counts. There are other, more advanced compression techniques, but these two are the most commonly used algorithms. Since integers in this system are usually not big and gamma coding is easy to implement, we picked gamma coding to compress the posting list.

After selecting gamma compression, we had to choose the best vector to use to compress our integers. To do this, we ran several experiments against the data shown in Figure 6 to compute the amount of space that would be required using several different vectors. The results of some of our experiments are shown in Table 3. As this table shows, optimal compression for n-gram counts and document gaps were achieved with different vectors. To simplify implementation, we chose the first vector in Table 3 as our compression scheme, although future versions of TELLTALE may use different vectors to compress different value sets. Even with our choice of a single vector, however, we were within 5.5% of the space required by the optimal two-vector compression scheme. Moreover, the optimal scheme for a given corpus can only be discovered by experimentation such as that we performed on our 1 GB corpus. Since the optimal choice for any particular corpus may be different, we chose a simple scheme that performed well.

Vector	N-gram counts (MB)	Document gap (MB)	Total (MB)
<0,1,2,3,4,5,6,7,8,9,10,11,12,14,16,18,20,28>	87.7	409.1	496.8
<0,1,2,3,4,6,8,10,12,14,16,18,28>	87.7	409.2	496.9
<0,0,1,2,3,4,5,6,7,8,10,12,14,16,18,28>	83.2	435.0	518.2
<0,0,0,1,2,3,4,5,6,7,8,10,12,14,16,18,28>	82.3	460.2	542.5
<0,0,0,0,1,2,3,4,5,6,7,8,10,12,14,16,18,28>	82.1	484.7	566.8
<0,0,2,4,6,8,10,12,14,16,18,28>	86.4	412.0	498.4
<0,2,4,6,8,10,12,14,16,18,28>	96.6	388.1	484.7
<0,3,6,9,12,15,18,21,28>	106.9	395.2	502.1
<0,2,3,6,9,12,15,18,21,28>	96.0	398.5	494.5
<0,1,2,4,6,9,12,15,18,21,28>	88.2	408.2	496.4

Table 3. Index sizes for various gamma compression vectors.

Using our gamma compression scheme with the vector in the first line of Table 3, we generated an on-disk index file covering ZIFF1, ZIFF2, AP1, and AP2. The raw text from these files was 960 MB, but the on-disk index consumed only 647 MB. While the table entry suggests that under 500 MB would be necessary, the table does not include additional data structures necessary to store document info and the n-gram headers; these structures make up the additional 150 MB.

Since they are only read in at startup, however, we decided not to attempt to compress them. Our experiment gave a compression ration of 0.67, which is nearly as good as `gzip`.

Gamma compression also improves query performance by reducing the amount of data that must be read for a single query. After loading the in-memory information for the 960 MB of text, TELLTALE can compute and sort about 300,000 documents' similarity result at the rate of 50 seconds per thousand characters in the query. This is considerably faster than our original compression scheme, as can be seen in Figure 8.

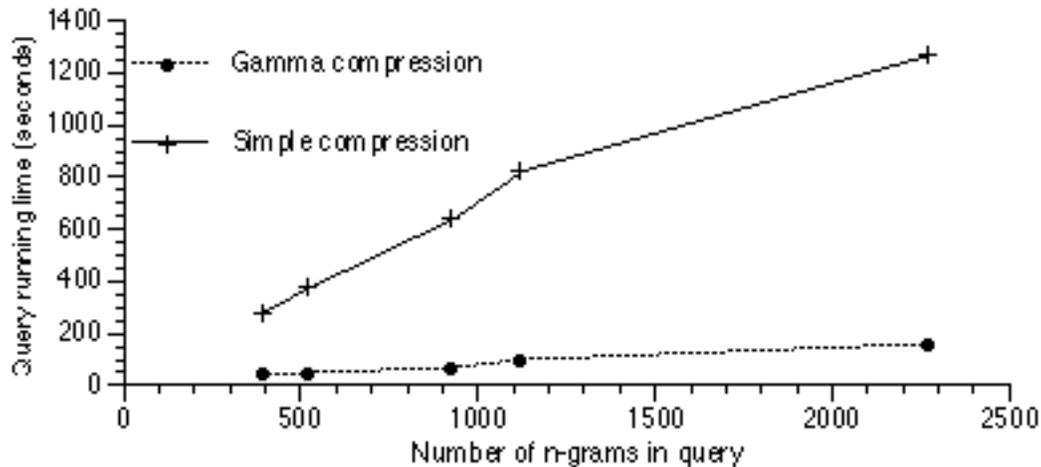


Figure 8. Retrieval time for different query sizes and compression methods.

3.4 Handling gamma compressed postings lists in memory

The dropping price of memory has made it possible to purchase large amounts of memory at a reasonable price. With the help of gamma compression, the compressed postings list is small enough to be loaded into main memory if allocation is handled intelligently. Since about 1 GB of raw text data can be indexed in under 700 MB, we can handle 1 GB of text data in less than 750 MB of main memory, allowing for a small amount of overhead. Doing so will improve the performance a great deal by eliminating disk I/O during a query.

The major difficulty with handling the compressed postings lists in memory is coping with the many bucket capacities necessary — some postings lists will be only a few bytes long, while others may require many kilobytes. Additionally, these buckets must grow dynamically as new documents are indexed. These requirements are best met using lists built from fixed size “chunks” of space connected in a linked list. The overhead for this scheme is relatively small — fixed size chunks that can hold 32 bytes require only four bytes of pointer overhead for an overhead of 12.5%. In addition, fixed size chunks waste some space because part of the last chunk is unused. On average, this will waste half of a chunk per n-gram, 16 bytes in our system. Thus, total overhead for a system that scanned in the 1 GB AP-ZIFF corpus would be 14 MB for unused chunk space and about 62.5 MB for pointers. In future versions of TELLTALE, we will attempt to reduce this overhead by allowing two or three different chunk sizes for maximum efficiency.

Operation using in-memory compressed postings lists is similar to that using uncompressed postings lists, but with the additional step that postings lists must be compressed before they are permanently stored. Also, postings lists must be uncompressed before they are used in similarity

calculations. While this technique uses less memory than uncompressed postings lists, it is somewhat slower because of the time needed to uncompress a postings list. A 200 MHz Pentium is capable of decompressing two million integers a second, and while this seems an impressive number, similarity calculations may need to process ten million postings or more. Thus, decompression time contributes significantly to similarity calculation time.

By using in-memory gamma compression rather than uncompressed postings lists, TELLTALE can reduce its memory usage by a factor of four or more, as Figure 9 shows. Both versions of TELLTALE keep the entire postings list as well as the other data structures, including the document hash table and n-gram hash table, in memory. The only difference between the two is that the gamma compressed version uses a great deal less memory for larger amounts of text. Note, however, that the original TELLTALE uses slightly less memory for small corpora; this occurs because the overhead for the gamma compression version of TELLTALE is slightly higher. However, this higher overhead is more than recovered as the amount of indexed text increases.

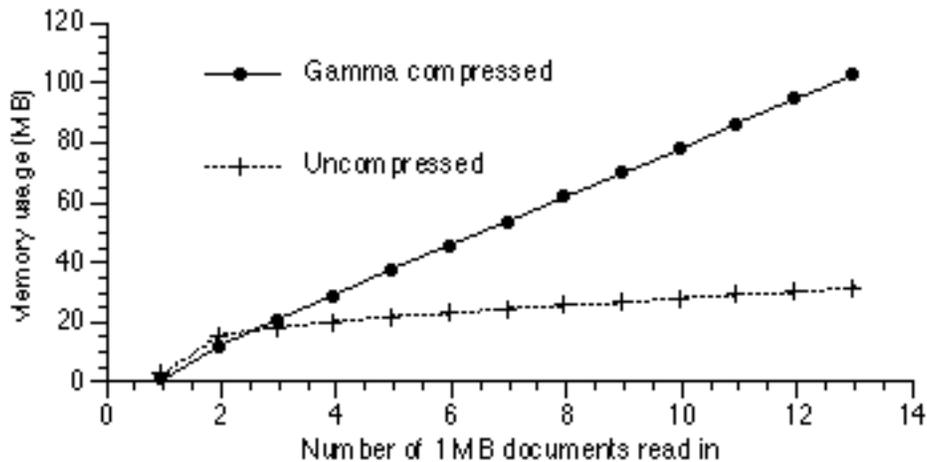


Figure 9. Memory usage for original and in-memory gamma compressed TELLTALE.

The in-memory gamma compressed version also provides increased speed compared to the on-disk version. Comparisons with the original, uncompressed in-memory version are less relevant because the original version can only handle very small corpora. Thus, we focused our attention on the relative performance of the gamma compressed postings lists on disk and in memory. We ran queries against the 257 MB ZIFF1 corpus, which contains 75,029 documents, 562,492 unique n-grams, and a total of 185,159,683 postings. We were limited to a corpus of this size because the machine on which the queries were run, a two processor SGI Origin 200, had only 256 MB of physical memory, and the use of virtual memory would have resulted in unacceptably slow performance due to excessive paging. When the entire ZIFF1 corpus was loaded into memory, it used 210 MB of memory, leaving the rest for operating system use. As shown in Figure 10, in-memory gamma compression is twice as fast as on-disk gamma compression, though the difference is not as large as we had expected. This may be due to the high performance XFS file system [18] used on the SGI server on which the experiments were run.

As the preceding experiments have shown, gamma compression performs well for n-gram-based IR. However, we had to adjust the gamma compression vectors to best compress the postings lists generated for n-grams because document gaps and, particularly, occurrence count distributions

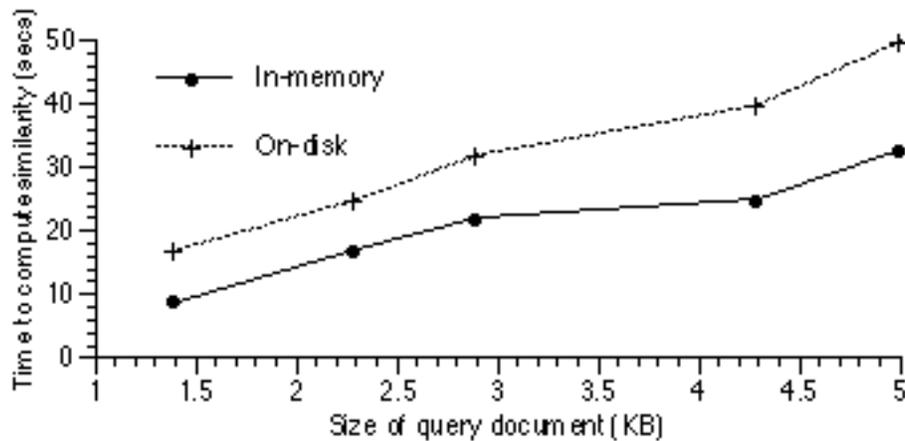


Figure 10. Performance comparison of in-memory and on-disk version with gamma compression.

are different between words and n-grams. Using these techniques, we expanded TELLTALE's capability from around 10 MB to over 1 GB while maintaining good query performance.

4 Exploring different similarity mechanisms

Because this is the first n-gram-based information retrieval system capable of handling a gigabyte of text, we were able to perform several experiments on using different document similarity schemes that were previously done only on relatively small corpora [19]. While our experiments were not extensive, they showed the effects of eliminating common n-grams from queries and the resulting performance gains. We also conducted some basic experiments on the effectiveness of tf.idf similarity using n-grams rather than words.

4.1 Incremental similarity calculations

We mentioned earlier that word-based IR systems remove stopwords because they don't help with retrieval and take up index space. Incremental similarity is based on the related idea that an n-gram which occurs in most documents is not important for retrieval. If an n-gram occurs in a large percentage of documents, it must be a common term and not helpful in distinguishing one document from another. At the same time, such n-grams have long postings lists, requiring TELLTALE to spend a relatively long time reading and uncompressing their postings lists. Can we ignore common n-grams without reducing retrieval performance? To test the effectiveness of ignoring common n-grams, we modified TELLTALE to include a threshold t ($0 \leq t \leq 1$) such that TELLTALE ignores n-grams that occur in more than $t \times numDocs$ documents, where $numDocs$ is the total number of documents in the corpus. Such n-grams may, but do not necessarily, correspond to stopwords. We therefore allow the TELLTALE user to adjust the threshold as they wish, trading off the risk of discarding meaningful n-grams and the expense of longer similarity computation time. The speedup from reducing the threshold is illustrated in Figure 11. As expected, lower thresholds require less computation time, but the improvement is not dramatic. Even for a threshold of 50%, the maximum improvement time is from 33 seconds to 28 seconds, or 85% of the original time. Because there are many more unique n-grams than words, there are fewer n-

grams that occur in most of the documents. Thus, omitting the most common ones does not result in a large reduction in memory usage or improvement in retrieval speed.

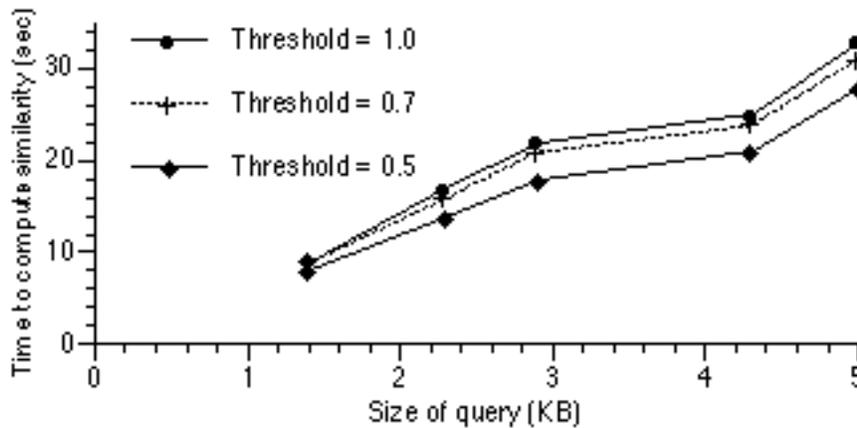


Figure 11. Performance for incremental similarity with different thresholds.

We next produced a simple precision-recall graph for reduced n-gram frequency thresholds. Since we did not have relevance judgments for queries on our corpus, however, we used approximate judgments based on the results from queries using a non-reduced corpus. Nonetheless, the graph in Figure 12 shows that, as expected, lower thresholds result in lower precision and recall. Given the relatively small improvement in retrieval speed, we believe that eliminating common terms from similarity computations may not be as effective for n-grams as it is for words.

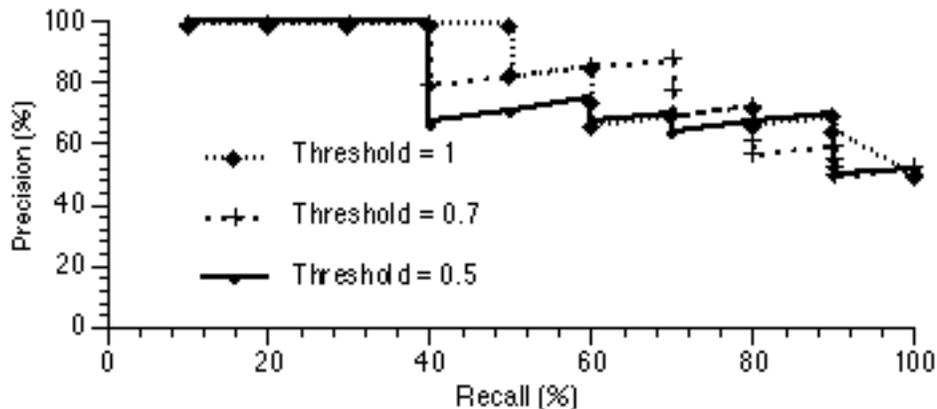


Figure 12. Recall-precision curve for different thresholds.

4.2 TF/IDF similarity

Another standard weighting system used in word-based systems is tf.idf (Term Frequency/Inverse Document Frequency) [19]. It uses the term frequency within a document times the log of the total number of documents over the number of documents containing the term. TELLTALE includes support for tf.idf weighting. Our experiments showed that tf.idf similarity can be calculated at about the same speed as the original TELLTALE similarity. Thus, if future work shows

that tf.idf yields more accurate similarity measures when using n-grams, as suggested by [19], TELLTALE will be able to support it with little loss in performance

5 Conclusions and future work

We greatly expanded TELLTALE's capacity and improved its performance, without compromising its ability to handle multilingual or garbled documents. It is these advantages, combined with an ability to perform retrieval using full documents rather than relatively short queries, that make TELLTALE a useful tool. However, there is still more work to be done. We hope to perform a more complete study of the tradeoffs between different similarity measures using n-grams rather than words. Because TELLTALE is the first system using n-grams that can handle a gigabyte of text, we hope to be able to show that n-grams are comparable to words as indexing terms in terms of retrieval effectiveness on large corpora. We are also performing experiments in using TELLTALE to index collections in non-English languages ranging from European languages such as French and Spanish to ideogram-based languages such as Chinese. Our preliminary results are promising, but more investigation needs to be done.

We used $n=5$ throughout this study. The larger the value of n , the larger the index required to store the corresponding (larger) number of unique n-grams. Longer n-grams tend to have more self-evident meaning, so ultimately the choice of n should be based on a trade-off between index size and meaning per n-gram. There are applications where smaller values of n are acceptable, and it remains to be shown how retrieval performance on a large corpus is affected as n increases.

We have demonstrated that it is possible to build an information retrieval engine using n-grams rather than words as terms that can handle gigabyte-sized corpora. The TELLTALE IR engine adapts techniques that have been used for word-based systems to n-gram-based information retrieval, making adjustments as necessary to account for the different term distributions exhibited by n-grams. Because there are many more unique n-grams than words in a document, TELLTALE must cope with many more unique terms and at least an order of magnitude more postings to allow indexing of a text corpus. By modifying standard techniques, we demonstrated a system that provides good performance on the large corpora that computers will be called upon to index. These techniques can also be used on other systems where performance and scalability are critical to better use of system resources and larger scale and faster processing.

Acknowledgments

The authors are grateful to the many people who contributed to this work by giving us feedback and suggestions. These include Claudia Pearce and Bill Rye at the Department of Defense and David Ebert at UMBC.

This work was supported in part by the U.S. Department of Defense.

References

- [1] Steve Lawrence and C. Lee Giles, "Searching the World Wide Web," *Science* **280**(3), 3 April 1998, pages 98 - 100.

- [2] Claudia Pearce and Ethan Miller, "The TELLTALE Dynamic Hypertext Environment: Approaches to Scalability," in *Advances in Intelligent Hypertext*, J. Mayfield and C. Nicholas, eds. Lecture Notes in Computer Science 1326, Springer-Verlag, October 1997, pages 109 - 130.
- [3] Ricardo Baeza-Yates and Berthier Ribeiro-Neto, *Modern Information Retrieval*, ACM Press, 1999.
- [4] Ian H. Witten, Alistair Moffat, and Timothy C. Bell, *Managing Gigabytes*, second edition, Morgan Kaufmann, 1999.
- [5] James P. Callan, W. Bruce Croft, and John Broglio, "TREC and Tipster experiments with INQUERY," *Information Processing and Management* **31**(3), 1995, pages 327 - 343.
- [6] D. Heckerman, "A tutorial on learning with Bayesian networks," Technical Report MSR-TR-95-06, Microsoft Research, March 1995 (revised November, 1996).
- [7] G. Salton and M.J. McGill, "The SMART and SIRE Experimental Retrieval Systems," in *Readings in Information Retrieval*, Karen Sparck Jones and Peter Willett, eds., Morgan Kaufmann, 1997, pages 381 - 399.
- [8] Claudia Pearce and Charles Nicholas, "TELLTALE: Experiments in a Dynamic Hypertext Environment for Degraded and Multilingual Data," *Journal of the American Society for Information Science*, April 1996, pages 263 - 275.
- [9] Marc Damashek, "Gauging Similarity with n-grams: Language-Independent Categorization of Text," *Science* **267**, 10 February 1995, pages 843 - 848.
- [10] C. E. Shannon, "Prediction and entropy of printed English," *Bell System Technical Journal* **30**, pages 50 - 64.
- [11] F. Cuna Ekmekcioglu, Michael F. Lynch, and Peter Willett, "Stemming and N-gram Matching For Term Conflation In Turkish Texts," available at <http://www.shef.ac.uk/uni/academic/I-M/is/lecturer/paper13.html#lovi68>.
- [12] The Unicode Consortium, *The Unicode Standard: World Wide Character Encoding*, Addison-Wesley, Redwood City, CA, 1992.
- [13] Brent B. Welch, *Practical Programming in Tcl and Tk*, 2nd edition, Prentice Hall, 1997.
- [14] Donna Harman, "The DARPA TIPSTER project," *ACM SIGIR Forum* **26**(2), Fall 1992, pages 26 - 28.
- [15] The Text Retrieval Conference. Information available at <http://trec.nist.gov>.
- [16] R. L. Rivest, "The MD5 Message-Digest Algorithm," RFC 1321, Network Working Group, April 1992.
- [17] *Secure Hash Standard*, FIPS-180-1, National Institute of Standards and Technologies, U.S. Department of Commerce, April 1995.

- [18] Adam Sweeney, Doug Doucette, Wei Hu, Curtis Anderson, Mike Nishimoto, and Geoff Peck, "Scalability in the XFS File System," *Proceedings of the Winter 1996 USENIX Conference*, San Diego, CA, pages 1 - 14.
- [19] James Mayfield and Paul McNamee. "N-gram vs. Words as Indexing Terms," available at URL <http://www.csee.umbc.edu/~mayfield/pubs/APL-TREC97prelim.ps>.