

Sentence-Forest Language Model: A Graph-theoretic Framework to Capture Local Term Dependencies

Ramesh Nallapati and James Allan
Center for Intelligent Information Retrieval
Department of Computer Science
University of Massachusetts
Amherst, MA 01003
{nmramesh, allan}@cs.umass.edu

ABSTRACT

We describe a new probabilistic graph-based language Model that captures adaptively the local term dependencies in documents. The new model works by boosting scores of documents that contain topic-specific local dependencies and exhibits the behavior of the unigram model in the absence of such dependencies. New features of the approach include modeling the syntactic structure of sentences in documents and a computationally efficient algorithm for capturing the most significant within-sentence term dependencies using a Maximum Spanning Tree approach, similar to Keith van Rijsbergen’s [13] modeling of document-level term dependencies.

We evaluated our model on three different tasks of the Topic Detection and Tracking (TDT) research program, namely Story Link Detection, Topic Tracking and New Event Detection. Our results show that the Local dependency language model consistently outperforms the basic unigram model on all three tasks.

Categories and Subject Descriptors

H.3.3 [Information Search and Retrieval]: Retrieval models - *language models, dependencies*

General Terms

Algorithms

Key Words

Term dependencies, language modeling, probabilistic approaches, graphical models, co-occurrences, sentences, maximum spanning tree, story link detection, topic detection and tracking

1. INTRODUCTION

Language Models have been found to be very effective in several information retrieval tasks. In the Language Modeling approach, we measure relevance of a document to a topic by the probability of its generation from the topic model [11]. One major assumption made in the unigram language modeling is the independence of all

terms with respect to one another. This allows us to compute the probability of generation of a document as the product of probabilities of generation of each term in the document, as shown in the following equation:

$$P(D|M) = \prod_{w \in D} P(w|M) \quad (1)$$

where D is the document in question, M is the topic model and w is a term in the document.

But to quote the famous probability theorist De Finetti, “dependence is the norm rather than the contrary” [4]. From our own understanding of natural language, we know that the assumption of term independence is a matter of mathematical convenience rather than a reality.

However, the ‘bag of terms’ approach of the unigram language model, as shown in equation 1, ignores all the dependencies and any other positional information of terms in the document. Hence, it seems desirable to have a more sophisticated model that is capable of capturing the semantics of documents rather than just the term distributions. A first step towards achieving this objective is capturing term dependencies, since dependencies establish associations between terms and may shed more light on the underlying semantics of the document than unigram term distributions alone.

The present work is an attempt to capture local term dependencies using a new variation of the language modeling approach.

The remainder of the report is organized as follows. Section 2 summarizes attempts made in the past in capturing dependencies. We present the methodology of the new *Sentence-Forest* language modeling approach in section 3. In this section, we describe the motivation behind the approach and several modeling issues involved. We evaluate the complexity of Sentence-Forest model and compare it with that of the unigram model in section 4. In section 5, we present a few interesting insights into why we think our new model works better than the unigram model. Section 6 describes some of the implementation details, while section 7 presents a brief description of Topic Detection and Tracking paradigm and two sub-tasks called Link detection and Topic tracking. In section 8, we describe the experiments performed and present the results obtained on two different tasks. Section 9 ends the discussion with a few observations and remarks on the performance of the Sentence-Forest model.

2. PAST WORK

A vast majority of IR models assume independence of terms because the assumption leads to a tractable representation and most IR systems practically have worked well under this assumption.

There have been very few successful attempts at modeling dependencies especially in the domain of formal probabilistic models.

Van Rijsbergen tried to capture document level term dependencies in his probabilistic modeling approach using Expected Mutual Information Measure (EMIM) scores between terms [13]. A maximum spanning tree is constructed, with the terms as the nodes and the EMIM scores as the weighted edges. The tree captures the maximum dependencies between terms in the document. These dependencies are used in computing the similarity score between two documents. However, the approach is computationally expensive and also unfortunately did not show promising results.

In other related work, Robertson and Bovey [14] tried including term pairs that have observable dependencies as separate terms with weights slightly different from the sum of weights or in some other way to allow for specific dependencies.

Turtle and Croft [16] investigated the use of an explicit network representation of dependencies by means of Bayesian inference theory. The use of such a network generalizes existing probabilistic models and allows integration of several sources of evidence within a single framework.

Fung and Crawford [5] worked on concept based information retrieval that captures dependencies between ‘concepts’ using a Bayesian inference network. One drawback of this approach is that the user has to identify the concepts manually in each document.

Attempts were also made to capture term dependencies using the vector space model. The generalized vector space model [17] is one such example which showed encouraging results.

In the area of language modeling, most attempts at capturing dependencies have been in the form of multigram language models [15]. Bigram and trigram models, though highly successful in the speech recognition task, have not met with great success in the domain of information retrieval systems.

In a very recent work, Nallapati and Allan [10] have shown that modeling sentences as maximum spanning trees, similar to Van Rijsbergen’s [13] modeling of documents, holds some promise in capturing local dependencies. Their *SenTree* model does not perform as well as unigram model, but they found that a linear mixture of *SenTree* and unigram models betters the unigram performance, but only in regions of low false alarm.

The present work, the *Sentence-Forest* model, is an improvement on the *SenTree* model. The *Sentence-Forest* model combines the features of both unigram and *SenTree* models into a single framework. We will show in our results section that the new model consistently outperforms the unigram model on two different TDT tasks.

3. METHODOLOGY

Our goal is to build a topic model from a story D_1 or a set of stories $\{D_1, \dots, D_n\}$ and then decide whether story D_{n+1} is predicted by that model. This section describes the methodology of the new model.

3.1 Exploiting sentence structure

A universal feature of all documents is the syntactic structure of sentences. Each sentence conveys a complete idea or a concept through a specific ordering of terms sampled from the language vocabulary. The sampling and ordering of terms depends on the intended concept and the underlying grammar of the language. The concepts or the semantics of the entire document are ultimately expressed as a grouping of such ordered samples of terms called sentences. In other terms, the semantics of a document are expressed through the syntax of sentences. Hence, we believe modeling sentences, rather than terms as individual entities, better captures the

underlying semantics of the document.

As we have seen earlier, unigram language models completely ignore the syntactic formation of sentences in documents. In the present approach, we attempt to capture it by modeling a document as a collection of sentences rather than as a ‘bag of terms’. We model each sentence graphically, as a forest of terms as we shall see later.

3.2 Probability of Sentence Generation

In the *Sentence Forest* approach, we assume each sentence to be independent of the other sentences. This assumption is certainly not valid but it is less stringent than the assumption of term independence. The assumption allows us to compute the probability of generation of a story from a topic model as follows:

$$P(D|M) = \prod_{S \in D} P(S|M) \quad (2)$$

where M is a topic model and S is a sentence in a story D . The tricky part is computing the probability of generation of each sentence. Ideally, one would compute the probability of generation of a sentence as follows:

$$P(S|M) = P(w_1, w_2, \dots, w_n|M) \quad (3)$$

where w_i is the i -th term and n is the number of terms in the sentence S . We also assume that all the terms in the sentence are unique, *i.e.*, we ignore any repeated occurrences of terms within a sentence. This assumption allows us to formulate our model in a clean way.

The joint probability in the right hand side of equation 3 is almost impossible to estimate due to the sparse nature of training data. To overcome this problem, we model the sentence as a forest in an approach motivated by Van Rijsbergen’s work. [13].

Given a sentence S and a model M , we first define the following undirected, weighted graph $G(V, E)$ as shown below.

$$V = \{w | w \in S\} \quad (4)$$

$$E = \{(w_i, w_j) \mid w_i, w_j \in V; P(w_i, w_j|M) > 0\} \quad (5)$$

and

$$\forall (w_i, w_j) \in E, \quad W(w_i, w_j) = P(w_i, w_j|M) \quad (6)$$

where $P(w_i, w_j|M)$ is the probability that w_i and w_j co-occur in the same sentence given the model M and $W(w_i, w_j)$ is the weight assigned to the edge (w_i, w_j) .

In other words, the set of nodes V in the graph G corresponds to the terms in the sentence S and an edge exists between terms w_i and w_j with a weight $P(w_i, w_j|M)$ if and only if they have a non-zero joint probability of within-sentence-co-occurrence with respect to the model. An example sentence-graph G is illustrated in figure 1. The graph has three connected components C_1, C_2 and C_3 as shown and corresponds to a sentence S that consists of the terms $\{w_1, \dots, w_8\}$. The thickness of the edges is indicative of the value of edge weight or the probability of within-sentence co-occurrence with respect to the model.

The first assumption we make is that the generation of terms in each connected component in the graph G is independent of terms in other connected components. If $\mathbf{w}_C = \{w | w \in C\}$ is the set of terms in a connected component C of G , then we can write

$$P(S|M) \approx \prod_{C \in G} P(\mathbf{w}_C|M) \quad (7)$$

where $P(\mathbf{w}_C|M)$ is the joint probability of all the terms in the connected component C . In the example of figure 1, the probability




Figure 1: A sentence Graph

of sentence generation can be approximated as shown below.

$$\begin{aligned}
 P(w_1, \dots, w_8 | M) &\approx P(\mathbf{w}_{C_1} | M) P(\mathbf{w}_{C_2} | M) P(\mathbf{w}_{C_3} | M) \\
 &= P(w_1, w_2, w_7, w_8 | M) P(w_3 | M) P(w_4, w_5, w_6 | M) \quad (8)
 \end{aligned}$$

Let's now consider how to approximate the joint probability of terms in a connected component C of G , given by $P(\mathbf{w}_C | M)$. In a nutshell, we approximate the joint probability of \mathbf{w}_C by the joint probability computed with respect to a Bayesian Network defined over the maximum spanning tree of C . In the following discussion, we first present the intuition behind such modeling and then present the actual model stating the assumptions made.

3.2.1 Intuition behind Maximum Spanning Tree Model

Let us assume that $\mathbf{w}_C = \{w_1, \dots, w_m\}$ for convenience. Using the chain rule, the joint probability in equation 7 can be expressed as

$$\begin{aligned}
 P(w_1, w_2, \dots, w_m | M) &= P(w_1 | M) \times P(w_2 | w_1, M) \times \\
 &P(w_3 | w_1, w_2, M) \times \dots \times P(w_m | w_1, w_2, \dots, w_{m-1}, M) \quad (9)
 \end{aligned}$$

As an approximation to this exact formula, we may ignore the higher order dependencies in each term in the right hand side of equation 9 and select from the conditioning variables, one particular variable that accounts for most of the dependence relation. In other words, we seek an approximation of the form

$$P(w_i | w_{i-1}, \dots, w_1, M) \approx \max_{0 \leq j < i} P(w_i | w_j, M) \quad (10)$$

Let $j(i)$ be the function that maps i into integers less than i such that $w_{j(i)}$ accounts most for the dependency of w_i . In other words, $j(i)$ is the value of j that maximizes the probabilities in equation 10. Then, the approximate probability of sentence generation is given by

$$P(\mathbf{w}_C | M) \approx \prod_{i=1}^m P(w_i | w_{j(i)}, M) \quad (11)$$

where w_0 is defined such that

$$P(w_i | w_0, M) = P(w_i | M) \quad (12)$$

If we imagine a dependency graph of relation 11 with terms as vertices and conditional probabilities as directed edges (directed from the conditioning variable to the conditioned variable), it is easy to

see that the relation represents a Bayesian Network in which each node has at most one parent. Moreover, if we ignore the direction of the edges, the network forms a spanning tree over the connected component C . These observations follow from the fact that there is at least one edge connected to each node and a node's dependence is defined only on the nodes preceding it, *i.e.*, $j(i)$ is always less than i , thus preventing the formation of any dependency cycles.

At this point, let us reconsider the chain rule expansion of the joint probability in equation 9. If we assume all orderings of terms in the sentence to be indistinguishable, then the chain rule can be defined in $m!$ possible ways. For instance, we could have expanded the joint probability as follows, too:

$$P(w_1, w_2, \dots, w_m | M) = P(w_m, w_{m-1}, \dots, w_1 | M) \quad (13)$$

$$\begin{aligned}
 &= P(w_m | M) P(w_{m-1} | w_m, M) P(w_{m-2} | w_m, w_{m-1}, M) \\
 &\dots P(w_1 | w_m, w_{m-1}, \dots, w_2, M) \quad (14)
 \end{aligned}$$

If we use a similar approximation as above on the new expansion in equation 14 ignoring the higher order dependencies and selecting the best conditioning variable in each term, the set of dependencies still forms a Bayesian network defined over a spanning tree of the graph. It is easy to see if we renumber $(w_m, w_{m-1}, \dots, w_1)$ as $(w'_1, w'_2, \dots, w'_m)$ and substitute in the chain rule expansion. We essentially end up with the same form as in equation 10.

Thus, the best approximation to the joint probability would be a Bayesian network defined over one of a maximum of $m!$ spanning trees, that best captures the dependencies. Clearly this is given by the Maximum Spanning Tree (MST) over the graph of C . The MST incorporates the most significant of dependencies between the terms of a sentence subject to the global constraint that the sum of them should be a maximum. Once the MST is computed, we can define a Bayesian network of dependencies by traversing breadth-wise or depth-wise on the MST starting from any leaf node and defining the direction of the edges to be same as the direction of traversal. Such definition of a Bayesian network ensures that every node has at most one parent as required by our approximation of first order dependencies. This follows from the fact that in depth-wise or breadth-wise traversal, each node is visited exactly once. Moreover, it can be easily shown using Bayes' rule that the joint probability remains the same irrespective of the starting node chosen in the definition of the Bayesian network over the spanning tree. Hence we can uniquely identify a Bayesian Network on each spanning tree.

This suggests modeling the joint probability of \mathbf{w}_C as follows. Form an undirected maximum spanning tree T_{Max} over C and define a directed Bayesian network on the tree as explained above. Then the joint probability $P(\mathbf{w}_C | T_{Max}, M)$ can be computed trivially over this network. The same idea is presented more formally below.

3.2.2 Maximum spanning tree model

We first introduce the following relation that equates the joint probability to a weighted summation over all spanning trees of the component C .

$$P(\mathbf{w}_C | M) \approx \sum_{T \in \text{spanning trees of } C} P(T | M) P(\mathbf{w}_C | T, M) \quad (15)$$

where $P(T | M)$ is the generative probability of tree T from model M and $P(\mathbf{w}_C | T, M)$ is the joint probability of terms in C given the tree structure of dependencies. We now assume that the probability mass of $P(T | M)$ is concentrated at the maximum spanning

tree T_{Max} of C . This assumption allows us to write:

$$P(\mathbf{w}_C|M) \approx P(\mathbf{w}_C|T_{Max}(C), M) \quad (16)$$

We estimate the right hand side of 16 by assuming a Bayesian network over T_{Max} obtained by traversing T_{Max} breadth-wise starting from any leaf node. We know that the joint probability thus defined is independent of the leaf node chosen, hence we can associate a unique Bayesian network with each T_{Max} . Given the Bayesian network, the joint probability is simply reduced to a product of first order conditional probabilities. For example, let's consider the graph in figure 1. For each component, we build a maximum spanning tree and define a Bayesian Network over the tree as shown in figure 2. As the figure indicates, we chose vertex w_2 as the starting leaf node in defining the Bayesian network over T_{Max} of component C_1 . Similarly, w_4 is the starting node of component C_3 . Component C_2 is a singleton node, hence there is no Bayesian network associated with it. Next, the joint probability of terms in each component is defined as follows.




Figure 2: Bayesian networks defined on each connected component's maximum spanning tree

$$\begin{aligned} P(\mathbf{w}_{C_1}|M) &\approx P(\mathbf{w}_{C_1}|T_{max}(C_1), M) \\ &= P(w_7|w_1, M)P(w_8|w_1, M) \times \\ &\quad P(w_1|w_2, M)P(w_2|M) \\ P(\mathbf{w}_{C_2}|M) &= P(w_3|M) \\ P(\mathbf{w}_{C_3}|M) &\approx P(\mathbf{w}_{C_3}|T_{max}(C_3), M) \\ &= P(w_5|w_6, M)P(w_6|w_4, M)P(w_4|M) \end{aligned} \quad (17)$$

Thus, we see that each sentence is modeled as a forest consisting of acyclic trees and singleton nodes, hence the name Sentence-Forest model.

3.3 Estimating the model parameters

As we have seen in section 3.2, to estimate the probability of sentence generation, all we need are the conditional probabilities of the form $P(w_i|w_j, M)$ for the dependencies in the Bayesian networks and unigram probabilities of the form $P(w_i|M)$ for the leaf nodes in the Bayesian nets and singleton nodes that have no dependencies associated with them. Typically these parameters of the model are estimated from a single document or a set of documents. Let D be the document from which we estimate the model parameters. We

use the maximum likelihood estimates for the model parameters as shown below:

$$P(w_i|M) \approx \lambda \frac{C(w_i, D)}{\sum_{w \in D} C(w, D)} + (1 - \lambda) \frac{C(w_i, GE)}{\sum_{w \in GE} C(w, GE)} \quad (18)$$

where $C(w, D)$ is the number of occurrences of term w in document D and GE is a large corpus of general English. We have used a smoothing parameter λ to prevent zero probabilities of unigrams.

The following maximum likelihood estimates are used to compute the joint and conditional probabilities:

$$P(w_i, w_j|M) \approx \frac{n_D(w_i, w_j)}{n_D} \quad (19)$$

$$P(w_i|w_j, M) \approx \frac{n_D(w_i, w_j)}{n_D(w_j)} \quad (20)$$

where n_D is the number of sentences in D and $n_D(w)$ is the number of sentences in document D that the argument w occurs in. Note that we have not used any smoothing in estimating the joint and conditional probabilities.

3.4 Likelihood ratio

Finally, the story's relevance score with respect to a model is given in terms of likelihood ratio which is defined by the following equation:

$$\begin{aligned} Score(D, M) &= \frac{P(D|M)}{P(D|GE)} \\ &= \frac{\prod_{S \in D} P(S|M)}{\prod_{w \in D} P(w|GE)} \end{aligned} \quad (21)$$

Note that in computing the probability of document generation with respect to the general English model $P(D|GE)$, we use the unigram model. Computing the likelihood ratio with respect to the general English model provides a sound basis for comparison of relevance and also serves as a normalizing factor to sentence and document lengths.

4. COMPUTATIONAL COMPLEXITY

In this subsection we discuss the time and space complexity of implementing the Sentence-Forest model with respect to a single story as the input. We will assume that the story contains N sentences and at most T terms per sentence.

The running time of each step in the algorithm is presented below:

1. Constructing the graph G of a sentence S with weighted edges: This requires computing the joint probability for all pairs of terms in the sentence, each of which can be done in constant time using hash tables. Thus, this step has a complexity of $O(T^2)$.
2. Building a maximum spanning tree for each of the component. The worst case arises when G is a fully connected graph. We use a greedy algorithm to build the MST. The running time of this step is $O(T^2 \text{Log}(T))$ if we use disjoint-set forest implementation with union by rank and path compression heuristics [2].
3. Defining a Bayesian network over the MST. This involves breadth-wise traversal starting from one of the leaf nodes.

Data set	Unigram	Sentence Forest
Train Set (6,361 pairs)	3	15
Test Set (27,538 pairs)	10	70

Figure 3: Comparison of average runtime of unigram and Sentence-Forest models on training and test sets in the link detection task: All values are in seconds.

The running time is $O(T + |E|)$. Since the number of edges in an MST spanning the entire graph is $T - 1$, breadth-wise traversal costs us only $O(T)$ time.

4. Computing the probability of occurrence of the sentence from the topic model and the background models: This requires generating the conditional probabilities of each edge in the MST and the unigram probabilities of the start node and any other singleton nodes, each of which can be done in constant time. Hence, this step has a complexity of $O(T)$, which is the size of the MST.

Thus, the overall running time per sentence is $O(T^2 \text{Log}(T))$. Thus, for the entire document, the complexity is simply given by $O(N) \times O(T^2 \text{Log}(T)) = O(T^2 N \text{Log}(T))$. In comparison, the unigram model has a time complexity of only $O(TN)$.

In practice, we notice that the run time of the Sentence-Forest model is only about 5-7 times higher than that of the unigram model as figure 3 indicates. We believe that despite the slightly lower speed, it is still fast enough for most real time applications, considering the fact that the Sentence-Forest model can process as many as 24,000 story pairs a minute.

Let us now look at the space-complexity of the model. Clearly, we need to index the joint and conditional probabilities for each pair of terms (w_i, w_j) in a sentence S in order to build the graph G (equation 5) and evaluate the joint probability of a sentence over the maximum spanning tree of each component in G (equation 16). Thus each sentence needs a space of $O(T^2)$, hence each document needs a space of $O(T^2 N)$. The unigram model, on the other hand, has a space-complexity of only $O(NT)$.

5. DISCUSSION

In this section we will try to justify modeling within-sentence dependencies and explain the intuition behind why we expect the new model to perform better than the traditional unigram model. Apart from effectiveness, we will also discuss the features of the algorithm that makes it efficient.

5.1 Why model sentences?

We believe that the most significant dependencies between terms occur within a sentence in the form of noun-phrases (e.g., gray-whale, balloon-adventure, Navy-sailor), proper names (e.g., Timothy-McVeigh, Bill-Clinton), associations (e.g., Glenn-NASA, asteroid-earth, Schindler-Oscar, Bush-Washington), etc. Dependencies do exist outside sentence boundaries too, but we posit that they do not tend to be as strong. This intuition explains several attempts in the past to incorporate noun phrases, proper name associations, etc. explicitly in retrieval models [3]. We believe our model includes all of those localized dependencies naturally into a unified probabilistic framework. From a computational viewpoint, modeling only within-sentence dependencies saves us considerable computational costs in comparison to modeling dependencies over the entire document (section 4), at the same time hopefully not compromising too much on effectiveness.

5.2 How might it help improve effectiveness?

Recall that our task is to compute the probability of generation of story D_2 with respect to a model $M(D_1)$ estimated from story D_1 . We will focus on the generative probability $P(S|M(D_1))$ of a particular sentence S in D_2 for illustration.

Consider examine the graph G of S defined with respect to the model of D_1 . Let w_i and w_j be two terms in S . As shown in equation 5, an edge (w_i, w_j) exists in G if and only if the probability of within-sentence-co-occurrence $P(w_i, w_j|M)$ is non-zero. Since we estimate $P(w_i, w_j|M)$ (see equation 19) entirely from D_1 's statistics, it follows that an edge (w_i, w_j) exists in G only if the terms w_i and w_j co-occur in at least one sentence in D_1 . Call all term pairs (w_i, w_j) of S that co-occur within a sentence boundary in D_1 'active' term pairs. Thus all edges in G correspond to active term pairs of S . If the edge (w_i, w_j) happens to be a part of the maximum spanning tree of one of the components C in G , then, the approximate generative probability of the sentence $P(S|M(D_1))$ will contain a term of the form $P(w_i|w_j, M)$ (see equation 17 for example). Thus in effect, the Sentence-Forest model searches for active term pairs in S and computes their conditional probabilities, while computing unigram scores for all other terms. In the absence of any active term pairs, the Sentence-Forest model smoothly collapses to a unigram generative model since the graph G would then only consist of independent singleton nodes.

While it is certainly not guaranteed by theory, it turns out that the maximum likelihood estimate of the conditional probability $P(w_i|w_j, M)$ is almost always an order of magnitude higher than the corresponding unigram estimate $P(w_i|M)$. This can be attributed to the fact that co-occurrence of any pair of terms is typically a much rarer event than each of their occurrences. Thus, if the sentence S contains many active term pairs, the Sentence-Forest model boosts the generative probability $P(S|M)$ via the conditional probabilities.

We expect that the sentence S contains considerable number of active term pairs if D_2 , the document that contains S , is on the same topic as that of D_1 and none otherwise. In such a scenario, if D_2 is on-topic, the score assigned by Sentence-Forest model is typically much higher than unigram score. If D_2 is off-topic with respect to D_1 , the generative probability computed by the Sentence-Forest model corresponds to unigram score. Thus we expect that on-topic pairs and off-topic pairs are well-separated by the Sentence-Forest model resulting in better performance.

5.3 What makes it efficient?

A simple unigram model has proven not only very effective in most IR tasks, but is also one of the most efficient algorithms in terms of its time and space complexity, mainly owing to its simplicity. Any model that incorporates more information than simple term distributions alone seems to invariably loses out to unigram model on efficiency. We believe that we have managed to improve on the effectiveness of the unigram model considerably while limiting the loss in efficiency within reasonable bounds for all practical purposes.

As we have seen in section 4, the time complexity of the Sentence-Forest model is only a factor of $O(T \text{Log} T)$ higher than that of the unigram model, while the space complexity is higher only by a factor of $O(T)$, where T is the maximum number of terms in a sentence.

One of the reasons for its efficiency being almost comparable with that of the unigram model is the fact that we take into account only within-sentence co-occurrences as dependencies, thereby reducing the number of first-order dependencies to evaluate from a maximum of $O((TN)^2)$ to just $O(T^2 N)$, a reduction by a factor

of $O(N)$.

Another important reason that makes the model suitable for most real-life systems is the fact that the model does not require one to estimate the conditional probabilities $P(w_i|w_j, M)$ over the entire space $|V| \times |V|$, where $|V|$ is the vocabulary size of the topic of M . Recall that a necessary (but not sufficient) condition that we compute conditional probability $P(w_i|w_j, M)$ for a pair of terms (w_i, w_j) is that they are active (see subsection 5.2) in document D from which we estimate the model M . Thus, for each document D from which we estimate a model $M(D)$, it is enough for us to index *a priori* the conditional probabilities of only those term pairs that co-occur in the same sentence in D . The number of these term pairs is much less than the entire space resulting in considerable savings.

There is another factor that keeps our model very efficient: as equation 21 indicates, we compute only unigram probabilities to estimate the generative probability of a document from the general English model GE . Note also that we do not use smoothing in estimating the conditional probabilities $P(w_i|w_j, M)$ (see equation 20). This effectively means that we never have to compute a conditional probability distribution with respect to the general English model, which results in a big saving in terms of both time and space. Note that since conditional probabilities are evaluated only for active term-pairs, there is no zero-probability problem and hence smoothing is rendered unnecessary. Nevertheless, smoothing is found to improve effectiveness in many applications. In our experiments, we have observed that smoothing the estimates of conditional probabilities with estimates from a general English model did not result in any significant improvement in effectiveness. Hence we have abandoned smoothing the conditional probability estimates in return for huge savings in time and space.

5.4 Comparison with the Sentree model

The current model is an improvement on the SenTree model [10]. In the SenTree model, for any sentence S , we construct a maximum spanning tree (MST) over the fully-connected sentence graph with respect to the statistics of the topic model M . Thus, even when topic-related dependencies may exist only among a small fraction of term-pairs in a sentence, the MST imposes a dependency structure over the entire sentence. As a result, most of the dependencies evaluated by the SenTree model turn out to be spurious, resulting in degraded performance [10]. The Sentence-Forest model seeks to overcome this drawback by adaptively modeling only topically relevant dependencies and exhibiting a unigram-like behavior by default. Thus it combines the best properties of both the SenTree and the unigram model into a single theoretical framework.

In terms of efficiency too, the Sentence-Forest model outperforms the Sentree model. The Sentree model requires smoothing of the conditional probabilities with the general English estimates to overcome the zero-probability problem, which involves huge time and space overheads. The Sentence-Forest model, on the other hand, computes conditional probabilities only for ‘active’ term pairs and hence does not require smoothing.

6. IMPLEMENTATION DETAILS

In this section, we discuss some of the system implementation details that are not covered in the discussion on methodology of the model.

During the process of training the model, we noticed that our expectation that active term pairs occur only in on-topic story pairs (see section 5) is not strictly valid. In fact, our data analysis revealed that many common term pairs such as ‘people say’, ‘officials stated’ co-occur within a sentence boundary in both stories

in many off-topic story-pairs. Such term co-occurrences do not carry much information about the topical content of the stories, but the Sentence-Forest model incorrectly assigns edges between such terms. These ‘spurious’ edges ultimately creep in as conditional probabilities and boost the scores of off-topic pairs. This results in a large number of false alarms resulting in a degradation in the performance of the model.

To suppress the effect of these spurious term pairs, we modified the definition of edges in equation 5 to include the following extra condition.

$$E = \{(w_i, w_j) \mid w_i, w_j \in V; P(w_i, w_j|M) > 0; L(w_i)L(w_j) > C\} \quad (22)$$

where

$$L(w) = \frac{P(w|M)}{P(w|GE)} \quad (23)$$

$L(w)$ is the likelihood ratio of the probability of term w given the model with respect to a general English model. It is similar to the familiar tf-idf weight and tells us the relative importance of the term with respect to the topic model. Thus, the modified definition of an edge in equation 22 requires the model to form an edge (w_i, w_j) only if the ‘relative importance’ of both the terms exceeds a threshold C . We empirically determined the best value of C to be 2500 and fixed it at this value on all our runs.

6.1 Tools

Clearly, the sentence-Forest model requires us to identify the boundaries of sentences in documents. We have used a simple heuristic-rule based sentence segmenter to detect sentence boundaries. We refer the reader to [10] for more information on the heuristic rules used.

Additionally, we have used a list of 423 most frequent words to remove stop words from stories. Stemming is done using the Porter stemmer [12] while the indexer and the Sentence-Forest model are implemented using Java.

7. TOPIC DETECTION AND TRACKING

The new model we present in this work is expected to address some of the issues in Topic Detection and Tracking (TDT) [1]. TDT is a research program investigating methods for automatically organizing news stories by the events that they discuss. It includes several evaluation tasks, each of which explores one aspect of that organization. In this section, we describe two tasks called Link Detection and Topic Tracking.

7.1 Link Detection task

Link Detection requires determining whether or not two randomly selected stories discuss the same topic. In the language modeling approach to link detection, we build a topic model $M(D_1)$ from one of the stories D_1 in the pair (D_1, D_2) . We then compute the generative probability of the second story D_2 from the model $M(D_1)$ as shown below.

$$Score(D_1, D_2) = P(D_2|M(D_1)) \quad (24)$$

Sometimes we may compute a *two-way score* to add symmetry to the formula, as shown below:

$$score(D_1, D_2) = \frac{1}{2}(P(D_2|M(D_1)) + P(D_1|M(D_2))) \quad (25)$$

If the score exceeds a pre-determined threshold, the system decides the two stories are linked. The system’s performance is evaluated using a topic-weighted DET curve [9] that plots miss rate

against false alarm over a large number of story pairs, at different values of decision-threshold. A Link Detection cost function C_{link} is then used to combine the miss and false alarm probabilities at each value of threshold into a single normalized evaluation score [18]. We use the minimum value of C_{link} as the primary measure of effectiveness and show DET curves to illustrate the error trade-offs.

7.2 Topic Tracking

The TDT topic tracking task is defined to be the task of associating incoming stories with topics that are known to the system. Each target topic is defined by a set of training stories. The tracking task is then to classify correctly all subsequent stories as to whether or not they discuss the target topic.

Similar to the link detection task, we estimate a model M_T for topic T from the set of its training stories. Given any subsequent story D , we measure the probability of its generation with respect to the model $P(D|M_T)$. The system decides that the story is on-topic if the score exceeds a pre-defined threshold. We use topic weighted DET curves and a minimum cost value to evaluate the system’s performance, as in link-detection task.

8. EXPERIMENTS AND RESULTS

In this section, we describe the experiments we performed on link-detection and topic-tracking and present the results we obtained.

8.1 Story Link Detection

The training set we used is a subset of the TDT2 corpus that includes six months of material drawn on a daily basis from six English-only news sources that are manually transcribed (when the source is audio). The collection comprises 6,361 story pairs and relevance judgments for each of them. The test set is the TDT3 corpus consisting of the 27,538 manually transcribed story pairs from multiple languages. If the source is non-English, we have used machine translations that are made available in the corpus. To derive the general English unigram model, we have used the same corpora as we performed experiments on.

We first trained the unigram model on the training set and used its best performance (at $\lambda = 0.2$) as the baseline for all the experiments. In our training experiments on the Sentence-Forest model, we performed a search for the best performing values of the parameters λ and L . The DET curve of the best performing values of $\lambda = 0.2$ and $L = 2500$ is shown in figure 4. It is clear from the plot that the Sentence-Forest model outperforms the unigram model. This is reflected in the fact that we were able to bring down the minimum cost C_{link} by 7% from that of the unigram model as shown in the same figure.

Having found the best performing values of various parameters, we now ran the system on the test set with the parameters set to the best training values. The performance of the system on the test set is shown in figure 5. Once again, we notice that the Sentence-Forest model consistently outperforms the unigram model resulting in a reduction of about 4% in the minimum cost.

8.2 Topic Tracking

We used the multi-lingual TDT3 corpus in our tracking experiments. The general English unigram model is derived from about 40,000 stories in the TDT2 corpus. We used manual transcriptions and machine translations to English wherever necessary. There were 38 topics to track and we provided one English training story per topic ($N_t = 1$). The topic model of each topic is computed




Figure 4: Performance of Sentence-Forest model on the training set (Link detection task)




Figure 5: Performance of Sentence-Forest model on the test set (Link detection task)

from the single training story and remains static (no adaptation is used) during the entire run. We used the best parameters obtained from training in the link detection task in our runs of unigram and Sentence-Forest models.

Figure 6 compares the performance of the Sentence-Forest model with that of the unigram model. We notice trends similar to the link detection task. The minimum cost of the Sentence-Forest model is about 6.4% lower than that of the unigram model. In both figures 5 and 6, we note that the Sentence-Forest DET curve dominates the unigram curve, so it’s even better than reducing the minimum cost: false alarm rate is reduced at all miss rates.

8.3 Comparison to state-of-the-art TDT systems

The experiments we have performed on the link detection task do not correspond to any official conditions, but the tracking run does correspond to the standard evaluation conditions of TDT 2001. The best system in that evaluation was Limsi [7] which achieved a minimum cost of 0.0959 (compared to the unigram cost of 0.2498 and Sentence-Forest minimum cost of 0.2339). We note that the state-




Figure 6: Performance of Sentence-Forest model on Tracking task

of-the-art TDT systems like Limsi [7] and Relevance Models [6] typically rely on elaborate boosting techniques such as document expansion and unsupervised adaptation while employing unigram model as their core technology.

The goal of the present work is to demonstrate that incorporating local dependencies within the language modeling framework improves effectiveness. Hence we have used the unigram model as our baseline in all our runs. Although the performance of our system is well short of the state-of-the-art systems, we believe our model is a contribution at a fundamental level, and it can be easily extended to accommodate several boosting techniques used in the state-of-the-art systems.

9. CONCLUSIONS AND FUTURE WORK

In this work, we have presented a new approach of modeling a document by capturing local dependencies in documents through adaptive Bayesian networks. Although this approach is built towards addressing a specific TDT task, we believe that the generality of the model permits one to apply it to any text classification or retrieval task.

Our experiments show that the Sentence-Forest model outperforms the baseline unigram model on different data sets and on two different TDT tasks. To the best of our knowledge, this is the only probabilistic model that performs better than unigram model using statistics from the document alone. As such, we think the most important contribution of this study is the evidence we have provided that the performance of the existing IR systems can be improved by employing more sophisticated language models that capture positional and other syntactic information in documents.

As part of our future work, we would like to test the performance of our model on other IR tasks such as ad-hoc retrieval. We also envision building a query or document expansion device using the Sentence-Forest model as the basic framework. Our hope is that the new expansion device may better other expansion models based on the unigram approach since our basic model outperforms the unigram model.

Acknowledgments

We would like to thank Andrew McCallum and Victor Lavrenko for their valuable comments. This work was supported in part

by the Center for Intelligent Information Retrieval and in part by SPAWARSYSCEN-SD grant numbers N66001-99-1-8912 and N66001-02-1-8903. Any opinions, findings and conclusions or recommendations expressed in this material are the author(s) and do not necessarily reflect those of the sponsor.

10. REFERENCES

- [1] Allan, J., Lavrenko, V. and Swan, R. Explorations Within Topic Tracking and Detection, *Topic Detection and Tracking: Event-based Information Organization*, James Allan, Editor, Kluwer Academic Publishers, 197-224, 2002.
- [2] Cormen, T. H., Leiserson, C. E. and Rivest, R. L. *Introduction to Algorithms*, MIT Press, 1990.
- [3] Croft, W. B., Turtle, H. R. and Lewis, D. D. The use of phrases and structured queries in information retrieval, *ACM SIGIR*, 32-45, 1991.
- [4] De Finetti, B. *Theory of Probability*, 1:146-161, Wiley, London 1974.
- [5] Fung, R. M., Crawford, S. L., Appelbaum, L. A. and Tong, R. M. An architecture for probabilistic concept-based information retrieval, *ACM SIGIR*, 455-467, 1990.
- [6] Lavrenko, V., Allan, J., DeGuzman, E., LaFlamme, D., Pollard, V. and Thomas, S. Relevance models for Topic Detection and Tracking, *Proceedings of the Conference on Human Language Technology (HLT)*, 2002.
- [7] Lo, Y. and Gauvain, J., The Limsi Topic Tracking System for TDT 2001, *Proceedings of TDT 2001 workshop*.
- [8] Manning, C. D. and Schütze, H. *Foundations of Statistical Natural Language Processing*, MIT Press, 1999.
- [9] Martin, A., Doddington, G., Kamm, T. and Ordowski, M. The DET curve in assessment of detection task performance, *EuroSpeech*, 1895-1898, 1997.
- [10] Nallapati, R. and Allan, J. Capturing Term Dependencies using a Language Model based on Sentence Trees, *ACM CIKM* 383-390, 2002.
- [11] Ponte, J. M. and Croft, W. B. A Language Modeling Approach to Information Retrieval, *ACM SIGIR*, 275-281, 1998.
- [12] Porter, M. F. An algorithm for suffix stripping, *Program*, 14(3):130-137, 1980.
- [13] Rijsbergen, V. *Information Retrieval*, Butterworths, 1979.
- [14] Robertson, S. E. and Bovey, J. D. Statistical Problems in the Application of Probabilistic Models to Information Retrieval, *Technical Report, Center for Information Science, City University*, 1982.
- [15] Song, F. and Croft, W. B. A General Language Model for Information Retrieval, *Information and Knowledge Management*, 1999.
- [16] Turtle, H. R. and Croft, W. B. *Inference Networks for Document Retrieval*, *ACM SIGIR*, 1-24, 1990.
- [17] Wong, S. K. M., Ziarko, W. and Wong, P. C. N. Generalized Vector Space Model in Information Retrieval, *ACM SIGIR* 18-25, 1985.
- [18] The Topic Detection and Tracking evaluation phase 2 plan, <http://www.nist.gov/speech/tests/tdt/tdt98/doc/tdt2.eval.plan.98.v3.7.pdf>.