

Discovering Ephemeral Associations among News Topics*

M. Montes-y-Gómez¹, A. Gelbukh¹, A. López-López²

¹ Centro de Investigación en Computación (CIC-IPN), México.
mmontesg@susu.inaoep.mx, gelbukh@cic.ipn.mx

² Instituto Nacional de Astrofísica, Óptica y Electrónica (INAOE), México.
alopez@inaoep.mx

Abstract

News reports are an important source of information about society. Their analysis allows to understand its current interests and to measure the social importance and influence of different events. In this paper, we focus on the study of a very common phenomenon of news: the influence of the peak news topics on other current news topics. We propose a text mining method to analyze such influences. We differentiate between the *observable* associations—those discovered from the newspapers—and the *real-world* associations, and propose a technique in which the real ones can be inferred from the observable ones. We argue that the discovery of the ephemeral associations can be translated into knowledge about interests of society and social behavior.

1 Introduction

The problem of analysis of large amounts of information has been solved to a good degree for the case of information that has fixed structure, such as databases with fields having no complex structure of their own. The methods for the analysis of large databases and the discovery of new knowledge from them are called data mining [Fayyad *et al.*, 1996]. However, this problem remains unsolved for non-structured information such as unrestricted natural language texts.

Text mining has emerged as a new area of text processing that attempts to fill this gap. It can be defined as data mining applied to textual data, i.e., as the discovery of new facts and world knowledge from large collections of texts that—unlike those considered in the problem of natural language understanding—do not explicitly contain the knowledge to be discovered [Hearst, 1999; Kodratoff, 1999]. Naturally, the goals of text mining are similar to those of data mining: for instance, it also attempts to uncover trends, discover

associations, and detect deviations in a large collection of texts.

In this paper, we focus on the analysis of a collection of news reports appearing in newspapers, newswires, or other mass media. The analysis of news collections is an interesting challenge since news reports have many characteristics different from the texts in other domains. For instance, the news topics have a high correlation with society interests and behavior, they are very diverse and constantly changing, and also they interact with, and influence, each other.

Montes-y-Gómez *et al.* [1999] have presented a method for the trend analysis of news. Here, we focus on the analysis of a very common phenomenon of news: the influence of the peak news topics over other current news topics.

We define a peak news topic as a topic with one-time short-term peak of frequency of occurrence, i.e., such that its importance sharply rises within a short period and very soon disappears. For instance, the visit of Pope John Paul II to Mexico City became a frequent topic in Mexican newspapers when the Pope arrived to Mexico and disappeared from the newspapers in a few days, as soon as he left the country; thus this is a peak topic.

Usually, these topics influence over the other news topics in two main ways: a news topic induces other topics to emerge or become important along with it, or it causes momentary oblivion of other topics.

The method we proposed analyzes the news over a fixed time span and discovers just this kind of influences, which we call *ephemeral associations*.

Additionally, we differentiate between the *observable* ephemeral associations, those immediately measured by the analysis of the newspapers, and the *real-world* associations. In our model, the real-world associations in some cases can be inferred from the observable ones, i.e., for some observable associations its possibility to be a real-world one is estimated.

* Work done under partial support from CONACyT grant 32003-A, REDII-CONACyT, and SNI, Mexico.




Figure 1. Ephemeral associations between news topics

The rest of the paper is organized as follows. Section 2 reviews some related work. Section 3 defines ephemeral associations and describes the algorithm for their detection. Section 4 introduces the distinction between the observable and the real-world associations and describes the general strategy for the discovery of the real-world associations. Section 5 presents some experimental results. Finally, section 6 draws some conclusions.

2 Related Work

The discovery of association rules is one of the most studied tasks of data mining and also text mining. In data mining, the discovery of associations rules has been defined as the discovery of rules of the form $X \Rightarrow Y (s/g)$ in a large set of transactions [Agrawal *et al.*, 1994]. Such a rule indicates that s percent of the cases contain both item sets ($X \cup Y$), and that g percent of the transactions containing the item set X also contain the item set Y . For instance, in a supermarket database an association rule could be *diaper* \Rightarrow *beer* (0.5 / 0.85) indicating that when people buy diapers, 85% of the times they also buy beer, and that 5% of all transactions include both items.

The discovery of this kind of association rules has also been used in the text analysis [Feldman and Hirsh, 1996; Feldman *et al.* 1997; Rajman and Besançon, 1997; Landau *et al.*, 1998; Rajman and Besançon, 1998]. In this case, the main goal is to uncover the hidden associations occurring among the “topics” mentioned in a text collection. For instance, for a medical corpus, the association rule *cancer* \Rightarrow *therapy* can emerge, indicating that texts talking about cancer tend to talk about therapy.

Usually, the discovery of this kind of association rules—in both, data and text mining—is done in two steps. First, sets of frequent items (topics) are built using an incremental strategy, and second, some association rules are statistically inferred from the frequent item (topic) sets.

In text mining, there is another group of methods related to the discovery of language patterns (collocations or co-occurrences of words at close positions) rather than facts from the real world [Ahonen *et al.*, 1997; Ahonen *et al.*,

1999; Lent *et al.*, 1997; Rajman and Besançon, 1997; Rajman and Besançon, 1998]. In this article, however, we are not interested in such kind of knowledge.

Such methods discover rules or patterns that indicate the co-occurrence of two sets of elements (topics or words appearing together in many texts, for the case of text mining), but they do not state anything about the temporal behavior or temporal relations between these elements.

Other methods of text mining focus on the temporal behavior of the text collections, i.e., on their trend analysis. Some of them allow analyzing the thematic distribution of a text

collection between two different moments [Feldman and Dagan, 1995; Montes-y-Gómez *et al.*, 1999]. Yet other methods allow identifying the topics having a special temporal behavior [Lent *et al.*, 1997].

The method described here is a combination of the latter approaches. It is concerned with the discovery of *ephemeral associations* among the news topics. This kind of associations not necessarily indicates the co-occurrence of two news topics, but mainly expresses the way these topics (particularly when one of them is a peak one) are related during a fixed time span.

3 Discovery of Ephemeral Associations

A common phenomenon in news reports is the influence of a peak news topic, i.e., a topic with one-time short-term peak of frequency, on the other news topics. This kind of influence is what we call an ephemeral association. An ephemeral association may mean an inverse or a direct temporal relation between two news topics (see the figure 1). An inverse association suggests that the peak topic probably produced the momentary oblivion of the other news topic, while a direct association indicates that the peak topic probably caused the momentary increase of the frequency of the other topic.

Given a peak topic and the surrounding data, we detect the ephemeral associations in two steps:

1. Construction of the probability distribution for each news topic over the time span around the peak topic.
2. Detection of the ephemeral associations in the observed data set (if any).

These steps are described in the next two subsections.

3.1 Construction of the Probability Distributions

Given a collection of news reports corresponding to the time span of interest, i.e., the period around the existence of the peak topic, we construct a structured representation of each news report, which in our case is a list of keywords or *topics*. Specifically, in our experiments we used a method

similar to one proposed by Gay and Croft [1990], where the topics are related to *noun strings*.

Once this procedure is done, a frequency f_k^i can be assigned to each news topic. The frequency f_k^i is calculated as the number of news reports for the day i that mention the topic k . It is more convenient, however, to describe each news topic k by a probability distribution $D_k = \{p_k^i\}$ by the days i , where for a given day i , p_k^i expresses the probability for a news topic randomly chosen from the reports of that day to be the topic k :¹

$$p_k^i = \frac{f_k^i}{\sum_{j \in \text{Topics}} f_j^i} \quad (1)$$

We will call the values p_k^i *relative probabilities*. A probability $p_k^i = 0$ indicates that the topic k was not mentioned on the day i . The relative probabilities p_k^i are advantageous for the discovery of the ephemeral associations mainly because they maintain a normalization effect over the news topics: for any day i ,

$$\sum_{k \in \text{Topics}} p_k^i = 1$$

This condition holds for the whole period of interest and means that the increase of the relative probability of one news topic always is compensated by the decrease of probability of some other topics, and vice versa.

3.2 Detection of Ephemeral Associations

The ephemeral associations express inverse or direct relations between the peak topic and some other current news topics (see the figure 1). Let the peak topic be, say, the topic $k = 0$ and the other one we are interested in be, say, the topic $k = 1$. The statistical method we use to detect the observable associations is based on the correlation measure r between the topics $k = 0$ and $k = 1$ [Freund and Walpole, 1990] defined as:

$$r = \frac{S_{01}}{\sqrt{S_{00}S_{11}}}, \text{ where} \quad (2)$$

$$S_{kl} = \sum_{i=1}^m (p_k^i p_l^i) - \frac{1}{m} \left(\sum_{i=1}^m p_k^i \right) \left(\sum_{i=1}^m p_l^i \right)$$

$$k, l = 0, 1.$$

Here, p_k^i are defined in the previous section, and m is the number of days of the period of interest.

¹ This roughly corresponds to the percentage of the space the newspapers devoted to the topic k on the day i .

The correlation coefficient r measures how well the two news topics are related to each other.² Its values are between -1 and 1 , where -1 indicates that there exists an exact inverse relation between the two news topics; 1 indicates the existence of an exact direct relation between the news topics, and 0 the absence of any relation at all.

Therefore, if the correlation coefficient between the peak topic and some other news topic is greater than a user-specified threshold u (i.e., $r < -u$) then there exists a *direct* ephemeral association between them. On the other hand, if the correlation coefficient is less than the threshold $-u$ (i.e., $r < -u$) then there exists an *inverse* ephemeral association between the two topics.

There are two reasons for introducing the user-specified threshold u . First, it softens the criterion so that we can approximate the way a human visually detects the association. Second, to cope with the relatively small data sets in our application: since few data are available (a peak topic persists over few days), random variations of topic frequencies unrelated to the effect in question can greatly affect the value of the correlation coefficient. A typical value recommended for u is around 0.5 .

4 Discovery of Real World Associations

Newspapers usually have a fixed size, and the editor has to decide what news to include in the day's number and what not to include. Thus, the frequencies of the news mentioned in a newspaper do not directly correspond to the number of events that happen in the real world on the same day.

In the next subsection, we explain this important difference between the real world news frequencies and the ones normalized by the fixed size of the newspaper. Then, we show how to estimate whether the observable associations are mainly due to the normalization effect or there is a possible real-world association component.

Note that some sources, such as newswires, do not present this problem: they do not have any fixed size, so the observed frequencies of the news reports are the real world ones. For such sources, the techniques discussed below do not make sense. An easier way to discover the same associations in this case is not to normalize the frequencies in the formula (1), using $p_k^i = f_k^i$ instead and then applying the formula (2).

However, if it is not clear or not known whether the news source presents the normalization problem, then the technique discussed below can be applied indiscriminately. This is because in the absence of normalization effect, the technique discussed below will give equally correct results, though with more calculations.

² The usual interpretation of the correlation coefficient is the following: $100 r^2$ is the percentage of the variation in the values of one of the variables that can be explained by the relation with the other variable.




Figure 2. Analysis of the peak topic “Visit of Pope”

4.1 The Notion of Real World Associations

Since our ultimate goal is to discover the associations that hold in the real world, it is important to distinguish between two different statistical characteristics of the topics appearing in the newspapers. One characteristic is the real-world frequency: the frequency with which the corresponding news comes from the information agencies, for instance. Another characteristic is the observable frequency, expressed as the pieces of news actually appearing in the newspapers.

To illustrate this difference, let us consider two sources of information: say, a journalist working in Colombia and another one working in Salvador. Let the first one send 30 messages each week and the second one send 30 messages in the first week and 70 messages in the second week. These are the real-world frequencies: 30 and 30 in the first week, and 30 and 70 in the second one (i.e., there was something interesting happening in Salvador in the second week). However, the newspaper has a fixed size and can only publish, say, 10 messages per week. Then it will publish 5 and 5 messages from these correspondents in the first week, but 3 and 7 in the second week. These are the observable frequencies, since this is the only information we have from the newspaper texts.

Our further considerations are based on the following two assumptions.

Assumption 1: The newspapers tend to have a constant “size”.³

Thus, the observable frequencies can be considered normalized, i.e., their sum is a constant, while the real world ones are not normalized. We assume that these two kinds of frequencies are proportional, being the proportion coefficient the normalization constant. Thus, we define a real-world ephemeral association as an association that holds between the topics in the real world and not only in the

³ The “size” of a newspaper not only depends on its physical size (for instance, the number of pages) but also on the number of the journalists, the time required for editing, printing, etc.

observable (normalized) data, and we consider that an observable ephemeral association is a combination of two sources: a (possible) real-world ephemeral association and the normalization.

The normalization effect is always an inverse correlation effect. This means that the increase of probability of the peak topic is always compensated by the decrease of probability of some other topics, and vice versa. Thus, we can conclude that any *direct* observable ephemeral association is, very probably, a real-world association.

Assumption 2: The peak topic proportionally takes away some space from each current news topic.

First, this assumption implies that the relative proportions among the rest of the news topics do not change if we take away the peak topic and its related topics. Second, no topic completely disappears only as a consequence of the normalization effect.⁴

4.2 Detection of Real World Associations

As we have noted, all direct associations should be considered real-world ones, so we only need to analyze the inverse ones. The idea is to restore the original distribution of the topics by eliminating the normalization effect, and check if this distribution still correlates with that of the peak topic.

The Assumption 2 allows us to estimate the probability distribution $D'_k = \{p'^i_k\}$ of the topic k as it would be without the normalization effect, where the probability p'^i_k expresses the relative probability of occurrence of the topic k on the day i after we take away the peak topic and its related topics. This probability is calculated as follows:

$$p'^i_k = \frac{f_k^i}{\sum_{j \in P} f_j^i} \quad (3)$$

Here the set P consists of the peak topic and its related topics (those with a direct association), while the frequency f_k^i indicates the number of news reports in the day i that mention the topic k .

Therefore, an inverse observable association between the peak topic and the news topic k is likely a real-world association if it remains after the normalization effect is eliminated from the topic k . In other words, if the correlation coefficient between the peak distribution and the corrected distribution D'_k is less than the user-specified threshold $-u$

⁴ Usually, newspaper editors design the newspaper format and contents in such a way that they expose all news of the day, even if briefly.

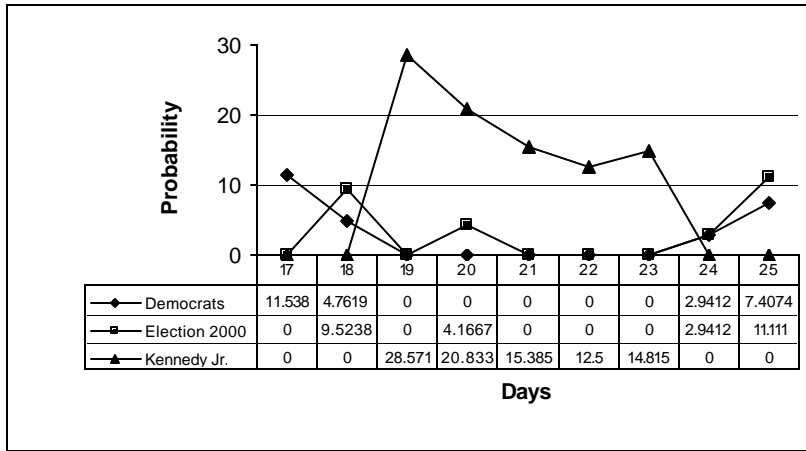


Figure 3. Analysis of the peak topic “Death of Kennedy Jr.”

(i.e., $r < -u$) then the inverse observable ephemeral association is likely a real-world one.

5 Experimental Results

To test these ideas, we used the Mexican newspaper *El Universal*.⁵ We collected the national news for the ten days surrounding the visit of Pope John Paul II to Mexico City, i.e., from January 20 to 29, 1999, and looked for some ephemeral associations between this peak topic and the other topics.

One of the associations detected with our method (using the threshold $u = 0.6$) was a direct ephemeral association between the peak topic and the topic *Virgin of Guadalupe*.⁶ The figure 2 illustrates this association. The correlation coefficient was $r = 0.959$ for the period between the 23 and 25 of January (stay of the Pope in Mexico), and $r = 0.719$ for the surrounding period between the 20 and 29 of January.

Since this association was a direct one, it had a high possibility for being a real-world one. This means that the topic *Virgin of Guadalupe* probably emerged because of the influence of the peak topic. Moreover, since this topic was the only one that had a direct association with the peak topic, we deduced that the visit of the Pope was strongly related with *Virgin of Guadalupe* (in fact, he has focused his discourse on this important Mexican saint).

Another interesting discovery was the inverse association between the peak topic and the topic *Raúl Salinas* (brother of the Mexican ex-president, Carlos Salinas de Gortari, sentenced in the 22 of January). The figure 2 also shows this association.

The correlation coefficient $r = -0.703$ between the 22 and 26 of January (period covering the *Visit of Pope* and the

sentencing of *Raúl Salinas*) indicates the existence of an inverse observable ephemeral association.

In order to determine the possibility of this association for being a real-world one, we analyzed the normalization effect. First, we built the probability distribution of the topic *Raúl Salinas* without considering the peak topic and its related topics (the topic *Virgin of Guadalupe* in this case). The new probability distribution was:

$$D'_{RaulSalinas} = \{ 5.88, 5.26, 25, 17.64, 6.25, 9.09, 11.11, 0, 0, 20 \}$$

Second, we recomputed the correlation coefficient between the peak topic and the topic *Raúl Salinas*. The new correlation coefficient $r = -0.633$ (between the 22 and 26 of January) indicated that it was very possible for this association to be real-world one. If this was true, then the topic *Raúl Salinas* went out of the attention because of the influence of the visit of the Pope to Mexico City.

As another example, we examined the peak topic *Death of Kennedy Jr.* This topic took place between the 18 and 24 of July of 1999. For the analysis of this peak topic, we used the news appearing in the national desk section of the newspaper *The New York Times*.⁷ Among our discoveries, there were two inverse ephemeral associations. One of them between the peak topic and the topic *Election 2000*, with $r = -0.68$, and the other one between the peak topic and the topic *Democrats*, with $r = -0.83$. The figure 3 shows these associations.

Since these associations were both inverse ones, we analyzed their normalization effect. First, we built their probability distributions without considering the peak topic:

$$D'_{Election2000} = \{0, 9.52, 0, 5.26, 0, 0, 0, 2.94, 11.11\}$$

$$D'_{Democrats} = \{11.53, 4.76, 0, 0, 0, 0, 0, 2.94, 7.4\}$$

Then, we recomputed their correlation coefficients. The probability distribution of the topic *Democrats* did not change (because of the zero probabilities of the topic *Democrats* during the peak existence). Thus, the correlation coefficient was again $r = -0.83$ and we concluded that this association had a high possibility for being a real-world one.

On the other hand, the new correlation coefficient between the topic *Elections 2000* and the peak topic, $r = -0.534$, was not less than the threshold $-u$ (we used $u = 0.6$), therefore, there was not enough evidence for this association to be a real-world one.

⁵ <http://www.el-universal.com.mx>

⁶ A Mexican saint whose temple the Pope visited.

⁷ The topics were extracted manually as opposed to the Spanish examples that were analyzed automatically.

6 Conclusions

We have analyzed a very frequent phenomenon in real life situations—the influence of a peak news topic on the other news topics. We have described a method for the discovery of this type of influences, which we explain as a kind of association between the two news topics and call *ephemeral associations*. The ephemeral associations extend the concept of typical associations because they not only reveal co-existence relations between the topics but also their temporal relations.

We distinguish between two types of ephemeral associations: the *observable* ephemeral associations, those discovered directly from the newspapers, and the *real-world* associations. We have proposed a technique with which the observable associations are detected by simple statistical methods (such as the correlation coefficient) and the real-world associations are heuristically estimated from the observable ones. Namely, our algorithm consists of the following steps:

1. Calculate the (observable) probabilities, formula (1);
2. Calculate the (observable) correlations between the peak topic and the other ones by the formula (2);
3. Select the topics that strongly correlate with the peak one, using a threshold $u \approx 0.5$;
4. Determine which associations are real-world ones:
 - a. All direct associations are real-world ones;
 - b. For the inverse associations,
 - i. Build the corrected distributions by the formula (3), using the knowledge obtained at the step 3;
 - ii. Calculate the (real-world) correlations between the peak topic and the other ones using the formula (2) and the corrected distributions;
 - iii. The topics for which this correlation is strong represent the real-world associations.

As future work, we plan to test these ideas and criteria under different situations and to use them to detect special circumstances (favorable scenarios and difficult conditions) that make the discovering process more robust and precise.

Finally, it is important to point out that the discovery of this kind of associations, the ephemeral associations among news topics, helps to interpret the relationships between society interests and discover hidden information about the relationships between the events in social life.

References

- [Agrawal and Srikant, 1994] Agrawal, Rakesh, and Ramakrishnan Srikant, *Fast Algorithms for Mining Association Rules*, Proc. 20th VLDB Conference, Santiago de Chile, 1994.
- [Ahonen-Myka *et al.*, 1997] Ahonen-Myka, Heinonen, Klemettinen, and Verkamo, *Applying Data Mining Techniques in Text Analysis*, Report C-1997-23, Department of Computer Science, University of Helsinki, 1997.
- [Ahonen-Myka *et al.*, 1999] Ahonen-Myka, Heinonen, Klemettinen, and Verkamo, Finding Co-occurring Text Phrases by Combining Sequence and Frequent Set Discovery, Proc. of the Workshop on Text Mining: Foundations, Techniques and Applications, IJCAI-99, Stockholm, 1999.
- [Fayyad *et al.*, 1996] Fayyad, Piatetsky-Shapiro, Smyth and Uthurusamy, *Advances in Knowledge Discovery and Data Mining*, Cambridge, MA: MIT Press, 1996.
- [Feldman and Dagan, 1995] Feldman and Dagan, *Knowledge Discovery in Textual databases (KDT)*, Proc. of the 1st International Conference on Knowledge discovery (KDD_95), pp.112-117, Montreal, 1995.
- [Feldman and Hirsh, 1996] Feldman and Hirsh, *Mining Associations in Text in the Presence of Background Knowledge*, Proc. of the 2nd International Conference on Knowledge Discovery (KDD-96), Portland, 1996.
- [Feldman *et al.*, 1997] Feldman, Klösigen, Yehuda, Kedar and Reznikov, *Pattern Based Browsing in Document Collections*, Proc. of the 1st Conference on Principles of Knowledge Discovery and data Mining (PKDD'97), Lecture Notes in AI, Springer Verlag, Norway, 1997.
- [Freund and Walpole, 1990] Freund and Walpole, *Estadística Matemática con Aplicaciones*, Cuarta Edición, Prentice Hall, 1990. (In Spanish)
- [Gay and Croft, 1990] Gay and Croft, *Interpreting Nominal Compounds for Information Retrieval*, Information Processing and Management 26(1): 21-38, 1990.
- [Hearst, 1999] Hearst, *Untangling Text Data Mining*, Proc. of ACL'99: the 37th Annual Meeting of the Association for Computational Linguistics, University of Maryland, 1999.
- [Kodratoff, 1999] Kodratoff, *Knowledge Discovery in Texts: A Definition and Applications*, Proc. of the 11th International Symposium on Foundations of Intelligent Systems (ISMIS-99), 1999.
- [Landau *et al.*, 1998] Landau, Feldman, Aumann, Fresko, Lindell, Liphstat and Zamir, *TextVis: An Integrated Visual Environment for Text Mining*, Proc. of the 2nd European Symposium on Principles of Data Mining and Knowledge Discovery, PKDD-1998.
- [Lent *et al.*, 1997] Lent, Agrawal, Srikant, *Discovering Trends in Text Databases*, Proc. of the 3rd Int'l Conference on Knowledge Discovery in Databases and Data Mining, Newport Beach, California, August 1997.
- [Montes-y-Gómez *et al.*, 1999] Montes-y-Gómez, López-López and Gelbukh, *Text Mining as a Social Thermometer*, Proc. of the Workshop on Text Mining: Foundations, Techniques and Applications, IJCAI-99, Stockholm, 1999.
- [Rajman and Besançon, 1997] Rajman and Besançon, *Text Mining: Natural Language Techniques and Text Mining Applications*, Proc. of the 7th IFIP 2.6 Working Conference on Database Semantics (DS-7), Chapman & Hall IFIP Proceedings serie. Leysin, Switzerland, 1997.
- [Rajman and Besançon, 1998] Rajman and Besançon, *Text Mining - Knowledge Extraction from Unstructured Textual Data*, 6th Conference of International Federation of Classification Societies (IFCS-98), Rome, 1998.