TrendNets: mapping emerging research trends from dynamic co-word networks via sparse representation

Marie Katsurai¹ · Shunsuke Ono²

Received: 23 May 2019 / Published online: 18 October 2019
© The Author(s) 2019

Abstract

Mapping the knowledge structure from word co-occurrences in a collection of academic papers has been widely used to provide insight into the topic evolution in an arbitrary research field. In a traditional approach, the paper collection is first divided into temporal subsets, and then a co-word network is independently depicted in a 2D map to characterize each period's trend. To effectively map emerging research trends from such a time-series of co-word networks, this paper presents TrendNets, a novel visualization methodology that highlights the rapid changes in edge weights over time. Specifically, we formulated a new convex optimization framework that decomposes the matrix constructed from dynamic co-word networks into a smooth part and a sparse part: the former represents stationary research topics, while the latter corresponds to bursty research topics. Simulation results on synthetic data demonstrated that our matrix decomposition approach achieved the best burst detection performance over four baseline methods. In experiments conducted using papers published in the past 16 years at three conferences in different fields, we showed the effectiveness of TrendNets compared to the traditional co-word representation. We have made our codes available on the Web to encourage scientific mapping in all research fields.

Keywords Emerging research trends · Dynamic co-word networks · Burst detection · Science mapping

Introduction

Due to the continuous increase of journals and conferences, as well as the wide spread of preprint servers such as arXiv¹, a massive amount of academic papers has accumulated. Revealing new research topics can bring significant benefits for people involved

¹ <https://arxiv.org/>.

✉ Marie Katsurai
katsurai@mm.doshisha.ac.jp
Shunsuke Ono
ono@c.titech.ac.jp

¹ Doshisha University, 1-3 Tatara Miyakodani, Kyotanabe-shi, Kyoto, Japan

² Tokyo Institute of Technology, 4259 Nagatsuta-cho, Midori-ku, Yokohama, Kanagawa, Japan

in the research environment, including researchers, research administrators, publishers, and funding bodies. However, it is hard even for experts to keep up with the large number of papers published.

To facilitate an understanding of an arbitrary research field, visualization of emerging research topics using papers published in the target field has attracted much attention (Callon et al. 1983; Assefa and Rorissa 2013; Muñoz-Leiva et al. 2012; Ronda-Pupo and Guerras-Martin 2012; Zhang et al. 2012; Hu et al. 2013; Li et al. 2010; Topalli and Ivanaj 2016). One widely used technique is *co-word analysis* (Callon et al. 1983), which maps the knowledge structure from word co-occurrences in the papers. In a traditional approach, the paper collection is first divided into temporal subsets (e.g., yearly or multiple years) according to the publication dates of the papers. Then, focusing on words in the content of papers such as titles and abstracts, the co-occurrence frequency of each word pair is calculated by counting the number of papers in which the two words appear together. Finally, the knowledge structure of each time period is visualized as a co-word network whose nodes correspond to words and whose edges reflect high co-occurrence frequencies between words, and in which the number of edges in the network is generally determined by thresholding edge weights. A series of co-word networks over time (called dynamic co-word networks) helps us to investigate how a target field has developed (Assefa and Rorissa 2013; Muñoz-Leiva et al. 2012; Ronda-Pupo and Guerras-Martin 2012; Zhang et al. 2012; Hu et al. 2013; Li et al. 2010; Topalli and Ivanaj 2016). However, the traditional approach does not have the means to highlight how the emerging research trends change from time period to period; it might depict edges, even if the corresponding word associations are stable or gradually become stronger over time. Such stable word pairs can disturb the finding of discriminative or emerging research topics that characterize the time period. Although there are several methods with which to visualize temporal changes of word clusters or topics (Wang et al. 2014; Song et al. 2014; Mei and Zhai 2005; Hall et al. 2008), these usually describe dominant research topics in a time period using general and frequent single words, rather than specific technical terms that suddenly attract attention.

To make it easier to understand the research topic evolution of a target field, this paper presents TrendNets, a novel emerging research trend visualization approach that detects rapid changes in edge weights in dynamic co-word networks. Our aim here is to find “bursty research topics”—topics that have been growing rapidly rather than topics that are merely popular. That is, a research topic is considered to be bursty in a given time period if it is heavily discussed in that time period, but not in the previous time periods. To measure the burstiness of a research topic, we calculate the difference in word co-occurrence frequency between two successive time periods. In the proposed approach, we first arrange the vectorized edge weights of co-word networks in columns of a single matrix in the order of time. Then, to detect bursty research topics, we formulate a convex optimization framework that decomposes the matrix into a smooth part corresponding to stationary topics and a sparse part corresponding to bursty topics in a computationally efficient manner. We term this framework *Fast Sparse–Smooth Matrix Decomposition*. After this optimization, the non-zero entries of the resulting sparse networks are considered to represent bursty topics for the corresponding time periods. Experiments on both toy data and real publication data show that TrendNets can effectively detect emerging research trends.

The main contributions of this paper can be summarized as follows:

- We introduce a novel bursty research topic detection scheme for a traditional co-word analysis for visualizing emerging research trends. The proposed method is computationally efficient, even for a large matrix calculated from dynamic co-word networks.
- We perform extensive experiments based on both synthetic and real-world datasets, which show that TrendNets effectively describes the bursty topics when compared with the conventional co-word representations.
- Our method can automatically detect technical phrases from a pool of title words without predefined vocabulary, which is useful because obtaining a set of titles is much easier than collecting keywords to papers provided by authors to papers.
- We have made our codes available to encourage science mapping in all disciplines.

The remainder of this paper is organized as follows: the next section provides a brief review of the related studies; the third section presents the details of TrendNets; the results of the experiments on the synthetic dataset and paper collections are presented in the fourth and fifth sections, respectively; finally, the paper is summarized and some possible directions for future work are suggested.

Related work

As scholarly big data is growing, science mapping has become more important for encouraging the activities of researchers, research administrators, and science policymakers. Many studies construct knowledge networks using the bibliographic metadata of papers to visualize the relationships between academic-related entities such as authors, papers, and technical terms. For example, co-authorship networks have been exploited to find influential researchers and research groups (Börner et al. 2005). Constructing citation or co-citation networks is also a popular approach for visualizing content-based relevance among research subfields (Shibata et al. 2008). However, it is difficult for citation-based methods to grasp the latest trend because papers usually require time to be cited by other papers. On the other hand, co-word networks have the advantage of being able to discover emerging technologies in a timely manner because they can be constructed as quickly as new papers are published. Using textual words can directly map the knowledge structure of a research field into a conceptual space. The visualization of word associations in a network form is known to provide intuitional understanding (Motter et al. 2002; Drieger 2013; Doerfel and Barnett 1999). Thus, co-word analysis has been widely used to reveal research advances in several fields, including science education (Assefa and Rorissa 2013), strategic management (Ronda-Pupo and Guerras-Martin 2012), patient adherence (Zhang et al. 2012), and information retrieval (Hu et al. 2013). It has also been used to characterize the publication trends of journals (Ravikumar et al. 2015) or conferences (Liu et al. 2014). In conventional visualization, the number of edges in a network is generally determined via the thresholding of edge weights (i.e., word co-occurrence frequency). However, this simple approach tends to produce a lot of word pairs that are stable over time.

To visualize the transition of research topics in dynamic co-word networks, there are methods that detect clusters in each co-word network separately and then track the clusters over time (Wang et al. 2014; Song et al. 2014). Specifically, Wang et al. (2014) applied community detection to a co-word network at each time period and connected communities based on their node-based or edge-based similarity across successive time periods. Song et al. (2014) clustered words via the Markov Random Field and tracked each cluster using

word similarity to show its development. These conventional methods can visualize the evolution of representative semantic clusters in a target research field. Because the clusters are usually labeled with general words, it is difficult to highlight specific technical terms that are emerging and likely to form a new technology.

Another line of topic evolution identification has leveraged probabilistic topic models such as Latent Dirichlet Allocation (LDA) (Blei et al. 2003). For example, Mei and Zhai (2005) applied a topic model to papers published in each time period separately and then linked the obtained topics across successive periods based on their similarities. Hall et al. (2008) applied LDA to an entire collection of papers to model research topics and then plotted the number of papers assigned to each topic within each time period. There are also advanced models that can directly treat papers' time stamps in topic modeling (Wang and McCallum 2006; Blei and Lafferty 2006; Wang et al. 2008). These topic model-based methods usually label resulting clusters with top frequent words within the clusters, which does not emphasize emergent technical terms. On the other hand, TrendNets considers how the co-occurrence frequency of a word pair suddenly increases from previous time periods, directly extracting temporally representative word pairs.

Several studies have regarded words displaying a sudden increase in usage to be indicative of emerging research trends (Mane and Börner 2004; Kim and Chen 2015). Traditional burst detection approaches include the thresholding of occurrences (e.g., the moving average (Vlachos et al. 2004)) and state transition modeling (e.g., a two-state automaton proposed by Kleinberg (2003)). One of the most famous science mapping tools is named CiteSpace (Chen 2006), and it implements the Kleinberg's burst detection method to show representative words in a specific time period. Kleinberg's method requires two parameters that are difficult to simultaneously tune. To increase the usefulness of science mapping, our method for constructing TrendNets is formulated to work with only a single parameter. In addition, the related works usually assume that research keywords of papers are obtainable for input data. However, author-provided keywords of conference papers are not always available as open bibliographic information. Thus, we opted to use paper titles only, which are much easier to collect than keywords. Because our method focuses on the burst of edge weights in co-word networks rather than single-word occurrences, we can automatically extract emergent research phrases/terms that are composed of multiple words in the network form. In our experiments, we demonstrate that our TrendNets can effectively organize meaningful phrases from single words within paper titles.

Proposed method

This section describes how to construct TrendNets to map emerging research trends from dynamic co-word networks. An overview of the proposed method is shown in Fig. 1. As shown, we first construct a co-word network for each time period and convert its edge weights into a single matrix whose rows and columns correspond to word pair indexes and time period indexes, respectively. Then, we decompose the matrix into a smooth part and a burst part. According to our definition of bursty research topics (see Introduction), the matrix decomposition step calculates the difference in word co-occurrence frequency between two successive time periods. The burst part corresponds to a large increase for a word pair, which is represented as a sparse matrix. Finally, the resulting sparse matrix is rearranged into a time-series of sparse co-word networks, in which each of the networks visualize bursty topics during the corresponding period. The details of each procedure are described below.

Fig. 1 An overview of the proposed method that constructs TrendNets

Co-word network construction

Let Ω be a collection of papers published in an arbitrary research field. We denote the number of unique words in Ω by M . According to the publication dates of the papers, we divide the paper collection Ω into exclusive temporal subsets $\Omega(1), \Omega(2), \dots, \Omega(T)$ such that $\Omega = \bigcup_{t=1}^T \Omega(t)$. To construct a co-word network at the t -th time period ($t \in \{1, 2, \dots, T\}$), following Liu et al. (2014), we calculate how often two words appear on the same paper on average during the period as follows:

$$\bar{W}_t(i, j) = \frac{n_t(i, j)}{|\Omega(t)|}, i \neq j, i, j \in \{1, 2, \dots, M\}, \quad (1)$$

where $n_t(i, j)$ is the number of papers in which the i -th and j -th words appear together during the t -th time period and $|\cdot|$ represents the cardinality of the set. This forms a symmetrical matrix $\bar{\mathbf{W}}_t \in \mathbb{R}^{M \times M}$, which corresponds to an adjacency matrix of a conventional co-word network. Then, as shown in Fig. 2, we turn the upper triangular part of $\bar{\mathbf{W}}_t$ into a column vector $\tilde{\mathbf{w}}_t \in \mathbb{R}^N$, in which $N = M(M - 1)/2$. Finally, we arrange all vectors over time in the columns to construct a single matrix as $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_T] \in \mathbb{R}^{N \times T}$.

Fast sparse–smooth matrix decomposition

In what follows, we extract sparse graphs from the time series of co-word networks by solving a newly formulated convex optimization problem. Specifically, we decompose a given $\mathbf{W}$ into the sum of a sparse matrix and a smooth matrix, in which the sparse matrix is expected to only contain rapidly changed entries that are closely related to bursty topics. The convex optimization problem is given as follows:

$$\min_{\mathbf{S}} \frac{1}{2} \|D(\mathbf{W} - \mathbf{S})\|_F^2 + \lambda \|\mathbf{S}\|_1, \quad (2)$$

where $\mathbf{S}$ contains extracted sparse graphs in its columns, D is a linear operator computing column differences, and $\|\cdot\|_1$ and $\|\cdot\|_F$ are the ℓ_1 and Frobenius norms, respectively. The operator D corresponds to the calculation of the difference in word co-occurrence frequency between two successive time periods. In Prob. (2), the first term plays a role in promoting the sparsity of $\mathbf{S}$ (NOTE: the ℓ_1 norm is the tightest convex relaxation of the ℓ_0 pseudo-norm, i.e., a reasonable sparsity measure). Meanwhile, the second term keeps the remaining matrix $\mathbf{W} - \mathbf{S}$ smooth in the column direction, implying that the matrix is expected to consist of entries that are not greatly changed during any periods.

Fig. 2 Matrix construction from a time series of co-word graphs

Because Prob. (2) is a convex but nonsmooth optimization problem and, indeed, has no closed-form solution, we have to use some iterative algorithms to solve it. In our framework, the fast iterative shrinkage–thresholding algorithm (FISTA) (Beck and Teboulle 2009) is adopted. FISTA can solve convex optimization problems in the following form:

$$\min_{\mathbf{x}} f(\mathbf{x}) + g(\mathbf{x}), \quad (3)$$

where f is a smooth convex function with a Lipschitzian gradient, and g is a possibly non-smooth convex function, and its proximity operator² (Moreau 1962) is available. The algorithm is given as follows: for an initial vector $\mathbf{y}_0$ and $z_0 = 1$, iterate

$$\begin{cases} \mathbf{x}_{k+1} := \text{prox}_{\frac{1}{L}g}(\mathbf{y}_k - \frac{1}{L}\nabla f(\mathbf{y}_k)), \\ z_{k+1} := \frac{1 + \sqrt{1 + 4z_k^2}}{2}, \\ \mathbf{y}_{k+1} := \mathbf{x}_k + \frac{z_k - 1}{z_{k+1}}(\mathbf{x}_k - \mathbf{x}_{k-1}), \end{cases} \quad (4)$$

where k is the iteration number and L is the Lipschitz constant of ∇f .

Let $f := \frac{1}{2}\|D(\cdot - \mathbf{S})\|_F^2$ and $g := \lambda\|\cdot\|_1$. Then, f is clearly differentiable, and its gradient:

$$\nabla f(\mathbf{S}) = -\mathbf{D}^T \mathbf{D}(\mathbf{W} - \mathbf{S})$$

has the Lipschitz constant L being equal to $\|\mathbf{D}\|_{op}^2$, where $\|\cdot\|_{op}$ is the operator norm. Meanwhile, since g is the ℓ_1 norm, its proximity operator boils down to the soft-thresholding operation:

$$\left[\text{prox}_{\frac{1}{L}g}(\mathbf{S}) \right]_{ij} = \text{sgn}(S_{ij}) \max\{S_{ij} - \frac{\lambda}{L}, 0\}. \quad (5)$$

As a result, FISTA can be readily applied to Prob. (2), which is summarized in Alg. 1.

² The proximity operator of index $\gamma > 0$ of a proper lower semicontinuous convex function h is defined by $\text{prox}_{\gamma h}(\mathbf{x}) := \arg \min_{\mathbf{y}} h(\mathbf{y}) + \frac{1}{2\gamma} \|\mathbf{y} - \mathbf{x}\|^2$.

Algorithm 1: Fast sparse-smooth matrix decomposition by FISTA

```

input :  $\mathbf{Y}_0 = \mathbf{W}$ ,  $z_0 = 1$ ,  $L = \|D\|_{op}^2$ ,  $\lambda$ ,  $k = 0$ 
1 while A stopping criterion is not satisfied do
2    $\mathbf{S}_{k+1} = \mathbf{Y}_k + \frac{1}{L} D^* D (\mathbf{W} - \mathbf{Y}_k)$ ;
3    $\mathbf{S}_{k+1} \leftarrow \text{soft-thresholding}(\mathbf{S}_k, \frac{\lambda}{L})$  by (5);
4    $z_{k+1} := \frac{1 + \sqrt{1 + 4z_k^2}}{2}$ ;
5    $\mathbf{Y}_{k+1} := \mathbf{S}_k + \frac{z_k - 1}{z_{k+1}} (\mathbf{S}_k - \mathbf{S}_{k-1})$ ;
6    $k \leftarrow k + 1$ ;
output:  $\mathbf{S}_k$ 

```

FISTA is a reasonable choice for solving Prob. (2), and the reasons are twofold: (1) its convergence rate is $\mathcal{O}(1/k^2)$ (Beck and Teboulle 2009), i.e., very fast, and (2) it uses only the gradient of f and the proximity operator of g , and, in our case, these can be computed efficiently as explained above (the computational cost is linear with the number of entries in $\mathbf{W}$). These properties are important in our framework because matrix $\mathbf{W}$ is large, and we must avoid expensive procedures, such as singular value decomposition. In all experiments, we set the Lipschitz constant L in Eq. (4) to 4.0.³

Visualization of bursty research topics

After solving (2), we transform the t -th column in the resulting sparse matrix $\mathbf{S}$ into a symmetrical matrix $\tilde{\mathbf{S}}_t \in \mathbb{R}^{M \times M}$. The non-zero, positive elements of $\tilde{\mathbf{S}}_t$ form a new co-word network for the t -th time period in which nodes represent words and edges associate the words to form bursty research topics. For the effective visualization of the research topics, nodes in the graph are generally grouped according to the graph structure (Liu et al. 2014; Silva et al. 2016), in which spectral clustering or community detection techniques are often exploited. For this paper, we use the Louvain method (Blondel et al. 2008) to find semantic clusters of words.

Experiment I: simulation on synthetic data

Synthetic dataset construction

We first quantitatively evaluated the performance of our matrix decomposition method that enables TrendNets. Because a ground truth set for bursty research topic detection is unavailable, we constructed a synthetic dataset using the following two steps:

- *Constructing stable co-word networks* We first generated 50 “stable” co-word networks, which are mostly similar to each other on the basis of a statistical model. Specifically, we used a subset of documents in the Reuters-21578 corpus⁴ that were tagged with “trade” to train a 5-gram language model. For each time period $t \in \{1, 2, \dots, 50\}$,

³ We conducted a preliminary experiment and observed that our matrix decomposition is not sensitive to L .

⁴ <http://www.daviddlewis.com/resources/testcollections/reuters21578/>.

Fig. 3 Illustrations of two types of bursts generated for the synthetic dataset. **a** Type-A bursts: word pairs that suddenly increase and disappear later; **b** Type-B bursts: word pairs that suddenly increase and keep the co-occurrence frequency after that. Deep blue bars indicate time periods that should be detected as bursts. (Color figure online)

the 5-gram language model produced a set of 5000 documents that were statistically identical to the given corpus in which the maximum number of words in each document was set to 12. The document set constructed for t is denoted by $\Omega(t)$. Discarding words that appear less than 30 times in $\{\Omega(t)\}_{t=1}^{50}$ resulted in a vocabulary of $M = 6696$ unique words. We counted the co-occurrence frequency of the i -th and j -th words in $\Omega(t)$, which is denoted by $G_t(i, j)$, producing a co-word network with adjacency matrix $G_t \in \mathbb{R}^{M \times M}$. The resulting time series of networks can be considered as being stable over time. For each word pair (i, j) , the mean and standard deviation values of $\{G_t(i, j)\}_{t=1}^{50}$ are denoted by $\mu(i, j)$ and $\sigma(i, j)$, respectively.

- *Generating synthetic bursts* Next, we artificially generated bursts on the time series of networks to construct ground truth data. As shown in Fig. 3, we considered the following two types of bursts: “Type-A bursts” (Fig. 3a) are word pairs that have a peak of frequency at a certain time and “Type-B bursts” (Fig. 3b) are word pairs that keep the high frequency after a sudden increase. To generate these bursts, for each time period t , we randomly picked a word pair (i', j') with probability $p = 0.05$ from all pairs such that $\mu(i', j') > 3$ and regarded the tuple (i', j', t) as a burst point. We further randomly classified each word pair into Type-A or Type-B with probability $p = 0.95$ and $p = 0.05$, respectively. For a burst point (i', j', t) labeled with Type-A, we increased its edge weight as follows:

$$G_t(i', j') \leftarrow \mu(i', j') + u_t(i', j')\sigma(i', j'), \quad (6)$$

where $u_t(i', j')$ is a random number drawn uniformly from $[3, 6]$. For a burst point (i', j', t) labeled with Type-B, we iteratively computed $G_{t'}(i', j')$ using Eq. (6) for $t' = t, t + 1, \dots, 50$. Finally, we ran the proposed method starting from Eq. (1), in which $n_t(i, j) = G_t(i, j)$ and $|\Omega(t)| = 50$, respectively.

Baselines

In the experiment, the proposed method was compared with the following four baseline methods:

- *Baseline 1: Thresholding* This approach identifies word pairs whose co-occurrence frequencies exceed a predetermined threshold value as bursts. That is, if $\tilde{W}_t(i, j)$ in Eq. (1)

is larger than threshold τ_1 , then the word pair (i, j) is detected as a burst at t . Baseline 1 corresponds to the traditional co-word network visualization.

- **Baseline 2: Thresholding of time derivative** This approach regards word pairs whose time derivatives of frequency exceed a predetermined threshold value as bursts. That is, if $\tilde{W}_t(i, j) - \tilde{W}_{t-1}(i, j)$ is larger than threshold τ_2 , then the word pair is marked as a burst at t .
- **Baseline 3: Thresholding of differences from the average over time** This approach detects word pairs if the difference between those frequencies and the average frequency over time exceeds a predetermined threshold value. That is, if $\tilde{W}_t(i, j) - \mu(i, j)$ is larger than threshold τ_3 , then the word pair is marked as a burst at t .
- **Baseline 4: Kleinberg's method** This approach exploits a well-known burst detection model presented by Kleinberg (Kleinberg 2003) and is implemented to visualize emerging research trends in conventional studies (Chen 2006; Katsurai 2017). It models a time series of a word pair's co-occurrence frequency as a finite-state automaton that consists of a base state and a burst state. The burst degree of the word pair at each time period was calculated from the state transition sequence. The parameters required in this model are s and γ . We used $s = 2.0$ and changed the value of γ .

Simulation results

After applying each method to synthetic data, we calculated the precision and recall measures as follows:

$$\text{Recall} = \frac{\text{\# of correctly detected bursts}}{\text{\# of synthetic bursts}}, \quad (7)$$

$$\text{Precision} = \frac{\text{\# of correctly detected bursts}}{\text{\# of detected bursts}}. \quad (8)$$

Figure 4 shows the Precision–Recall (PR) curve for each method, while Table 1 shows the Area Under the PR Curve (AUC). As shown, we found that the proposed method achieved the best burst detection performance. Baseline 1 provided the worst performance, which implies that the traditional co-word visualization might not highlight an emerging trend in a certain time period. Baseline 4, an application of Kleinberg's burst detection, achieved a relatively better performance; however, there is unfortunately no automatic method for determining its two parameters s and γ . Our burst detection method is more practical because it works with only a single parameter λ , which changes the sparsity of networks. Baseline 2, that calculates the time derivative, is the most similar to our method because both aim to detect a sudden increase in edge weights of dynamic co-word networks. However, Baseline 2 often gives false detection of word pairs associated with large values of $\mu(i, j)$ because the time derivative of such word pairs tends to become relatively large. Our method solved this problem by considering the smoothness along the time direction for each word pair.

Fig. 4 PR curve for burst detection results on synthetic data

Table 1 AUC measured using burst detection results on synthetic data

Proposed method	Baseline 1	Baseline 2	Baseline 3	Baseline 4
0.876	0.148	0.845	0.720	0.427

Experiment II: TrendNets construction using conference papers

This section presents the results of the experiments on real-world datasets to verify the effectiveness of TrendNets. For dataset construction, we chose the following international conferences: CVPR, INFOCOM, and ICASSP. These conferences are considered prominent in the fields of computer vision, networking, and signal processing, respectively. For each conference, we collected titles presented over the past 16 years from DBLP.⁵ DBLP is an open bibliographic database in the fields of computer science and engineering, which provides titles, author names, publication dates, and conference/journal names, but not abstracts, for papers. Then, from the set of words extracted from each title, we removed stop words (e.g., “a”, “the”, “from”) and unnecessary symbols (e.g., “:”, “?”, “!”). Finally, we performed word stemming (Porter 1980). Table 2 summarizes the details of the three datasets.

We divided each conference’s dataset into $T = 8$ temporal subsets so that each time period consists of two years. Our method was computationally efficient: for example, matrix decomposition with $\lambda = 4.0 \times 10^{-4}$ on the ICASSP dataset took approximately 2.31 seconds running on a workstation with a 3.1 GHz Intel Xeon E5-1680 Processor.

Figure 5 shows TrendNets constructed for the latest four periods of the CVPR dataset, which used $\lambda = 8.5 \times 10^{-6}$, considering the limited space of the paper. In the figure, nodes of the same color belong to the same cluster, as found via Louvain method. To make it easier to understand research topics, each usage of words in the stemmed forms was replaced with a non-stemmed word whose co-occurrences with other words on the map were the highest. As shown in Fig. 5, TrendNets effectively captured the characteristics of each time period in CVPR. During 2013–2014, the appearance of “human action recognition” and a burst of “pose tracking” are captured (Fig. 5b). Figure 5c

⁵ <https://dblp.uni-trier.de/>.

Table 2 Details of each conference's dataset

Conference	#Papers	#Unique words
CVPR	7957	5468
INFOCOM	5305	5199
ICASSP	22,431	12,534

(a) TrendNets for CVPR 2011–2012.

(b) TrendNets for CVPR 2013–2014.

(c) TrendNets for CVPR 2015–2016.

(d) TrendNets for CVPR 2017–2018.

Fig. 5 TrendNets constructed using the CVPR dataset ($\lambda = 8.5 \times 10^{-6}$). **a** CVPR 2011–2012, **b** CVPR 2013–2014, **c** CVPR 2015–2016, **d** CVPR 2017–2018

contains a meaningful cluster consisting of “convolutional neural networks” and “object detection,” which characterizes the period from 2015 to 2016. We can see from Fig. 5d that recent emerging trends in learning are “adversarial learning” and “reinforcement learning”. A new emerging term, “visual question,” is also shown as an emerging trend of CVPR 2017–2018. The results suggest that using TrendNets helps us understand topic evolution in the target conference.

Comparison with traditional co-word representation

Finally, we show how TrendNets can be different from the original co-word networks using INFOCOM and ICASSP as examples. We performed the proposed method on

Fig. 6 TrendNets ($\lambda = 1.7 \times 10^{-3}$) and their original co-word networks for the INFOCOM dataset in a comparison of the proposed method and the traditional method. **a** TrendNets for 2005–2006, **b** original co-word network for 2005–2006, **c** TrendNets for 2017–2018, **d** original co-word network for 2017–2018. The numbers of edges of **(b)** and **(d)** were set to the same as **(a)** and **(c)** via the thresholding of edge weights, respectively

the INFOCOM dataset using $\lambda = 1.7 \times 10^{-3}$. Figure 6 shows the two dynamic networks corresponding to TrendNets and the original co-word networks for INFOCOM during 2006–2007 and 2017–2018. For fair comparison, the number of edges in the original networks were set to the same of those of the corresponding TrendNets via thresholding. Compared with the original co-word networks in Fig. 6b, d, TrendNets found multiple semantic clusters to explain the details of the bursty research topics, as shown in Fig. 6a, c. Specifically, the conventional visualization (Fig. 6b, d) linked a lot of words to the central term “networks,” which makes it hard to capture the research topics. Our method solved this conventional co-word visualization’s problem and made it possible to show popular clusters such as “scheduling control” and “learning,” as shown in Fig. 6a, c, respectively.

Interestingly, as shown in Figs. 5c, d and 6a, b, the word “networks” commonly appears in the two conferences TrendNets. Thanks to the network form-based visualization, we can infer the meaning of each conferences “networks” on the basis of surrounding words. In our future work, we will detect and analyze such a difference in word usage across conferences in other fields.

As compared to the INFOCOM results, the original co-word networks seem to successfully visualize emerging trend topics. Unfortunately, a static term “source separation” exists in both Fig. 7b, d, and more temporally representative terms can be seen in TrendNets. Since these results of the ICASSP dataset show only a slight difference, we can assume that there

are differences between conferences in how research topics evolve in terms of time, and it would be valuable to investigate the possibility of the proposed method to estimate it.

In summary, the proposed method enables bursty research trend visualization using a single parameter λ only, even if they are inconspicuous in the original co-word networks. To encourage science mapping in all research fields, the codes for the proposed method are available on the Web.⁶ In addition, our TrendNets effectively revealed emerging research topics that are composed of multiple words, even though we only used the titles of papers without manually assigned keywords. Note that using abstracts in addition to the titles has the potential to cover additional words relevant to the papers that are absent from the titles.

Conclusion and future work

We have developed a novel emerging research trend mapping approach that detects rapid changes in the edge weights of dynamic co-word networks. We have formulated a convex optimization problem and provided an efficient solution to provide sparse networks named TrendNets. Our algorithm works in time linear with the number of words in original networks. In addition, the proposed method can easily work with a single parameter λ , which is practical and user friendly. According to experiments using a synthetic dataset, the proposed method achieved better burst detection performance compared to other baseline methods. Experiments conducted on three conference datasets (CVPR, INFOCOM, and ICASSP) showed the advantage of TrendNets in detecting hot topics that are inconspicuous in the traditional co-word representation. Interestingly, the degree of effectiveness of TrendNets differs among the conferences: we can assume that the temporal stability of a word distribution depends on the research area. Quantifying this type of characteristic is a part of our future work.

Further room for investigation and improvement exists in our work. We currently considered only the rapid changes of the edge weights, ignoring the node degrees of original dynamic networks. To fully make use of the original network structure for research trend visualization, more sophisticated models need to be developed. It is also necessary to design more effective visualization of TrendNets. For example, we should examine how to set the colors and sizes of network components such as nodes, edges, and labels. We will continually improve the proposed method and conduct additional experiments using a larger dataset consisting of multiple conferences/disciplines.

Acknowledgements This research was partly supported by JSPS KAKENHI Grant Number 17K12794.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Assefa, S. G., & Rorissa, A. (2013). A bibliometric mapping of the structure of STEM education using co-word analysis. *Journal of the American Society for Information Science and Technology*, 64(12), 2513–2536.

⁶ <http://mm.doshisha.ac.jp/sci2/TrendNets.html>

- Beck, A., & Teboulle, M. (2009). A fast iterative shrinkage–thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2, 183–202.
- Blei, D. M., & Lafferty, J. D. (2006). Dynamic topic models. In *Proceedings of international conference on machine learning (ICML)* (pp. 113–120).
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *The Journal of Machine Learning Research*, 3, 993–1022.
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 10, P10008.
- Börner, K., Dall'Asta, L., Ke, W., & Vespignani, A. (2005). Studying the emerging global brain: Analyzing and visualizing the impact of co-authorship teams. *Complexity*, 10(4), 57–67.
- Callon, M., Courtial, J.-P., Turner, W. A., & Bauin, S. (1983). From translations to problematic networks: An introduction to co-word analysis. *Social Science Information*, 22(2), 191–235.
- Chen, C. (2006). Citespace II: Detecting and visualizing emerging trends and transient patterns in scientific literature. *Journal of the American Society for Information Science and Technology*, 57(3), 359–377.
- Doerfel, M. L., & Barnett, G. A. (1999). A semantic network analysis of the international communication association. *Human Communication Research*, 25(4), 589–603.
- Drieger, P. (2013). Semantic network analysis as a method for visual text analytics. *Procedia—Social and Behavioral Sciences*, 79, 4–17.
- Hall, D., Jurafsky, D., & Manning, C. D. (2008). Studying the history of ideas using topic models. In *Proceedings of the conference on empirical methods in natural language processing (EMNLP)* (pp. 363–371).
- Hu, C.-P., Hu, J.-M., Deng, S.-L., & Liu, Y. (2013). A co-word analysis of library and information science in China. *Scientometrics*, 97(2), 369–382.
- Katsurai, M. (2017). Bursty research topic detection from scholarly data using dynamic co-word networks: A preliminary investigation. In *Proceedings of IEEE international conference on big data analysis (ICBDA)* (pp. 115–119).
- Kim, M. C., & Chen, C. (2015). A scientometric review of emerging trends and new developments in recommendation systems. *Scientometrics*, 104(1), 239–263.
- Kleinberg, J. (2003). Bursty and hierarchical structure in streams. *Data Mining and Knowledge Discovery*, 7(4), 373–397.
- Li, L., Li, X., Cheng, C., Chen, C., Ke, G., Zeng, D. D., et al. (2010). Research collaboration and ITS topic evolution: 10 years at T-ITS. *IEEE Trans Intelligent Transportation Systems*, 11(3), 517–523.
- Liu, Y., Goncalves, J., Ferreira, D., Xiao, B., Hosio, S., & Kostakos, V. (2014). CHI 1994–2013: Mapping two decades of intellectual progress through co-word analysis. In *Proceedings of the ACM conference on human factors in computing systems (CHI)* (pp. 3553–3562).
- Mane, K. K., & Börner, K. (2004). Mapping topics and topic bursts in PNAS. *Proceedings of the National Academy of Sciences*, 101(suppl 1), 5287–5290.
- Mei, Q., & Zhai, C. (2005). Discovering evolutionary theme patterns from text: An exploration of temporal text mining. In *Proceedings of ACM SIGKDD international conference on Knowledge discovery in data mining (KDD)* (pp. 198–207).
- Moreau, J. J. (1962). Fonctions convexes duales et points proximaux dans un espace hilbertien. *Comptes Rendus de l'Académie des Sciences de Paris*, 255, 2897–2899.
- Motter, A. E., de Moura, A. P. S., Lai, Y.-C., & Dasgupta, P. (2002). Topology of the conceptual network of language. *Physical Review E*, 65, 065102.
- Muñoz-Leiva, F., Viedma-del, J. M. I., Sánchez-Fernández, J., & López-Herrera, A. G. (2012). An application of co-word analysis and bibliometric maps for detecting the most highlighting themes in the consumer behaviour research from a longitudinal perspective. *Quality & Quantity*, 46(4), 1077–1095.
- Porter, M. F. (1980). An algorithm for suffix stripping. *Program*, 3(14), 130–137.
- Ravikumar, S., Agrahari, A., & Singh, S. N. (2015). Mapping the intellectual structure of scientometrics: A co-word analysis of the journal scientometrics (2005–2010). *Scientometrics*, 102(1), 929–955.
- Ronda-Pupo, G. A., & Guerras-Martin, L. A. (2012). Dynamics of the evolution of the strategy concept 1962–2008: A co-word analysis. *Strategic Management Journal*, 33(2), 162–188.
- Shibata, N., Kajikawa, Y., Takeda, Y., & Matsushima, K. (2008). Detecting emerging research fronts based on topological measures in citation networks of scientific publications. *Technovation*, 28(11), 758–775.
- Silva, F. N., Amancio, D. R., Bardosova, M., Costa, L. F., & Jr, O. N. O. (2016). Using network science and text analytics to produce surveys in a scientific topic. *Journal of Informetrics*, 10(2), 487–502.
- Song, M., Heo, G. E., & Kim, S. Y. (2014). Analyzing topic evolution in bioinformatics: investigation of dynamics of the field with conference data in DBLP. *Scientometrics*, 101(1), 397–428.

- Topalli, M., & Ivanaj, S. (2016). Mapping the evolution of the impact of economic transition on Central and Eastern European enterprises: A co-word analysis. *Journal of World Business*, <https://doi.org/10.1016/j.jwb.2016.06.003>.
- Vlachos, M., Meek, C., Vagena, Z., & Gunopulos, D. (2004). Identifying similarities, periodicities and bursts for online search queries. In *Proceedings of the ACM SIGMOD international conference on management of data (SIGMOD)* (pp. 131–142).
- Wang, C., Blei, D. M., & Heckerman, D. (2008). Continuous time dynamic topic models. In *Proceedings of the international conference on uncertainty in artificial intelligence (UAI)* (pp. 579–586).
- Wang, X., Cheng, Q., & Lu, W. (2014). Analyzing evolution of research topics with NEViewer: A new method based on dynamic co-word networks. *Scientometrics*, 101(2), 1253–1271.
- Wang, X., & McCallum, A. (2006). Topics over time: A non-Markov continuous-time model of topical trends. In *Proceedings of the international conference on knowledge discovery and data mining (KDD)* (pp. 424–433).
- Zhang, J., Xie, J., Hou, W., Tu, X., Xu, J., Song, F., et al. (2012). Mapping the knowledge structure of research on patient adherence: Knowledge domain visualization based co-word analysis and social network analysis. *PLoS ONE*, 7, 1–7.