

# Linear Scale Semantic Mining Algorithms in Microsoft SQL Server's Semantic Platform

Kunal Mukerjee  
Microsoft  
Redmond, WA, USA  
+1-425-703-6009  
kunal@microsoft.com

Todd Porter  
Microsoft  
Redmond, WA, USA  
+1-425-703-5337  
tporter@microsoft.com

Sorin Gherman  
Microsoft  
Redmond, WA, USA  
+1-425-538-4725  
soring@microsoft.com

## ABSTRACT

This paper describes three linear scale, incremental, and fully automatic semantic mining algorithms that are at the foundation of the new Semantic Platform being released in the next version of SQL Server. The target workload is large (10 – 100 million) Enterprise document corpuses. At these scales, anything short of linear scale and incremental is costly to deploy. These three algorithms give rise to three weighted physical indexes: Tag Index (top keywords in each document); Document Similarity Index (top closely related documents given any document); and Semantic Phrase Similarity Index (top semantically related phrases, given any phrase), which are then query-able through the SQL interface. The need for specifically creating these three indexes was motivated by observing typical stages of document research, and gap analysis, given current tools and technology at the Enterprise. We describe the mining algorithms and architecture, and also outline some compelling user experiences that are enabled by the indexes.

## Categories and Subject Descriptors

H.2.8 [Database Management]: Database Applications – *data mining, statistical databases.*

**General Terms:** Algorithms.

## Keywords

Semantic Platform, semantic mining, linear scale, incremental, keyword extraction, document similarity

## 1. INTRODUCTION

Managing unstructured and structured data in separate systems leads to many problems, such as consistency control, synchronizing backups, and supporting multiple systems. Increasingly, this pain drives users to move unstructured data into relational databases. Such unstructured data includes files (e.g., office documents, PDFs), large text fragments (e.g., emails, wiki pages, forum comments), etc.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

KDD '11, August 21–24, 2011, San Diego, California, USA.  
Copyright 2011 ACM 978-1-4503-0813-7/11/08...\$10.00.

As users store more documents in databases, it is critical that databases help manage them. Doing so effectively requires more than storing and retrieving bits; it requires providing access to the information in documents. To illustrate, suppose a consumer products company keeps corporate documents in a database. These documents describe which employees work on which products, how products are related, how they compare to the competition, etc. However, if the database treats these documents merely as opaque bits, users cannot query this information. To address this, we introduce the Semantic Platform for SQL Server, which mines and then exposes structured concepts within unstructured data to database queries, by building three physical indexes: 1) Tag Index (TI), which can return the top keywords, given a document; 2) Document Similarity Index (DSI), which can return the most closely related documents, given any document; and 3) Semantic Phrase Similarity Index (SPSI), which returns the top semantically related phrases, given any phrase.

The view from one level up is that in SQL Server, we are building a tightly streamlined and integrated Document Understanding Platform, comprising Integrated Full Text Search (IFTS) and Semantic Platform for documents, under a unified functional surface of SQL language extensions. We avoid creating an entirely separate syntax for Semantic Platform, by merely extending the IFTS syntax that is already familiar to many customers. The processing flow and algorithms are highly optimized so that documents are crawled, parsed and mined in a unified, 1-pass algorithm. The algorithm is incremental, and efficiently updates the index as the corpus evolves over time (increments may be triggered both manually and automatically). The Full-Text Index (FTI) is one output, in order to satisfy the needs of document search. Simultaneously, document level semantics are mined into the three new indexes: TI, DSI and SPSI, without needing to re-crawl and re-parse the original documents.

Additionally, the user does not have to go through any special configuration steps in order to tell the system about their data, e.g. set up special classifiers, etc. Everything involved in mining these indexes, figuring out what are the semantic concepts contained in the document corpus, how they cluster, etc., is fully automatic and inherent to the system. This means that we have created a very low barrier to entry for users of the document understanding system.

We conducted experiments to verify two main properties: linear scalability and quality. Results show that building the FTI, TI, and DSI is linear and highly scalable. Also, experimental results show that both TI and DSI produce very competitive precision numbers when compared to other methods.

The rest of the paper is organized as follows: The rest of Section 1 describes related products and algorithms, and then describes how we used customer and scenario focused design to help guide our design. In Section 2, we describe the overall architecture and components. Experimental results are described in detail in Section 3. We describe some potential applications in Section 4. We conclude with future work in Section 5.

### 1.1 Related Products and Algorithms

Starting in February 1990, the NSF has been publishing a set of goals for database research and commercialization every 5 years. The NSF prospectus [12] clearly identifies the need to cover IR over unstructured and semi-structured data and document processing, as a large emerging need for future database product lines. Fifteen years later we have Full-Text Search trending towards mainstream in multiple database product lines including Microsoft SQL Server; Oracle acquiring OpenText in 2008, etc.

We are currently at the cusp of widespread adoption of document IR in the database, as predicted by leading NSF researchers. In the Enterprise space, Oracle, PostgreSQL, Microsoft and others are competing in order to become entrenched players in the space of Document Understanding. Oracle's Open Text [10] has created integration with Share Point and Microsoft Duet (Office + SAP). Oracle Open Text offers a broad based document strategy, providing good support for web docs, via their Red Dot component, and over semi-structured data, by way of LiveLink. Semi-structured search is also supported with LiveLink ECM [10]. Microsoft's primary products in the space of document understanding include SharePoint and FAST, and also in-database IFTS, which has been offered for a few releases now.

At the algorithmic level, two semantic mining classics are Latent Semantic Indexing (LSI) [5] and Latent Dirichlet Allocation (LDA) [2]. These have been extensively used and built upon, in both Enterprise and Web based search. There are a number of approaches for extracting relevant keywords from a document, e.g. [3].

The main advantages of our semantic mining algorithms over the existing algorithms are that our algorithms scale linearly with the size of data, as we show in the results section, and they are incremental. All of these are essential ingredients for mining large (e.g. 10 to 100 million) Enterprise document corpuses and keeping the indexes "fresh" as the corpus evolves over time.

In [4] the authors propose a language-independent means of gauging topical similarity in unrestricted text, using character level n-grams and a simple vector-space technique to infer categorization and similarity. By contrast, we bootstrap our system with fairly large and complete statistical language models (LM), and work with ngrams over LM words that are output by the word breakers in our system [6], to cheaply inject high level structural and contextual constraints into our feature space very early on (almost at the outset). This saves compute time for our system.

In a related work also based on character level n-grams [3], the authors are motivated to construct a highlight/abstract extraction system that does not rely on any language-specific resources such as stemmers and stop lists, whereas the scenario is exactly opposite for IFTS and Semantic Platform in SQL Server, where the language is explicitly identified, and space constraints are not as tight. So we do employ stop lists, stemmers and sophisticated LMs as relatively cheap components from our perspective, to reduce CPU load, which aligns better with the more imperative aspects of our typical customer workload.

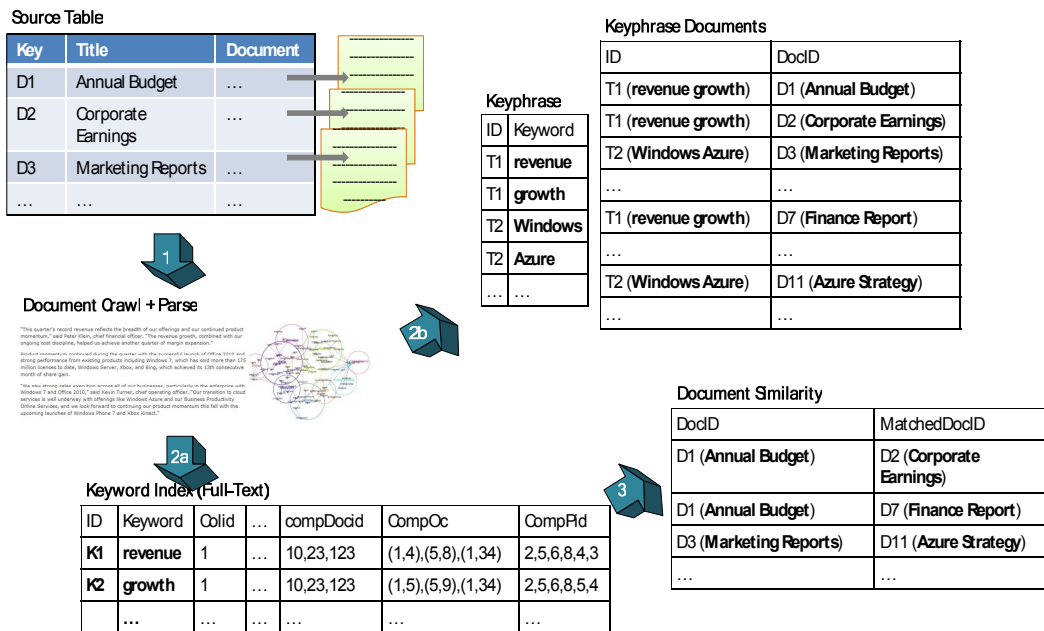


Figure 1. System design directly mirrors typical customer activities in researching large document corpuses: 1) FT1 enables full text search, returning references to the most relevant documents; 2) The next step is to query top documents for most relevant keywords and phrases – this is facilitated by T1; 3) The subset of interesting documents is used to query all related documents, using the DSI; and 4) Search terms are refined by way of SPSI (not shown here) – this takes us back to step 1.

In [7] the authors do inject language specific resources; they go one step further and use a knowledge model based on training a space of concepts on Wikipedia. Their Explicit Semantic Analysis (ESA) system reports good correlation with subjective evaluation, and also has the advantage of working over a vector space that makes intuitive sense compared to the above character-based vector space methods. However, at the Enterprise (the primary domain for SQL Server), a lot of the information tends to be very domain-specific, technical, and full of jargon and acronyms. For this reason we construct a system that does use a vector space of language ngrams, but stops short of constructing a vector space over a generic knowledge model, because we have found that these tend to be too generic. An adaptive knowledge model is more interesting for our purpose, however there are several engineering as well as Test challenges to this, given that database users tend to value determinism and reproducibility very highly.

For document similarity, our approach is close to the classic vector space model [11]. We use cosine-similarity over the documents represented in the vector space of weighted keywords, and use LM entropy to boost semantic salience to correlate well with subjective evaluation.

Our approach to phrase similarity is close to [8]. Whereas their objective is to ultimately use the phrase graph to infer document similarity, ours is to infer a semantic thesaurus that incorporates transitive relationships between phrases. We also focus on scalability aspects and linear approximation whilst generating a transitively closed phrase graph, because linear scale is pretty much a hard pre-requisite for SQL Server’s customer base.

Most important from our perspective is the fact that our mining algorithms are tightly integrated as a cohesive system such that the documents are parsed or incrementally scanned in a single pass, and all the data flows linearly through the entire system to create the three new indexes, plus the already existing FTI. Additionally, the one pre-existing and three new indexes are created and maintained to directly support four very fundamental customer needs and interactions, which are described next. This makes the end-to-end system highly optimized, and ensures that all functionality and end outputs are directly motivated by customer requirements.

## 1.2 Customer and Scenario Focused Design

Customer focus groups reveal that the following is a common work flow of events and actions when perusing and researching a document corpus, as depicted in Figure 1: 1) First, the user, by now very accustomed to using search engines, starts by searching the corpus for an idea, a word or a phrase; 2) Next, when the search returns a number of “hits” comprising Enterprise documents, the user needs to quickly skim the top results for their content without reading the entire document – this is an important differentiator for the Enterprise user as opposed to the user of web search, because Enterprise documents may be hundreds of pages long, whereas web pages are short enough to be scanned rapidly; 3) When a small number of candidate documents have been identified, the user wishes to round off their search by querying for all documents that are related to those exemplars – without this the research would not be complete; and 4) Finally, the user may wish to re-submit a search by refining the search terms, typically by seeking out *semantically related* words or phrases. For instance, “Google” may translate into {search engine, relevance} by way of ordinary thesaurus lookup; however, a semantically related term for “Google”, when mined over a corpus of Microsoft documents, might result in “competition”.

With the results of step 4, the user returns once more to step 1, and the process refines and repeats until the research is completed. The fundamentals of system design of Semantic Platform follow almost directly from the above work flow. Internally, we provide four physical indexes to support the above queries. 1) FTI, already available in past releases of SQL Server, enables searching of words and phrases; 2) TI summarizes each document into its top phrases; 3) DSI finds the top related documents to a given document; and 4) SPSI provides semantically related phrases to a given phrase.

The family of mining algorithms that processes a large enterprise document corpus linearly and incrementally to produce the above four indexes, constitutes a significant advance in the state of art, because they may be applied to unify the search and semantic inference processing suites of a significantly large number of Enterprise products. With modifications to handle spam and account for web-specific relevance characteristics such as hyperlink count (page rank), they may become interesting for search engines. Our algorithms can provide substantial savings of capital expenditure as well as reduce document crawl- to-index availability, i.e. end-to-end throughput and latency of these products may be improved.

In this paper, we present the Semantic Platform being released in the next version of SQL Server – and encourage people to try it out. This paper specifically describes the underlying mining algorithms of Semantic Platform, and provides a brief outline of the SQL extensions and new user experiences that are easy to build by querying these indexes.

## 2. SEMANTIC MINING ALGORITHMS AND ARCHITECTURE

In this section, we first show the system level block diagram, which shows how all the algorithmic pieces of IFTS and Semantic Platform fit, and how data flows through the tightly combined IFTS and Semantic Platform system end-to-end. Next, we will provide details on the three new Semantic mining algorithms.

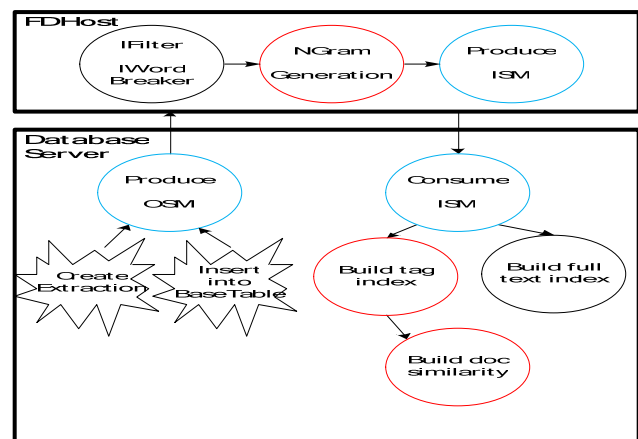


Figure 2. Overall system architecture.

### 2.1 Top Level System Diagram

Figure 2 shows how all the pieces fit at the system level. We start with a table in SQL database, where one column contains references (e.g. URIs) to actual documents. If the source of the data resides in a relational database, the primary key in the relational table or view goes in the main document index key

field. Once the extraction process is created and launched by the user, the documents are crawled and parsed in a separate process called FDHost, which runs as a service, hosting third party DLLs such as IFilters and word breakers, which crack open the document formats, e.g. \*.doc, \*.pdf, and extract words and word sequences (ngrams). Shared memory is used for inter-process communications between the database server processes. Once the ngrams have been received back in the database server process, the TI and FTI are built and/or incremented concurrently. After the TI has been built, it is in turn used as the source to construct the DSI. After DSI is built or incremented, the whole process is complete.

Due to the close integration of the Semantic mining platform with the pre-existing Full Text Search pipeline, the cost/effort to develop and maintain this architecture is moderate. We took 8 months with a team of 6 Developers, 5 Testers and 1 Program Manager to productize the system from start to finish.

## 2.2 Mining the Semantic Phrase Similarity Index (Graph)

**N.B.:** This index is not being shipped in the next version of SQL Server, but in a subsequent release to be determined by Microsoft.

In this section we first describe the top level objectives of the information we are trying to mine, to clarify for the reader, fundamental differences with related approaches like [8]. Whereas the end objective in [8] is to infer document similarity, our objective is to infer a phrase similarity graph that is approximately closed under transitive closure, and has near-linear time computational requirement.

### 2.2.1 Top Level Objective: Semantic Thesaurus

The overall objective of mining the semantic SPSI is to derive a semantic thesaurus which embodies the inferred semantic space over phrases used at the Enterprise.

For example, given a Microsoft corpus of documents, I would like “Google” to translate into things like “search engine”, “ad revenue”, etc., that I would normally get from a Thesaurus lookup, but additionally, also get “competition”. In case the corpus is mainly discussing the competition between these companies, I want “competition” to be heavily weighted. I also wish to discover so-called latent relationships [5] between words and phrases, such that if  $A \rightarrow B$  with a high weight and  $B \rightarrow C$  with a high weight, then I wish to infer the  $A \rightarrow C$  relationship. Most importantly, I wish to construct the mining algorithm to scale linearly, be incremental, and possible to integrate tightly into the overall FTI and Semantic Platform architecture outlined in Section 2.1.

The SPSI algorithm consists of eight stages, numbered S1 – S8 in




Figure 3. SPSI mining algorithm: Stages S1-S8 and data structures that go with them.

Figure 3. Corresponding to each stage is a data structure abstraction, as shown in the figure.

To facilitate understanding this algorithm, we will re-use a worked example from [15]. Consider that the document corpus we are working with contains the following three documents, with ids:  $D_1, D_2$  and  $D_3$ .

$D_1$ : “Cat ate cheese”

$D_2$ : “Mouse ate cheese too”

$D_3$ : “Cat ate mouse too”

In S1, which does tokenization and lemmatization into base forms, we transform the corpus thus:

$D_1$ : “cat eat cheese”

$D_2$ : “mouse eat cheese too”

$D_3$ : “cat eat mouse too”

In S2 we simply substitute LM ids for each word (or ngram), so that the rest of the algorithm works with integers. This makes the processing more efficient than working with strings. However, for ease of explanation, we will revert back to using strings for the remainder of this worked example.

$D_1$ : 2022, 788, 4077

$D_2$ : 16788, 788, 4077, 345

$D_3$ : 2022, 788, 16788, 345

In S3, where we use an LM to filter out the low entropy words in the language, we end up with the following:

$D_1$ : “cat **eat** cheese”

$D_2$ : “mouse **eat** cheese **too**”

$D_3$ : “cat **eat** mouse **too**”

We preprocess the text using LMs to derive entropy sliced pyramids (ESP). The top levels of these pyramids contain the most salient semantic entities. Lower levels of the ESP contain progressively more terms, and the lowest level contains everything and can be used for full text search. ESP essentially gives us a multi-resolution decomposition of the data, and imposes the priority/scale constraints that the top levels of the pyramid contain both the most semantically valuable as well as least quantity of data, which means we can “skim off the top” and compute only the most salient information, entities and




Figure 4. Entropy sliced pyramid structure. Each level contains all terms at or above the entropy threshold for that level. We may decide to truncate the pyramid at some level if we only need it for semantic indexing, otherwise we go all the way to the bottom for full text indexing.

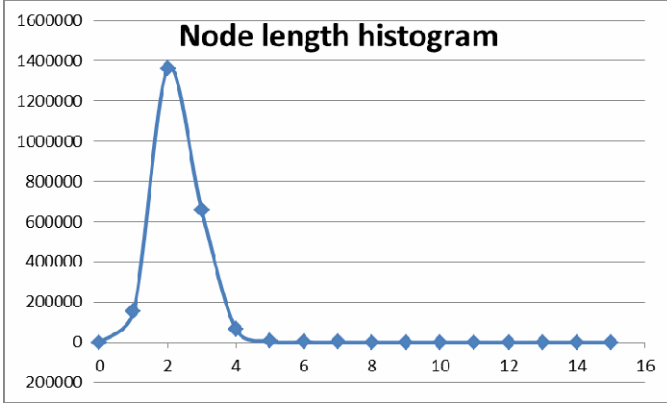
relationships faster than the rest. This mechanism enables us to deal very effectively with information at scale. Figure 4 shows an example of using an ESP on the incoming text (presumably from a resume): “Skill set: I have used C++ and MATLAB”.

In S4, we create the inverted ngram index, as shown below:

“cat”: { $D_1, D_3$ }

“cheese”: {D<sub>1</sub>, D<sub>2</sub>}  
 “mouse”: {D<sub>2</sub>, D<sub>3</sub>}  
 “cat cheese”: {D<sub>1</sub>}  
 “mouse cheese”: {D<sub>2</sub>}  
 “cat mouse”: {D<sub>3</sub>}

Here it’s worth mentioning that we only used 1, 2, and 3-grams.



**Figure 5. Ngram node length histogram mined from a 3000 doc corpus**

There is a big space/disk cost to using high order ngrams, and we repeatedly found evidence of diminishing returns for ngrams longer than 3 in a number of large (Enterprise) corpuses. This is shown in Figure 5.

In S5 we filter out sets with cardinality less than some threshold. For example, in our running example, we may filter out the sets: “cat cheese”, “mouse cheese” and “cat mouse” because they are all singletons. Sometimes singleton sets are indicative of noise, and they are rarely useful in linking semantically similar content.

In S6, we create weighted adjacency lists, using the Jaccard

$$\text{metric: } J(D_1, D_2) = \frac{|D_1 \cap D_2|}{|D_1 \cup D_2|}. \text{ This provides a notion of}$$

semantic nearness between any two ngrams, as the count of documents they are both present, over those where at least one is present. In our running example, the pairwise Jaccard metric applied to the surviving ngrams produces the following output, which is an Adjacency Matrix:

cat: {(cheese, 0.33), (mouse, 0.33), (cat cheese, 0.5), (cat mouse, 0.5)}

cheese: {(cat, 0.33), (mouse, 0.33), (cat cheese, 0.5), (mouse cheese, 0.5)}

mouse: {(cat, 0.33), (cheese, 0.33), (mouse cheese, 0.5), (cat mouse, 0.5)}

cat cheese: {(cat, 0.5), (cheese, 0.5)}

mouse cheese: {(cheese, 0.5), (mouse, 0.5)}

cat mouse: {(cat, 0.5), {mouse, 0.5}}

For S6, we compute the Jaccard graph that relates all pairs of ngram nodes incrementally, with each incoming document, as described below.

*Induction Hypothesis:* Graph is Jaccard complete (i.e. all pairs have distance given by the Jaccard metric) after we have seen the (i-1)<sup>th</sup> document.

*Induction Step:* We only partially update the graph to change numerators and denominators of only the “active ngrams” – these are the ngrams that occur in the i<sup>th</sup> document. After the update step the graph is once more Jaccard complete.

In the following we outline the steps of the fast Jaccard update algorithm on receiving the i<sup>th</sup> document.

1. Represent the incoming document as a set of ngrams:  $Doc(i) = \{ng_{i1}, ng_{i2}, \dots, ng_{in}\};$
2. Mark previously unseen and previously seen ngrams, creating the set partitions:  
 $Doc(i) = Seen(i) \cup Unseen(i);$
3. Insert the unseen set into the Adjacency Matrix, the unseen rows become:  $Doc(i)$ , all numerators and denominators = 1;
4. For each ngram in the Seen set, append the Unseen set, and initialize all numerator = denominator = 1;
5. For each ngram in Seen set, update the old row with denominator++; and if the old word is in  $Doc(i)$  then numerator++.

Step S7 is actually an optimization/pre-processing step to speed up S8, and so we first describe S8, which is the transitive closure over the graph obtained from S6.

At the output of S6, we have a pairwise connected weighted graph represented by the adjacency matrix. From this it is possible to compute connected components and semantic clusters and neighborhoods. Indeed, that is what we do in S7. However, in order to capture the latent semantic relationships, it is necessary to consider transitive relationships between ngram nodes. If we fail to do that, the semantic inference available from the entire process tends to be locked into “vocabulary silos”.

We compute the transitive closure over this graph using the classic Floyd-Warshall algorithm, by defining the relaxation step to be the maximum over edge weight between vertices  $i, j$ , and intermediate vertex  $k$ :  $e(i, k) \times e(k, j)$ . This exposes the latent semantic relationships, and gives us functional equivalence with LSI. The end result is an entity relatedness graph with edge weights representing relationship strengths between any pair of vertices. Note that our edge weights are symmetrical, and so we only need to deal with the upper triangular half of the adjacency matrix, discounting diagonal entries (each node is trivially related to itself).

*Theorem:* Given a Graph  $G$ , defining the relaxation step on an edge connecting vertices  $V_i$  and  $V_j$  as:  $e(i, j) = \text{Max}\{e(i, j), e(i, k) \times e(k, j)\}$  where the base edge weights represent Jaccard similarity, is necessary and sufficient to induce transitive closure over  $G$ , exposing latent semantic relationships.

*Proof:* If Jaccard similarity is considered as an approximation of probability that nodes  $V_i$  and  $V_j$  are related, and if the relationship between any two (non-identical) pairs of vertices (e.g.  $\{V_i, V_j\}$  and  $\{V_i, V_k\}$ ) is independent, then their joint probability is the result of multiplying independent probabilities. But this is the same as multiplying edge weights stemming from Jaccard similarity.

To complete the running example, here are the outputs of the Transitive Closure step. Latent relationships that emerge after transitive closure are highlighted.

cat: {(cheese, 0.33), (mouse, 0.33), (cat cheese, 0.5), (cat mouse, 0.5), (**mouse cheese, 0.165**)}

cheese: {(cat, 0.33), (mouse, 0.33), (cat cheese, 0.5), (mouse cheese, 0.5), (**cat mouse, 0.165**)}

mouse: {(cat, 0.33), (cheese, 0.33), (mouse cheese, 0.5), (cat mouse, 0.5), (**cat cheese, 0.165**)}

cat cheese: {(cat, 0.5), (cheese, 0.5), (**mouse, 0.165**), (**mouse cheese, 0.25**), (**cat mouse, 0.25**)}

mouse cheese: {(cheese, 0.5), (mouse, 0.5), (**cat, 0.165**), (**cat cheese, 0.25**), (**cat mouse, 0.25**)}

cat mouse: {(cat, 0.5), {mouse, 0.5), (**cheese, 0.165**), (**cat cheese, 0.25**), (**mouse cheese, 0.25**)}

### 2.2.2 Transitive Closure Optimizations (S7)

Transitive Closure on a graph is an expensive  $O(N^3)$  operation. We outline a series of optimizations that reduce the end-to-end complexity.

First, we *enforce sparseness* on the graph from S6 by clamping weights,  $e(i,j)$ , that are less than a *sparseness threshold*,  $T$ , to zero: *If  $(e(i,j) < T) e(i,j) = 0$ , then  $e(i,j) = 0$ .*

We next run connected components algorithm on the sparse graph. Connected components is  $O(N)$ . Once the  $M$  components are identified, the transitive closure on the entire graph reduces to running transitive closure on the  $M$  components. Thus, we

transform the overall complexity from  $O(N^3)$  to  $O(\sum_{i=0}^M n_i^3)$ .

A spin-off data structure of connected components is the *Component Index*, which is actually a hierarchy of embedded sub-components, corresponding to different values of sparseness threshold. This enables us to work with components and clusters later on.

We now state without formal proof, a conjecture based on substantive empirical evidence, that actually results in completely eliminating the expensive transitive closure operation for large corpuses. ***Either transitive closure is necessary and the document set is small, in which case it is cheap to compute; or transitive closure is asymptotically unnecessary*** (and can therefore be entirely eliminated), beyond filling in the pair-wise Jaccard similarities, when the document set is large.

*Intuitive Explanation:* Consider documents relating Obama and the US presidency. These are likely to be many. Consider documents relating Obama and basketball. These are likely to be few. In a small collection there may not be any single document commenting on relationships between the office of the President of USA and basketball. However, if a sufficiently large document set is constructed, then there is a high chance that some document does talk about how presidents relate to the game of basketball.

This intuitively explains why the net contribution of transitive closure, or latent relationships, in general, may be expected to diminish as we approach very large document sets.

*Experimental Evidence:* We now present the experimental evidence (Figure 6) which shows that the contribution of the weight deltas due to the transitive closure step seems to follow roughly an exponential decay curve as the corpus size increases. We have found that this generalizes over a variety of different Enterprise (e.g. SharePoint design document repository) and academic corpuses (e.g. New England Journal of Medicine).

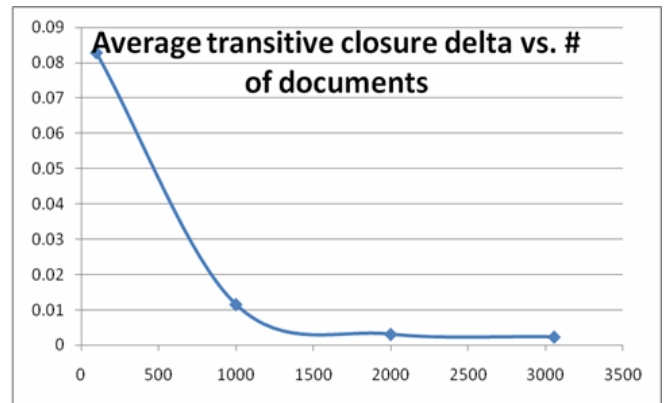


Figure 6. Contribution of Transitive Closure diminishes for large document corpuses.

## 2.3 Mining the Tag Index

The primary objective of mining the tag index (TI) is the need to summarize long enterprise documents (sometimes 200 – 500 pages long) into their top few key words and phrases. This is very important in order to support the search experience, because the user might otherwise find search results, which are links to long documents, very time consuming to go through. Good tag indexing therefore directly results in productivity gains for the user of an enterprise search based system such as IFTS.

The TI mining algorithm is based on a LM and weight and threshold functions. TI is the output of the algorithm: given a document id, the TI returns a list of its top N key phrases, each with a weight signifying relative importance in the document.

The data flowing into the TI mining algorithm is shown in the system block diagram above (Figure 2). The TI mining algorithm is kept simple for scalability purposes, and is based on cross-entropy: we compare the word distributions from the LM (which contains expected distributions) to the word distribution from the current document. We score the ngrams based on how much more frequently they appear in a document, compared to their expected frequency as computed based on LM. This can be viewed as an approximation to Term Frequency Inverse Document Frequency (TFIDF) weighting, but we are replacing the inverse document frequency portion with the LM frequency. This enables our algorithm to be corpus independent, which is essential for Enterprise document corpuses that are getting new documents daily – we do not want to re-compute the TI every time this occurs. Also, our algorithm is not machine learning based, so there is no training phase, it can be used without the need of a good corpus representative sample data set for initial training.

A caveat is that for short documents the word distributions are not very reliable. Therefore, we take into account document length as well. For all ngrams in any document that are considered as a possible tag, the algorithm requires: 1) Ngram frequency in the document; and 2) Ngram frequency in the language. This is either directly available in the LM, or is inferred from the LM, by treating words as i.i.d and adding their log probabilities.

The pseudo-code for extracting keywords is provided below.

```
ExtractTopNKeyPhrases (doc)
{
  languageModel <- SelectLanguageModel(doc)
  candidates <- Empty topN priority queue
  for each ((ngram, locations) in doc) do
    score = CrossEntropy(
```

```

doc.GetSize(),
locations.Length,
    // ngram frequency in doc
languageModel.Frequency(ngram)
);
Candidates.Add(ngram, score);
end for
return topN candidates by descending score
}

ComputeCrossEntropy(docSize, docFrequency,
                    LMFrequency)
= (docFrequency/docSize) x (BASELog10(LMFrequency))

```

## 2.4 Mining the Document Similarity Index

The DSI mining algorithm works with a populated and stable TI as its input/data source. The end output of the algorithm is an index which can be queried by any given document id, and returns the top N other documents which share common highly ranked key phrases with the given document, with connection weights for each related document.

We only consider up to top K' candidate documents that we consider for similarities (to avoid the  $O(N^2)$  CPU issue), and only top K (where  $K \leq K'$ ) of those similarities will be stored. This allows us to avoid the  $O(N^2)$  space issue. This greedy heuristic can be summarized as: given the target doc T, we have the list of keywords in descending order by their relative weight within T given by the TI; we start exploring the keywords list top-down and we find the documents where these keywords are also highly weighted; we stop when we have found K' such documents, or we finished exploring the keywords list for doc T; those K' documents are our candidates.

Given two documents,  $doc_1$  and  $doc_2$ , the TI gives us the list of tags and corresponding weights for each:

$Doc_1: \{(tag_{11}, weight_{11}), (tag_{12}, weight_{12}), \dots, (tag_{1k}, weight_{1k})\}$

$Doc_2: \{(tag_{21}, weight_{21}), (tag_{22}, weight_{22}), \dots, (tag_{2k}, weight_{2k})\}$

We use cosine-similarity to compute the similarity between two documents, as shown below:

$$CosineSimilarity(doc_1, doc_2) = \frac{DotProduct(doc_1, doc_2)}{(Norm(doc_1) * Norm(doc_2))}$$

where:

$$DotProduct(doc_1, doc_2) = \sum \{Weight(tag, doc_1) * Weight(tag, doc_2)\}$$

for each tag that is common to  $doc_1$  and  $doc_2$ , and

$$Norm(doc_i) = \sqrt{\sum Weight(tag, doc_i)^2} \text{ for each tag in } doc_i.$$

The pseudo-code for the DSI mining algorithm is provided below.

```

ExtractTopKDocumentSimilarities(TagIndex TI)
{
for each (unprocessed docId) do
// Step 1: Find top K' candidates
topK' Candidates=SelectSimCandidates(TI,
                                     docId);

// Step 2: Find similarities,
// relative to candidates
resultsHeap <- EmptyHeap(size=K);
for each (candidateId in topK') do
similarity=CosineSimilarity(T, docId,
                           candidateId);
resultsHeap.Add(candidateId,
                similarity);
end for
// Step 3: Select top K results
}

```

```

for each (candidateId, similarity) in
resultsHeap do
DSI.WriteRow(docId, candidateId,
             similarity);
end for
end for
}

```

The details of the SelectSimCandidates function are provided in pseudo-code.

```

SelectSimCandidates(TI, docId)
{
resultSet = Empty(Map<int, double>);
for each (tag, weight1) in TI(docId) do
for each ((candidateId, weight) in
(select topK' docs from TI[tag]) do
if (candidateId ∉ resultSet)
resultSet.Add(candidateId, 0.0);
end if
resultSet[candidateId] += weight1 * weight2;
end for
end for
return top K items in resultSet;
}

```

## 3. EXPERIMENTAL RESULTS

Traditional accuracy metrics for full text search and web search are Precision and Recall [1]. Whereas the precision of semantic search is good, the recall is expected to be somewhat poor in comparison to full text search products, because we follow transitive links to other documents.

For our tagging experiment, we used the social tagging dataset culled from Delicious that was used in [9]. There are two versions of the TI algorithm that we tested - the normal algorithm and a filtered version (Filtered TI). The filtered version takes the keywords produced by the TI and then filters out keywords that do not appear in the document a minimum number of times normalized by the document length. We compare our results to corpus dependent machine learning methods: KEA [14], Linear Support Vector Machine (SVM) and Ranking Support Vector Machine (Ranking SVM) [9]. For details on parameter settings, test/training splits, etc., please see [9]. We also compare our results to the standard Term Frequency (TF) and TFIDF methods. Table 1 shows our results based on the Precision at 5 (P @ 5) and Precision at 10 (P @ 10) metrics.

The results are quite promising. Without filtering, TI produces results that are approximate to TFIDF, but in a corpus independent way. Once we filter TI, we achieve results that surpass TF and TFIDF and, in fact, are closing in on KEA. The advantage is that our method scales linearly. One interesting result is that TF performed better than TFIDF, which could be due to the small corpus size (600). Overall, the results show that TI provides a solid foundation upon which DSI can be built.

**Table 1. Precision results for our algorithm (TI and Filtered TI) compared to others**

	P @ 5	P @ 10
TI	<b>0.289</b>	<b>0.221</b>
TFIDF	0.295	0.23
TF	0.34	0.267
Filtered TI	<b>0.38</b>	<b>0.27</b>
KEA	0.437	0.31
SVM	0.469	0.323
Ranking SVM	0.495	0.349

For document similarity precision performance, we constructed a corpus of 2500 documents scraped from Wikipedia. The corpus consisted of 25 topics, each with 100 documents each. Sample topics include “american actors”, “american singers”, and “national football league players”. Note, some of the documents belonged in two categories, such as an actor that was also a singer. This negatively affects precision results for all methods including DSI. To test the precision, we took a single document from each of the topics and used it as the query document, and then we obtained the top 10 most similar documents. Then precision numbers at 1, 3, 5, and 10 were calculated. We compared our results to the traditional Cosine similarity measure (CosSim) based on TFIDF [1].

The results in Table 2 show that DSI provides a similar precision level to CosSim, but does so in linear time instead of in  $O(N^2)$  time, and is completely corpus independent, unlike CosSim. This is a big win for customers whose workloads are in the millions of documents, anything other than linear will simply not be feasible at such scale.

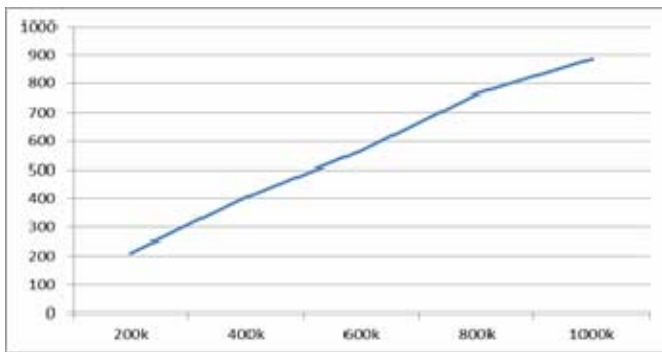


Figure 7. Total execution time in seconds versus total number of documents.

Table 2. Precision results for DSI and CosSim

	P @ 1	P @ 3	P @ 5	P @ 10
DSI	0.8	0.8	0.79	0.78
CosSim	0.8	0.82	0.83	0.82

Figure 7 shows how FTI, TI and DSI mining stack up in terms of end-to-end execution time (i.e. sum total of computing all three indexes) on a real customer workload sampled at 200k document increments, with an average sample size of 40KB per document in plain text format. The test machine had 32GB of RAM and a DOP of 32. The results clearly show linear scaling. Pre-release bits of this product were released to a restricted set of users, and have been well received, with one Microsoft internal and one external customer showing intent to adopt.

#### 4. APPLICATIONS

Based on mining the SPSI (Section 2.1), we propose to enable a set of concept-centric information navigation experiences (Figure 8). This enables the user to browse a document collection by following links through concepts, rather than through document names or a sequence of search queries – especially useful in large document collections where individual documents are large, e.g. at the Enterprise. In Figure 9, we show how the semantic graph may be used to enhance the Windows Explorer user experience.




Figure 8. Exploring a concept graph

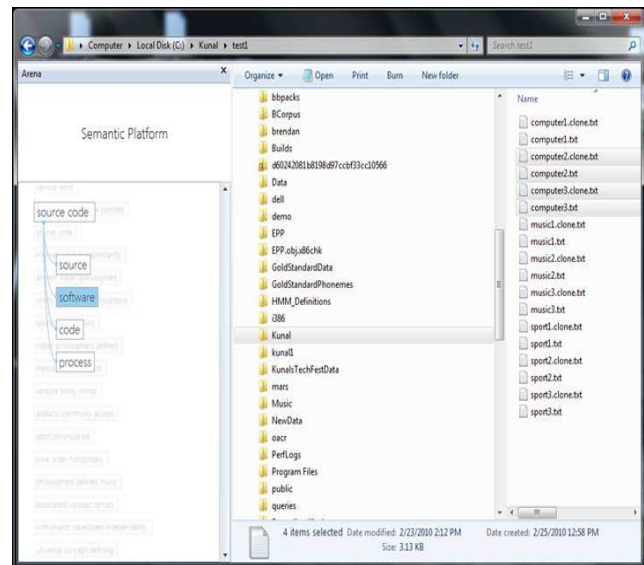


Figure 9. Exploring mined concepts in a file system. When a concept is browsed on the left, the corresponding files are highlighted on the right.

#### 5. CONCLUSIONS

We have described the semantic mining algorithms that lie at the core of Semantic Platform, which is being released in the next version of SQL Server, and shown that they scale linearly with large document corpuses. The combination of the four indexes facilitates the process and work flow of browsing, searching and researching information at the Enterprise, where documents can be large compared to web pages. It also enables new and intuitive user experiences along the lines of browsing concepts that are contained in documents instead of the documents themselves, e.g. using filenames.

Future work includes keyword precision improvements for the TI by utilizing additional corpus independent features such as first occurrence and phrase distribution [9]. Additionally, for the DSI, exploring different candidate document selection strategies and document similarity functions being researched; these could yield gains in performance and precision. We also plan to expose more tuning knobs to better control the degree of completeness of DSI results. Adaptive language modeling is yet another promising future direction when mining Enterprise document corpuses.

## 6. REFERENCES

- [1] Baeza-Yates, R., and Ribeiro-Neto, B., *Modern Information Retrieval*, Addison-Wesley, 1999.
- [2] Blei, D. M., Ng, A. Y., and Jordan, M., Latent Dirichlet allocation, *Journal of Machine Learning Research* No. 3, pp. 993 – 1022.
- [3] Cohen, J. D., Highlights: Language- and domain independent automatic indexing terms for abstracting, *Journal of the American Society of Information Science*, Volume 46, Issue 3, pp. 162-174, April 1995.
- [4] Damashek, M., Gauging similarity with N-grams: Language-independent categorization of text, *Science* 267, Feb. 1995.
- [5] Deerwester, S., et al., Improving Information Retrieval with Latent Semantic Indexing, *Proceedings of the 51st Annual Meeting of the American Society for Information Science* 25, 1988, pp. 36–40.
- [6] Full Text Search Overview (<http://msdn.microsoft.com/en-us/library/ms142571.aspx>).
- [7] Gabrilovich, E., Markovitch, S., Computing semantic relatedness using Wikipedia-based explicit semantic analysis, in *Proceedings of the 20<sup>th</sup> International Joint Conference on Artificial Intelligence*, 2007.
- [8] Hammouda, K., M., Kamel, M., S., Document Similarity using a Phrase Indexing Graph Model, *Journal of Knowledge and Information Systems*, Vol. 6, Issue 6, November 2004.
- [9] Jiang, X., Hu, Y., Li, H., A Ranking Approach to Keyphrase Extraction, *Microsoft Research Technical Report*, 2009.
- [10] McNabb, K., Moore, C., and Levitt, D., Open Text Leads ECM Suite Pure Plays, *The Forrester Wave Vendor Summary*, Q4 2007.
- [11] Salton, G., Wong, A., Yang, C. S., “A Vector Space Model for Automatic Indexing”, *Communications of the ACM*, vol. 18, nr. 11, pages 613-620, 1975.
- [12] Silberschatz, A., Stonebraker, M., and Ullman, J., *Database Research: Achievements and Opportunities into the 21<sup>st</sup> Century*. Technical Report, Stanford, 1996.
- [13] Tan, P-N., Steinbach, M., and Kumar, V., *Introduction to Data Mining*, 2005.
- [14] Witten, I.H., Paynter, G. W., Frank, E., Gutwin, C., and Nevill-Manning, C., G., KEA: Practical automatic keyphrase extraction, *Proc. DL '99*, pp. 254-256.
- [15] Zamir, O., and Etzioni, O., Web Document Clustering: A Feasibility Demonstration, in *Proc. ACM SIGIR '98*, 1998.