

EvoRiver: Visual Analysis of Topic Coopetition on Social Media

Guodao Sun, Yingcai Wu, *Member, IEEE*, Shixia Liu, *Senior Member, IEEE*,
Tai-Quan Peng, Jonathan J. H. Zhu, and Ronghua Liang, *Member, IEEE*

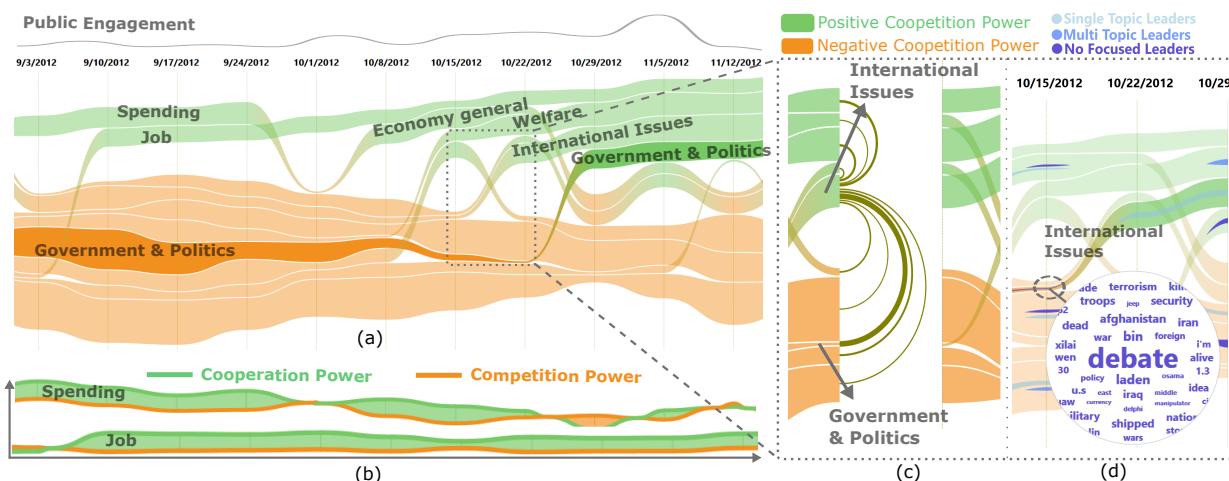


Fig. 1. (a) Topic coopetition dynamics during the 2012 U.S. presidential election with EvoRiver, showing most of the topics were transitioning from competition to cooperation during that time; (b) playfair-style chart of *spending* and *job* to unfold their coopetition power; (c) pairwise similarity between *international issues* and other topics with connected arcs; (d) word cloud of *international issues*.

Abstract— Cooperation and competition (jointly called “coopetition”) are two modes of interactions among a set of concurrent topics on social media. How do topics cooperate or compete with each other to gain public attention? Which topics tend to cooperate or compete with one another? Who plays the key role in competition-related interactions? We answer these intricate questions by proposing a visual analytics system that facilitates the in-depth analysis of topic coopetition on social media. We model the complex interactions among topics as a combination of carry-over, coopetition recruitment, and coopetition distraction effects. This model provides a close functional approximation of the coopetition process by depicting how different groups of influential users (i.e., “topic leaders”) affect coopetition. We also design EvoRiver, a time-based visualization, that allows users to explore coopetition-related interactions and to detect dynamically evolving patterns, as well as their major causes. We test our model and demonstrate the usefulness of our system based on two Twitter data sets (social topics data and business topics data).

Index Terms—Topic coopetition, information diffusion, information propagation, time-based visualization.

1 INTRODUCTION

Social media allow millions of users to consume, produce, and disseminate huge volumes of highly diverse information on social networks [20, 27]. This information may concern different topics, such as celebrity news and personal updates, which can reach many users as soon as they are uploaded on the Internet. Different topics may interact with one another during their propagation, which can lead to complex dynamics of information diffusion. These topics may compete or cooperate with one another to gain attention from social media users [30, 40]. For instance, Nokia and HTC are direct competitors in the smartphone

market and have created their respective Twitter accounts to promote their products. These companies compete with each other by tweeting various topics to attract the attention of potential customers and to improve their brand awareness. These companies may also cooperate with each other to some extent. For example, HTC and Nokia reportedly mocked the launch of the Samsung Galaxy S5 by tweeting “Buyers remorse: Coming soon to S5 owners” and “Not the SameSung” with an attached image of a Windows Phone, respectively [28].

Understanding topic competition and cooperation (jointly called *coopetition*) presents many useful applications. For example, social marketing specialists express their great interest in understanding topic cooperation, such that they can effectively insert their desired topics into appropriate trending topics on social media. The inserted topics can then leverage the trending topics to become more salient among users. Understanding topic competition is also valuable. Social marketing specialists identify the most competitive topics to tweet about in order to gain attention from the audience of other topics. Therefore, insights can be obtained by detecting and analyzing topic coopetition despite the difficulty of such a problem [12, 30, 40, 44].

Understanding topic coopetition is hindered by two major challenges, namely, the quantitative measurement of dynamic topic coopetition and the interactive visualization of measured coopetition. Previous studies have modeled the competition and cooperation among multiple propagating contagions [30] or memes [12] on social media. However, these models can only determine whether two contagions or memes are

- Guodao Sun and Ronghua Liang are with Zhejiang University of Technology. This work was done when Guodao Sun was an intern at Microsoft Research. E-mail: godoor.sun@gmail.com, rhliang@zjut.edu.cn.
- Yingcai Wu and Shixia Liu are with Microsoft Research. Shixia Liu is the correspondence author. E-mail: {yingcai.wu, shliu}@microsoft.com.
- Tai-Quan Peng is with Nanyang Technological University. E-mail: winsonpeng@gmail.com.
- Jonathan J.H. Zhu is with City University of Hong Kong. E-mail: j.zhu@cityu.edu.hk.

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014; date of publication xx xxx 2014; date of current version xx xxx 2014.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

competing or cooperating with each other without directly quantifying the power (or strength) of such processes. These models ignore the vital role of influential social media users (i.e., *opinion leaders*) in information dissemination [20, 42]. Therefore, the co-competition power for a topic and the influence of opinion leaders remain unquantifiable.

Topic co-competition can be regarded as a dynamic and a fundamentally distinctive topic behavior. The effective visualization of a dynamic topic co-competition remains a major challenge. Previous studies have merely visualized topic co-competition by using a stacked graph [44], which cannot be used to display topic co-competition in a single coherent view. An effective and comprehensive co-competition model can produce dynamic and multidimensional data with special features by measuring competition and cooperation powers, as well as the influence of opinion leaders on each topic. These features, which include contextual information such as tweets, hinder the visualization of topic co-competition.

We propose a new visual analytic system that can assist analysts in exploring and analyzing the dynamics of co-competition among multiple topics on social media. We propose a new model that characterizes topic co-competition as a combination of carry-over, co-competition recruitment, and co-competition distraction effects. We introduce the concept of *topic publics*, which is borrowed from political science [22], to provide additional information that can help in the identification of opinion leaders (see Section 4.1). This concept can distinguish three types of opinion leaders, namely, *single-topic*, *multi-topic*, and *no-focused leaders*. We identify these three groups of leaders collectively as topic leaders for the sake of simplicity. Our model explicitly shows the contributions of topic leaders to the dynamics of topic co-competition, which provides new insight into topic co-competition mechanisms. Although our model can comprehensively capture the characteristics of topic co-competition, our model generates complex data with special features, as we discussed in the previous paragraph.

We resolve this problem by proposing EvoRiver, a time-based visualization that transforms time-varying, multidimensional information into an interactive visualization. EvoRiver employs a river metaphor and represents a topic as a strip. By moving the river strips up and down, EvoRiver can visualize the overall topic co-competition trend over time. Different groups of topic leaders are visually encoded as threads that are then overlaid to the strips to visualize the co-evolutionary relationship between the topics and the topic leaders. However, the huge solution space hinders the design of a legible, compact, and aesthetically appealing EvoRiver layout with minimum crossings and wiggles. We address this problem by converting the layout generation problem into discrete and continuous optimization problems. These issues are then solved by using an adapted two-level *Directed Acyclic Graph* (DAG) algorithm and a quadratic optimization algorithm, respectively. The layout method automatically enhances important patterns (e.g., prominent diverging and converging patterns) in the layout. EvoRiver supports visual exploration and sense-making through a rich set of user interactions, which allows users to interact with the visualization to locate and to investigate interesting patterns for new insights.

- We propose a new model that can quantitatively characterize the dynamic topic co-competition-related interactions, as well as the influence of topic leaders on such interactions.
- We design EvoRiver, a visual system that can assist analysts in their investigation of the complex dynamics of topic co-competition-related interactions on social media.
- We offer profound insights into the dynamics of topic co-competition and the influence of topic leaders by using two large-scale social media data sets that cover the areas of business and politics.

2 RELATED WORKS

This section reviews related works on stacked graphs and visual analysis of information diffusion on social media.

Stacked Graphs. Stacked graphs are widely used in different applications [14, 24, 26, 34, 36, 39]. Stacked graphs can simultaneously visualize time series individually and collectively [9]. Havre et al. [19] proposed ThemeRiver, a seminal technique, for creating a smooth stacked graph. Byron and Wattenberg [9] introduced Streamgraphs, which improves the legibility and aesthetics of stacked graphs. Stream-

graphs have been used in several text visualization systems [13, 26, 35] to support topic-based analysis and exploration. Cui et al. [13] introduced TextFlow, which captures and visualizes the temporal split and merge behaviors of topics. Xu et al. [44] employed stacked graphs to display the time-varying competitiveness of topics on social media. Both approaches [13, 44] employ a composite visual design that draws threads above the strips of a streamgraph. This design visualizes the co-occurrence between keywords (threads) and topic clusters (strips) [13], as well as demonstrates the influence of opinion leaders (threads) on the saliency of topics (strips) [44]. We use the same composite visual design to demonstrate the influence of topic leaders (threads) on co-competition power (strips). However, co-competition power may turn positive or negative at each time stamp which makes co-competition power semantically different from the encoded information in other approaches [13, 44]. The traditional stacked graph layouts cannot visualize co-competition power because adding the positive values with the negative values does not generate logical results. Therefore, we propose EvoRiver to visualize the evolving relationships between the co-competition power of topics and topic leaders.

Visual Analysis of Information Diffusion. Many information diffusion models assume that multiple memes flow in isolation on social media [18]. The assumption has been contradicted in previous studies [7, 40, 41], which suggest that memes primarily compete for the attention of social media users [7, 40, 41]. Recent findings reveal that memes may also collaborate to attract attention [12, 30]. Myers and Leskovec [30] proposed a probabilistic model to approximate the probability for a user to adopt a new contagion after being exposed to a sequence of contagions. Coscia [12] proposed an empirical approach that could compute a set of conditional probabilities for each pair of memes on Quickmeme.com, which could directly detect meme co-competition. Compared with previous studies [12, 30], our model characterizes both competition and cooperation among temporal salient topics (e.g., education and economics) rather than among highly volatile memes. Our model can quantitatively measure the co-competition power for each topic, whereas existing models can only detect whether two memes are competing or cooperating. Our model can also evaluate the influence of topic leaders on the dynamics of topic co-competition. New visualizations have been created to display and to analyze the information flow on social media [5, 6, 10, 37, 43, 44]. Whisper [10] used a sunflower metaphor to visualize the spatiotemporal diffusion process of information. However, these methods cannot easily visualize the relationships among multiple topics through the dynamics of information diffusion.

Our proposed visualization approach can demonstrate the temporal relationships among topics and the influence of topic leaders on the diffusion process. Our approach can be distinguished from existing approaches [44] in three aspects. First, we propose a highly comprehensive model that can estimate both the competition and cooperation powers among topics. Second, given that topics may have a positive or a negative co-competition power at each time stamp, the stacked graph of [44] cannot visualize co-competition power because adding the positive values with the negative values cannot generate valid results. We address this issue by designing EvoRiver to provide a visual summary of topic co-competition. Third, the opinion leaders that are investigated in [44] are defined according to quantitative measures (the number of tweets forwarded) without considering any content measure. Therefore, the ecological validity of such an approach remains questionable. We address this issue by combining quantitative and content measures to analyze the influence of opinion leaders on topic co-competition.

3 SYSTEM OVERVIEW

Our system comprises three major parts, namely, data preprocessing, data analysis, and interactive visualization. The data preprocessing part uses LIBSVM [11], a well-known library for support vector classification, to extract relevant tweets. The extracted tweets are further indexed by a high-performance text search engine (Apache Lucene [1]). The data analysis part is fed with various time series data. Our co-competition model can perform various quantitative measurements to characterize dynamic co-competition among topics, as well as determine the influence of topic leaders. The visualization component is fed with the output of

the data analysis part. EvoRiver is the core component for visualizing the cooperation power to illustrate the topical transition between competition and cooperation. EvoRiver also supports a detail-on-demand and in-place visualization for investigative exploration.

4 TOPIC LEADERS AND COOPERATION MODEL

This section presents the background information on topic leaders, briefly introduces an existing competition model [44], and further describes our proposed cooperation model.

4.1 Topic Leaders

We introduce the concept of “*topic publics*” to facilitate the modeling of topic cooperation. Adopted from the concept of “*issue publics*” in political science [22], *topic publics* refer to those individuals who always focus on a small number of public issues (i.e., topics) exclusively and intensively. This focused topic orientation has been attributed to motivational factors, such as self-interests, social identifications, and cherished basic values [8] as well as resource constraints, such as limitations in time, cognitive capacity, and emotional spending [46]. Empirical studies of public opinion survey data provide consistent evidence for the widespread and enduring existence of topic publics. In general, around 10% to 20% of the American public who are passionately concerned about a single topic (single-topic publics) [22, 29]. Despite their small size, topic publics are far more persistent and vocal than the other members of the society. Therefore, topic publics often carry a greater political clout [23]. The concept of topic publics provides additional and crucial information to identify opinion leaders. In previous studies of Twitter and other similar media, opinion leaders have been traditionally defined by quantitative measures (e.g., number of followers) without considering any content measure. We use information on topic publics to identify three types of opinion leaders, namely, *topic leaders*, which are described below.

Single-topic leaders are the most active, popular, and influential members of the single-topic publics who focus on one topic exclusively.

Multi-topic leaders are the most active, popular, and influential members of the multi-topic publics, and simultaneously attend to several topics (empirically determined to be 2-5 in the current study).

No-focused leaders are the most active, popular, and influential members of the non-focus publics, and are concerned about nearly all topics, which is essentially the same as a lack of focal topics.

We use the quantitative measures of activity, popularity, and influence to separate opinion leaders from their followers. Other content measures (i.e., type of topics) can be used to separate single- and multi-topic leaders and no-focused leaders from one another, while assuming that each leader is largely influential among his/her followers. In this way, we establish a crucial, yet often missing, linkage between opinion leaders and followers. Section 6.1 explains our classification scheme.

4.2 Competition Model

The competition model [44] assumes that various topics exist in an environment that offers limited public attention. Therefore, these topics compete for public attention and media coverage [40]. Given k topics and n groups of topic leaders, the model is defined as follows.

$$p_i^t = \alpha_i p_i^{t-1} + \sum_{g=1}^n m_{i,g}^{t-1} \sum_{j=1, j \neq i}^k \beta_{i,j,g} p_j^{t-1} - p_i^{t-1} \sum_{j=1, j \neq i}^k \sum_{g=1}^n \beta_{j,i,g} m_{j,g}^{t-1} \quad (1)$$

where α and β are regression coefficients. The model accounts for the saliency of topic i at time t (i.e., p_i^t) with a combination of three effects: carry over effect ($\alpha_i p_i^{t-1}$) from time $t-1$ to time t , and effects of

- **Competition Recruitment** ($\sum_{g=1}^n m_{i,g}^{t-1} \sum_{j=1, j \neq i}^k \beta_{i,j,g} p_j^{t-1}$) means that topic i can attract followers from other topics (e.g., j) by the topic leader groups (e.g., g) advocating topic i on the followers of other topics (e.g., j).
- **Competition Distraction** ($p_i^{t-1} \sum_{j=1, j \neq i}^k \sum_{g=1}^n \beta_{j,i,g} m_{j,g}^{t-1}$) means that other topics (e.g., j) can distract followers from topic i by the topic leader groups (e.g., g) advocating other topics (e.g., j) on the followers of topic i .

The dependent variable, p_i^t , denotes the saliency of topic i as perceived by the public at time t . It is defined as the ratio between the number of the tweets of topic i at time t and the number of all the tweets at time t . The independent variables include the public saliency of topic i (i.e., p_i^{t-1}), and the coverage of topic i by opinion leaders g (i.e., $m_{i,g}^{t-1}$) at time $t-1$. The coverage is obtained by dividing the number of tweets of topic i by the total number of tweets posted by g at time $t-1$.

4.3 Cooperation Model

The competition model [44] has been proven useful in revealing the competitive relationships among different topics during the 2012 US presidential election. However, this model assumes that competition is the sole form of relationships that can exist among topics, and thus, neglecting other forms of relationships, such as cooperation and independence. Some studies [12, 30] suggest that topics (or contagions, memes) may compete and cooperate with one another to gain public attention. Cooperation can outweigh competition as the primary form of relationship among topics [12]. Therefore, the competition model must be refined by simultaneously considering the dynamics of competitive and cooperative interactions among topics.

We refine the competition model [44] by incorporating the similarity between topics into the current cooperation model. Selective exposure theory [17, 33] in communication and social psychology research argues that as cognitive misers, individuals tend to expose themselves to information that they are familiar with or are concerned about. These individuals focus on similar topics to mitigate the dissonance and cognitive load that are required in their information processing. In other words, highly similar topics are expected to cooperate with one another to divert public attention from other topics. In other words, highly similar topics are expected to cooperate with one another to recruit public attention from other topics. Topical similarity $\theta_{i,j}$ is measured by the semantic similarity that is weighted by the temporal correlation between two topics. The similarity is computed as $\theta_{i,j} = \mu_{i,j} * v_{i,j}$ where $\mu_{i,j}$ denotes the Pearson product-moment correlation coefficient between the time series sequence of topic i and j , and $v_{i,j}$ denotes the semantic similarity between topic i and j .

Temporal correlation refers to the relationship between the dynamic time series sequences of two topics. For time series sequence generation, we first split the study period into a number of time segments (i.e., 4-hour window in our paper). In each segment, the number of the tweets under a topic is counted, and further normalized to zero mean and unit-variance. A fine-grained time series sequence will provide more observations for linear regression in model analysis. Our choice of 4-hour window is not an arbitrary decision. The 4-hour time window can help our model capture the temporality of users’ activities[21], and can also provide us with adequate observations within each time segment. The computation of the semantic similarity between two topics is challenging. It may be affected by word ambiguity, stop words removal, etc. We adopt the cosine similarity method, a widely used measurement in data mining domain, to measure the semantic similarity of two topics. We use term frequency (i.e., the number of occurrence of a word in a tweet) to generate the feature vector of a topic. The semantic similarity of two topics at a certain time point is defined as the cosine similarity of the two topics at that time.

With $\theta_{i,j}$, we propose the following cooperation model that captures the competitive and cooperative interactions among topics.

$$p_i^t = \alpha_i p_i^{t-1} + \sum_{g=1}^n m_{i,g}^{t-1} \sum_{j=1, j \neq i}^k \beta_{i,j,g} p_j^{t-1} - p_i^{t-1} \sum_{j=1, j \neq i}^k \sum_{g=1}^n \beta_{j,i,g} m_{j,g}^{t-1} \quad (2a)$$

$$+ \sum_{g=1}^n \sum_{j=1, j \neq i}^k (m_{i,g}^{t-1} + m_{j,g}^{t-1}) \theta_{i,j} \sum_{h=1, h \neq i, h \neq j}^k \gamma_{j,h,g} p_h^{t-1} \quad (2b)$$

$$- p_i^{t-1} \sum_{g=1}^n \sum_{j=1, j \neq i}^k \sum_{h=1, h \neq i, h \neq j}^k (m_{j,g}^{t-1} + m_{h,g}^{t-1}) \theta_{j,h} \gamma_{h,i,g} \quad (2c)$$

where α , β , and γ are regression coefficients. In the proposed model, the first three terms (Equation (2a)) are adopted from the competition

model described in Equation (1), which correspond to carry over, competition recruitment, and competition distraction effects, respectively. The remaining two terms are as follows:

- * **Cooperation Recruitment** (Equation (2b)) means that topic i and i 's cooperative topics (e.g., topic j that is similar to topic i) can attract followers from other topics (e.g., h) by the topic leader groups (e.g., g) advocating topic i and i 's cooperative topics on the followers of other topics. The term stands for the cooperative recruitment effect of two topics (i.e., topic i and j) on other remaining topics with the contribution from different groups of topic leaders.
- * **Cooperation Distraction** (Equation (2c)) means that topic j and j 's cooperative topics (e.g., topic h that is similar to topic j) can distract followers from topic i by the topic leader groups (e.g., g) advocating topic j and j 's cooperative topics on the followers of topic i . The term (Equation (2c)) highlighted in the second rectangle denotes the cooperative distraction effect caused by any other two topics.

4.4 Measuring Coepetition Power

In this study, the **competition/cooperation power** of topic i is defined as the magnitude of how competitive/cooperative the topic is in recruiting attention from the public. Therefore, we employ the competition and cooperation recruitment effects to develop two measurements for topic competition and cooperation powers. Standard linear regression is employed to solve Equation (2), in which the product terms of the independent variables are treated as individual independent variables. However, we cannot treat the estimated coefficients (such as $\beta_{i,j,g}$ and $\gamma_{j,h,g}$) as the competition and cooperation powers for our analysis because they are not intuitively interpretable. We need a stringent method to measure the competition and cooperation powers of the topics and to reveal the recruitment effect that can be attributed to different topic leader groups. Therefore, we adopt the squared semi-partial correlation sr^2 to estimate the cooperation and competition powers of different topics. The unique-effect of sr^2 is additive, normalized, and comparable within an equation or over all equations of a system. This measurement enables an analytical comparison among n topic leader groups within each topic or among k topics across all topic leader groups over time.

A stepwise regression is applied to identify the competition and cooperation powers. The competition power of a topic is defined as the sr^2 of the competition recruitment term (i.e., $m_{i,g}^{t-1} p_j^{t-1}$) with same subscript i , whereas the cooperation power is defined as the sr^2 of the cooperation recruitment term (i.e., $(m_{i,g} + m_{j,g}) \theta_{i,j} p_h$) with same subscript i .

Coepetition Power. We propose a composite measure called *coepetition power* to describe the rise and fall of the competition and cooperation powers of each topic. This measure identifies the difference between the two powers of each topic. A topic tends to be cooperative when its coepetition power is greater than 0 and vice versa. The contributions of topic leaders on cooperation and competition powers can be obtained by summing the sr^2 regarding the same group of topic leaders. The measurement of coepetition power enables our system to present the dynamics of topic cooperation and competition.

5 VISUAL DESIGN

This section describes the user requirements collected from our domain experts and derives a set of design goals. After our discussion, we introduce our visualization techniques based on the design goals.

5.1 User Requirements

We collaborated with two scholars from Communication and Media Studies for this project. We aimed to explore the complex dynamics of interactions among the topics and issue publics on social media. With the domain experts, we worked on our research problems, which were to define and refine the coepetition model, and to design visualization techniques iteratively. We also derived research questions from our domain experts, which are described as follows:

Q1 How and when do topics cooperate or compete with one another to gain public attention over a long period? How does the overall competition power vary over time?

- Q2 What are the topics that tend to cooperate or compete with one another? What are the similarities and differences of competing and cooperating topics?
- Q3 Who plays the key role in various coepetition-related interactions? Who exerts the greatest influence on a highly cooperative topic that has changed from a highly competitive topic?
- Q4 What are the similarities and differences in the roles played by different groups of issue publics through the dynamics of topic coepetition-related interactions?
- Q5 How do the correlations among topics and issue publics co-evolve over time? Do issue publics always focus on a few topics? If not, how often do they divert their attention to other topics?
- Q6 How can our visual analysis system assist in the formation and validation of the hypotheses when an interesting pattern emerges? For instance, is the pattern triggered by a breaking news event?

These requirements helped us derive the appropriate design principles and make judicious decisions on our visual design.

5.2 Design Goals

We defined the following design principles based on the user requirements to guide our visual design:

- G1 **Summarize dynamic topic coepetition.** The design must provide a clear and compact visual summary of temporal topic coepetition. A time-oriented design is employed in EvoRiver by considering the importance of time in addressing questions on temporal patterns (Q1) or on the co-evolutionary patterns of topics and topic leaders (Q4 and Q5). This design can also facilitate the connection of external events (i.e., breaking news events) with topic coepetition patterns (Q6). Other analysis tasks (Q2 and Q3) can also benefit from this design as the time attribute usually serves as important contextual information for the analysis. Domain experts wish to identify the individual coepetition power for each topic and the overall coepetition power for all topics over time (Q1). Therefore, EvoRiver also uses a visual design that is similar to stacked graphs.
- G2 **Provide a visual metaphor.** Our collaborators prefer simple and intuitive designs. An intuitive visual representation can help enhance their understanding of topic coepetition (Q1 to Q6). Such a design should also be self-explaining. EvoRiver uses a river metaphor and encodes a topic as a river thread to display the natural flow of topics from being highly competitive to highly cooperative over time, or vice versa. The visual encoding of convergence and divergence patterns has been inspired by river confluence and bifurcation, in which the former occurs when two or more streams merge into a river and the latter occurs when a river flows into two or more streams. This visual metaphor allows us to visualize important patterns where topics converge to cooperation (or competition) or diverge from cooperation (or competition). The river metaphor is particularly helpful in addressing Q1 and Q2.
- G3 **Compare topics pair-wisely.** The design must allow a user to compare topics pair-wisely (Q2). Domain experts prefer to identify the similarities among topics in the context of topic coepetition dynamics. Therefore, we vertically split EvoRiver and fill the gap with an in-place view to illustrate the pairwise similarities among topics. The in-place view can provide an occlusion-free visualization that is close to the topic coepetition context.
- G4 **Relate topic leaders to topics.** The time series data of topic leaders must be visually related to the topics to facilitate the detection and analysis of their correlation patterns (Q3 to Q5). EvoRiver visualizes the related data in one coherent view to facilitate the identification of correlation patterns (Q3 to Q5). The system employs a composite design by drawing a set of threads, which represent the topic leaders on the river strips that are associated with the topics.
- G5 **Reduce visual clutter.** The crossings and wiggles of the threads and strips in EvoRiver may generate visual clutter that can hinder users from seeking and analyzing information (Q1 to Q6). An effective design must have minimal clutter and clearly visualize data. We devised an optimization method to optimize our layout.
- G6 **Highlight and unfold patterns.** The visualization must enhance the prominent patterns to reveal such patterns immediately upon

their emergence. It must also provide additional details to help users investigate patterns and formulate hypotheses (Q6).

5.3 Visualization Techniques

This section introduces our interactive visualization techniques that are designed based on the aforementioned design goals.

5.3.1 EvoRiver Visualization

EvoRiver is the core visual component of our visual analytics system that aims to provide a comprehensive visual summary of how the cooperation power of various topics evolves over time (G1) with a familiar river metaphor (G2). EvoRiver provides an in-place view to compare topic pair-wisely (G3). The system visually relates topic leaders to topics by using a composite visual design (G4). We design an optimization algorithm to optimize the EvoRiver layout, minimize visual clutter (G5), and highlight prominent patterns (G6). A rich set of user interactions is supported to reveal patterns (G6).

Visual Encodings. We design EvoRiver to visualize the evolutionary cooperation power for each topic through a river metaphor. G1 and G2 require a time-oriented visual design with a familiar metaphor, which is similar to stacked graphs. However, we could not directly use a stacked graph to provide a meaningful visual summary of cooperation power. This power can become positive or negative at any time, and traditional stacked graphs cannot deal with such variations. Therefore, we introduce EvoRiver to overcome this problem.

Figure 4 shows an EvoRiver layout, which shows the flow of topics that changes from cooperative to competitive over time, or vice versa. Each river strip visually encodes a topic, and the height of each strip reflects the cooperation power of the topic. EvoRiver comprises two parts, namely, the top cooperative strips with a positive cooperation power and the bottom competitive strips with a negative cooperation power. Competitive and cooperative strips are stacked to display the overall positive and negative cooperation powers, respectively.

We use a composite design to relate the topic leaders to the topics (G4). Different topic leader groups are visually represented as distinct threads with distinguishable colors. Two types of threads, namely, continuous and transition threads, are directly overlaid onto the river strips. Such a composite view could help perceive the contribution of topic leaders to the dynamics of topic cooperation immediately. A continuous thread, which is represented as a solid line, is drawn on a strip if the topic leaders contribute to the cooperation power of the related topic (see Figure 3 (b)). The contribution is represented by the thickness of the thread. A transition thread, which is represented as a dotted line, indicates how a group of topic leaders diverts their attention from one topic to another (see Figure 3 (b)). The density of the dots reflects the strength of such a transition, which is derived from using a soft matching approach to estimate the focus transition [44].

Layout Generation. In EvoRiver, a strip can move between the top and bottom parts where the strips in any part are stacked. The threads that are overlaid on the strips may also switch among different strips. This visual design may produce undesirable visual effects, such as wasted screen space, wiggles, and edge crossings, which are caused by the arbitrary placement of strips and threads. The simultaneous optimization of these effects may result in a large search space. We adopt an optimization strategy from storyline visualization [25] to generate an effective and aesthetically appealing layout (G2 and G5). This strategy optimizes the highly undesirable effects and ensures that subsequent optimizations do not affect the previous optimization results. Previous studies identify crossings and wiggles as the first and second most undesirable effects, respectively [31]. Thus, our strategy divides the optimization process into two parts, namely, discrete (to reduce crossings) and continuous (to reduce wiggles and to boost symmetry).

Discrete optimization reduces the crossings by using a two-level barycenter method [16] to adjust the ordering of strips and threads. This algorithm involves two steps, namely, a DAG for optimizing the ordering of all strips, and a DAG for optimizing the ordering of the threads within each strip. Figure 2 shows the ordering process. Figure 2 (a) shows an initial ordering with crossings among strips and threads, and two crossings between threads. Strips A, B, and C are

sorted according to their barycenter scores (2, 3, and 1.27, respectively). After that, we sort the threads in each strip (see Figure 2 (c)).




Fig. 2. Sorting of Strips and threads: (a) initial ordering; (b) result after sorting the strips; (c) result after sorting the threads.

The continuous optimization process aims to produce a symmetrical layout with minimum wiggles for river strips. The strips are stacked on the top and bottom parts of EvoRiver, and the number of strips may vary at each time stamp. This design can generate a strange zigzag-shaped layout that occupies much screen space (see Figure 3 (a)). We define two layout metrics, namely, *symmetry* and *alignment*, to address this problem. Following Gestalt theory [38], we use the symmetry metric to produce a symmetrical layout that can strengthen the perceptiveness of EvoRiver. The alignment metric reduces the undesirable effects of wiggles. We utilize an objective function to transform a layout problem into an optimization problem. Given n strips during a period with T time stamps, we define the objective function that optimizes the top part of EvoRiver as follows.

$$\alpha \sum_{t=1}^{T-1} \sum_i^n (y_{i,t} - y_{i,t+1})^2 + \beta \sum_{t=1}^{T-1} (S_t - S_{t+1})^2 \quad (3)$$

$$\text{Subject to } y_{i+1,t} - y_{i,t} = \frac{c_{i,t}}{2} + \frac{c_{i+1,t}}{2}, \text{ if } o_{i,t} < o_{i+1,t} \quad (3a)$$

$$0 < y_{i,t} < B \quad (3b)$$

$$y_{k,t} - y_{j,t} > H, \forall k \in \text{Top part}, \forall j \in \text{Bottom part} \quad (3c)$$

The function has a wiggle term on the left and a symmetric term on the right. The wiggle term adds the wiggle distances of every strip between adjacent time stamps. We only compute the distances between adjacent wiggles of a strip when these wiggles are located in the same part (top or bottom). The wiggle term must be minimized to create a smooth strip with fewer wiggles. The symmetric term adds the difference of the central positions of the top (or bottom) part between the adjacent time stamps. This term aligns the central positions of the neighboring regions of the top (or bottom) part to achieve a symmetrical layout.

- α and β are the weights for the two terms. We set $\alpha = 1$ and $\beta = 20$ after the experiments to create excellent layouts.
- m_t is the number of strips in the top (or bottom) part at time t .
- $y_{i,t}$ is the vertical position of strip i at time t (see Figure 3 (d)).
- B is the boundary of the top (or bottom) part (see Figure 3 (d)).
- S_t is the average vertical position of the strips in the top (or bottom) part, $S_t = \frac{1}{m_t} \sum_i^{m_t} y_{i,t}$ (see Figure 3 (d)).
- c_i is the cooperation power of strip i .
- $o_{i,t}$ is the order of strip i at time t .

Equation (3) is a quadratic convex optimization equation with three linear constraints, which are described as follows:

- (3a) Order constraint refers to the ordering of strips that is determined during discrete optimization, which must be preserved.
- (3b) Boundary constraint ensures that the produced layout is within a bounding region with the size of B .
- (3c) Gap constraint separates the top part of the layout from the bottom part. H is a constant that adjusts the gap.

We use Mosek [4] to find the global optimum in polynomial time.

Pattern Enhancement A highly compact and symmetrical layout can be produced by minimizing the wiggle and asymmetry between adjacent time. However, this process may also “hide” some prominent patterns. For example, the converging and diverging behaviors of strips may not be easily perceived in a layout (See Figure 3 (b)). Both patterns reflect important analysis scenarios in which important

topics converge into or diverge from cooperation (or competition). Our flexible optimization method can automatically enhance the patterns by extending the distance between two converging/diverging river strips, which enables users to perceive important patterns (G6). This method enlarges and highlights the patterns by adding a constraint to the optimization process (i.e., Equation (3)). We determine the location of the convergence or divergence of large strips, and we then add a constraint equation to Equation (3) to enlarge and highlight such patterns.

$$|y_{i,t} - y_{i,t+1}| = \gamma \sum_i^{m_i} h_i \quad (4)$$

where $y_{i,t+1}$ and $y_{i,t}$ denote the vertical positions of strip i at two adjacent time stamps, in which $h_i = c_{i,t+1}$ is for the convergence pattern and $h_i = c_{i,t}$ is for the divergence pattern. We set $\gamma = 0.5$ in our optimization system. Figure 3 (e) shows the right term of the above equation. Figures 3 (b) and (c) show the comparison of the optimized layouts without and with pattern enhancement, respectively.

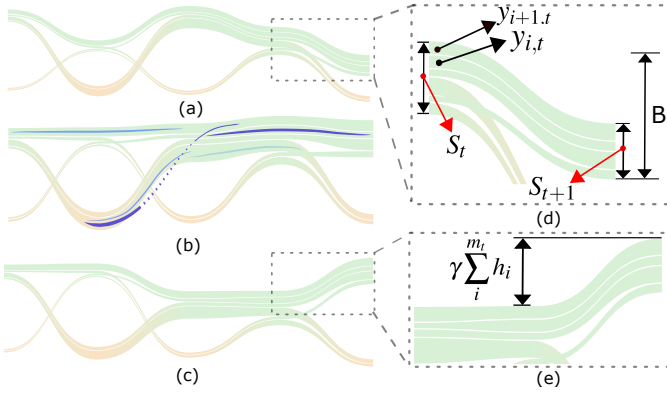


Fig. 3. (a) Initial Layout without optimization; (b) optimized layout; (c) optimized layout with enhanced divergence and convergence patterns; (d) and (e) parameters for Equations (3) and (4), respectively.

5.3.2 User Interactions

EvoRiver supports various advanced and basic interactions, such as panning and zooming (G6).

Show overview first and details-on-demand. EvoRiver shows a visual summary of topic cooperation at the beginning to keep things simple and short. Users can view the evolution of relationships between topics and topic leaders by simply pressing the “Space” key, which leads them to the threads that represent the topic leaders strips.

Compare topics pair-wisely. An in-place view offers an occlusion-free visualization that facilitates the exploration of pairwise similarities between topics (G3). Figure 1 (c) shows an example of the in-place view. A user can click on a river strip, and EvoRiver is split vertically. The empty space is expanded to show the in-place view. Pairwise similarities between the focus topic and other topics is drawn as an arc in the split region, and the width of the arc encodes the similarity.

Examine conversations of topic leaders. A user can click on a thread, and a word cloud appears to show a visual summary of the keywords that are being used by the group of topic leaders. The tweets of these topic leaders on the corresponding topic are listed in the tweet list view on the right side of EvoRiver. Clicking a keyword can update the tweet list by automatically filtering out the tweets that do not contain such a keyword. When a user hovers over a word in one word cloud, the background of the same word in other word clouds will be highlighted. Therefore, for less salient words, it will be easier to see the difference of the words in different word clouds.

Unfold cooperation power. A user may want to unfold the cooperation power of a topic and examine its competition and cooperation powers in details (G6). Our system displays a line chart similar to the famous Playfair’s trade-balance chart that shows the competition and cooperation powers. Figure 1 (b) shows the chart with two lines in orange and green representing the competition and cooperation powers, respectively. The regions between the two lines intuitively encode cooperation power. A line chart displaying the number of tweets posted about the topic is also displayed above EvoRiver (see Figure 1 (a) top).

6 EVALUATION AND CASE STUDIES

We conduct an experiment to validate our cooperation model and use two case studies to demonstrate the usefulness and effectiveness of our system. We also invite two domain experts to evaluate our system.

6.1 Data Preparation

Two large-scale Twitter data sets are collected. The first one is called **business topic data** collected from January 1 to December 20, 2013, and covers nine large IT companies that produce consumer devices: *Amazon, Apple, Blackberry, Google, HTC, Microsoft, Nokia, Samsung, and Sony*. A total of 436,791,811 tweets about the company are collected using 1,539 keywords and hashtags, such as #iPad, Xbox, and Nokia. The second is called **social topics data** that cover the most important general topics, including *law and order, health care, government and politics, welfare, job, general economy, environment and energy, money, spending, and international issues*. These topics, regarded as those that the public is most concerned about, are derived by our collaborators from Gallup [2] using 449,835,519 tweets extracted by 3,336 keywords. The tweets are grouped into the corresponding topics according to the keywords and hashtags used. According to the 4-hour time window, time series sequence of the topics in above two data are divided into 2,124 and 2,192 time points, respectively.

Considering the extensive list of general keywords, the retrieved tweets contain spams and irrelevant messages. An SVM classification technique [11] is used to clean up the data; two professional coders are hired to manually label 1,000 tweets as relevant or irrelevant for each topic. Inter-coder reliability is measured using the Krippendorffs alpha, a popular measure of inter-coder reliability in communication and other social scientific research. The alpha value is 0.83 ($p < 0.01$), suggesting that the coding process is reliable. A model is trained for each topic using 600 tweets; the remaining 400 tweets are used to test the model. Common words in tweets are removed before classification. We choose term frequency (i.e., the number of occurrence of a word in a tweet) to generate feature vectors considering tweets are always short [45]. The average precision and recall rates are 0.80 and 0.84, respectively. The trained models are used to classify the tweets in each topic, and approximately 70% of the tweets are classified as relevant.

Classification scheme of topic leaders. A user is considered a member of k -topic publics if 0.75 k or more of his/her tweets (with a minimum = $2k$) focus on each of the k topics. For example, 1-topic publics devote 75% or more of their tweets to one topic alone. Of the total number of users, 15% focus on one topic, 3% on two topics, 1% on three topics, 0.2% on four topics, 0.1% on five topics, and even fewer on six or more topics. To simplify, one-topic users are placed into single-topic publics, two- to five -topic users are placed under multi-topic publics, and the rest are placed under no-focus publics.

In each category, the topic leaders are extracted based on their Klout scores [3], a popular measure of influence on social media. For single-topic publics, we select the top 50 users from each of the 10 topics to arrive at 500 leaders. For multi-topic and no-focused publics, we select the top 500 users from each group, respectively. Most single-topic leaders are political figures, organizations, and interest groups (for social topics data), and official accounts (for business data), whereas most no-focus leaders are from the mass media. The multi-topic leaders are a mix of the previous two groups.

6.2 Model Evaluation

Time series models estimated in the study are evaluated by three measures: the overall goodness of fit (R^2) of the regression model, the standard error of the estimates ($se_{\hat{y}}$), and the presence of autocorrelation (Durbin-Watson d [15]).

R^2 indicates the explanatory power of the model, as measured by the proportion of the variance in the dependent variable explained by the vector of the independent variables included in the equation. As shown in the upper panel of Table 1, the mean value of the R^2 , averaged from the 2,196 date points of the time series for each of the 10 equations, ranges from 0.95 to 0.99, suggesting that more than 90% of the fluctuations in public attention to the 10 issues are explained by the model. As such, the model appears to be highly effective and robust.

	Economy	Environment	Government	Health	Int'l Issues	Job	Law	Money	Spending	Welfare
R^2	0.96 (0.005)	0.97 (0.005)	0.98 (0.003)	0.97 (0.005)	0.97 (0.005)	0.99 (0.001)	0.96 (0.009)	0.98 (0.004)	0.95 (0.006)	0.96 (0.01)
$se_{\hat{y}}$	0.01 (0.001)	0.002 (0.0005)	0.05 (0.01)	0.03 (0.001)	0.01 (0.002)	0.02 (0.001)	0.02 (0.008)	0.01 (0.0005)	0.002 (0.0003)	0.006 (0.002)
DW - d	1.85 (0.14)	2.01 (0.11)	1.93 (0.014)	1.83 (0.14)	1.85 (0.07)	2.00 (0.05)	1.98 (0.07)	1.89 (0.11)	2.12 (0.08)	2.04 (0.05)

Table 1. Average and standard deviation of the three measures, and the evaluation result shows that our model is highly effective and robust.




Fig. 4. Visualization of the dynamics of topic cooperation and competition from August 13, 2012 to March 25, 2013.

$se_{\hat{y}}$ describes the predictive power of the model, as measured by the precision of predictions based on the model. As shown in the middle panel of Table 1, the mean values of $se_{\hat{y}}$ for most issues vary in a narrow range of 0.002-0.02 on a scale of 0-1. The mean values of $se_{\hat{y}}$ for government and health issues are 0.05 and 0.03, respectively, which are both greater than those for other eight issues but still acceptable.

DW-d tests the autocorrelation between the adjacent residuals of the model, with d ranging from 0 (perfectly positive autocorrelation) to 4 (perfectly negative autocorrelation). As shown in Table 1, the d values are all close to 2 (i.e., absence of autocorrelation), suggesting that the residuals of the 10 equations are essentially white noise.

6.3 Case Studies

This section demonstrates the use of the system for exploring cooperation among social topics and business topics.

Cooperation among Social Topics The first case study is conducted to show the use of EvoRiver to help analysts interactively explore the dynamic cooperation and competition among various topics.

The EvoRiver allows users to gain a quick overview of the dynamics of topic cooperation. Figure 4 shows orange strips in the bottom part which are more and larger, than the green strips in the top part of the EvoRiver visualization; thus, most of the topics tend to be competitive over a long period of time, suggesting that these topics are inclined to recruit public attention from other topics. Figure 4 shows some strips switching between the top and bottom parts of the visualization, indicating that the corresponding topics may transit from competition to cooperation or vice versa. This pattern is more apparent and frequent from September 24, 2012 to November 5, 2012 (in Figure 4 (a) during the 2012 US Presidential election season). Several topics switch from competition to cooperation one by one (i.e., *spending*, *welfare*, *economy general*, *international issues*, and *government and politics*).

EvoRiver allows the unfolding of the pattern and investigation of a phenomenon. For each switching topic, the topic is selected as focus (e.g., *international issues*) and the EvoRiver is split for the in-place view near the transition time interval (Figure 1 (c)); from the in-place view, *government and politics* is always the most similar topic to the focus topic. For instance, the arc that connects the focus topic (*international issues*) to *government and politics* is the thickest one among all other arcs in Figure 1 (c). Next, we examine the detailed relationship between *government and politics* and every switching topic when the switching topic transits. Figure 1 (d) highlights the strip of *international issues* with the keywords when it switches from competition to cooperation. The largest keyword in the word cloud is “debate” indicating that it is the most popular of the keywords used. After examining the tweets with “debate” in the tweet list view, we find that the third presidential debate took place on Oct 22, 2013. This event is a part of the topic of *government and politics*. The debate mostly focused on the international issues of America, which suggests that *government and politics* cooperates with the *international issues* to some extent. Similar observations are noted for the relationships between *government and politics* and other remaining switching topics.

The topic *government and politics* is considered the driving force that pushes the topics to switch and be more cooperative. The tendency

of the topic for cooperation gradually decreases its cooperation power (highlighted in dark orange). Nevertheless, Figure 1 (a) shows that *government and politics* remains more competitive than cooperative until October 29, 2012. It is interesting to examine why the topic finally switches. We check the temporal public engagement of the topic on the top of EvoRiver (see Figure 1 (a)). The curve shows that the public engagement of *government and politics* increases suddenly after Oct 29, 2012. We speculate that the sudden increase may be related to the switching behavior of the topic.

EvoRiver enables us to easily relate topic leaders to topics over a period of time and explore interesting patterns. Around December 17, 2012, *law and order* transited from competition to cooperation (see Figure 4 (b)). We explore the relationship among *law and order* and other topics by examining the pairwise similarities between *law and order* and other topics in the in-place view (see the enlarged region in the dashed line rectangle in Figure 5 (a)). *Law and order* is most similar to *government and politics* which was highly cooperative during that time. We further enable the composite display mode of EvoRiver to inspect the relationship between the topic leaders and the topics (Figure 5 (b)). We can see that no-focused leaders (in dark blue) exert continuous influence over a long period of time on *government and politics*, and start to engage into *law and order*.

We click on the threads representing the topic leaders to explore the tweets posted by these no-focused leaders on the two topics (see Figure 5 (b)). In the word clouds, the keyword “gun” is the most salient for both topics. When we examine the related tweets that contain “gun”, we find that those no-focused leaders with respect to *law and order* are mostly talking and criticizing the gun control laws after the gun shooting at Sandy Hook Elementary School in Connecticut on December 14, 2012. The no-focused leaders with respect to *government and politics* are tweeting about President Obama’s action towards gun control. One example of such tweet is: “@Piers Morgan: Never seen President Obama so upset. And he just ordered the gun control debate to happen...”. The event is considered the reason for the high similarity and the resulting cooperation between the two topics. The fact of both topics cooperating recruits more public salience from other topics. Figure 5 (b) also shows another interesting pattern within a topic. *No-focused leaders* play a major role in promoting the cooperation power of *health care* together with *single-topic leaders* and *multi-topic leaders*. An examination of the word cloud and relevant tweets of these *no-focused leaders* shows that most of the users are not only discussing gun control laws but also emphasizing providing mental health services for people with ongoing mental health conditions (e.g., “@Michael Moore: The way to honor these dead children is to demand strict gun control, free mental health care, and an end to violence...”).

Tweets sent by *single-topic leaders*, the interest group that take advantage of opportunities to promote their intention for increasing the budget on mental health care, mainly call for more funding (for example, “@Andy Borowitz: Politicians: if, as you say, this is about mental health and not guns, why are you cutting funds for mental health?”). Clearly, the *single-topic leaders* contribute significantly to recruiting public attention from gun control to mental health services, resulting in the high competition power of *health care*.

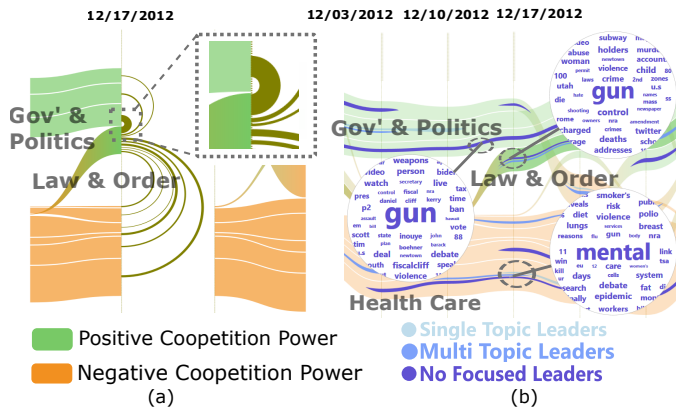


Fig. 5. (a) In-place view using *law and order* as the focus topic; (b) detailed analysis with word clouds in a composite display of strips and threads.

Moreover, EvoRiver is useful for comparing the various roles entailed by different groups of topic leaders through the dynamics of topic competition. Figure 6 shows an overall distribution of different groups of topic leaders over a long period of time. Two obvious patterns are notable. First, there are more dark blue and light blue threads than the normal blue threads, suggesting that no-focused leaders and single-topic leaders play a more significant role in shaping the dynamics of topic competition. Second, more light blue threads than dark blue threads appear in Figure 6 (a) before the week of November 6, 2012 (the date of the U.S. presidential election), whereas more dark blue threads appear in Figure 6 (b) after the week. This pattern indicates that non-topic leaders exert more significant influence on the dynamics of topic competition before the week, but single-topic leaders dominate the influence after that week. Collaborators of the present study are particularly interested in this pattern, and hypothesized the observation may be largely related to the presidential election. In the week of November 12, 2012, we can observe two transition lines representing no-focused leaders were transitioning between *government and politics* and *international issues* (see Figure 6 (c)). Two word clouds are then examined for the two topics, and the keyword “gaza” is most salient for both topics. Further examination of “gaza” in the tweet list view reveals no-focused leaders with respect to *international issues* are mostly talking about the attack on Gaza and Israel’s call for U.S.’s support on Gaza crisis, and no-focused leaders with respect to *government and politics* are tweeting about President Obama’s attitude towards the Gaza crisis: we are fully supportive of Israel’s right to defend itself. This event can be considered the reason that the two topics continued to be cooperative after the third U.S. presidential debate.




Fig. 6. Visual summary of the co-evolutionary relationships among topics and topics leaders for the social topics data.

EvoRiver also allows users to examine the competition and cooperation power of a topic. Figure 1 (a) shows that some topics, such as *job* and *spending*, appear to have an equal amount of competition power. The Playfair-style chart is used to unfold the competition power of the topics. In Figure 1 (b), the competition and cooperation powers of *job* and *spending* are displayed using the orange and green curves, respectively. The cooperation power of *spending* is much higher than that of *job*, indicating that *spending* is more cooperative than *job*. This pattern demonstrates that the same competition power may have rather different implications and thus the Playfair-style charts are needed to explore the competition power.

Cooperation among Topics of IT companies To demonstrate the system’s capability, it is applied to business data analysis. Figure 7 presents the findings of the analysis, which shows that some patterns are similar to those in the social topics data as observed in Figure 4. Whereas most topics still tend to be competitive, some topics may switch between cooperation and competition. *Nokia*, the topic with largest cooperation power, switches from competition to cooperation under the influence of no-focused leaders on September 2, 2013. The public engagement curve above EvoRiver indicates that the public engagement of *Nokia* peaked during that time. A word cloud is used to reveal the pattern through digging into the data and to examine what no-focus leaders are talking about; the keyword “deal” is given considerable attention. Word clouds of the topics that also switched during the same time interval are then examined for any relationship between *Nokia* and the topics. The examination has shown that Microsoft is as prominent as the keyword “deal”. Interestingly, the no-focused leaders likewise exert a certain influence on *Microsoft*; an examination of tweets containing the keyword “deal” shows that the exhibited pattern is related to *Microsoft’s* acquisition of *Nokia’s* Devices and Services Division in a \$7.17 billion deal. *Microsoft* and *Nokia* cooperated with each other to attract public attention immediately after the deal.

EvoRiver allows us to visually compare different behaviors of distinct topic leader groups. Figure 7 further shows that different topic leader groups play different roles in shaping topic competition dynamics on Twitter. During the time range in the grey ellipse, single-topic leaders exert considerable influence over *Microsoft* and *Samsung*, whereas no-focused leaders play a more significant role in *Amazon*. By examining the tweet lists (see the tweets shown in Figure 7), we find that the single-topic leaders are the official accounts of the companies who promote their services or products. In contrast, the no-focused leaders in *Amazon* are the media accounts such as CNN and Wall Street Journal, who mainly report the breaking news that Amazon founder Jeff Bezos bought Washington Post. From this case study, we can see the recruitment behaviors of different topic leader groups are complex and may change over time. EvoRiver is valuable and enables analysts to detect and analyze the behaviors quickly.

6.4 User Feedback

Two professors (PA and PB) in Communication and MCedia studies from two universities were asked to work on this study, identify research problems, and collect design requirements. The system was iteratively improved throughout the frequent meetings with the domain experts. The case studies were conducted when the system was ready. The experts provided interesting insights into the research findings. Their feedback is summarized as follows:

Visualization Design. The visual design of EvoRiver was received very well by both PA and PB. They agreed that the tool is intuitive, engaging, and easy to use, and were very impressed by the interactive features. PA said that the user interactions are smooth and very helpful for data analysis and exploration. He mentioned that the river metaphor of EvoRiver based on river confluence and bifurcation greatly helps his understanding of the design, and added that the consistent use of the color in the system further enhances his understanding of the visual encodings. PB was impressed by the visualization, and particularly liked the composite design. He said that the design is “great” and that such a great design “allows me to easily connect topic leaders to the corresponding topics.” However, despite the experts’ appreciation of the overall design, they found the Playfair-style chart below EvoRiver difficult to understand because they have never seen this kind of chart before. Their inability to find similar metaphors also led to the difficulty. Nevertheless, the design is accepted and liked for its feature for unfolding cooperation patterns after the design is explained.

Usability. Both users confirmed the usefulness and effectiveness of the system and wanted to use the system in teaching and research. PA said that “The system is a great tool. I can use it to quickly find interesting competition patterns among topics and then drill down to the patterns to see more details.” He especially liked the feature of automatic pattern enhancement, which allows him to see the diverging and converging patterns easily. PB noted that the system is not only

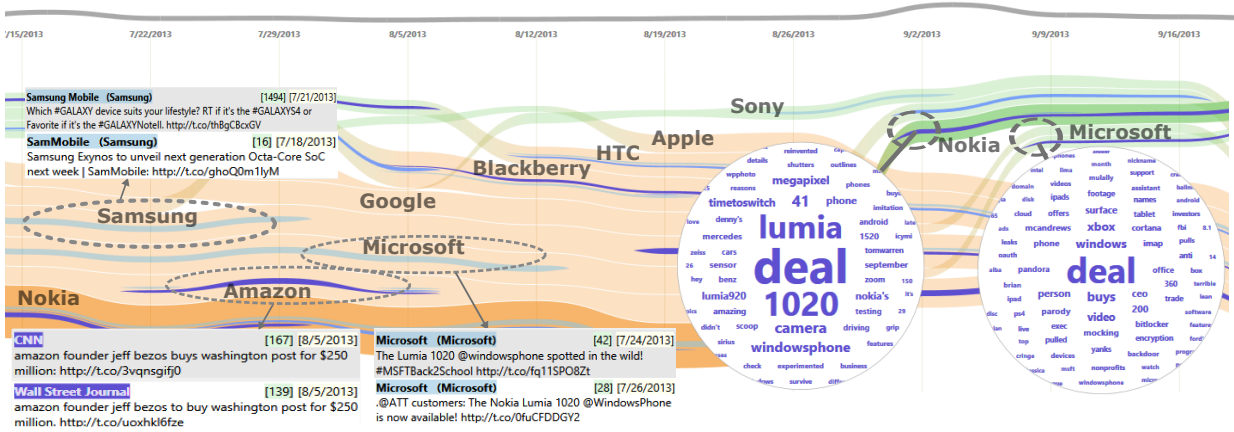


Fig. 7. Visualization of the dynamics of topic competition for the business data from July 17, 2013 to September 16, 2013.

useful for data analysis but also helpful for easily communicating their findings to colleagues or a wider audience.

Suggestion. The users provided valuable suggestions to improve the system. PA suggested the showing of a tweet volume curve above EvoRiver to display the popularity of a topic over time, and proposed the novel classification of users from the perspective of “topic publics”. PB suggested that the design be kept as simple as possible and that too many views should not be introduced. PB also suggested that the system should support exporting the numerical values to an Excel file.

6.5 Discussion

The experiments, case studies, and user evaluation confirm the usability and effectiveness of the system. This study has two important implications. First, the study reveals some very interesting patterns and insight into the dynamics of topic competition on social media: (1) topics tend to compete with other topics in most cases, as seen in large-scale data sets covering both politics and business scenarios over a long period of time; and (2) cooperative topics are usually semantically similar to each other, which has been rarely reported and discussed in the field.

Second, the classification scheme of topic leaders makes it possible gleaning of insight into various roles entailed by different groups of topic leaders. This study shows that multi-topic leaders or no-focus topic leaders may exert greater impact on the dynamics of topic competition. The classification of topic leaders also leads to a special circumstance that our results show fewer transition lines than the ones in [44]. For example, we find that the single topic leaders mostly focus on just one topic, thus the corresponding threads usually do not transit.

Our system can be applied to other domains. For example, in the business marketing domain, different products compete with one another to attract more customers (similar to public salience of topics in our paper), when the coverage (such as coverage in advertising and news reports) of one product by topic leaders (or other types of influential entities such as TV channels or newspapers) may have a recruitment effect on the customers of other products. This is a very common competition scenario between commercial products in daily lives. On the other hand, some products may also cooperate with each other to gain more customers. For example, the sales growth of iPhone may lead to the increasing sales of Mac to some extent, and vice versa. This is also a popular cooperation scenario between products. Our visual analysis system can be used to model and analyze dynamic cooperativeness and competitiveness between commercial products from different industries, which will offer practical insights for marketing professionals.

The EvoRiver visualization is also of great usefulness in a scenario where stacked graphs must be split into two parts. For instance, EvoRiver can be used to analyze the performance of different soccer teams over time. Each strip represents a team and the height of the strip at each time point represents the number of goals of that team at that time point. The strip can transit between two states (Win or Loss). Famous soccer players can be overlaid on the strips as threads to highlight their contribution to the number of goals of the team.

The system is also intended to support data streaming and detect interactions among topics in real time. When new topics emerge, we first compute the independent variables such as the public salience of the new topics and coverage of the topics by topic leaders. Then the competition power of the new topics can be estimated similarly. However, visualization of new topics in runtime brings new challenges. We need to further develop a sophisticated layout algorithm to not only reduce the crossings between the strips of new topics and those of the existing ones, but also maintain the stability of the layout and preserve mental map of existing visualization. A future plan is the development of a web-based system. The data preprocessing, model analysis, and optimization of the visualization layout could be managed at the server side, and the result will be rendered at the client side.

The present work, however, has some limitations. First, the classification scheme of topic leaders is derived based on the statistical analysis and the suggestion of domain experts. The classification method works well in the case studies herein but it may not be perfect to classify topic leaders for other data sets. A systematic study of the classification method for topic leaders is worth further study. Second, in SVM classification, new models needed to be trained for new data sets, which are different from the politics and IT company data. Collecting sufficient and good training samples is time-consuming and challenging, thus the possibility of using clustering methods to remove irrelevant tweets needs to be studied [32]. Third, the competition power is derived using only the recruitment effect in the model. In the future, the distraction effect needs to be studied and whether this effect could also implicitly affect the dynamics of topic competition needs to be examined. We would also like to investigate the visual design alternatives. For example, in the in-place view, we can use the transparency/brightness of strips while excluding the focus strip to represent the similarity, or connect the arc to the right side of the split area.

7 CONCLUSION

This paper introduces EvoRiver, a visual analysis system for interactive analysis and understanding of the cooperation and competition among topics on social media via a seamless integration of a novel model and a set of new visualization techniques. The model characterizes the dynamic interactions among topics and their co-evolving relationships with topic leaders on social media. The interactive visualization techniques are designed to display the complex dynamics of topic interactions captured by the model intuitively. The complicated details of the model are hidden from analysts but the model can be updated implicitly through interactions with the visualizations to better characterize the topic interactions.

ACKNOWLEDGMENTS

The work is partly supported by the National Natural Science Foundation of China (No. 61379076), the Program for New Century Excellent Talents in University (No. NCET-12-1087), Zhejiang Provincial Natural Science Foundation (No. LR14F020002), HKRGC GRF (CityU 154412), and Microsoft (NTU M4061358.060).

REFERENCES

- [1] Apache lucene. <http://lucene.apache.org/core/>, Mar. 2014.
- [2] Gallup. <http://www.gallup.com>, Mar. 2014.
- [3] Klout. <http://klout.com>, Mar. 2014.
- [4] Mosek. <http://www.mosek.com/>, Mar. 2014.
- [5] Project cascade. <http://nytlabs.com/projects/cascade.html>, Mar. 2014.
- [6] Revisit. <http://moritz.stefaner.eu/projects/revisit>, Mar. 2014.
- [7] S. Asur, B. A. Huberman, G. Szabo, and C. Wang. Trends in social media: Persistence and decay. In *Proceedings of the International AAAI Conference on Weblogs and Social Media*, pages 434–437, 2011.
- [8] D. S. Boninger, J. A. Krosnick, and M. K. Berent. Origins of attitude importance: Self-interest, social identification, and value relevance. *Journal of Personality and Social Psychology*, 68(1):61–80, 1995.
- [9] L. Byron and M. Wattenberg. Stacked graphs - geometry & aesthetics. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1245–1252, 2008.
- [10] N. Cao, Y.-R. Lin, X. Sun, D. Lazer, S. Liu, and H. Qu. Whisper: Tracing the spatiotemporal process of information diffusion in real time. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2649–2658, 2012.
- [11] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):27:1–27:27, 2011.
- [12] M. Coscia. Competition and success in the meme pool: a case study on Quickmeme.com. In *Proceedings of the International Conference on Weblogs and Social Media*, 2013.
- [13] W. Cui, S. Liu, L. Tan, C. Shi, Y. Song, Z. Gao, H. Qu, and X. Tong. TextFlow: Towards better understanding of evolving topics in text. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2412–2421, 2011.
- [14] W. Dou, L. Yu, X. Wang, Z. Ma, and W. Ribarsky. HierarchicalTopics: Visually exploring large text collections using topic hierarchies. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2002–2011, 2013.
- [15] J. Durbin and G. S. Watson. Testing for serial correlation in least squares regression. *Biometrika*, 37(3/4):409–428, 1950.
- [16] P. Eades and S. Whitesides. Drawing graphs in two layers. *Theoretical Computer Science*, 131(2):361–374, 1994.
- [17] D. Frey. Recent research on selective exposure to information. *Advances in experimental social psychology*, 19(1):41–80, 1986.
- [18] A. Guille, H. Hacid, C. Favre, and D. A. Zighed. Information diffusion in online social networks: A survey. *ACM SIGMOD Record*, 42(2):17–28, 2013.
- [19] S. Havre, E. Hetzler, P. Whitney, and L. Nowell. ThemeRiver: Visualizing thematic changes in large document collections. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):9–20, 2002.
- [20] M. Hu, S. Liu, F. Wei, Y. Wu, J. Stasko, and K.-L. Ma. Breaking news on twitter. In *Proceedings of ACM CHI*, pages 2751–2754, 2012.
- [21] F. T. Juster and F. P. Stafford. The allocation of time: Empirical findings, behavioral models, and problems of measurement. *Journal of Economic Literature*, 29(2):471–522, 1991.
- [22] J. A. Krosnick. Government policy and citizen passion: A study of issue publics in contemporary america. *Political Behavior*, 12(1):59–92, 1990.
- [23] J. A. Krosnick and S. Telhami. Public attitudes toward israel: A study of the attentive and issue publics. *International Studies Quarterly*, 39(4):535–554, 1995.
- [24] S. Liu, W. Cui, Y. Wu, and M. Liu. A survey on information visualization: recent advances and challenges. *The Visual Computer*, 2014, to be appeared.
- [25] S. Liu, Y. Wu, E. Wei, M. Liu, and Y. Liu. StoryFlow: Tracking the evolution of stories. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2436–2445, 2013.
- [26] S. Liu, M. X. Zhou, S. Pan, Y. Song, W. Qian, W. Cai, and X. Lian. Tiara: Interactive, topic-based visual text summarization and analysis. *ACM Transactions on Intelligent Systems and Technology*, 3(2):25:1–25:28, 2012.
- [27] A. Marcus, M. Bernstein, O. Badar, D. Karger, S. Madden, and R. Miller. Twitinfo: aggregating and visualizing microblogs for event exploration. In *Proceedings of ACM CHI*, pages 227–236, 2011.
- [28] C. Matyszczyk. Nokia and htc mock samsung galaxy s5 launch (badly). http://news.cnet.com/8301-17852_3-57619437-71/nokia-and-htc-mock-samsung-galaxy-s5-launch-badly/, Feb. 2014.
- [29] M. McCombs and J. H. Zhu. Capacity, diversity, and volatility of the public agenda: Trends from 1954 to 1994. *Public Opinion Quarterly*, 59(4):495–525, 1995.
- [30] S. Myers and J. Leskovec. Clash of the contagions: Cooperation and competition in information diffusion. In *Proceedings of IEEE International Conference on Data Mining*, pages 539–548, 2012.
- [31] H. C. Purchase. Which aesthetic has the greatest effect on human understanding? In *Proceedings of the International Symposium on Graph Drawing*, pages 248–261, 1997.
- [32] B. Renoust, G. Melançon, and M.-L. Viaud. Measuring group cohesion in document collections. In *Proceedings of IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, pages 373–380, 2013.
- [33] D. O. Sears and J. L. Freedman. Selective exposure to information: A critical review. *Public Opinion*, 31(2):194–213, 1967.
- [34] C. Shi, W. Cui, S. Liu, P. Xu, W. Chen, and H. Qu. RankExplorer: Visualization of ranking changes in large time series data. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2669–2678, 2012.
- [35] L. Shi, F. Wei, S. Liu, L. Tan, X. Lian, and M. X. Zhou. Understanding text corpora with multiple facets. In *Proceedings of IEEE Symposium on Visual Analytics Science and Technology*, pages 99–106, 2010.
- [36] G. Sun, Y. Wu, R. Liang, and S. Liu. A survey of visual analytics techniques and applications: State-of-the-art research and future challenges. *Journal of Computer Science and Technology*, 28(5):852–867, 2013.
- [37] F. Viégas, M. Wattenberg, J. Hebert, G. Borggaard, A. Cichowlas, J. Feinberg, J. Orwant, and C. Wren. Google+ ripples: A native visualization of information flow. In *Proceedings of the international conference on World Wide Web*, pages 1389–1398, 2013.
- [38] C. Ware. *Information Visualization, Third Edition: Perception for Design (Interactive Technologies)*. Morgan Kaufmann, 3rd edition, 2012.
- [39] M. Wattenberg and J. Kriss. Designing for social data analysis. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):549–557, 2006.
- [40] L. Weng, A. Flammini, A. Vespignani, and F. Menczer. Competition among memes in a world with limited attention. *Scientific reports*, 2:335, Jan. 2012.
- [41] F. Wu and B. A. Huberman. Novelty and collective attention. *Proceedings of the National Academy of Sciences*, 104(45):17599–17601, 2007.
- [42] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts. Who says what to whom on twitter. In *Proceedings of the international conference on World Wide Web*, pages 705–714, 2011.
- [43] Y. Wu, S. Liu, k. Yan, M. Liu, and F. Wu. OpinionFlow: Visual analysis of opinion diffusion on social media. *IEEE Transactions on Visualization and Computer Graphics*, 20(12), 2014.
- [44] P. Xu, Y. Wu, E. Wei, T.-Q. Peng, S. Liu, J. J. H. Zhu, and H. Qu. Visual analysis of topic competition on social media. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2012–2021, 2013.
- [45] J. Yin, A. Lampert, M. Cameron, B. Robinson, and R. Power. Using social media to enhance emergency situation awareness. *IEEE Intelligent Systems*, 27(6):52–59, 2012.
- [46] J. H. Zhu. Issue competition and attention distraction: A zero-sum theory of agenda-setting. *Journalism Quarterly*, 69(4):825–836, 1992.