

TopicTracker: A Platform for Topic Trajectory Identification and Visualisation

Yong-Bin Kang^{a,*}, Timos Sellis^{b,**}

^a*Department of Media and Communication, Swinburne University of Technology, Australia*

^b*Facebook, USA*

Abstract

Topic trajectory information provides crucial insight into the dynamics of topics and their evolutionary relationships over a given time. Also, this information can help to improve our understanding on how new topics have emerged or formed through a sequential or interrelated events of emergence, modification and integration of prior topics. Nevertheless, the implementation of the existing methods for topic trajectory identification is rarely available as usable software. In this paper, we present **TopicTracker**, a platform for topic trajectory identification and visualisation. The key of **TopicTracker** is that it can represent the three facets of information together, given two kinds of input: a time-stamped topic profile consisting of the set of the underlying topics over time, and the evolution strength matrix among them: evolutionary pathways of dynamic topics, evolution states of the topics, and topic importance. **TopicTracker** is a publicly available software implemented using the R software.

Keywords: TopicTracker, topic trajectory, topic evolution, topic tracking

*Corresponding Author

**Work done while at Swinburne University of Technology

Email addresses: ykang@swin.edu.au (Yong-Bin Kang), tsellis@fb.com (Timos Sellis)

Code metadata

C1	Current code version	v1.0.0-beta
C2	Permanent link to code/repository used for this code version	https://github.com/Yongbinkang/topicTracker
C3	Code Ocean compute capsule	
C4	Legal Code License	MIT License
C5	Code versioning system used	git
C6	Software code languages, tools, and services used	R
C7	Compilation requirements, operating environments & dependencies	R version ≥ 3.6 , and R packages (igraph, hash, plotrix)
C8	If available Link to developer documentation/manual	https://github.com/Yongbinkang/topicTracker
C9	Support email for questions	ykang@swin.edu.au, yongbin.kang@gmail.com

Table 1: Code metadata

1. Motivation and significance

Topic trajectory identification is a research area that has attracted significant attention from scientific institutions and innovation-industry sectors. In this area, a fundamental is the use of *topic modeling* to discover latent thematic topics (or concepts) from a document collection, where each topic consists of its representative terms extracted from the collection [1]. More recently, dynamic topic modeling has also been utilised to identify topics and their evolution over a time period [2, 3, 4]. Identifying topic trajectories provides precious insights into the dynamics of topics over time. For example, in scientific and patented innovation domains, such trajectories can significantly help to distinguish outstanding research or technological topics, and discover their evolutionary pathways reflecting how new topics have been emerged or formed through a sequential of the events of the emergence, modification and integration of past topics [5, 6, 7, 8]. We view a *trajectory of topics* as a main stream or an evolutionary pathway of topics over time¹.

¹To simplify the presentation, we do not distinguish between ‘trajectory’ and ‘evolutionary pathways’ and use them interchangeably.

Study	Software availability	Domain
He et al. (2009) [9]	X	
Jo et al. (2011) [10]	X	
Song et al. (2014) [6]	X	Scientific literature (topic: technology/knowledge concept)
Zhou et al. (2017) [5]	X	
Zhang et al. (2017) [8]	X	
Jung et al. (2020) [11]	X	
Yoon et al. (2011) [7]	X	
Zhong et al. (2016) [12]	X	
Lee et al. (2017) [1]	X	
Park et al. (2017) [13]	X	Patents (topic: technology/knowledge concept)
Triulzi et al. (2020) [14]	X	
Huang et al. (2020) [15]	X	
Qiu et al. (2020) [16]	X	
Zhang et al. (2015) [2]	X	E-commerce (topic: market brand)
Greene et al. (2016) [3]	X	Politics (topic: political agenda)
Song et al. (2016) [4]	X	History (topic: historical event)
Gaul et al. (2017) [17]	X	Online news (topic: online news)

Table 2: Previous studies on topic trajectory identification and software availability.

Also, we define an *evolutionary pathway* as a series of evolutionary relationships between older topics and newer topics. Consequently, due to this merit, many studies were conducted in recent years for designing different methods for topic trajectory identification (see Table 2).

Unfortunately, the implementation of these methods to encourage their use by the wider scientific community that are interested in topic trajectory identification has been still remained limited. Primarily, a large body of the methods is not available as readily usable software as seen in Table 2. In topic modeling and its application areas, existing software tools were developed with little consideration for preparing and formatting the data for topic trajectory identification in a simple and easy way to use. This requires the users to directly implement their algorithms for topic trajectory identification from the results of topic modeling. Further, this leads to unnecessary time spent for data preparation and limits effective comparison of results produced by different topic trajectory identification models as well.

To address these issues, we have developed **TopicTracker**, a platform for topic trajectory

identification and visualisation. TopicTracker is a software implemented using the R software.

2. Software description

We present the architecture of TopicTracker and its main functionalities.

2.1. Software Architecture

The architecture of TopicTracker’s code design is depicted in Fig. 1. It is designed to distinguish the inferential module for building a *Topic Evolution Tree* (TET) of topics (Phase 1) from the code for visualising their topic trajectories (Phase 2). One key idea in designing TopicTracker is to discover evolutionary pathways (i.e., topic trajectories) between non-contemporary topics by constructing their most likely genealogy tree (i.e., TET) over a given time. Because of this flexibility, one can easily customise the code to modify or define a new shape of a TET.




Figure 1: The architecture of TopicTracker.

The following summarises the workflow within the architecture:

- First, two kinds of input data must be provided: a *temporal topic profile* comprising of the profile of the underlying time-stamped topics, and a $N \times N$ *Topic Evolution Strength* (TES) matrix, where N denotes the number of the topics in the profile.

- Second, TopicTracker infers the TET of the topics in the profile with a user-specified parameter, `min_tes` (Phase 1), the minimum TES among the topics to find their possible ancestors. The topics with TESs less than `min_tes` will not be considered in finding their ancestors. Given a topic v in the TET, its pathway to the most ancestral topic indicates the trajectory of v . Thus, TET is a backbone reflecting evolutionary pathways of topics.
- Third, TopicTracker visualises the inferred TET with the five central modules (Phase 2) using three parameters: `min_reborn` is the minimum time period of topics which has been elapsed between the moment of their emergence until their evolved topics appear; `min_dead` is the minimum time period of topics that goes into unobserved; and `min_tes`.

An example TET is presented in Fig. 2, where each node denotes a topic. The evolutionary




Figure 2: An example TET generated by TopicTracker. The parameters are set as: `min_reborn` = 2 years, `min_dead` = 1 year, and `min_tes` = 0.2. The explanation about this example is provided in Section 3.

relationship between two topics is denoted by a directed edge, and its TES is represented by a different edge color. The evolution state of each node is marked with a different node color. The y-axis shows the importance of topics normalised in $[0,1]$. We emphasize that the integration of *three information facets* in TET provides insight about topic trajectories: the evolutionary pathways of topics with their TESs over time, evolution states of topics over time, and topic

importance. Below we elaborate the construction process of a TET.

2.2. Software Functionalities

We now present the detailed descriptions about the input data and the two phases embodied in TopicTracker.

2.2.1. Input data format

The two kinds of input data² must be given to run TopicTracker: a temporal topic profile \mathcal{P} , and a TES matrix \mathbf{M} . \mathcal{P} contains the descriptions about the target topics whose trajectories will be generated. An example of \mathcal{P} is given in Table 3, and it has the following fields for each topic v :

- **id** is the unique identifier of v .
- **index** is the unique integer index of v (starting from 0).
- **weight** is the weight of v given each year.
- **year** (in format yyyy) is the year that v was generated.
- **words** are the top- N words representing v . The information about **weight**, **year**, and **words** can be generated by a topic model.³

The \mathbf{M} matrix is a $N \times N$ matrix, where N is the number of the topics presented in \mathcal{P} , where:

- The i -th row and j -th column of \mathbf{M} represents the TES of the i -th row topic (old) towards the j -th column topic (new).
- A TES only exists between a pair of two *non-contemporary* topics as we only estimate the TES between topics on different time slots. Thus, we set TESs between *contemporary topics* as 0. Also, by default, we set the diagonal entries to be 1.
- All the topics are sorted in ascending order according to their time slots. Thus, the first is the most ancient topic and the last the most recent topic.

²Discussion of generating this input data is beyond the scope of this paper.

³For example, a dynamic topic model [18] can generate such information. Discussion about generating such information is beyond the scope of this paper.

- \mathbf{M} is a non-symmetric matrix, where all entries below the main diagonal do not hold any values, as we are only interested in the calculation of the TES between two topics x and y , if only if $time(x) < time(y)$, where $time$ is the function returning the time slot of a given topic.
- All entries in \mathbf{M} are normalised in $[0,1]$, where the higher the more important.

An example of \mathbf{M} is given in Table 4. To estimate TESs in \mathbf{M} , the concept of similarity has been widely used in most related works [9, 5, 11, 16, 1, 3]. The fundamental is to formulate a similarity measure between x and y using the aggregated similarities between x 's top- k descriptive terms and y 's top- k descriptive terms (often $K=10$) [19, 3].

2.2.2. Creating TET layout

We construct a TET from the given \mathcal{P} and \mathbf{M} . The goal of constructing a TET is to identify the most likely genealogy of topics over a given time period. Technically, this problem is formulated as finding an optimum branching in a directed tree, where each direct edge connects a topic and the topics evolved from it. Let $TET = (V, E)$ be a directed, weighted tree, where $V = \{v_1, \dots, v_n\}$ is the set of nodes that correspond to the n topics in \mathcal{P} . These topics are time-stamped meaning that these are collected according to the k different time slots, where k is the carnality of the years observed in \mathcal{P} . $E \subseteq V \times V$ represents the set of directed edges that reflect evolutionary relationships in V , such that $(v_i, v_j) \in E$ if only if $time(v_i) < time(v_j)$, where the $time$ function specifies the time slot of a given topic. An edge (v_i, v_j) is an ordered pair of two nodes v_i and v_j , and is interpreted as there is a directed dependency from an ancestry topic v_i to a descendent topic v_j . Each edge is associated with its TES, drawn from \mathbf{M} . A TES reflects how strong each ancestry-descendent relationship is between two topics.

The TET algorithm is implemented in the function `buildTES()` in `topicTracker.R`, and has the following bases:

1. We assume that there is the dummy *root* node in a TET to merely make a tree. Thus, the most ancient topic for each topic are connected to the root node.
2. Ancestry topics always precede their descendant topics in time.
3. A topic v can have no ancestor, meaning that v is newly emerged. On this occasion, v is connected to the root node.

4. Each topic v can have multiple ancestors as more than one topics in the past can be assembled together into the emergence of v in a later time.
5. In a TET, evolutionary transitivity relationships are logically inferred by navigating the edges between past and new topics. Suppose that there are two evolutionary relationships: an older topic ‘**wireless technology**’ influences on the generation of a recent topic ‘**mobile devices**’; and a topic ‘**mobile devices**’ influences on the generation of a more recent topic ‘**Android phone**’. Then, we infer the relationship that ‘**wireless technology**’ can also influence on the generation of ‘**Android phone**’.
6. A complexity of a TET is determined by the number of edges connected in the TET. This number is determined by `min_tes`. That is, we only include all edges whose TESs are equal to and greater than `min_tes`. Also, in a TET, cycles are not observed as ancestries cannot go back in time.
7. An evolutionary pathway of a topic v is defined as a sequence of connected topics from v to the root node. Given the same evolutionary pathway, among all possible parents $\text{Par}(v)$ of v , some are more likely than others. A parent is the immediate ancestor of a topic. This likelihood is estimated based on the TESs, from \mathbf{M} , between v and $\text{Par}(v)$. We choose the parent with the highest TES with v . To illustrate, in Fig. 3 (a), given a topic **F**, there are three possible evolutionary paths: (1) $\{(B, F)\}$, (2) $\{(B, C), (C, F)\}$, and (3) $\{(B, D), (D, F)\}$. Given the paths (1) and (2), there are two possible parents for **F**, **B** and **C**, as these exist on the same pathway. In this case, we choose **C** as its TES with **F** is higher than that of **B**. Thus, we do not connect **F** to **B**. The TESs are indicated by different edge colors as presented in the legend. In the same manner, given the paths (1) and (3), we choose **D** not **B** as the parent of **F**. The result is given in Fig. 3 (b).

2.2.3. Topic trajectory visualiser

Topic trajectory visualiser is implemented in the function `visualiseTET()` in `topicTracker.R`. The function has the following modules:

- *Create TET layout*: We create a TET using `layout_as_tree()` (the Reingold-Tilford graph layout algorithm) in the R `igraph` package. We also define the positions of topics by their weights given by `weight` in \mathcal{P} and time information given by `year` in \mathcal{P} within the layout.
- *Determine evolution states of TET nodes*: We determine the evolution states of topics in the TET, and mark them using different colors (see below).




Figure 3: An example of the trajectories.

- *Define R plotting function:* We define the `plot()` function: the properties of nodes and edges within the defined `layout_as_tree()` to visualise them in the TET.
- *Adjust overlapped labels of TET nodes:* The overlapped node labels in the TET are separated to avoid their overlaps using `thigmophobe.labels` in the package `plotrix`.
- *Define x-axis and y-axis labels and visualize TET:* We finally define the ticks of x-axis and y-axis and draw their labels in the TET.

TopicTracker can identify the trajectories of the underlying topics, detecting their evolution states: (a) when and how a new topic is born? (b) how a topic can influence on the generation of newer topics? (c) what is a topic continually flourishing? (d) how a new topic can emerge implanted by what past topics? (e) what are the dead topics no longer observed? and (f) what are reborn topics distinguishable again in later time? To address them, we measure the evolution states of underlying topics using the TESs given in **M**:

- *born:* Topics born are the topics newly emerging. These are the topics emerging without any ancestors.
- *split:* Topics split are the topics split into more than two topics in the next generations.
- *fused:* Topics fused are the topics whose emergence has been made by more than two topics in the previous generations.

- *reborn*: Topics reborn are the topics that re-emerge after a user-specific time period, `min_reborn`. This state indicates that the topic reborn has attention after `min_reborn`, while during `min_reborn` it had been unnoticed.
- *dead*: Topics dead are the topics that go into unobserved (being unpopular) during `min_dead`. During `min_dead`, if a topic rarely influences the generation of topics in the next generation, we regard it as dead topics.
- *flourishing*: Topics flourishing are the topics continually and actively being used or influencing on the generation of some topics in the next generations.

Note also that each topic can be marked by two evolutionary states: its emergence reason from the past generation, and its influence on the next generation. The former state is called ‘emerging-state’ as it captures how it emerges. The latter is called ‘evolving-state’ as it reveals its evolving state for topics in the next generations. The evolution states are integrated with the TET, and this unified knowledge identifies the trajectories of the underlying topics. Finally, the underlying topics are positioned considering their weights, given by the `weight` field in \mathcal{P} , on the y-axis (see also Fig. 1). This helps us to identify which topics more important than the others.

3. Illustrative examples

We present illustrative examples of `TopicTracker` which can also be reproduced in the provided Github URL. Assume that the following temporal topic profile \mathcal{P} in Table 3 and the TES matrix \mathbf{M} in Table 4 are given to `TopicTracker`. In \mathcal{P} , suppose that the `words` field contains the top-4 words for each topic.

As seen in Tables 3 - 4, there are 11 topics whose time slots are from 2001 to 2005. Given \mathcal{P} and \mathbf{M} , `TopicTracker` creates the TES shown in Fig. 2. We now discuss what insight we observe from this TET:

- The label of each topic comes from the `label` string in Table 3. If the user wants to display `id` and/or `words` as a label, the user can modify the code in `visualiseTES()`⁴.
- The topics at the same time slot are aligned together by the y-axis. The timeline is split by years 2001-2005. assuming that topics are captured and identified using a year.

⁴We have commented which block needs to be modified in the code

id	index	label	weight	year	words
t1	0	A	0.75	2001	['opinion', 'computer', 'lab', 'user', 'human']
t2	1	B	0.5	2001	['time', 'lab', 'opinion', 'human', 'computer']
t3	2	C	0.55	2002	['time', 'application', 'interface', 'user', 'computer']
t4	3	D	0.4	2002	['lab', 'user', 'abc', 'interface', 'application']
t5	4	E	0.1	2002	['application', 'interface', 'abc', 'opinion', 'computer']
t6	5	F	0.8	2003	['interface', 'computer', 'application', 'response', 'system']
t7	6	G	0.5	2003	['lab', 'abc', 'survey', 'opinion', 'time']
t8	7	H	0.7	2004	['machine', 'lab', 'system', 'response', 'human']
t9	8	I	0.5	2004	['opinion', 'interface', 'time', 'application', 'lab']
t10	9	J	0.65	2005	['opinion', 'human', 'user', 'survey', 'system']
t11	10	K	0.45	2005	['application', 'interface', 'system', 'survey', 'abc']

Table 3: An example temporal topic profile \mathcal{P} . See also `softwarex_example_topic_profile.csv` under the `data` directory of the Github URL.

	A	B	C	D	E	F	G	H	I	J	K
A	1	0	0.1	0.1	0.1	0.9	0.1	0.7	0.1	0.1	0.1
B		1	0.3	0.5	0.1	0.1	0.1	0.1	0.1	0.1	0.1
C			1	0	0	0.1	0.1	0.1	0.1	0.1	0.1
D				1	0	0.9	0.1	0.1	0.1	0.1	0.1
E					1	0.2	0.3	0.1	0.1	0.1	0.1
F						1	0	0.1	0.1	0.1	0.1
G							1	0.1	0.75	0.1	0.1
H								1	0	0.9	0.9
I									1	0.1	0.1
J										1	0
K											1

Table 4: An example TES \mathbf{M} . The labels A-K are provided for the illustration only. In the actual input TES, only a $N \times N$ matrix without the labels should be given, where N is the number of the topics in \mathcal{P} . See also `softwarex_example_tes_matrix.csv` under the `data` directory of the Github URL.

- The TESs are represented by the ‘Evolutionary strength’ legend. For example, A influences the generation of both F and H, evidenced by $A \rightarrow F$ and $A \rightarrow H$. The TES of A associated with F and H are assigned to the corresponding edges, respectively, using different colors. There is no meaning of the length of the edges.
- The evolutionary states of the topics are represented by different colors in the ‘Evolution

states' legend. Each topic has two evolution states: emerging-state and evolving-state. For example, the emerging-state of A is born as it newly emerges at 2001 without any ancestors. Its evolving-state is split as it influences the generation of both F and H as indicated by the directed edges. H's emerging-state is reborn, with $\text{min_reborn} = 2$ years, meaning that H had been unnoticed during min_reborn from 2001 to 2003, but noticed again by being influenced by A after min_reborn at 2004. Its evolving-state is split as it influences on the generation of both J and K.

Uncolored topics represent flourishing. Both emerging-states and evolving-states of topics D, G, I, J and K are flourishing. D and G are influencing the generation of other topics in the next generations: $D \rightarrow F$, $G \rightarrow I$. I, J and K are under the incubation period which is less than min_reborn from the latest observed time slot, 2005.

The emerging-state of C is flourishing as generated by the influence of B. Its evolving-state is **dead** as its influence had not been observed for the past 3 years from the latest time slot 2005, where the 3 years are greater than min_dead . The emerging-state of F is fused as generated by the co-influence of both A and D. Its evolving-state is dead with the same reason of that of C.

- The y-axis shows relative importance of topics. Referring to $B \rightarrow D$, we see a declining trend of B (the weight decreases from 0.5 to 0.4) during 2001 - 2002, but that topic is getting more popular during 2002 - 2003 as indicated by $D \rightarrow F$ (the weight increases from 0.4 to 0.8).

4. Impact

Through a TET, we can observe the integrated, precious information about topic trajectories: evolutionary pathways of dynamic topics with evolution strengths indicated by the directed edges with different edge colors, evolution states of the topics indicated by different node colors, and the topic importance indicated by topic weight on the y-axis. The goal of building a TET is to build the most likely genealogy of non-contemporary topics over time. This TET models topic trajectory information: the evolution pathways between non-contemporary topics based on their evolution strengths.

TopicTracker enlightens credible evolutionary relationships between non-contemporary topics. For example, in science and technology domains, TopicTracker could contribute to uncovering how technological or knowledge topics can change and influence the generation of newer topics

through a series of evolution events over time. In e-commerce domains, **TopicTracker** could also help to track the evolution of market-competitive product-related topics in a product market over time.

TopicTracker is designed to generate topic trajectory identification and visualisation with few parameters. It can run with two simple data formats as explained in Section 2: a temporal topic profile, and a TES matrix. For any models that can provide these formats, their ability to identify topic trajectories can be easily analysed and visualised by **TopicTracker**. We believe that **TopicTracker** could help the user to allow a greater focus on methodological developments of their evolution strength matrix between time-stamped topics, rather than their implementation for topic trajectory identification. To the best of our knowledge, **TopicTracker** is the first generation of a framework that can be used in broader communities interested in identifying and visualising topic trajectory information.

5. Conclusions

In this paper, we presented a platform, **TopicTracker**, that can identify and visualise topic trajectory information. **TopicTracker** is equipped with the capability of addressing the issues still remained in the information retrieval community: how to identify trajectories of evolving topics over time? how to represent evolution states of the underlying topics at a particular time slot? and how to visualise topic trajectory information? As the backbone of topic trajectory information, we presented that **TopicTracker** uses TET which aims to induce a most likely genealogy tree for evolving topics. We presented the formal definition of TET, its constituent elements, and detailed descriptions about how to construct a TET from two kinds of input data: a temporal topic profile and a TES matrix showing the evolution strengths among the topics in the profile. Another key strength of **TopicTracker** is its ability to visualise three facets of useful trajectory information in a TET together: the evolutionary pathways of dynamic topics with their inter-evolution strengths, their evolution states at a particular time, and their relative importance. We believe that **TopicTracker** can provide a platform for topic trajectory analysis that can promote ease-of-use.

6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by the LP170100416 project, funded by an Australian Research Council's Linkage grant. We thank Professor Beth Webster, Director of the Centre for Transformative Innovation at Swinburne University of Technology, for providing construct advice and support in developing TopicTracker. We also thank both Dr. Don Klinkenberg and Dr. Thibaut Jombart who are the creators of the 'phybreak' and 'seqtrack' packages, respectively, for inspiring the foundation of TopicTracker. Special thanks to Dr. Don Klinkenberg for providing valuable suggestions for implementing TopicTracker.

References

- [1] W. S. Lee, S. Y. Sohn, Identifying Emerging Trends of Financial Business Method Patents, *Sustainability* 9 (9) (2017) 1–21.
URL <https://ideas.repec.org/a/gam/jsusta/v9y2017i9p1670-d112597.html>
- [2] H. Zhang, G. Kim, E. P. Xing, Dynamic topic modeling for monitoring market competition from online text and image data, in: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '15*, Association for Computing Machinery, New York, NY, USA, 2015, p. 1425–1434. doi:10.1145/2783258.2783293.
URL <https://doi.org/10.1145/2783258.2783293>
- [3] D. Greene, J. P. Cross, Exploring the political agenda of the european parliament using a dynamic topic modeling approach., *CoRR* abs/1607.03055 (2016).
URL <http://dblp.uni-trier.de/db/journals/corr/corr1607.html#GreeneC16>
- [4] J. Song, Y. Huang, X. Qi, Y. Li, F. Li, K. Fu, T. Huang, Discovering hierarchical topic evolution in time-stamped documents, *Journal of the Association for Information Science and Technology* 67 (4) (2016) 915–927. arXiv:<https://asistdl.onlinelibrary.wiley.com/doi/pdf/10.1002/asi.23439>, doi:10.1002/asi.23439.
URL <https://asistdl.onlinelibrary.wiley.com/doi/abs/10.1002/asi.23439>
- [5] H.-k. Zhou, H.-m. Yu, R. Hu, Topic discovery and evolution in scientific literature based on content and citations, *Frontiers of Information Technology & Electronic Engineering* 18 (10) (2017) 1511–1524. doi:10.1631/fitee.1601125.
URL <https://app.dimensions.ai/details/publication/pub.1099692062>

- [6] M. Song, G. E. Heo, S. Y. Kim, Analyzing topic evolution in bioinformatics: investigation of dynamics of the field with conference data in dblp, *Scientometrics* 101 (1) (2014) 397–428. doi:10.1007/s11192-014-1246-2.
URL <https://doi.org/10.1007/s11192-014-1246-2>
- [7] J. Yoon, K. Kim, Identifying rapidly evolving technological trends for r&d planning using sao-based semantic patent networks, *Scientometrics* 88 (1) (2011) 213–228. doi:10.1007/s11192-011-0383-0.
URL <https://doi.org/10.1007/s11192-011-0383-0>
- [8] Y. Zhang, G. Zhang, D. Zhu, J. Lu, Scientific evolutionary pathways: Identifying and visualizing relationships for scientific topics, *Journal of the Association for Information Science and Technology* 68 (8) (2017) 1925–1939. arXiv:<https://asistdl.onlinelibrary.wiley.com/doi/pdf/10.1002/asi.23814>, doi:10.1002/asi.23814.
URL <https://asistdl.onlinelibrary.wiley.com/doi/abs/10.1002/asi.23814>
- [9] Q. He, B. Chen, J. Pei, B. Qiu, P. Mitra, L. Giles, Detecting topic evolution in scientific literature: How can citations help?, in: *Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM '09*, Association for Computing Machinery, New York, NY, USA, 2009, p. 957–966. doi:10.1145/1645953.1646076.
URL <https://doi.org/10.1145/1645953.1646076>
- [10] Y. Jo, J. E. Hopcroft, C. Lagoze, The web of topics: Discovering the topology of topic evolution in a corpus, in: *Proceedings of the 20th International Conference on World Wide Web, WWW '11*, Association for Computing Machinery, New York, NY, USA, 2011, p. 257–266. doi:10.1145/1963405.1963444.
URL <https://doi.org/10.1145/1963405.1963444>
- [11] S. Jung, W. C. Yoon, An alternative topic model based on common interest authors for topic evolution analysis, *Journal of Informetrics* 14 (3) (2020) 101040. doi:<https://doi.org/10.1016/j.joi.2020.101040>.
URL <http://www.sciencedirect.com/science/article/pii/S1751157719303517>
- [12] S. Zhong, B. Verspagen, The role of technological trajectories in catching-up-based development: An application to energy efficiency technologies, MERIT Working Papers 013, United Nations University - Maastricht Economic and Social Research Institute on Innovation and

Technology (MERIT) (Mar. 2016).

URL <https://ideas.repec.org/p/unm/unumer/2016013.html>

- [13] H. Park, C. L. Magee, Tracing technological development trajectories: A genetic knowledge persistence-based main path approach, *PLOS ONE* 12 (1) (2017) 1–18. doi:10.1371/journal.pone.0170895.
URL <https://doi.org/10.1371/journal.pone.0170895>
- [14] G. Triulzi, J. Alstott, C. L. Magee, Estimating technology performance improvement rates by mining patent data, *Technological Forecasting and Social Change* 158 (2020) 120100. doi:<https://doi.org/10.1016/j.techfore.2020.120100>.
URL <http://www.sciencedirect.com/science/article/pii/S0040162520309264>
- [15] Y. Huang, F. Zhu, A. L. Porter, Y. Zhang, D. Zhu, Y. Guo, Exploring technology evolution pathways to facilitate technology management: From a technology life cycle perspective, *IEEE Transactions on Engineering Management* (2020) 1–13.
- [16] Z. Qiu, Z. Wang, Technology forecasting based on semantic and citation analysis of patents: A case of robotics domain, *IEEE Transactions on Engineering Management* (2020) 1–21.
- [17] W. Gaul, D. Vincent, Evaluation of the evolution of relationships between topics over time, *Advances in Data Analysis and Classification* 11 (1) (2017) 159–178. doi:10.1007/s11634-016-0241-2.
URL <https://doi.org/10.1007/s11634-016-0241-2>
- [18] D. M. Blei, J. D. Lafferty, Dynamic topic models, in: *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, Association for Computing Machinery, New York, NY, USA, 2006, p. 113–120. doi:10.1145/1143844.1143859.
URL <https://doi.org/10.1145/1143844.1143859>
- [19] D. O’Callaghan, D. Greene, J. Carthy, P. Cunningham, An analysis of the coherence of descriptors in topic modeling, *Expert Systems with Applications* 42 (13) (2015) 5645 – 5657.