

Recovering Implicit Thread Structure in Newsgroup Style Conversations

Yi-Chia Wang, Mahesh Joshi

Language Technologies Institute
Carnegie Mellon University
Pittsburgh, PA 15213
{yichiaw,maheshj}@cs.cmu.edu

William W. Cohen

Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213
wcohen@cs.cmu.edu

Carolyn Rosé

Language Technologies Institute/
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15213
cprose@cs.cmu.edu

Abstract

On-line discussions are composed of multiple inter-woven threads, regardless of whether that threaded structure is made explicit in the representation and presentation of the conversational data. Recovering the thread structure is valuable since it makes it possible to isolate discussion related to specific subtopics or related to particular conversational goals. In prior work, thread structure has been reconstructed using explicit meta-data features such as “posted by” and “reply to” relationships. The contribution of this paper is a novel approach to recovering thread structure in discussion forums where this explicit meta-data is missing. This approach uses a graph-based representation of a collection of messages where connections between messages are postulated based on inter-message similarity. We evaluate three variations of this simple baseline approach that exploit in different ways the temporal relationships between messages. The results show that the three proposed approaches outperform the simple threshold-cut baseline.

Introduction

As the language technologies community focuses more and more on the analysis of social media rather than primarily focusing on expository text, the characteristics that set conversational data apart must be considered explicitly. Whereas expository text is organized around the information that is presented, typically in a single coherent argument, conversation is organized around a shared task where multiple participants act and react to one another as the scenario unfolds. Because the interaction is not, and indeed cannot, be fully planned out ahead of time, the structure emerges from the complex and intertwined separate intentions of each of the actors as they participate in their joint task. It has been widely noted that the temporal ordering of contributions to an ongoing conversation can be exploited in the analysis of social media (Kleinberg, 2002; Kumar et al., 2003; Dubinko et al., 2006). Nevertheless, with multiple actors, each with their own goals and discourse intentions, it cannot be assumed that the conversation proceeds with a single focus at a

time, which would result in a simple, treelike structure (Grosz and Sidner, 1986). Rather, conversational data may have a complex, multi-threaded structure (Rose et al., 1995; Wolf and Gibson, 2005). Thus, limiting the structural representation of conversational data to a singular sequential ordering is overly simplistic and limiting.

Isolating the individual threads that make up a complex, multi-party interaction plays a pivotal role in making sense out of the individual contributions on multiple levels. For example, from a group knowledge building perspective, teasing out the individual threads is valuable for characterizing the consensus building style of the actors (Wang et al., 2007). On the level of conversation acts, identifying which contributions fit together on a single thread is useful for identifying adjacency pairs and thus leveraging the notion of sequential relevance for disambiguating the actions performed (Stolcke et al., 2000). Thread structure can also be used to identify particularly influential contributions by analyzing backward links between messages to identify which contributions were responsible for stimulating the longest discussion threads. Polarity prediction within the area of sentiment analysis may benefit from the identification of thread structure. For example, if Participant C disagrees with B’s disagreement with A’s opinion, then C agrees with A. This cannot be correctly resolved from the characteristics of individual contributions in the absence of the discourse structure.

Consistent with this view, thread structure has been advantageously applied to many different research problems in the area of social media analysis, such as email summarization (Carenini et al., 2007), text classification task in discussion forums (Wang et al., 2007), and newsgroup search (Xi et al., 2004). In this prior work, meta-data that makes the thread structure explicit played an important role in identifying the parent child relationships between conversational contributions. Unfortunately, there are many popular conversation streams that are not organized explicitly into threads. The lack of explicit thread representation may limit the exploitation of thread structure in the analysis of social media. In contrast to this prior work, we develop an algorithm for recovering thread structure where it is not made explicit in associated meta-data that accompanies conversational contributions.

As an example of such a data set, for our analysis, we use a data set collected within a multi-party educational game environment called LegSim (Wilkerson and Fruland, 2006). The LegSim environment is a simulated legislature where students play the roles of state representatives, who are responsible for proposing and examining bills. In this environment, students engage in debates possessing a complex, threaded structure. However, that thread structure is neither explicit in the stored representation of the data nor from the interface that the students use to engage in the multi-party conversation.

The contribution of this paper is a novel approach to recovering thread structure in discussion forums where explicit meta-data for reconstructing the thread structure is missing. We refer to this task as the *thread structure recovery task*. In particular, thread structure recovery is the process whereby a parent message is explicitly linked to one or more responding child messages. Our approach consists of two steps. First, we use a connectivity matrix constructed using a shallow message similarity measure to represent the implicit relationships between messages within the same conversation stream. Each element in the matrix is a weighted link between a message pair, which measures the content similarity between the messages. The second step is to determine parent-child relationships using the connectivity matrix. A simple baseline approach merely uses a threshold to identify which similarity measures are strong enough to constitute an explicit link. We propose three variations of this baseline approach that leverage in different ways the temporal relationships between messages.

In the remainder of the paper, we will first review some important related work in section. Next, we will describe the formal definition of the task and dataset as well as the details of our algorithms. Then, we will present our evaluation by comparing the effectiveness of the alternative versions of our basic approach. We conclude with some directions for our continued investigation.

Related Work

As discussed in the previous section, some recent work has demonstrated that explicit thread structure can be advantageously utilized in the analysis of social media. First, it can be used for boosting the performance of automatic analysis. For example, Carvalho and Cohen (2005) exploited the sequential ordering among emails within the same conversation thread to improve the performance of email-act classification. Also, Xi et al. (2004) introduced an effective ranking function for newsgroup search. In that work, Xi and colleagues extracted a set of meta features related to the structure of newsgroup discussion threads, such as the number of children of the current message, the position of the current message in the thread tree, and so on. Similarly, Wang and colleagues (2007) developed a novel approach to leveraging context for classifying newsgroup style discussion segments. Their solution was to use the explicit

thread structure to identify pairs of potentially related segments of text. They then used shallow semantic similarity metrics to compute a feature that indicated the maximum similarity value between a segment contributed by one participant and other potentially related segments contributed by other participants in the discussion thread. This added feature significantly improved performance at identifying the consensus building style exhibited by the individual discussion contributions over a baseline using only word level features.

Recently, some research related to thread structure has been done in the context of email summarization work, where recovering thread structure is necessary for constructing coherent summaries of email conversations. To that end, Yeh et al. (2006) studied how to use quoted messages, string matching metrics, and email headers to reassemble the treelike thread structure of email conversations. Inspired by Yeh et al.'s work, Carenini et al. (2007) proposed a new email summarization framework, ClueWordSummarizer. They used what they refer to as a fragment quotation graph to organize conversations and subsequently construct lexical chains. They utilized quotation markers (e.g., ">") as well as quotation depth (the number of quotation marks in the beginning of every line) to build the quotation graph.

In addition to facilitating the automatic analysis of social media, identification of thread structure can be used to aid in the work that conversation analysts do with the data by hand. For example Trausan-Matu et al. (2007) presented a conversation analysis support tool that provided a visualization of a threaded conversation. Topics were identified by means of expressions such as "Let's talk about" or "What about", and then the WordNet ontology was used for recovering threads by means of automatically constructed lexical chains. The visualization assisted analysts in identifying meaningful subsets of messages to examine together as a coherent conversational unit.

Our work is most related to Shen et al.'s work (2006), in which they introduced a partial solution to the thread reconstruction task for synchronous chat data. In their work, a single-pass clustering algorithm was modified to exploit the temporal information in the streams by computing centroids that updated dynamically as the analysis proceeded sequentially through the data stream. The purpose of this analysis was to identify subtopic clusters of contributions. Our work is similar in that we also exploit temporal relationships between contributions. However, their approach fell short of a complete solution in that it did not identify the explicit parent-child relationships between contributions, as we do in our work. Also, Shen et al. employed a different data structure and search technique. Specifically, they employed an incremental clustering algorithm in which they dynamically computed centroids for their clusters based on a combination of semantic similarity and temporal proximity. In our work, we introduce an approach to recover thread structure using a graph-based connectivity matrix, which similarly uses both similarity and temporal

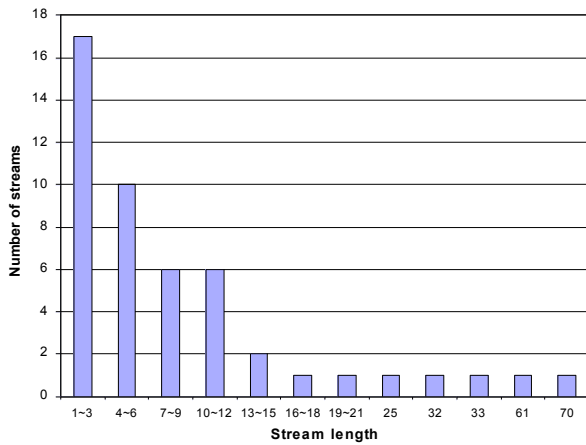


Figure 1. Distribution of Conversation Stream Size.

information. Our evaluation demonstrates the importance of both types of information, and contrasts alternative approaches to combining temporal cues with similarity based predictions.

Problem Definition

In this section, we introduce our dataset and offer an example to illustrate the thread reconstruction problem.

LegSim Corpus

The LegSim¹ corpus was provided to us by John D. Wilkerson at University of Washington in Seattle. That dataset contains a collection of messages contributed within an educational multi-player web-based teaching tool called LegSim. LegSim has been in instructional use since 2000, and was used by approximately 1400 high school, community college and university students during the 2005-06 school year. It includes features that allow students to participate in the form of online discussions. It simulates a legislative environment where students in a government or public policy course assume representational roles and advance legislative proposals that reflect their own political priorities.

We chose the data from one college level Civics course from Fall 2006. 93 students participated in the course, 71 of which actively participated in the on-line discussion. They were able to post comments and opinions regarding the proposed legislation (a.k.a. a bill). Over the course of the semester, 478 messages were posted. In this data, the messages were organized into 48 separate streams, as determined by which bill they referred to. The histogram in Figure 1 shows the detailed distribution of number of messages per stream. The longest stream is composed of 70 messages, while the shortest one contains only 1

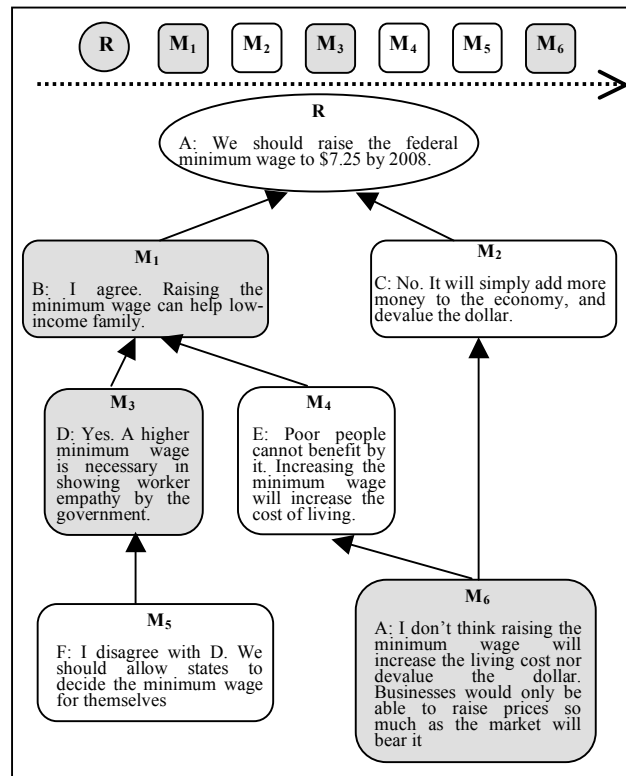


Figure 2. A thread example in LegSim and the goal of thread recovery.

message. The average number of messages per stream is 9.98.

A Simple Thread Example

Figure 2 displays a simplified conversation example from the LegSim dataset. The nodes above the dashed arrow denote a discussion stream, arranged from left to right in the same order in which they were posted. Every discussion stream consists of discussion related to a specific topic, in this case a proposed bill that is under debate. For example, the topic of this example is whether the government should raise the minimum wage. Participant A posted the root message R to initiate the discussion. Following that is a sequence of reply messages (M1 to M6), which respond to previous messages. The graph illustrates the thread relationships among the messages in the discussion stream. A directed link in the graph represents a parent-child relationship, with arrows that point from child to parent (e.g., M3→M1). In conversational data, it is possible for a message to have multiple parents (i.e., one message can reply to several previous messages at the same time.) Take the message M6 as an example. It expresses a disagreement with both M2 and M4. Similarly, a parent may have more than one child. For example, both M3 and M4 respond to M1.

Because of the lack of explicit thread structure information in the dataset, before we began any of our experimentation, we manually reconstructed the thread structure by assigning links between messages based on

¹ <http://www.legsim.org/>

their observed semantic relationship and explicit discourse markers. For example, the following is part of a discussion example from the dataset in which two “representatives” are discussing whether they should pass a bill that proposes “Local Law Enforcement Celebration Day”.

Representative X said “Every day our Police Officers put their life in danger to protect and serve the people. I feel that this bill is the very least that we could do to express our gratitude. It is non-controversial so this should pass quickly.”

Representative Y said “I agree... it is the least we can do to show our appreciation... I think it should easily pass also... Thanks Rep Tung for proposing it.”

It is clear semantically that the message by Representative Y comments on the message by Representative X. First, explicit discourse markers, namely “I agree” as well as “also”, mark that this message is a response to an earlier message. The statement “It should pass” echoes part of the earlier message. Also, its sequential position, adjacent to the earlier message, offers a clue. So, during our manual annotation process, we assigned a parent-child link between these two message based on these types of clues.

We verified that our annotation of parent-child relationships was reliable. Since we expected that the thread structure would tend to be more complicated for longer threads, we evaluated our agreement over three randomly selected threads from those with longer lengths (within the upper most quartile). Altogether these three threads (of lengths 10, 19, and 32 respectively) constituted 15% of the corpus. Two humans annotated these three threads using the definition outlined above. The Cohen’s kappa is 0.87 for the three threads. This shows that the annotators achieved a high level of consensus for the link annotation task.

The two human annotators were trained to use the following discourse clues to identify parent-child links when doing annotation. *Response markers*, namely “I agree with X” as well as “In response to X”, mark that a message is a response to an earlier message. *Addition markers*, such as “We also have to remember...” and “That’s still wrong”, mark that the message has some additional comments for the earlier discussion or arguments against a previous post. *Echo markers* echo part of the earlier messages. *Question and answer* in a candidate parent and the current message respectively can link two messages together, too. Finally, *subtopic phrases* have the potential to chain messages together. For example, there are two subtopics (helping low-income families versus devaluing the dollar) in the discussion of the bill “whether the government should raise the minimum wage.” The phrases “low-income family” and “poor people” mark that a message should belong to the former subtopic discussion. A phrase such as “devalue the dollar” or “economy”, on the other hand, suggests that the message is talking about the later subtopic.

Among the 48 streams mentioned above, 39 streams have two or more messages. After manual reconstruction of the thread structure, 28 of these 39 streams contained at least one parent-child link. We use these 28 streams in the experiments reported below.

Definition of Thread Recovery Task

Given a discussion stream in which messages are sorted by the posting time, the thread recovery task is to construct the thread graph (as in the example in Figure 2). The links in the graph identify which of the previously contributed message the current message replies to. In our problem, three assumptions are made:

Assumption 1: The default parent of each reply message is the initial message of the stream. Thus, that is the assignment made by the algorithm if no parent is identified further down the thread. As we have mentioned above, all messages in the same stream are all about one single topic, and the root message is the one that starts the discussion.

Assumption 2: We assume each message can only link back to those messages that were posted before it in time.

Assumption 3: Except for the posting time of messages, no meta information is available, such as author names, quotation marks, and message titles. We make this assumption since this information is not always available, and the purpose of our experimentation is to determine what can be used to recover thread structure in the absence of this meta-information.

These assumptions take the form of constraints that are enforced in all versions of our algorithm described below.

Alternative Approaches to Thread Recovery

In this section, we first explain how to build a connectivity matrix to reflect semantic similarity between messages. A threshold is then applied to the weighted links in the matrix in order to identify the thread links. To leverage the sequential nature of conversation streams, we describe three variations to the baseline algorithm, which are able to take temporal information into account.

Data Representation

We utilize a vector space model to represent messages. A message is a vector of terms², which are weighed by the term frequency (TF) and inverse document frequency (IDF). Here, TF stands for the number of times a term occurs in a message. IDF is the log of the number of messages in the stream divided by the number of messages in the stream that contain the term. TF is normalized, as displayed in the formula below. There are many variations available for computing TF.IDF. Following is the one we used, which is the UMass version of the leading Okapi algorithm (Robertson et al., 1994):

² In our experiment, we used non-stemmed unigrams as features.

$$TF.IDF = \frac{tf}{tf + 0.5 + 1.5 \cdot \frac{msglength}{avg_msglength}} \cdot \log\left(\frac{N}{df}\right)$$

where N is the number of messages in the conversation stream.

Intuitively, we selected TF.IDF for our term weights in order to emphasize the most salient terms in the messages in our message similarity measure. This is based on the intuition that what makes a message distinctive within its stream is what is most likely to elicit a response. However, there are instances where this may not be the case. Thus, while using TF.IDF term weights may have some desirable consequences, we cannot rule out at this point the possibility that a different term weight scheme may have been more appropriate. We leave experimentation along these lines for future research

Message-to-Message Graph Building

Given a time-ordered sequence of messages $M = \{ m_i \mid 1 \leq i \leq n \}$ ³ in a conversation stream, we view each message as a node in a graph. We then build a directed graph by creating edges from m_j to all messages before it in time (i.e. m_i where $i < j$). The associated weight of the edge is computed with the cosine similarity measure. We use a connectivity matrix \mathbf{W} to represent the graph, where each element w_{ij} in \mathbf{W} denotes the weighed edge from m_i to m_j in the graph. The formal definition of similarity matrix $\mathbf{W} = [w_{ij}]_{n \times n}$ is given below:

$$w_{ij} = \begin{cases} \frac{\vec{m}_i \cdot \vec{m}_j}{\|\vec{m}_i\| * \|\vec{m}_j\|}, & \text{if } i > j \\ 0, & \text{otherwise} \end{cases}$$

where \vec{m}_i is the term vector of message m_i . Note that according to the definition, \mathbf{W} is a lower triangular matrix where the main diagonal is zero.

Thread Structure Recovery

Once we have the message-to-message matrix \mathbf{W} , which represents semantic similarity between messages, we are ready to reconstruct the thread structure for M . First, we describe our baseline algorithm (GR_B). In order to decide whether two messages should be linked by a parent-child relationship, a *threshold* is employed as the baseline. Only when the cosine similarity between messages, i.e., w_{ij} in \mathbf{W} , is larger than *threshold* we will assign a link for them.

³ Note that m_j is the first message immediately following the start message. The reason that we didn't take the root message into account is because all messages in the same stream are supposed to discuss the same topic initiated by the root. Thus, it is likely that the similarity score between root and any of its child messages will normally be high relative to that of other pairs of child messages, which could interfere with thread recovery.

Considering each entry in \mathbf{W} , we get the baseline thread graph $\mathbf{G} = [g_{ij}]_{n \times n}$.

$$g_{ij} = \begin{cases} 1, & \text{if } w_{ij} > \text{threshold} \\ 0, & \text{otherwise} \end{cases}$$

where g_{ij} indicates a link from message m_i to m_j and is set to 1 only if w_{ij} is larger than *threshold*. By adjusting the threshold, we can explore different degrees of similarity between messages.

After informally observing some thread instances in our dataset, we observed that a child message is usually not far away from its parent. As a consequence, we introduced three penalizing functions for \mathbf{W} , which can differently leverage the temporal ordering and distance between messages. We consider the relative position of messages instead of using real time-stamps.

Penalizing \mathbf{W} Using Fixed Window Size (GR_F). The central idea behind the first variation on the baseline algorithm is that only messages within a certain window size are considered to be the potential parents of some message. More specifically, a fixed time window is employed to transform the similarity matrix \mathbf{W} . This yields a new similarity matrix $\mathbf{A} = [a_{ij}]_{n \times n}$:

$$a_{ij} = \begin{cases} w_{ij}, & \text{if } |i - j| < s \\ 0, & \text{otherwise} \end{cases}$$

where s is the window size we consider; and $|i - j|$ is the distance between message m_i and m_j . The algorithm for detecting links here (and also in the other two variations) is the same as the baseline except that we substitute the baseline similarity matrix \mathbf{W} with \mathbf{A} . Note that there are two parameters – s and *threshold* – that need to be specified in this penalizing strategy.

Penalizing \mathbf{W} Using Dynamic Window Size (GR_D). Considering that the length of different conversation streams may vary, it might be beneficial if the window size can dynamically reflect the length of the stream. Thus, we use a dynamic window size and use it to compute the new weighted matrix $\mathbf{B} = [b_{ij}]_{n \times n}$ as shown below:

$$b_{ij} = \begin{cases} w_{ij}, & \text{if } |i - j| < p \cdot n \\ 0, & \text{otherwise} \end{cases}$$

where n is the number of messages belonging to the conversation stream M ; p denotes the percentage of previous messages we want to consider when determining the parents of the current message. This approach also involves two parameters – p and *threshold*.

Penalizing \mathbf{W} Using Time Distance (GR_T). Finally, instead of using an indicator function to indicate a subset of elements in \mathbf{W} , we penalize the similarity by the time distance between the two messages. The result similarity matrix $\mathbf{C} = [c_{ij}]_{n \times n}$ is defined as following:

$$c_{ij} = \begin{cases} \left(\frac{1}{|i - j|} \right) \cdot w_{ij}, & \text{if } i \neq j \\ 0, & \text{otherwise} \end{cases}$$

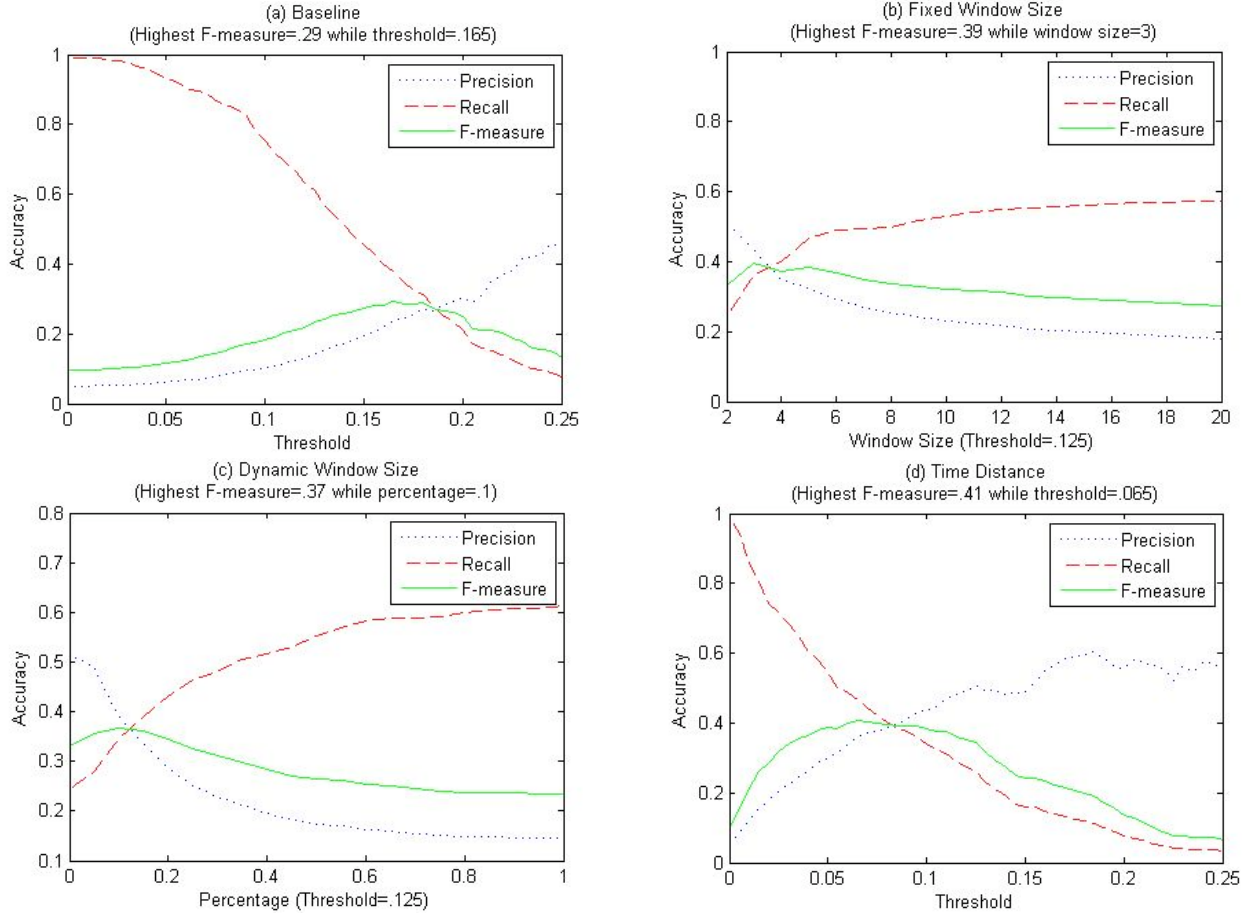


Figure 3. Performance of GR_B , GR_F , GR_D , and GR_T while parameters vary.

where the reciprocal of $|i - j|$ is used to decrease the cosine similarity value in term of the distance between the two messages. For this variation, we only need to consider the similarity *threshold*.

Experiments

We have described the baseline algorithm for recovering thread structure as well as the three variations of it that make use of temporal information. In this section, we compare these algorithms in order to evaluate the contribution of temporal information to the thread structure recovery task. We first introduce the evaluation metrics and then the experimental results. The final part of this section shows the results of parameter tuning.

Evaluation Methods

We use precision, recall, and F-measure to evaluate our results. We explain the definition for each of the three metrics in terms of the thread recovery task as follows:

$$Precision = \frac{|(real\ links) \cap (predicted\ links)|}{predicted\ links}$$

$$Recall = \frac{|(real\ links) \cap (predicted\ links)|}{real\ links}$$

$$F-measure = \frac{2 * Precision * Recall}{Precision + Recall}$$

where we employ the balanced *F-measure* in which precision and recall are evenly weighted.

Results

In our evaluation, we compare the three variant algorithms with the baseline algorithm that only takes similarity into account. The performance curves are shown in Figure 3.

Before we present a statistical comparison of tuned performance across conditions, we will first illustrate how our four approaches vary as we manipulate the parameters that are associated with them. Both Time/Distance and Similarity are used to narrow the range of previous messages that are considered as potential parent messages. Thus, as the threshold on Time or Distance is made more

