

## RESEARCH ARTICLE

# Microblog topic identification using Linked Open Data

Ahmet Yıldırım<sup>1</sup>\*, Suzan Uskudarli<sup>1</sup>

Complex Systems Research Lab, Department of Computer Engineering, Bogazici University, Istanbul, Turkey

\* [ahmet.yil@boun.edu.tr](mailto:ahmet.yil@boun.edu.tr)

## Abstract

Much valuable information is embedded in social media posts (microposts) which are contributed by a great variety of persons about subjects that of interest to others. The automated utilization of this information is challenging due to the overwhelming quantity of posts and the distributed nature of the information related to subjects across several posts. Numerous approaches have been proposed to detect topics from collections of microposts, where the topics are represented by lists of terms such as words, phrases, or word embeddings. Such topics are used in tasks like classification and recommendations. The interpretation of topics is considered a separate task in such methods, albeit they are becoming increasingly human-interpretable. This work proposes an approach for identifying machine-interpretable topics of collective interest. We define topics as a set of related elements that are associated by having posted in the same contexts. To represent topics, we introduce an ontology specified according to the W3C recommended standards. The elements of the topics are identified via linking entities to resources published on Linked Open Data (LOD). Such representation enables processing topics to provide insights that go beyond what is explicitly expressed in the microposts. The feasibility of the proposed approach is examined by generating topics from more than one million tweets collected from Twitter during various events. The utility of these topics is demonstrated with a variety of topic-related tasks along with a comparison of the effort required to perform the same tasks with words-list-based representations. Manual evaluation of randomly selected 36 sets of topics yielded 81.0% and 93.3% for the precision and F1 scores respectively.

## OPEN ACCESS

**Citation:** Yıldırım A, Uskudarli S (2020) Microblog topic identification using Linked Open Data. PLoS ONE 15(8): e0236863. <https://doi.org/10.1371/journal.pone.0236863>

**Editor:** Syed Ahmad Chan Bukhari, Yale University, UNITED STATES

**Received:** January 26, 2019

**Accepted:** July 8, 2020

**Published:** August 11, 2020

**Copyright:** © 2020 Yıldırım, Uskudarli. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** The keywords used to query the Twitter API are provided in the manuscript. The authors abided by Twitter's, Wikipedia's, DBpedia's, and Wikidata's Privacy Policy and Terms of Use and Service. With respect to the restricted data stemming from the TOC of Twitter's data sharing policy, we have provided all tweet IDs, which can be fetched by researchers for purposes of replicability. This manner of sharing abides by the TOC of Twitter API use. Additionally, all relevant data related to this work are uploaded to Figshare and publicly available via the following DOI and URL:

## Introduction

Microblogging systems are widely used for sharing short messages (microposts) with online audiences. They are designed to support the creation of posts with minimal effort, which has resulted in a vast stream of posts relating to issues of current relevance such as politics, product releases, entertainment, sports, conferences, and natural disasters. Twitter [1], the most popular microblogging platform, reports that over 500 million tweets are posted per day [2]. Such systems have become invaluable resources for learning what people are interested in and how they respond to events. However, making sense of such large volumes of posts is far from

10.6084/m9.figshare.7527476 [https://figshare.com/articles/S-BounTI\\_Semantic\\_Topic\\_Identification\\_approach\\_from\\_Microblog\\_post\\_sets/7527476](https://figshare.com/articles/S-BounTI_Semantic_Topic_Identification_approach_from_Microblog_post_sets/7527476). The implementation of this work, including its source code, is also uploaded to Figshare and publicly available via the following DOI and URL: 10.6084/m9.figshare.5943211 [https://figshare.com/articles/S-BounTI\\_Semantic\\_Topic\\_Identification\\_approach\\_from\\_Microblog\\_post\\_sets\\_An\\_application/5943211](https://figshare.com/articles/S-BounTI_Semantic_Topic_Identification_approach_from_Microblog_post_sets_An_application/5943211).

**Funding:** AY is supported by the Turkish Ministry of Development under the TAM Project, number DPT2007K120610. This work is researched and developed in the SosLab at Boğaziçi University, Istanbul, Turkey.

**Competing interests:** The authors have declared that no competing interests exist.

trivial, since posts tend to be limited in context (due to their short length), informal, untidy, noisy, and cryptic [3]. Furthermore, content related to the same topic is typically distributed over many contributions posted by numerous users.

Various approaches have been developed to gain insight into the topics that emerge on microposts. Some of the most popular topic detection approaches are based on latent Dirichlet allocation (LDA) [4–6], latent semantic analysis (LSA) [7, 8], and non-negative matrix factorization (NMF) [9, 10]. These methods capture topical keywords from post sets to represent topics that can be utilized in classification, recommendation, and information retrieval tasks. Topics are represented with bags-of-words, along with weights indicating the strength of their association with the microposts. Alternative approaches are based on the change in the frequency of terms of interest (i.e., words and hashtags) [11–16] to capture trending topics. More recently, word-embeddings have been used to capture both the semantic and the syntactic features of words to improve the relevancy of the words representing the topics [5, 6, 17].

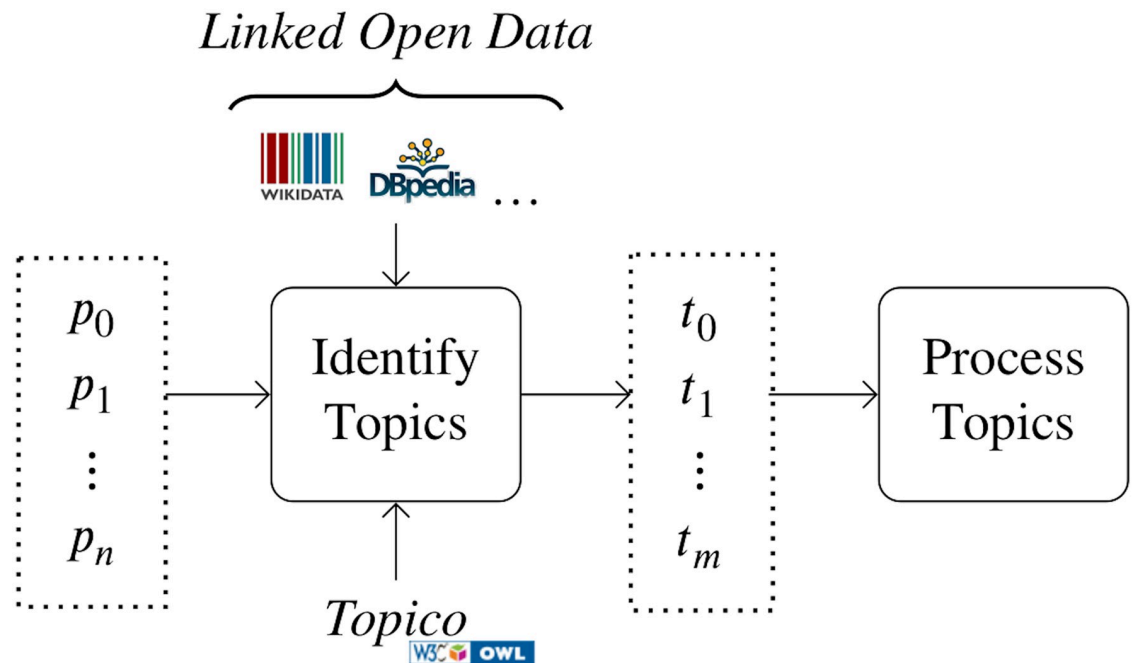
The determination of topics related to people or groups in social media facilitates content-specific user recommendation as opposed to the more familiar friend of friend recommendations obtained from follower networks [18, 19]. For this purpose, approaches that process the content and the user behavior to recommend users have been proposed that extracts content [18] or social network analysis of co-occurring content [19].

Natural language processing (NLP) techniques are utilized to yield more human-readable topics. Sharifi *et al.* propose a reinforcement-based method on consecutively co-occurring terms to summarize collections of microposts [20]. BOUN-TI [21] also produces human-readable topics using cosine similarity among collections of tweets and Wikipedia articles. The up-to-date nature of Wikipedia pages successfully captures topics of current relevance with highly readable titles. While such topics are easily human-interpretable, they are less suitable for automated processing.

Some approaches identify topics within single posts by linking meaningful fragments within them to external resources such as Wikipedia pages [22–25]. In the context of microposts, the detection of topics for single posts is not very effective due to their limited context. This work focuses on topics that have gained traction within the crowd by aggregating contributions relevant to topics from numerous posts.

This work proposes an approach, S-BOUN-TI (Semantic-Boğaziçi University-Topic Identification) that produces machine-interpretable (actionable) semantic topics from collections of microposts (Fig 1). A topic is considered to be a collection of related *topic elements*. The *topic elements* are Linked Open Data (LOD) [26] resources that are linked to fragments of posts. LOD is an ever-growing resource (1,255 datasets with 16,174 links as of May 2020) of structured linked data regarding a wide range of human interests made accessible with the semantic Web standards. Such linking enables capturing the meaning of the content that is expressed in alternative manners such as the terms “FBI”, “feds” and “Federal Bureau of Investigation” all link to the web resource [http://dbpedia.org/resource/Federal\\_Bureau\\_of\\_Investigation](http://dbpedia.org/resource/Federal_Bureau_of_Investigation), and “guns n’ roses”, and, “gunsn’roses” to the web resource [http://dbpedia.org/resource/Guns\\_N'\\_Roses](http://dbpedia.org/resource/Guns_N'_Roses) in DBpedia. These resources are further linked to other resources through their properties. For example, the resource of the *FBI* provides information about *law enforcement*, *its director*, and the resource of *Guns N’ Roses* provides information about their *music genre* and *group members*. The linking of fragments within posts to such resources greatly expands the information at our disposal to make sense of microposts. For example, the topics that are recently talked about regarding *law enforcement* or *rock concerts* could be easily retrieved. To represent topics, an ontology (*Topico*) is specified under the W3C semantic Web [27, 28] standards.

The main goal of this work is to explore the feasibility of linking informal conversational content in microposts to semantic resources in LOD to produce relevant machine-interpretable



**Fig 1. Overview of identifying semantic topics from a set of microposts.** Entities within microposts ( $p_i$ ) are linked to semantic entities in *Linked Open Data*, which are processed to yield semantic topics ( $t_j$ ) expressed with the *Topico* ontology.

<https://doi.org/10.1371/journal.pone.0236863.g001>

topics. First, the potential elements of topics are determined by processing linked entities in a collection of posts. Then, the elements are assigned to topics by processing a co-occurrence graph of the entities. Finally, the topics are created by processing the elements and representing them with the *Topico* ontology. To the authors' knowledge, this is the first approach that utilizes semantic Web and LOD to identify topics within collections of microposts.

To assess the viability of the proposed approach, we developed a prototype to generate topics for 11 collections of tweets gathered during various events. The utility of these topics (totaling 5248) is demonstrated with a variety of topic-related tasks and a comparison of the effort required to perform the same tasks with words-list-based (wLB) representations. An evaluation of randomly selected 36 sets of topics yielded 81.0% and 93.3% for the precision and  $F_1$  scores respectively.

The main contributions of this work are:

- an approach for identifying semantic topics from collections of microposts using LOD,
- the *Topico* ontology to represent semantic topics,
- an analysis of semantic topics generated from 11 datasets with over one million tweets, and
- a detailed evaluation of the utility of semantic topics through tasks of various complexities.

To enable the reproducibility of our work and to support the research community, we contribute the following:

- a prototype to generate semantic topics [29].
- the semantic topics generated from the datasets and the identifiers of the tweets of 11 datasets [30, 31],

- a demonstration endpoint for performing semantic queries over the generated topics (<http://soslab.cmpe.boun.edu.tr/sbounti>), and
- the manual relevancy-annotations of S-BOUN-TI topics from 36 sets corresponding to approximately 5760 tweets [31].

The remainder of this paper is organized as follows: The *Related work* section provides an overview of topic identification approaches. The key concepts and resources utilized in our work are presented in the *Background* section. The proposed approach is described in the *Approach to identifying semantic topics* section. An analysis of the topics generated from various datasets and their utility is detailed in the *Experiments and results* section. Our observations related to the proposed approach and the resulting topics are presented along with future directions in the *Discussion and future work* section. Finally, in the *Conclusions* section, we remark on our overall takeaways from this work.

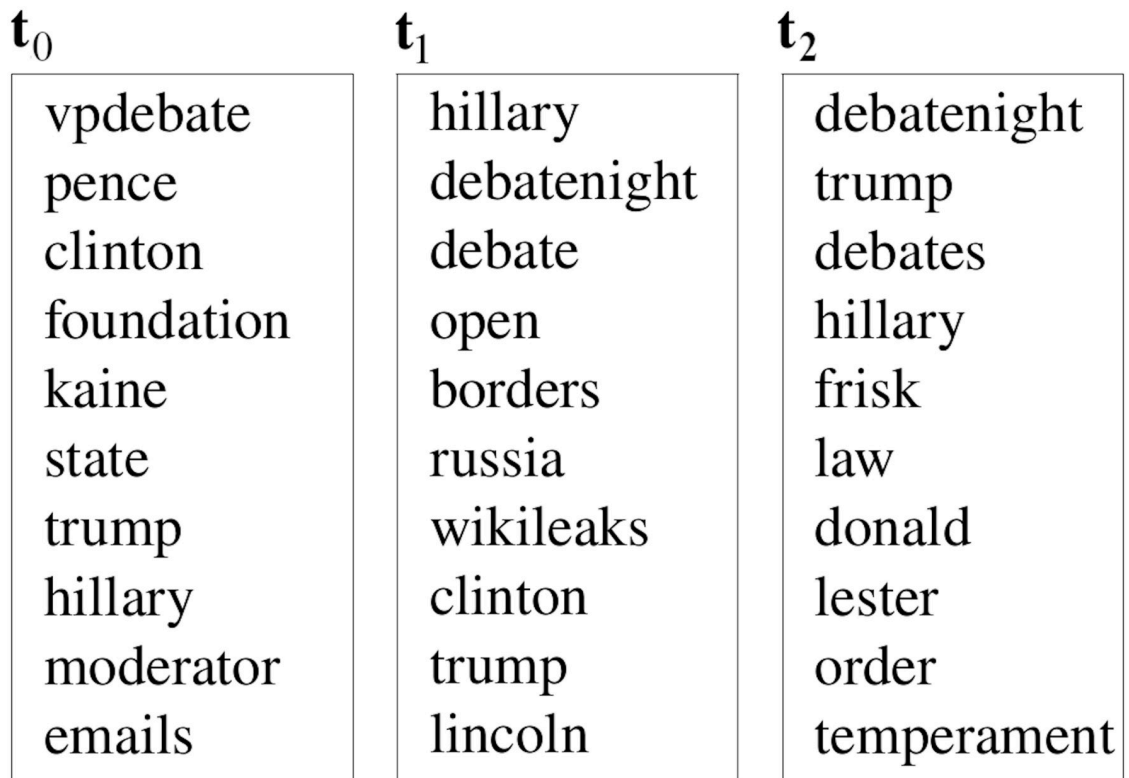
## Related work

The approaches to making sense of text can be characterized in terms of their input (i.e., sets of short, long, structured, semi-structured text), their processing methods, the utilized resources (i.e., Wikipedia, DBpedia), and how the results are represented (i.e., summaries, words, hashtags, word-embeddings, topics).

Various statistical topic models, such as latent semantic analysis (LSA) [7, 8], non-negative matrix factorization (NMF) [9, 10], and latent Dirichlet allocation (LDA) [32], aim to discover topics in collections of documents. They represent documents with topics and topics with words. The topics are derived from a term-document matrix from which a document-topic and a topic-term matrix are produced. LSA and NMF methods achieve this with matrix factorization techniques. LDA learns these matrices with a generative probabilistic approach that assumes that documents are represented with a mixture of a fixed number of topics, where each topic is characterized by a distribution over words. The determination of the predefined number of topics can be difficult and is typically predetermined based on domain knowledge or experimentation. The sparseness of the term-document matrix stemming from the shortness of the posts presents challenges to these approaches [33–35]. The topics produced by these approaches are represented as lists of words and will be referred as words-lists-based approaches (WLB). The interpretation of the topics is considered a separate task.

LDA has been widely utilized for detecting topics in microposts. Some of these approaches associate a single topic to an individual post [36, 37], whereas others consider a document to be a collection of short posts that are aggregated by some criteria like the same author [38], temporal or geographical proximity [39], or content similarity that indicates some relevance (i.e., hashtag or keyword) [40]. Fig 2 shows the top ten words of some LDA topics resulting from tweets that we collected during the 2016 U.S. presidential debates (produced by *TwitterLDA* [41]). Some approaches use the co-occurrence of words in the posts to capture meaningful phrases and use these bi-terms in the generative process [4, 42, 43]. The determination of the predefined number of topics for large collections of tweets that are contributed by numerous people can be quite difficult.

More recently, word-embeddings learned from the Twitter public stream and other corpora (i.e., Wikipedia articles) have been used to improve LDA based [5, 6] and NMF based [44] topics. In some cases, word-embeddings are used to enhance the posts with semantically relevant words [17, 45]. Another utility of word-embeddings is in assessing the coherence of topics by determining the semantic similarity of their terms [46].



**Fig 2. The top 10 words of topics generated (using TWITTER-LDA) from a collection of tweets gathered during the 2016 U.S. presidential debate.**

<https://doi.org/10.1371/journal.pone.0236863.g002>

Alternatively, in micropost topic identification, some approaches consider topics as a set of similar posts, such as those based on the fluctuation in the frequency of the terms of interest (i.e., words and hashtags) [11, 12]. The evolution of such trending topics can be traced using the co-occurrences of terms. Marcus *et al.* [14] created an end-user system that presents frequently used words and representative tweets that occur during peak posting activity which they consider as indicators of relevant events. Petrović *et al.* [47] use term frequency and inverse document frequency (*tf-idf*) to determine similar posts and select the first post (temporally) to represent topics. An alternative similarity measure is utilized by Genc *et al.* [48] who compute the semantic distances between posts based on the distances of their linked entities in Wikipedia's link graph. In these approaches, the interpretation of what a topic represents is also considered a separate task.

Sharifi *et al.* [20] produce human-readable topics in the form of a summary phrase that is constructed from common consecutive words within a set of posts. BOUN-TI [21] represents topics as a list of Wikipedia page titles (that are designed to be human-readable) which are most similar (cosine similarity of *tf-idf* values) to a set of posts. While these topics are human-comprehensible, they are less suitable for automated processing.

The approaches mentioned thus far have been domain-independent, however, in some cases domain-specific topics may be of interest. In the health domain, Prieto *et al.* explore posts related to specific sicknesses to track outbreaks and epidemics of diseases [49] by matching tweets with illness-related terms that are manually curated. Similarly, Parker *et al.* utilize sickness-related terms which are automatically extracted from Wikipedia [50]. Eissa *et al.* [51] extract topic related words from resources such as DBpedia and WordNet to identify topics

related to user profiles. As opposed to machine learning-based approaches that generate topics, these approaches map a collection of posts to pre-defined topics.

Entity linking approaches [52] have been proposed to identify meaningful fragments in short posts and link them to external resources such as Wikipedia pages or DBpedia resources [22–25]. These approaches identify topics related to single posts. Such approaches may not adequately capture topics of general interest since they miss contextual information present in crowd-sourced content.

Semantic Web technologies [53] are frequently utilized to interpret documents with unique concepts that are machine-interpretable and interlinked with other data [54–56]. For example, events and news documents have been semantically annotated with ontologically represented elements (time, location, and persons) to obtain machine-interpretable data [57–59]. Parliamentary texts are semantically annotated with concepts from DBpedia and Wikipedia using look-up rules and entity linking approaches [60, 61]. Biomedical references have been normalized with an unsupervised method that utilizes the OntoBiotope ontology [62] that models microorganism habitats and word-embeddings [63]. In this manner, they can map text like *children less than 2 years of age* to the concept *pediatric patient* (OBT:002307) that bears no syntactic similarity.

Our work utilizes semantic Web technologies to identify topics from domain-independent collections of microposts and to express them. Like many of the other approaches, we aggregate numerous posts. The ontology specification language OWL [64] is used to specify *Topico* to represent topics. The elements of topics are identified via entity linking using LOD resources. The collective information gathered from sets of posts is utilized in conjunction with the information within LOD resources to improve the topic elements. The elements of topics are related based on having co-occurred in several posts. In other words, numerous posters have related these elements by posting them together. This can result in topics that may seem peculiar, such as the *FBI* and the U.S. presidential candidate *Hillary Clinton*, which became a hot subject on Twitter as a result of public reaction. A co-occurrence graph is processed to determine the individual topics. The topic elements are the URIs of web resources that correspond to fragments of posts. Various fragments may be associated with the same resource since our approach aims to capture the meaning (i.e., “FBI”, “feds” and “Federal Bureau of Investigation”) to [http://dbpedia.org/resource/Federal\\_Bureau\\_of\\_Investigation](http://dbpedia.org/resource/Federal_Bureau_of_Investigation)). Semantically represented topics offer vast opportunities for processing since short unstructured posts are mapped to ontologically represented topics consisting of elements within a rich network of similarly represented information.

## Background

This section describes the basic concepts and tools related to semantic Web and ontologies, entity linking, and Linked Open Data that are used in this work.

Ontology is an explicit specification that formally represents a domain of knowledge [28, 65]. It defines inter-related concepts in a domain. The concepts are defined as a hierarchy of *classes* that are related through *properties*. Ontologies often refer to definitions of concepts and properties in other ontologies, which is important for reusability and interoperability. Resource Description Framework Schema (RDFS) definitions are used to define the structure of the data. Web ontology language (OWL) [64] is used to define semantic relationships between concepts and the data. Ontology definitions and the data expressed with ontologies are published on the Web and referred to with their unified resource identifiers (URI). To easily refer to the ontologies and data resources, the beginning of URIs are represented with namespace prefixes. For example *dbp:* is often used to refer to <http://dbpedia.org/resource/>. A specific

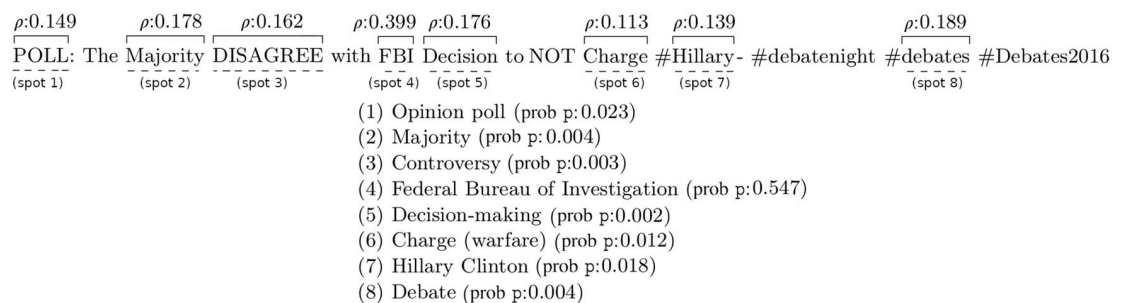


entity is referred to with its namespace prefix and the rest of the URI following it such as *dbr:Federal\_Bureau\_of\_Investigation* for [http://dbpedia.org/resource/Federal\\_Bureau\\_of\\_Investigation](http://dbpedia.org/resource/Federal_Bureau_of_Investigation) (which is the definition of FBI in DBpedia).

We use OWL language to define the *Topico* ontology to express microblog topics. Other ontologies that *Topico* refers to are DBpedia to express encyclopedic concepts, (FOAF) [66] to express agents (with emphasis on people), w3c basic Geo vocabulary [67, 68] and Geonames [69] to express geolocations, *Schema.org* [70, 71] to express persons and locations, and w3c time [72] to express intervals of topics. The namespace prefixes that are referred to in this paper are given in S1 Table.

*Entity linking* is used to identify fragments within text documents (surface forms or spots) and link them to external resources that represent real-world entities (i.e., dictionaries and/or encyclopedias such as Wikipedia) [52]. Entity linking for microposts is challenging due to the use of unstructured and untidy language as well as the limited context of short texts. We use TagMe [24] for this purpose as it offers a fast and well-documented application programming interface (API) [73]. TagMe links text to Wikipedia articles to represent entities. This is suitable for our purposes since the articles are cross-domain and up-to-date. Given a short text, TagMe returns a set of linked entities corresponding to its spots. For each result, TagMe provides a goodness value ( $\rho$ ) and a probability (p) which are used to select those that are desirable. The chosen entities are treated as candidate topic elements. Fig 3 illustrates the response of TagMe for a short text. Here, the spot *FBI* is linked to [https://en.wikipedia.org/wiki/Federal\\_Bureau\\_of\\_Investigation](https://en.wikipedia.org/wiki/Federal_Bureau_of_Investigation) with  $\rho = 0.399$  and  $p = 0.547$ . We use  $\rho$  and p to determine the viability of a topic element.

Linked Data [26, 74] specifies the best practices for creating linked knowledge resources. Linked Open Data (LOD) refers to the data published using Linked Data principles under an open license. It is an up-to-date collection of interrelated web resources that spans all domains of human interests, such as music, sports, news, and life sciences [75, 76]. LOD contains 1,255 datasets with 16,174 links among them (as of May 2020) [77]. With its rich set of resources, LOD is suitable for representing the elements of topics, such as [http://dbpedia.org/resource/Federal\\_Bureau\\_of\\_Investigation](http://dbpedia.org/resource/Federal_Bureau_of_Investigation) to represent “FBI”. Among the most widely used data resources in LOD are DBpedia [78] with more than 5.5 million articles derived from Wikipedia (as of September 2018 [79]) and Wikidata [80, 81] with more than 87 million items (as of June 2020 [82]). DBpedia is a good resource for identifying entities such as known persons, places, and events that often occur in microposts. For example, in the short text: “*POLL: The Majority DISAGREE with FBI Decision to NOT Charge #Hillary—#debatenight #debates #Debates2016*” the entity linking task identifies the spots *Hillary* and *FBI* that are linked to *dbr:Hillary\_Clinton* and *dbr:Federal\_Bureau\_of\_Investigation* respectively. Both Wikidata and DBpedia support



**Fig 3. Entity linking results from TagMe for a short text.**

<https://doi.org/10.1371/journal.pone.0236863.g003>

semantic queries by providing SPARQL endpoints [83, 84]. SPARQL [85] (the recursive acronym for SPARQL Protocol and RDF Query Language) is a query language recommended by W3C for extracting and manipulating information stored in the Resource Description Framework (RDF) format. It utilizes graph-matching techniques to match a query pattern against data. SPARQL supports networked queries over web resources which are identified with URIs [86]. This work uses LOD resources in SPARQL queries to demonstrate the utility of the proposed approach and to identify if topic elements are persons or locations.

## Approach to identifying semantic topics

This work focuses on two main aspects related to extracting topics from collections of microposts: their identification and their representation. More specifically, the determination of whether LOD is suitable for capturing information from microposts and if semantically represented topics offer the expected benefits. The key tasks associated with our approach are (1) identifying the elements of topics from collections of microposts, (2) determining which elements belong to which topics, and (3) semantically representing the topics. This section presents a topic identification approach and describes its prototype implementation which is used for evaluation and validation purposes. First, we describe the ontology developed to represent topics since it models the domain of interests and, thus, clarifies the context of our approach. Then, we present a method for identifying topics from micropost collections, which will be represented using this ontology. While describing this method, aspects relevant to the prototype implementation are introduced in context. Finally, various implementation details are provided at the end of this section.

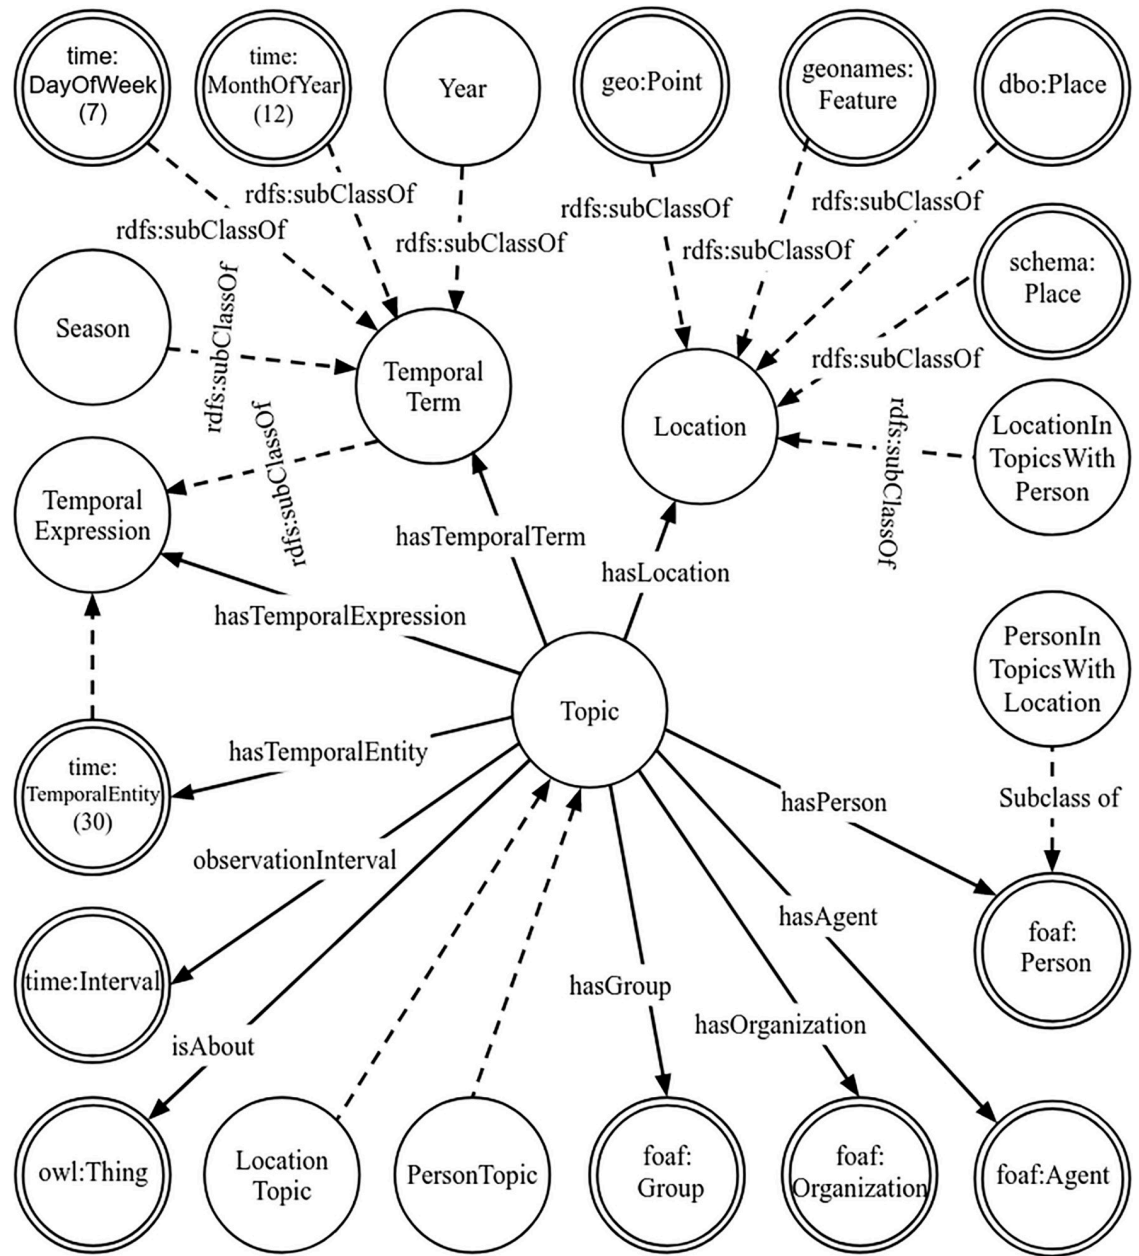
## Topico ontology for representing topics

In the context of this work, a topic is considered to be a set of elements that are related when numerous people post about them in the same context (post). Here, we focus on the elementary aspects of the topics that are most common in social media. For this purpose, we define an ontology called *Topico* that is specified with the Web Ontology Language (OWL) [64] using Protégé [87] according to the Ontology 101 development process [88]. The main classes and object relations of the ontology (*Topico*) are shown in Fig 4. Further object properties are shown in S1 Appendix. This section describes some of the design decisions and characteristics relevant to *Topico*. In doing so, the references to the guidelines recommended for specifying ontologies are shown in italic font. The prefix *topico* is used to refer to the *Topico* namespace.

The first consideration is to *determine the domain and scope of the ontology*. Representing topics that emerge from collections of microblog posts is at the core of our domain. Thus, the ontology must reflect the concepts (classes) and properties (relations) common to microblogs. It aims to serve as a basic ontology to represent general topics that could be extended for domain-specific cases if desired. The simplicity is deliberate to create a baseline for an initial study and to avoid premature detailed design.

To *enumerate the important terms in the ontology*, we inspected a large volume of tweets. We observed the presence of well-known people, locations, temporal expressions across all domains since people seem to be interested in the “who, where, and when” aspects of topics. What the topic is about varies greatly, as one would expect. As a result we decided to focus on the agents (persons or organizations), locations, temporal references, related issues, and meta-information of topics. Based on this examination, *a definition of classes and a class hierarchy* was developed. The main class of *Topico* is *topico:Topic* which is the domain of the object properties *topico:hasAgent*, *topico:hasLocation*, and *topico:hasTemporalExpression* which relate a topic to people/organizations, locations, and temporal expressions. To include all other kinds





**Fig 4. The fundamental classes and object properties of *Topico*.** The classes defined in other ontologies are shown with double circles and labeled with the prefixes of the namespaces such as FOAF. *rdfs:subClassOf* relationships are represented with dashed lines.

<https://doi.org/10.1371/journal.pone.0236863.g004>

of topic elements we introduce the *topico:isAbout* property (i.e., *Topic1 topico:isAbout dbr:Abortion*). We defined several temporal terms as instances of the *topico:TemporalExpression* class. Also, since the subjects of conversation change rapidly in microblogs, the required property *topico:observationInterval* is defined which corresponds to the time interval corresponding a collection (timestamps of the earliest and latest posts). This information enables tracking how topics emerge and change over time, which is specifically interesting for event-based topics like political debates and news. A topic may be related to zero or more of elements of each type. Topics with no elements would indicate that no topics of collective interest were identified. Such information may be of interest to those tracking the topics in microblogs. In

```

<owl:NamedIndividual rdf:about="http://soslab.cmpe.boun.edu.tr/sbounti/
  topics_algorithmV0.5.owl#2016_first_presidential_debate_50_52_topic_23">
  <rdf:type rdf:resource="http://soslab.cmpe.boun.edu.tr/ontologies/topico.owl#Topic"/>
  <foaf:maker rdf:resource="http://soslab.cmpe.boun.edu.tr/sbounti/topics_algorithmV0.5.owl#algorithmV0.5"/>
  <rdfs:label xml:lang="en">2016 First Presidential Debate 50-52 minutes topic: 23</rdfs:label>
  <topicCreatedAt rdf:datatype="http://www.w3.org/2001/XMLSchema#dateTimeStamp">
    2017-06-26T13:05:37Z</topicCreatedAt>
  <observationInterval rdf:resource="http://soslab.cmpe.boun.edu.tr/sbounti/
    topics_algorithmV0.5.owl#interval_2016-09-27T01-50-00Z-2016-09-27T01-51-59Z"/>
  <hasPerson rdf:resource="http://dbpedia.org/resource/Lester_Holt"/>
  <hasPerson rdf:resource="http://dbpedia.org/resource/Donald_Trump"/>
  <isAbout rdf:resource="http://dbpedia.org/resource/Racial_profiling"/>
  <isAbout rdf:resource="http://dbpedia.org/resource/Terry_stop"/>
  <hasTemporalTerm rdf:resource="http://soslab.cmpe.boun.edu.tr/sbounti/
    topics_algorithmV0.5.owl#year_2016"/>
</owl:NamedIndividual>

```

**Fig 5. A semantic topic extracted from 50-52nd minutes of the first presidential debate of U.S. 2016 elections.** This topic is related to Lester Holt (Journalist) and Donald Trump (presidential candidate) regarding racial profiling and terry stopping (stop and frisk) in the U.S. in 2016. Automatic enumeration gave the topic number 23 to this topic.

<https://doi.org/10.1371/journal.pone.0236863.g005>

our prototype, however, an approach that yields topics with at least two elements was implemented since we were interested in the elements of the topics.

With respect to the *consider reusing existing ontologies* principle, we utilize the classes and properties of existing ontologies whenever possible such as w3c OWL-Time ontology [72], FOAF, Schema.org, and Geonames. FOAF is used for agents and persons. The classes *schema:Place*, *dbo:Place*, *geonames:Feature*, and *geo:Point* are defined as subclasses of *topico:Location*. Temporal expression of w3c OWL-Time ontology [72] are grouped under *topico:TemporalExpression*. The temporal expressions of interest which were not found are specified in *Topico*.

A topic related to the first 2016 U.S. presidential debate (27 September 2016) is shown in Fig 5. This topic is related to the 2016 U.S. presidential candidate Donald Trump, the journalist Lester Holt (*topico:hasPerson*), racial profiling, and terry stopping (*topico:isAbout*) in the United States (*topico:hasLocation*) in 2016 (*topico:hasTemporalTerm*). The subject of racial profiling and terry stopping (stop and frisking mostly of African American men) frequently emerged during the election.

Further information about *Topico* may be found at [89] and the ontology itself is published at <http://soslab.cmpe.boun.edu.tr/ontologies/topico.owl>.

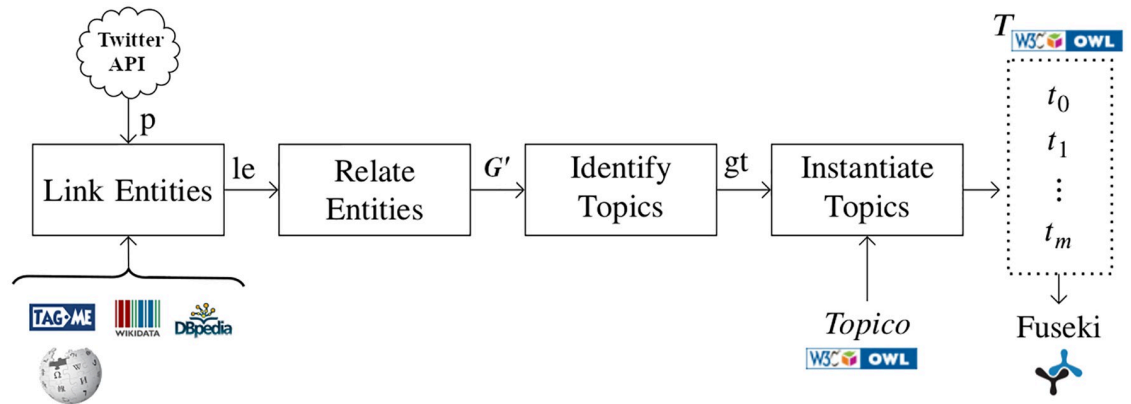
## Identifying topics

The task of topic identification consists of identifying significant elements within posts and determining which of them belong to the same topic. An overview for identifying topics is shown in Fig 6, which takes a set of microposts and results in a set of topics represented with *Topico* (S-BOUN-TI topics). Semantic topics are stored in RDF repositories to facilitate processing. Algorithm 1 summarizes the process of generating semantically represented topics given a collection of microposts. It has three phases: the determination of candidate topic elements, topic identification, and topic representation.

### Algorithm 1 Topic extraction from microposts

1: Input: $P$	▷ micropost set
2: Output: $T$ : OWLDocument	▷ semantic topics
3: $unlinkedSpots \leftarrow []$	▷ unlinked spots
4: $elements \leftarrow []$	▷ candidate elements
5: $types \leftarrow []$	▷ types of elements
6: $le \leftarrow []$	▷ linked entities
7: $G, G'$ : graph	▷ initial and pruned graphs
8: $gt$ : {}	▷ set of sub-graphs
9: $sTopic$ : OWL	▷ S-BOUN-TI topic

*Phase 1—Identify Candidate Elements*



**Fig 6. An overview of identifying semantic topics from collections of tweets.** An overview of identifying semantic topics from collections of tweets where  $p$  is the set of tweets;  $le$  is the set of linked entities (candidate elements);  $G$  is the co-occurrence graph of candidate topic elements;  $gt$  is the set of sub-graphs of  $G$  whose elements belong to the same topic; and  $T$  is the set of semantic topics represented in OWL. *TagMe*, *Wikidata*, *DBpedia*, and *Wikipedia* are external resources used during entity linking. *Topico* is the ontology we specified to express semantic topics. All topics are hosted on a Fuseki SPARQL endpoint.

<https://doi.org/10.1371/journal.pone.0236863.g006>

```

10: for each  $p$  in  $P$  do
11:    $elements[p] \leftarrow entities(p) \cup temporalExpressions(p)$ 
12:    $unlinkedSpots[p] \leftarrow unlinkedSpots(p)$ 
13: end for
14: for each  $p$  in  $P$  do
15:    $le[p] \leftarrow reLink(elements[p], elements)$ 
16:   # Add the newly linked entities
17:    $le[p] \leftarrow le[p] \cup linkSpots(unlinkedSpots[p], elements,$ 
     $unlinkedSpots)$ 
18: end for
    Phase 2—Identify topics
19:  $G \leftarrow relate(le)$  ▷ construct co-occurrence graph
20:  $G' \leftarrow prune(G, \tau_e)$  ▷ prune the graph using  $\tau_e$ 
21:  $gt \leftarrow subgraphsOfRelatedNodes(G', \tau_{sc})$  ▷ sub-graphs of related
    nodes of  $G'$ 
22:  $observationInterval \leftarrow getObservationInterval(P)$  ▷ timestamps
    of the earliest and latest posts
23: for each  $v$  in  $G'$  do
24:    $types[v] \leftarrow getType(v, P, \tau_{loc})$  ▷ determine the type of
    element ( $v$ )
25: end for
    Phase 3—Represent topics with Topico
26: for each topic in  $gt$  do
27:    $sTopic \leftarrow sem-topic(topic, types, observationInterval)$  ▷ rep-
    represent as s-BOUN-TI topic
28:    $T.add(sTopic)$  ▷ Add to topics for collection  $P$ 
29: end for
30: return  $T$ 

```

The first phase determines the candidate topic elements that are extracted from each post (Lines 10-13).  $P$  is a set of posts. The function  $entities(p)$  returns the entities within a post  $p$ . We denote a post and its corresponding entities as  $\langle p, l \rangle$  where  $l$  are linked entities. Determining the candidate elements entails the use of an entity linker that links elements of microposts to external resources and a rule-based temporal term linker. We defined temporal term linking rules [31] to detect frequently occurring terms like the days of the week, months, years, seasons, and relative temporal expressions (i.e., *tomorrow*, *now*, and *tonight*) to handle the various

ways in which they are expressed in social media. We denote linked entities as [spots]  $\mapsto$  [URI], where all the spots that are linked to an entity are shown as a list of lowercase terms and the entities are shown as URIS. For example, [north dakota, n. dakota]  $\mapsto$  [dbr:North\_Dakota] indicates the two spots *north dakota* and *n. dakota* that are linked to *dbr:North\_Dakota* where some posts refer to the state North Dakota with its full name and others have abbreviated the word north as “n.”. The entity linking process may yield alternatively linked spots or unlinked spots. For example, the spot “Clinton” may be linked to any of *dbr:Hillary\_Clinton*, *dbr:Bill\_Clinton*, *dbr:Clinton\_Foundation* or not at all (unlinked spot). Such candidate topic elements may be improved by examining the use of patterns within the collective information for agreement among various posters. Thus, the linked entities retrieved from all the posts are used to improve the candidate elements attempting to link previously unlinked spots or altering the linking of a previously linked spot (Lines 14-18). In our prototype, entities are retrieved using TagMe which links spots to Wikipedia pages. We map these entities to DBpedia resources which are suitable for the semantic utilization goals of our approach (see the [Background](#) section). At the end of this phase, any remaining unlinked spots are eliminated, yielding the final set of candidate topic elements.

The second phase decides which elements belong to which topics. We consider the limited size of microposts to be significant when relating elements since the user chose to refer to them in the same post. The more often a co-occurrence is encountered the more significant that relation is considered since the aim is to capture what is of collective interest. In this work, the term co-occurring elements/entities is defined to be the co-occurrence of the spots within a post to which these elements are linked.

To identify topics, we construct a co-occurrence graph of the candidate topic elements.

Let  $LEP = \{ \langle p, l \rangle \mid p \in P \wedge l = entities(p) \}$ . Let the co-occurrence graph  $G = (V, E)$  where  $V = \bigcup_{p \in P} entities(p)$  and  $E = \{ (v_i, v_j) \mid v_i, v_j \in V \wedge v_i \neq v_j \wedge \{v_i, v_j\} \subseteq l \wedge \langle p, l \rangle \in LEP \}$ .

Let  $w: E \rightarrow \mathbb{R}_{[0,1]}$  be a function that returns the weight of an edge and is defined as:

$$w(e) = \frac{|\{ p \mid e = (v_i, v_j) \in E \wedge v_i \neq v_j \wedge \{v_i, v_j\} \subseteq l \wedge \langle p, l \rangle \in LEP \}|}{|P|} \tag{1}$$

[Fig 7](#) shows an example co-occurrence graph constructed from four micropost texts. The linked entities obtained from these posts are *dbr:Donald\_Trump*, *dbr:Lester\_Holt*, *dbr:Social\_Profiling*, *dbr:Terry\_Stop*, *dbr:Constitutionality*. Within four posts, the co-occurrence between some of these entities ranged between 0.25 to 0.75. Thus, we have extracted a significantly rich set of information from the posts in terms of relating them to web resources which themselves are related to other resrouces via data and object properties.

To represent collective topics (those of interest to many people) the weak elements are eliminated prior to identifying the topics (Line 20). The weak edges ( $w(e) < \tau_e$ ) are removed. All vertices that become disconnected due to edge removal are also removed. The following equations describe how  $G = (V, E)$  is pruned to obtain the final co-occurrence graph  $G' = (V', E')$ :


$$E' = \{ e \mid e \in E \wedge w(e) \geq \tau_e \} \tag{2}$$

$$V' = \{ v \mid \exists x [ (x, v) \in E' \vee (v, x) \in E' ] \} \tag{3}$$

The co-occurrence graph  $G'$  represents all related topic elements (Line 20). [S2 Fig](#) shows a co-occurrence graph obtained at this step.

Post text	Entities
Lester reminds Trump that Stop-and-Frisk was ruled unconstitutional. #debatenight	<i>dbr:Lester_Holt, dbr:Donald_Trump</i> <i>dbr:Terry_stop, dbr:Constitutionality</i>
Lester Holt got the receipts on this stop and frisk s.t. #debatenight	<i>dbr:Lester_Holt, dbr:Terry_stop</i>
Add stop and frisk to the mix of racial profiling ...that's a great idea #Debates.	<i>dbr:Terry_stop, dbr:Racial_profiling</i>
Donald Trump tells Lester Holt he is wrong on #stopandfrisk in regards to it being a form of racial profiling	<i>dbr:Donald_Trump, dbr:Lester_Holt</i> <i>dbr:Terry_stop, dbr:Racial_profiling</i>

Co-occurrence Graph




**Fig 7. A sample co-occurrence graph (G) of topic elements.** Nodes are candidate topic elements and edge labels represent weights. The table at the top shows the linked entities. Here, the entities are DBpedia resources, shown using the namespace *dbr* is <http://dbpedia.org/resource> (i.e., *dbr:Terry\_stop* is [http://dbpedia.org/resource/Terry\\_stop](http://dbpedia.org/resource/Terry_stop)). The spots within the microposts are encircled with a box. The entities are shown in the second column. Note that alternative forms of terry stopping have been linked to the same entity (*dbr:Terry\_stop*). The co-occurrence graph captures the elements of topics within posts. Its nodes represent the entities and the edges represent the degree to which their corresponding spots co-occurred in the posts.

<https://doi.org/10.1371/journal.pone.0236863.g007>

$G'$  is processed to yield sets of related topic elements, each of which will represent a topic (Line 21). The criteria for determining the topics (sub-graphs of  $G'$ ) are (1) an element may belong to several topics since it may be related to many topics, (2) topics with more elements are preferable as they are likely to convey richer information, (3) topics with few elements are relevant if their relationships are strong (i.e., topics of intense public interest such as the death of a public figure).

The maximal cliques algorithm [90] is used to determine the sub-graphs, where for a graph  $G = (V, E)$ ,  $C \subseteq V$  is a clique  $\Leftrightarrow \forall_{v,v' \in C} (v, v') \in E, v \neq v'$ . Maximal cliques are sub-graphs that are not subsets of any larger clique. They ensure that all elements in a sub-graph are related to each other. An examination of the maximal cliques obtained from co-occurrence graphs revealed that many of them had a few (two or three) vertices. This is not surprising since it is unlikely that many elements become related through many posts. Another observation is that some elements (vertices) occur with a far lower frequency than the others. Since topics with few very weak elements are not likely to be of great interest, they are eliminated with the use of the  $\tau_{sc}$  threshold:  $\frac{freq(v)}{|P|} < \tau_{sc}$  where  $v \in V$  and  $freq(v) = |\{(p, l) | v \in l \wedge (p, l) \in LEP\}|$ . Fig 8 shows an example of how topics are identified from graphs.





**Fig 8. An example entity co-occurrence graph and its corresponding topics.** A co-occurrence graph, where  $v_i$  are candidate topic elements and  $w_{ij} = w(e_{ij})$ . The edge  $e_{ij}$  is eliminated for being weak ( $w_{13} < \tau_e$ ). Three maximal cliques (topics) emerge:  $\{v_0, v_1, v_2\}$ ,  $\{v_0, v_3, v_4, v_5\}$ , and  $\{v_0, v_3, v_4, v_5\}$ . If  $freq(v_0) < \tau_{kc}$  or  $freq(v_6) < \tau_{kc}$  then  $\{v_0, v_6\}$  will also be eliminated.

<https://doi.org/10.1371/journal.pone.0236863.g008>

After the elements of the topics are determined, additional information necessary to represent the final topics is obtained. The observation interval is computed using the posts with the earliest and the latest timestamps in  $P$  (Line 22). Since our *s-BOUN-TI* topics represent *persons*, *locations*, and *temporal expressions* the entity types of elements are resolved (Lines 23-25). The type of *temporal expressions* is determined while they are being extracted. The types of other elements are identified with semantic queries. For example, if the value of the *rdf:type* property includes *foaf:Person* or *dbo:Person* its type is considered to be a person. To determine if an entity is a location, first, the value of *rdf:type* is checked for a location indicator (*schema:Place*, *dbo:PopulatedPlace*, *dbo:Place*, *dbo:Location*, *dbo:Settlement*, *geo:SpatialThing* and *geonames:Feature*). Locations are quite challenging as they may be ambiguous and be used in many different manners. Then, the contexts of the spots corresponding to entities are inspected for *location indicators* (succeeding the prepositions *in*, *on*, or *at*) within the post-collection. Again, we employ a threshold ( $\tau_{loc}$ ) to eliminate weak elements of location type. Finally, an entity  $v$  is considered a location if  $\frac{|location-prepositions(v.LEP)|}{|P|} > \tau_{loc}$ . For example, the entity *FBI* in “*FBI reports to both the Attorney General and the Director of National Intelligence.*” is considered as an agent, whereas in “*I’m at FBI*” it is considered as a location.

At this stage, we have the sub-graphs of  $G'$  (topics) along with the types of topic elements. The final phase is to represent these topics with the *Topico* ontology. Each  $t \in T$  is mapped to an instance of *topico:Topic* (Lines 26-29). The topics are related with their elements in accordance in accordance with their types. For example, elements of type *person* are associated with



a topic with the *topico:hasPerson* property. The property *topico:isAbout* is used for all elements of type other than the types of *person*, *location*, and *temporal expression*. The observation intervals are associated with the *topico:observationInterval* property. The instantiated topics are referred to as *S-BOUN-TI* topics and are ready for semantic processing.

Note that other graph algorithms could be used to obtain topics. Also, alternative pre- and post-processing steps could be utilized. For example, it may be desirable to eliminate or merge some topics to yield better results. For illustration purposes, let's consider the consequence of using the maximal-cliques algorithm on pruned graphs. The maximal-cliques algorithm requires all of its elements to be related. The pruning of weak bonds introduces the potential of severing the relations necessary to be identified as a topic. In such cases, very similar topics may emerge, such as those that differ in only a single element. This conflicts with our desire to favor a variety of topics with higher numbers of elements. Not pruning the graph to prevent such cases would unreasonably increase the cost of computation since the original graphs are very large and consist of many weak relations. A post-processing step could be introduced to merge similar topics with the use of two thresholds:  $\tau_c$  for topic similarity and  $\tau_{e_{min}}$  for an absolute minimum edge relevancy weight. Let  $T_0$  be the set of cliques ( $T_0 \subseteq P(V)$ ). The set of merged cliques,  $T$ , is obtained by:

$$T = \begin{cases} (T_0 \setminus \{t_i, t_j\}) \cup \{t_k\} & \text{if } \text{jaccard}(t_i, t_j) > \tau_c \wedge \forall v_x, v_y [w((v_x, v_y)) > \tau_{e_{min}}] \\ T_0 & \text{otherwise} \end{cases} \quad (4)$$

where  $t_k = t_i \cup t_j$ ,  $t_i, t_j \in T_0$ ,  $v_x \in t_i$ ,  $v_y \in t_j$ . Higher values for  $\tau_c$  or  $\tau_{e_{min}}$  lead to more topics that are similar to one another.

### Prototype infrastructure

The services used to acquire external information are: the TagMe API for entity linking suggestions, the DBpedia and Wikidata for fetching semantic resources to be used as topic elements, and the Phirehose Library [91] for continuously fetching posts from the Twitter streaming API filter endpoint [92]. TagMe and Twitter have granted us access tokens to make API requests. DBpedia makes resources available under the Creative Commons Attribution-ShareAlike 3.0 License and the GNU Free Documentation License. Wikidata resources are under CC0 license which is equivalent to public domain access (both of which provide a public SPARQL endpoint). Our implementation has complied by all the terms and conditions related to the use of all services. The prototype was deployed on a virtual machine based on VMWare infrastructure running on Intel Xeon hardware with 2 GBS of RAM and Linux operating system (Ubuntu). The implementation of the maximal-cliques algorithm [90] is run within the R [93] statistical computation environment. In all of the processes mentioned above, we use a local temporary cache to reduce unnecessary API calls to reduce network traffic.

Finally, all topics are represented as instances of *topico:Topic*, serialized into OWL, and stored in a Fuseki [94] server (a SPARQL server with a built-in OWL reasoner) for further processing.

### Experiments and results

The main focus of this work is to examine the feasibility of using LOD resources to identify useful processable topics. Accordingly, our evaluation focuses on the examination of the characteristics and the utility of the generated *S-BOUN-TI* topics. We considered it important to generate topics from real data, which we gathered from Twitter. The quality and utility of the resulting topics are examined by:

- inspecting the characteristics of their elements (*Semantic topic characteristics* section),

- comparing the effort required to perform various tasks in comparison to topics generated by words-list-based (WLB) approaches (*The utility of semantic topics in comparison to WLB topics* section), and
- manually assessing their relevancy (*Topic relevancy assessment* section).

Furthermore, to gain insight into the similarity of topics with topics generated by other methods we compared them with human-readable (*Comparison with human-readable topics* section) and WLB topics (*Comparison with WLB topics* section).

## Datasets

For evaluation purposes S-BOUN-TI topics were generated from 11 datasets consisting of 1, 076, 657 tweets collected during significant events [31]. The Twitter Streaming API was used to fetch the tweets via queries which are summarized in Table 1. The first four sets were fetched during the 2016 U.S. election debates, which are significantly greater than the others. We expected that there would be a sufficient quantity of interesting tweets during the debates, which was indeed yielded plenty of diverse tweets (~48 tweets/second) resulting in very large datasets. The remainder of the datasets (except [PUB]) were collected during other notable events. These are focused due to a particular person such as Carrie Fisher or a concept such as *concert*. The debates related sets were collected for the duration of the televised debates and the remainder were collected until they reached at least 5000 posts. The [PUB] dataset was collected to inspect the viability of topics emerging from tweets arriving at the same time but without any query (*public stream*). Note that the Twitter API imposes rate limits on the number of tweets it returns during heavy use. Although they do not disclose their selection criteria, the tweets are considered to be a representative set.

**Table 1. The queries to fetch the datasets from Twitter and information about the collections.**

Dataset	Explanation	Twitter Query	Start time (UTC)	$\Delta t$ (m)
[pd <sub>1</sub> ]	2016 First presidential debate	election2016, 2016election, @HillaryClinton, @realDonaldTrump, #trump, #donaldrump, #trump Pence2016, hillary, hillaryclinton, hillarykaine, @timkaine, @mike_pence, #debates2016, #debatenight	2016-09-27T01:00:00Z	90
[pd <sub>2</sub> ]	2016 Second presidential debate	same as <i>pd<sub>1</sub></i>	2016-10-10T01:00:00Z	90
[pd <sub>3</sub> ]	2016 Third presidential debate	same as <i>pd<sub>1</sub></i>	2016-10-20T01:00:00Z	90
[vp]	2016 Vice presidential debate	keywords in <i>pd<sub>1</sub></i> , #vpdebate2016, #vpdebate	2016-10-05T01:00:00Z	90
[BA]	The divorce of Angelina Jolie and Brad Pitt	#Brangelina	2016-09-20T23:38:38Z	21
[CF]	The death of Carrie Fisher	Carrie Fisher	2016-12-28T13:59:50Z	15
[CO]	Tweets related to the keyword <i>concert</i>	concert	2016-12-02T19:00:00Z	60
[ND]	North Dakota demonstrations	north dakota	2016-12-03T06:59:48Z	14
[TB]	Toni Braxton became trending	Toni Braxton	2017-01-08T07:08:56Z	765
[IN]	Inauguration of President Trump	#inauguration, Trump, @realDonaldTrump	2017-01-21T20:41:44Z	6
[PUB]	A sample of public English tweets	(no keyword)	2016-12-02T20:29:53Z	8

<https://doi.org/10.1371/journal.pone.0236863.t001>

**Table 2. The datasets used to create S-BOUN-TI topics.**

Dataset	Posts (#)	Distinct-Poster	
		#	(%)
[pd <sub>1</sub> ]	259200	206827	79
[pd <sub>2</sub> ]	259203	187049	72
[pd <sub>3</sub> ]	258227	181436	70
[VP]	256174	135565	52
[BA]	5900	4777	79
[CF]	7932	6753	85
[CO]	5326	4743	89
[ND]	7466	6231	83
[TB]	5948	4506	75
[IN]	5809	5425	93
[PUB]	5472	5365	98

<https://doi.org/10.1371/journal.pone.0236863.t002>

Table 2 shows the number of posts and the ratios of distinct posters. The number of posts during the debates ( $pd_1, pd_2, pd_3$ , and [VP]) are fairly similar. For all datasets, the percentage of unique contributors is generally greater than 70%, which is desirable since our approach aims to capture topics from a collective perspective.

S-BOUN-TI topics are generated from collections of tweets. The debate datasets were segmented into sets of tweets posted within a time interval to capture the temporal nature of topics. Throughout the remainder of this paper, a collection of posts will be denoted as  $[ds_{id}][t_s, t_e]$  where  $ds_{id}$  is the name of the dataset, and  $t_s$  and  $t_e$  are the starting and ending times of a time interval. For example,  $pd_1 [10, 12)$  refers to the tweets in  $pd_1$  that were posted between the 10<sup>th</sup> to the 12<sup>th</sup> minutes of the 90-minute long debate. The earliest tweet is considered to be posted at the 0<sup>th</sup> minute, thus  $t_s = 0$  for the first collection of a dataset.

## Experiment setup

The first consideration is to determine the size of the collections. Streams of posts can be very temporally relevant, as is the case during events of high interest (i.e., natural disaster, the demise of a popular person, political debates). Furthermore, the subjects of conversation can vary quite rapidly. Short observation intervals are good at capturing temporally focused posts. Processing time is also significant in determining the size of the collections. When the rate of posts is high, the API returns approximately 5800 tweets per two minutes. Under the best of circumstances (when all required data is retrieved from a local cache), the processing time required for a collection of this size is approximately four minutes. Whenever API calls are required the processing time increases. We experimented with generating S-BOUN-TI topics with collections of different sizes and decided on limiting the size of collections to about 5000–8000 posts. This range resulted in meaningful topics with reasonable processing time. During heavy traffic, it corresponds to approximately 2–3 minutes of tweets, which is reasonable when topics tend to vary a lot.

As described earlier, our approach favors topics with a higher number of elements of significant strength. Table 3 shows the values of the thresholds we used to generate the topics where all values are normalized by the collection size.

The thresholds  $\tau_\rho$  and  $\tau_p$  are confidence values used to link entities as defined by the TagMe API, are set to the recommended default values. Higher values yield fewer candidate topic elements, thus fewer topics.

**Table 3. The values of thresholds for generating topics.**

	Value	Description
$\tau_p$	0.15	entity link confidence
$\tau_{\rho}$	0.35	spot confidence
$\tau_e$	0.001	weak edge pruning weight
$\tau_{sc}$	0.01	small clique removal
$\tau_c$	0.8	clique merge similarity
$\tau_{e_{min}}$	0.0005	minimum edge weight for clique merge
$\tau_{loc}$	0.01	weight of location entities with preposition

<https://doi.org/10.1371/journal.pone.0236863.t003>

Crowd-sourcing platforms typically exhibit long tails (a few items having relatively high frequencies and numerous items having low frequencies), which is also observed for entities we identified. For example, highly interconnected and dominant six entities in a co-occurrence graph extracted from  $pd_1$  are: *Debate*, *Donald\_Trump*, *Hillary\_Clinton*, *year:2016*, *Tonight*, and *Now* with weights of 0.12, 0.11, 0.11, 0.10, 0.07, and 0.03 respectively (maximum edge weight is 0.12). Similar distributions are observed in other collections.

All the thresholds are set heuristically based on experimentation. The threshold for eliminating weak edges ( $\tau_e$ ) is set to 0.001 based on the desire to capture entities with some agreement among posters. The threshold  $\tau_{sc}$  that is used to eliminate small cliques consisting of weak elements is set to 0.01. The threshold for clique similarity  $\tau_c$  is set to 0.8. Similar cliques were merged as a post-processing step (as explained in the *Approach to identifying semantic topics* section) where  $\tau_{e_{min}}$  is set to 0.0005 ( $\tau_c/2$ ). Finally,  $\tau_{loc} = 0.01$  to decide whether an entity is collectively used as a location. The cliques that remain after applying these thresholds are considered as collective topics.

To examine the impact of pruning applied before identifying topics, we traced the topic elements to the posts from which they were extracted. The percentage of posts that end up in the topics vary according to the dataset, with an average of 58% for vertices and 43% for edges (see [S2 Table](#)). Since the remaining vertices and edges are relatively strong, the resulting topics are considered to retain the essential information extracted from large sets of tweets.

## Semantic topic characteristics

[Table 4](#) summarizes the type of elements of the generated  $s$ -BOUN-TI topics. Most of them have persons, which is not surprising since tweeting about people is quite common. Topics with persons emerged regardless of whether the query used to gather the dataset included persons. Temporal expressions occurred more frequently in topics that were generated from datasets that correspond to events where time is more relevant (i.e., [CO]).

The viability of our method requires the ability to link posts to linked data, thus the linked entities must be examined. From our datasets, some of the spots and its corresponding topic element that we observe are: [donald, trump, donald trump, donald j. trump, donald j.trump]  $\rightarrow$  [*dbr:Donald\_Trump*], [stop and frisk, stopandfrisk, stop-and-frisk]  $\rightarrow$  [*dbr:Terry\_stop*], and [racial divide, racial profiling, racial profile, racial segment, racial violence]  $\rightarrow$  [*dbr:Racial\_profiling*]. As seen, the topic captures the intended meaning regardless of how it was expressed by the contributors.

Tweets can be very useful in tracking the impact of certain messaging since people tend to post what is on their mind very freely. Considering political campaigns, much effort is expended on deciding their talking points and how to deliver them. The ability to track the impact of such choices is important since it is not easily observable from the televised event or

Table 4. The frequencies of the types of topic elements.

Set	Topic		Person		Location		Temp.		isAbout	
	#	#	%	#	%	#	%	#	%	
[pd <sub>1</sub> ]	1221	1121	91	8	0.6	808	66	1129	92	
[pd <sub>2</sub> ]	1120	1068	95	32	2	559	49	1,010	90	
[pd <sub>3</sub> ]	1214	1130	93	11	0.9	265	21	1118	92	
[vp]	1511	1377	91	50	3	395	26	1380	91	
[BA]	9	6	66	0	0	7	77	7	77	
[CF]	35	34	97	0	0	18	51	27	77	
[CO]	31	7	22	2	6	19	61	29	93	
[ND]	43	5	11	40	93	11	25	43	100	
[TB]	46	46	100	0	0	1	2	43	93	
[IN]	18	18	100	8	44	9	50	17	94	

<https://doi.org/10.1371/journal.pone.0236863.t004>

its transcripts. To give an example, during the 86<sup>th</sup> minute of the 2016 U.S. vice presidential debate, the candidates were talking about abortion and its regulation. The topics of [vp] [86, 88] relate to the vice-presidential candidates Tim Kaine and Mike Pence and the subjects of law, faith, and religion. This shows that the audience resonated with the debate at that time. Counterexamples can be seen in some of the topics generated from the [IN] dataset, which were gathered with query terms the inauguration of Donald Trump as the U.S. President. It so happened that the *Women's March* event that was to take place the subsequent day was trending and was related to the inauguration through those who posted tweets during this time. As a result, the topics included people such as *Madonna* and *Michael Moore* who were very active in the Women's March. Besides, the locations London, France and Spain appeared in topics from tweets expressing support for the Women's March. As a result, the topics that were captured rather accurately reflected what was on the mind of the public during the inauguration.

Finally, we examined whether public streams ([PUB]) would yield topics, which we expected they would not since there would not be sufficient alignment among posts. Indeed no topics were generated, however, some entities were identified. We speculate that in public datasets collected during major events, such as earthquakes and terrorist attacks, the strength of ties could be strong enough to yield topics, although this must be verified.

### The utility of semantic topics in comparison to WLB topics

The main purpose of this work is to produce topics that lend themselves to semantic processing. To demonstrate the utility of S-BOUN-TI topics, we provide a comparative analysis with topics represented as lists of terms (WLB) in terms of the effort required to perform various topic related tasks. The effort is described with the use of the helper functions shown in Table 5.

Table 5. Helper functions for performing topic related tasks.

Description	Abbreviation
Entity identification	EI
Type resolution	TR
External resource utilization	EX
Time of contribution	TI
Rule definition	RD
Query optimization	QO
Semantic analysis	SA

<https://doi.org/10.1371/journal.pone.0236863.t005>

This section describes how various tasks are achieved with *S-BOUN-TI* topics and what would need to be done if *WLB* topics were used. For readability purposes, the queries are described in natural language. The *SPARQL* queries may be found in the supporting material.

**Task 1: Who occurs how many times in the topics related to Hillary Clinton:**

This is a simple task that can be achieved by querying the topics for persons who co-occur with Hillary Clinton (and counting them) as shown in Listing 1.

**Listing 1 Query: Who occurs how many times in the topics related to Hillary Clinton?**

```
SELECT ?person (COUNT(?topic) AS ?C)
WHERE {
  ?topic topico:hasPerson dbr:Hillary_Clinton;
  topico:hasPerson ?person.
  FILTER (?person NOT IN (dbr:Hillary_Clinton))
}
GROUP BY ?person
ORDER BY DESC(?C)
```

The first three results (out of 41) are:

Person	Count
<i>dbr:Donald_Trump</i>	"2205"^^xsd:integer
<i>dbr:Bill_Clinton</i>	"338"^^xsd:integer
<i>dbr:Tim_Kaine</i>	"314"^^xsd:integer

Other persons include Barack Obama, Ruth Bader Ginsburg, Al Gore, Ronald Reagan, Anderson Cooper, Abraham Lincoln, and George Washington. Since *S-BOUN-TI* topics represent persons (with *topico:hasPerson*) retrieving those who co-occur with Hillary Clinton is trivial.

To achieve the same results with *WLB* representations, it is necessary to determine the words that represent the person Hillary Clinton, other people, and count them. Thus, a type-resolution (*TR*) task to identify persons is needed which requires entity identification (*EI*) using an external resource (*EX*) with information about people.

**Task 2—When do the topics related to women’s issues occur?**

This type of query is useful to know when certain topics emerge and to track whether they trend, persist, or diminish within streaming content. The time of observation of subjects is significant since they tend to rapidly change on social media.

The query in Listing 2 identified 166 topics in 66 time intervals related to topics about *abortion*, *rape*, and *women’s health*. Retrieving the time intervals is straightforward, since the *topico:observationInterval* property captures this information. An inspection of the linked entities revealed that the posters used different terms related to these concepts, such as [rape, raped, rapist, rapists, raping, sexual violence, serial rapist] → [*dbr:Rape*].

**Listing 2 Query: When do the topics related to women’s issues occur?**

```
SELECT DISTINCT ?startTime ?endTime WHERE {
  ?topic topico:observationInterval ?interval.
  ?interval time:hasBeginning ?begin.
  ?interval time:hasEnd ?end.
  ?begin time:inXSDDateTime ?startTime.
  ?end time:inXSDDateTime ?endTime.
  {?topic topico:isAbout dbr:Rape.}
  UNION {?topic topico:isAbout dbr:Abortion.}
  UNION {?topic topico:isAbout dbr:Women\'s_health.}}
```

To achieve the same results with *WLB* representations, the terms related to the women’s issues and the time intervals (*TI*) corresponding to the time they appeared must be determined.

**Task 3—When do the top 50 issues related to topics including Hillary Clinton and/or Donald Trump appear?**





```

?topic topico:hasPerson ?person.
?topic topico:isAbout ?about.
?topic topico:observationInterval ?interval.
?interval time:hasBeginning ?intervalStart.
?intervalStart time:inXSDDateTime ?time.
FILTER(?person IN
  (dbr:Hillary_Clinton, dbr:Donald_Trump))}
GROUP BY ?time ?about ?person}
FILTER (?about=?issueOfInterest)}

```

To gain some insight regarding how the resulting issues corresponded with the actual debates we inspected them along with their transcripts [95–98]. For example, racism was an issue that was mostly discussed by the candidates during the second half of the first presidential debate and the first half of the third debate. The topics we identified also revealed that racism was mostly posted during the same time (see the rows labeled *White\_people* and *Black\_people* in Fig 9). Furthermore, an inspection of the tweets posted during the same time also included tweets related to racism that were posted in pro-Republican and pro-Democratic contexts. For the topics related to *Tax* and Donald Trump only ([vp][48, 50) and pd<sub>3</sub>[80, 82)) the corresponding tweets were indeed related to only Donald Trump. On the other hand, while the candidates were talking about ISIS, Iraq, and the position of the United States in the Middle East (pd<sub>3</sub>[68, 70)) the identified topics were related to illegal immigration and income tax. In this case, we observe a lack of resonance between what was transpiring during the debate and the topics of interest to the posters (who preferred to post about other matters).

This query demonstrates the use of the inverse relationships *topico:isAPersonOf* and *topico:inTopic*. These relationships are inferred through reasoning according to the definitions in *Topico* ontology (see descriptions of object relationships in Description Logic in S1 Appendix).

To achieve the same result with *wlb* topics, the terms indicating *Hillary Clinton*, *Donald Trump*, and the terms corresponding to the top 50 issues must be identified (EI) which requires reference to external resources (EX). Finally, when the issues emerged must be identified (TI).

#### Task 4—Which politicians occur in the topics?

This task requires determining the occupation of persons, which may be of interest to known people. While *s-BOUN-TI* topics include the persons, the DBpedia entities that represent them often do not include the *dbo:occupation* or worse are related to incorrect entities. However, Wikidata entities utilize *wdt:P106* (occupation) property with persons quite systematically. Since DBpedia refers to equivalent Wikidata entities with the *owl:sameAs* property, the occupation of a person can be retrieved from Wikidata. The query in Listing 4 fetches the politicians in topics extracted from the debates by: (1) fetching all persons in the *s-BOUN-TI* topics from our endpoint, (2) retrieving the Wikidata identifiers of these persons from the DBpedia endpoint, and (3) identifying the persons whose occupation (*wdt:P106*) is a *Politician* (*wd:Q82955*) from the Wikidata endpoint. Among the results are: *dbr:Abraham\_Lincoln*, *dbr:Bill\_Clinton*, *dbr:Colin\_Powell*, *dbr:Bernie\_Sanders*, and *dbr:Saddam\_Hussein*. Query optimization (QO) is performed to reduce the search space by prioritizing the sub-queries according to their expected response sizes. This example shows the benefits of using LOD in our topics, where the links within the entities lead to a multitude of options.

**Listing 4 Query: Which politicians occur in the topics?** This query performs three queries to the *s-BOUN-TI* DBpedia, and Wikidata endpoints. *wdt:P106* refers to the occupation property. *wd:Q82955* refers to the *politician* concept.

```

SELECT DISTINCT ?person WHERE {
  ?topic topico:hasPerson ?person}
SELECT ?DbPediaPerson ?wikidataPerson WHERE {

```

```

?DbPediaPerson owl:sameAs ?wikidataPerson.
FILTER (?DbPediaPerson IN
(<http://dbpedia.org/resource/Donald_Trump>,
<http://dbpedia.org/resource/Lester_Holt>,
...)).
FILTER regex(str(?wikidataPerson), "^.*wikidata\\.org.*$")
SELECT DISTINCT ?person WHERE {
?person wdt:P106 wd:Q82955.
FILTER (?person IN (wd:Q22686, wd:Q3236790,
...)) }

```

Performing this task with WLB topics requires the identification of persons who are politicians (EI, EX, TR).

#### Task 5—Which rock music band is performing where?

Social media is often used to get information about events. The query shown in Listing 5 fetches the bands and locations of rock music concerts. This requires fetching the names of the bands performing rock concerts and the locations of these concerts. The type of concert is determined with the use of DBpedia resources (EX). This query returns results like *dbr:Guns\_N\_Roses* and *dbr:Mexico\_City*. Country music concerts are found by appropriately revising the query (*dbc:Country\_music\_genres*), which returns results like *dbr:Luke\_Bryan* and *dbr:Nashville,\_Tennessee*. The locations of the concerts are retrieved from the tweets, whereas the genres of bands are determined from DBpedia. This query demonstrates the use of *topico:LocationTopic* which is a subclass of *topico:Topic*. The individuals of this type are determined through reasoning according to the definitions in *Topico* ontology.

#### Listing 5 Query: Which rock music band is performing where?

```

SELECT ?musicGroup ?location {
SERVICE <http://193.140.196.97:3030/topic/sparql>{
  SELECT ?topic ?musicGroup ?location WHERE {
    ?topic topico:isAbout dbr:Concert.
    ?topic a topico:LocationTopic.
    ?location topico:isLocationOf ?topic.
    {?topic topico:isAbout ?musicGroup.}
  UNION
  {?topic topico:hasPerson ?musicGroup.}}
SERVICE <http://dbpedia.org/sparql>{
  SELECT ?musicGroup2 WHERE {
    ?musicGroup2 a schema:MusicGroup.
    ?musicGroup2 dbo:genre ?musicGenre.
    ?musicGenre dct:subject dbc:Rock_music_genres }}
FILTER (?musicGroup = ?musicGroup2)}

```

To achieve this with WLB topics, the terms referring to the bands and locations (EI and TR) are needed. External resources (EX) are needed to identify bands, locations, and the genre of the bands. Also, as explained in the *Identifying topics* section, the context of the location terms must be examined to determine if they were indeed used as locations.

#### Task 6—Which of the issues related to Barack Obama during the 2012 and 2016 U.S. election debates are the same?

In politics, it is useful to know which issues persist over time. The query shown in Listing 6 queries the topics identified during the 2012 and the 2016 U.S. election debates. For this purpose, topics were generated from the tweets gathered during the 2012 U.S. Presidential debates [99]. The resulting elements are: *dbr:Debate*, *dbr:President\_of\_the\_United\_States*, *dbr:Debt*, *dbr:Question*, *dbr:Tax*, *dbr:Tax\_cut*, *dbr:Golf*, *dbr:Economy*, *dbr:Black\_people*, *dbr:Racism*, *dbr:Violence*, *dbr:Birth\_certificate*, *dbr:Lie*, *dbr:Muslim*, *dbr:Barack\_Obama\_presidential\_campaign\_2008*, *dbr:Russia*, *dbr:Iraq*, *dbr:Immigration*, *dbr:Blame*, and *dbr:*

*Central\_Intelligence\_Agency*. There are many issues one would expect to see in a presidential debate such as taxes, violence, and the economy. Among other issues that appear in both years are racism, immigration, Muslim, Iraq, and Russia. One might be surprised to see golf (*dbr: Golf*) in this list; alas, the amount of golf played by candidates seems to be a matter of public interest. An inspection of the tweets confirms that the amount of golf that Barack Obama played became a topic of discussion.

**Listing 6 Query: Which of the issues related to Barack Obama are the same in the 2012 and 2016 U.S. election debates?** This is a federated query that queries two endpoints, one for the debates in 2012 and one for the debates in 2016.

```
SELECT ?about1 {
SERVICE <http://193.140.196.97:3031/topic/sparql>{
  SELECT DISTINCT ?about1 WHERE {
    ?topic1 topico:isAbout ?about1.
    ?topic1 topico:hasPerson dbr:Barack_Obama}
SERVICE <http://193.140.196.97:3032/topic/sparql>{
  SELECT DISTINCT ?about2 WHERE {
    ?topic1 topico:isAbout ?about2.
    ?topic1 topico:hasPerson dbr:Barack_Obama}
FILTER(?about1=?about2)}
```

To obtain a similar result with *wlb* representations, words common to topics of 2012 and 2016 must be retrieved. The results would be terms rather than concepts. For conceptual results, entity identification (EI and EX) could be used.

**Task 7—Which religions and ethnicity were mentioned during the 2012 and 2016 debates?**

*Topico* explicitly represents only persons, location, and temporal elements. To detect other types, external resources must be used. The query shown in Listing 7 utilizes knowledge about religions and ethnicities in Wikidata to retrieve related topics in the 2012 and the 2016 U.S. elections' debate topics with Query 1 that retrieves all religions from the Wikidata endpoint and Query 2 that retrieves the topics that include any of the items fetched in Query 1. A program that optimizes this query by feeding the output of Query 1 to Query 2 is used for this task (OO). The same process is repeated for ethnic groups. The tweets themselves refer to specific religions or ethnicities (i.e., Christian and Mexican). This query enables retrieving information about religions and ethnicities independent of any specific instance.

**Listing 7 Query: Which religions were mentioned during the 2012 and 2016 debates?**

This query is issued using two queries. Query 1: Get the religions from Wikidata, where the property *P279\** means all subclasses and *Q9174* is the identifier for the religion class. Query 2: Get the topics that include religions.

```
PREFIX wdt: <http://www.wikidata.org/prop/direct/>
PREFIX wd: <http://www.wikidata.org/entity/>
SELECT DISTINCT ?item ?article
WHERE {
  ?item wdt:P279* wd:Q9174.
  ?article schema:about ?item.
FILTER (
  SUBSTR(str(?article),9,17)="en.wikipedia.org/") .}
SELECT ?about (COUNT(?about) as ?C)
WHERE {
  ?topic topico:isAbout ?about.
  FILTER (?about IN (
    dbr:Buddhism, dbr:Jainism,
    ...
    dbr:Tapa_Gaccha, dbr:Zen))}
GROUP BY ?about
ORDER BY DESC(?C)
```

The query for 2012 returned only *dbr:Catholicism*, whereas for 2016 it returned *dbr:Islam\_in\_the\_United\_States*, *dbr:Islam*, and *dbr:Sunni\_Islam*. A manual inspection of tweets confirms the difference in tweeting about religion. In 2012, Catholicism was a subject of concern related to abortion and in 2016 Islam became an issue in the context of the Iraq War and the 9/11 terrorist attacks.

The issues regarding ethnicity in 2012 were *dbr:African\_Americans*, *dbr:Russians*, *dbr:Egyptians*, *dbr:Jews*, *dbr:Mexican\_Americans*, *dbr:Arabs*, and *dbr:Israelis*. Ethnic references were also present during 2016, however with differing emphasis: *dbr:Russians*, *dbr:Hispanic*, *dbr:Asian\_Americans*, *dbr:Chinese\_Americans*, *dbr:Hispanic\_and\_Latino\_Americans*, *dbr:Mexican\_Americans*, and *dbr:Mexicans*.

Furthermore, we observed that the topic elements that co-occur with *dbr:African\_Americans* also varied. With the support and opposition to the *black lives matter* movement, the elements *dbr:Police* and *dbr:Racism* were observed in 2016.

To accomplish this task with WLB representations, the identification of religions and ethnic groups are needed (TR) that requires entity identification (EI) using an external resource (EX).

#### Task 8—Which people are related to the same topics?

Semantic representation enables inference from present information, such as introducing the *vcard:hasRelated* property that specifies relationships among people and organizations (see vCard ontology [100]). This property can be used to relate people that occur in the same topic using the Semantic Web Rule Language (SWRL) [101] by defining the rule (RD):

```
Topic(?topic) ^ hasPerson(?topic, ?person1) ^ hasPerson(?topic, ?
person2)
-> vcard:hasRelated(?person1, ?person2)
```

A reasoner can be used to relate all people who are in the same topic with the *vcard:hasRelated* relation, such as *dbr:Donald\_Trump*, *dbr:Hillary\_Clinton* and *dbr:Lester\_Holt*. Such rules support the introduction of subjective inquiries of interest. Also, software agents aware of the vCard ontology may reason about this information.

In the WLB case, persons in topics must be identified (TR), which requires entity identification (EI) using external resources (EX). There must be some way of expressing this relation so it can be referenced.

#### Task 9—What are the categories of topics?

Topic enrichment [102–104] through external resources (EX) may greatly enhance the utility of topics. A useful enrichment for S-BOUN-TI topics would be to relate them to their DBpedia subject categories through the *dct:subject* property. For example, the category of the *dbr:Job* is *dbc:Employment*, which indirectly relates all S-BOUN-TI topics having *dbr:Job* as a topic element to *dbc:Employment*. The following SWRL rule (RD) enriches S-BOUN-TI topics with *topico:isAbout* relations to the categories of their elements:

```
Topic(?topic) ^ isAbout(?topic, ?element) ^ dct:subject(?element, ?
category)
-> isAbout(?topic, ?category)
```

Thus, all topics related to the DBpedia category *dbc:Employment* can be fetched with:

```
SELECT ?topic WHERE {
  ?topic topico:isAbout dbc:Employment}
```

In this query, when the DBpedia category *dbc:Employment* is replaced with *dbc:Law\_enforcement\_operations\_in\_the\_United\_States* the results include topics with the element *dbr:Stop-and-frisk\_in\_New\_York\_City*. Therefore, with this simple rule definition, it becomes possible to relate the topics with their categories and query the topics according to these categories.

A similar enrichment for WLB topics requires external resources (EX) and functionality such as semantic analysis (SA) of topics that could require considerable programming.

**Table 6. Subtasks to perform *Task-1-Task-9* for WLB vs S-BOUN-TI topics.**

Task	WLB	S-BOUN-TI
1	TR, EX, EI	–
2	TI	–
3	EI, TI, EX	–
4	TR, EX, EI	EX (2), QO
5	EI(2), TR (2), EX	EX
6	–	–
7	TR, EX, EI	EX, QO
8	TR, EX, EI	RD
9	SA, EX	RD, EX

<https://doi.org/10.1371/journal.pone.0236863.t006>

Table 6 summarizes the effort to perform *Task-1* through *Task-9* for S-BOUN-TI and WLB approaches in terms of the subtasks that must be performed. The subtasks are described in terms of the helper functions, where those that must be performed numerous times are indicated with a subsequent parenthesized number (i.e., TR(2), for two type resolutions). For the sake of brevity, we assume the existence of primitive functions (i.e., string, set, and list operations) and query support, which are not indicated in the comparison.

Semantically represented topics offer many opportunities when they are utilized in conjunction with resources and ontologies within LOD. The utility of S-BOUN-TI topics is most apparent when it yields results that are not directly accessible in the source content. The use of semantic rules enables enriching topics with general or highly domain-specific information. The latter being quite lucrative for domain-specific applications.

### Topic relevancy assessment

A comparative evaluation of the relevancy of S-BOUN-TI topics is difficult since the proposed approach has no precedence and produces topics that are significantly different from other approaches. The effort required for manual evaluation is complex, highly time-consuming, and error-prone since it involves the simultaneous examination of large sets of tweets (approximately 5800 per collection) and many semantic resources for every topic. The level of effort and diligence required to evaluate topics through surveys or services such as Amazon Mechanical Turk [105] that rely on human intelligence was deemed prohibitive. However, to gain insight regarding the relevancy of the topics, a meticulous evaluation was performed by the authors of this work with the assistance of a web application we developed for this purpose (see S1 Fig). This tool presents a set of topics to be annotated as *very satisfied*, *satisfied*, *minimally satisfied*, *not satisfied*, or *error* (when URIs are no longer accessible) along with optional comments to document noteworthy observations. During annotation, an evaluator may view the tweets from which the topics were generated as well as a word cloud that presents the words in proportion to their frequency. Also, the linked entities and temporal expressions extracted from the tweets can be inspected.

For evaluation purposes, 10 topics from randomly selected 36 intervals (9 from each debate) were annotated (S3 Table). Two annotators evaluated 24 intervals, 12 of which were identical to compute the inter-annotator agreement rate. The topics to be evaluated were selected based on a higher number of topic elements since they result from higher levels of alignment among posters. As such, they were deemed more significant to evaluate. Of the topics shown to annotators, there were 3 of size 8, 13 of size 7, 66 of size 6, 162 of size 5, 147 of size 4, 87 of size 3, and 2 topics of size 2.



The topics are presented per interval since they are all identified from the same collection. The evaluator is expected to inspect each topic to determine if it is related to tweet collection from which it was generated (by also inspecting the tweets). Each element of each topic is inspected by visiting their DBpedia resources to determine their relevancy to the collection in the *context* of the other elements of the topic. An element that is related to the tweet set, but not in the context of the other elements is considered irrelevant. Each topic is labeled as: *very satisfied* only if all of the topic elements are valid; *satisfied* only if one of the topic elements is incorrect; *minimally satisfied* if more than one element is incorrect while retaining significantly valuable information; and *not satisfied* if several topic elements are incorrect (i.e., the relative temporal expression may be true but does not convey sufficiently useful information). Note that the evaluation was performed in a strict manner, where a penalty is given for any kind of dissatisfaction—regardless of the source of the error. For example, if a web resource on DBpedia has incorrect information (which happens), the annotation of that topic is penalized. This was done to avoid subjective and relative evaluation as well as to assess the viability of the resources being used. Furthermore, since *S-BOUN-TI* topics are produced for machine interpretation the accuracy of topic elements is quite significant. It is also easier to identify mistaken elements in contrast to assessing a whole document as an error.

The results are examined in two ways: for topics marked either *Very satisfied* or *Satisfied* (assuming general satisfaction) and for topics annotated exclusively as *Very satisfied*. The evaluation resulted in the precision and  $F_1$  scores of 74.8%, 92.4% when considering only those marked as *Very satisfied*, and 81.0%, 93.3% when *Very satisfied* or *Satisfied*. The  $F_1$  scores (computed as defined by Hripcsak and Rothschild [106]) indicate a high degree of agreement among annotators.

### Comparison with human-readable topics

In an earlier work (*BOUN-TI* [21]), we identified human-readable topics from collections of microblogs (Wikipedia page titles). *BOUN-TI* models collections of posts as bags of words and compares their *tf-idf* vector with the content of Wikipedia pages to identify a ranked list of topics. The titles of the pages represent topics that are easily human-interpretable. *BOUN-TI* topics are satisfactory for human consumption, especially since they are descriptive titles produced by the prolific contributors of Wikipedia. Misleading topics can result when several subjects are posted about with similar intensities, such as the topic *Barack Obama citizenship conspiracy theories* derived from the words *Barack* and *citizen* whereas context of *citizen* was in *Hillary is easily my least favorite citizen in this entire country*—clearly not related to Barack Obama. Such cases occur as a consequence of using bag-of-words to model the documents. *S-BOUN-TI* overcomes this issue by considering both the wider context of the collection and the local context of posts while identifying topics. The context of individual tweets is used to determine potential topic elements, while the context of collections to capture the collective interest and patterns of use.

We inspected and compared *BOUN-TI* and *S-BOUN-TI* topics by deriving them from the same datasets (see [S1 Fig](#)). To give example, for  $pd_1$  [26–28], some of *BOUN-TI* topics are: *Donald Trump*, *Hillary Clinton*, *Bill Clinton*, *Barack Obama's Citizenship*, and *Laura Bush*. Since *S-BOUN-TI* topics include many elements, we will suffice by mentioning some of the topic elements: persons *dbr:Hillary\_Clinton*, *dbr:Donald\_Trump*, *dbr:Lester\_Holt* (the moderator of the debate) and other elements *dbr:Debate*, *dbr:ISIS*, *dbr:Fact*, *dbr:Interrupt*, *dbr:Watching*, and *dbr:Website*. These elements are identified because people were talking about ISIS, a high level of interruptions during the debate, and Hillary Clinton's fact-checking website.

The evaluation of BOUN-TI topics yielded 79.3% and 89.0% for precision and  $F_1$  scores for those marked *Very Satisfied* only and 88.9% and 94.0% if annotated as *Very Satisfied* or *Satisfied*. The scores for BOUN-TI are higher, which is largely influenced by two factors. Firstly, BOUN-TI topics are single titles that tend to be high level, thus they tend to be relevant even when not very specific. For example, the *Debate* topic would be considered relevant to a collection of posts about a specific debate at a specific time, even though it is very general. In S-BOUN-TI, topics are more granular with more elements, thus the evaluator scrutinized each element to determine relevancy in a manner that penalizes mistakes. Since S-BOUN-TI topics are intended for machine processing, a harsher judgment is called for.

In summary, we have found some similarities between S-BOUN-TI and BOUN-TI topics. In some cases, the corresponding DBpedia resources of BOUN-TI topics (Wikipedia pages) were elements of S-BOUN-TI topics that indicate a similarity between the results of BOUN-TI and S-BOUN-TI. Both approaches produce relevant topics, while S-BOUN-TI produces a greater variety and more granular topics in comparison to BOUN-TI topics. In general, BOUN-TI captures higher-level human-readable (encyclopedic) topics, while S-BOUN-TI picks up on lower-level elements that provide conceptual information that lend themselves to a greater variety of machine-interpretation such as *Barack Obama is a person and was a president*.

### Comparison with WLB topics

Latent Dirichlet allocation (LDA) is one of the most popular topic models, which makes a comparison with S-BOUN-TI interesting. To perform a comparison, LDA topics are generated with TwitterLDA [41] with the two-minute intervals of the datasets  $pd_1$ ,  $pd_2$ ,  $pd_3$ , and [VP] that are used for generating S-BOUN-TI topics (using the default values of LDA  $\alpha = 0.5$ ,  $\beta = 0.01$ , number of iteration = 100). Topics were generated for alternative values of LDA parameters for the expected number of topics:  $N = 2 - 10, 20, 30, 40, 50$ .

The topic representations of S-BOUN-TI and LDA are very different where LDA topics capture terms expressed by contributors as words-list-based (WLB) topics, S-BOUN-TI topics map original content to instances in LOD which are expressed with LOD resources and the OWL language [64]. With S-BOUN-TI a set of alternative words that are contributed may be mapped to the same semantic entity, capturing the intended meaning rather than how it was articulated.

To get a rough idea about the similarity of topics, we utilized the label (*rdfs:label*) of the S-BOUN-TI topic elements. The union of the lowercase form of words in the labels of all elements is compared using Jaccard similarity with topmost 10 terms of LDA topics (according to their distributions). We observed that there are cases that LDA and S-BOUN-TI topic elements are the same but not matching due to some syntactic difference. For example, if an LDA topic element is “emails”, and S-BOUN-TI topic element is *dbr:Email*, the strings “emails” and “email” do not match which results in lower similarity scores. To address similar issues, we assumed that the cases that one term is a substring of the other are matching. Each S-BOUN-TI topic is element-wise compared with LDA topics that are generated for the same input set.

The maximum similarity of an S-BOUN-TI topic in an interval is considered its similarity. The average of such similarities in an interval is the similarity measure obtained from that interval. And, the average similarity for all intervals is the average for a dataset, which ranges between 60-70% with a maximum of 77%. Since the comparison of the topics is performed on elements of different levels, the results give a very rough idea. The semantic similarity is expected to be higher. We would have been concerned if the comparisons resulted in very low values since that would indicate a significantly different relation among topic elements. As a result, we observe considerable coverage between the topics identified by these approaches, which is interesting for future work towards alternative methods for identifying topic elements. For

both methods, and in the case of using any other comparison methods to compare *S-BOUN-TI* topics with words-list-based topics, there is still the issue of word versus entity comparison. Automatically assessing the relevancy of topics without a gold standard is a challenging issue that requires domain knowledge and understanding of “topics” in the domain. We address these operations for future work.

### Evaluation summary


To assess the proposed approach, *S-BOUN-TI* topics were generated from sets of tweets and examined by inspecting their characteristics, using them in processing tasks, and comparing them with topics generated from *BOUN-TI*. Our main inquiry was to assess the viability of generating topics from collections of microposts with the use of resources on *LOD*. We found that considerable links between tweets and *LOD* resources were identified and that identifying topics from the constructed entity co-occurrence graph yielded relevant topics. With semantic queries and reasoning, we saw that it was possible to reveal information that is not directly accessible in the source (tweets), which could be very useful for those (i.e., campaign managers, marketers, journalists) who are following information from social media.

The proposed approach is a straightforward one aimed at gaining a basic understanding of the feasibility of mapping sets of tweets to semantically-related entities. If possible, this would facilitate a vast number of applications that harvest the richly connected web of data. Our observations lead us to believe that this is possible. Furthermore, this approach would improve by enhancing the techniques used to identify and relate topic elements, refining the topic representation, and with the increasing quality of data on *LOD*, which have been improving in terms of quantity and quality during the span of this work, a most encouraging prospect. Potential improvements are elaborated in the following section.

### Discussion and future work

In this section, we discuss some of our observations regarding the approach we proposed and present some future directions. The main objective of this work was to examine the feasibility of linking informal, noisy, and distributed micropost content to semantic resources in *LOD* to produce relevant machine-interpretable topics. We specifically focused on subjects of significant interest from a collective perspective. Topics of general interest lead to vast numbers of microposts. We generated semantic topics from a variety of tweets collections and represented them with an ontology that we developed for this purpose. The semantic topics were subjected to various tasks to examine their utility. The results show that relevant topics were identified for a diverse set of subjects. In the *Experiments and results* section, we presented the semantic topics generated from collections of tweets with emphasis on a complete set collected during the four major debates of U.S. elections (a total of 1036800 tweets). The utility of the resulting topics (respectively 1221, 1120, 1214, 1511 number of topics) was demonstrated through various tasks that facilitated the understanding of the issues relevant to the debate watchers, such as the persons, the locations, the temporal and other aspects of interest. Furthermore, issues at higher conceptual levels such as violence, ethnicities, and religions were revealed.

In our experiments, we observed that our approach produces relevant topics for diverse contexts. The topics of an entirely different context can be observed in a subject that is of great interest during the final preparations of this article, namely the coronavirus 2 (SARS-CoV-2) pandemic (a.k.a. COVID-19) that is widely reflected on social media. A preliminary exploration of topics generated from collections of tweets related to COVID-19 also yielded relevant topics. In this case collections of tweets posted during the same time for 53 consecutive days were inspected to get a general sense of the issues of relevance. There were 140 people, 32 locations,



**Fig 10. Topic elements for *topico:isAbout* in topics generated from COVID-19 dataset.** The more yellow the color is the more topics exist including that topic element.

<https://doi.org/10.1371/journal.pone.0236863.g010>

46 temporal expressions, 1097 issues that distinctly occurred in the topics. Among the occupations of people are politicians (several heads of states), journalists, singers, and athletes. The locations were dominated by China, Wuhan, Italy. While many locations were across the globe (i.e., Germany, West Bengal, and London) others were regions within the U.S. (i.e., Texas, Michigan, and Louisiana). This is reasonable since the time intervals of the tweets correspond to midday in the United States and COVID-19 cases were spiking in various parts of the country. There were also numerous temporal references, the most frequent ones being *now*, *today*, *tonight*, the months of January through May. The occurrence and the frequency of the specific temporal terms are significantly different from those encountered in the debate related sets, which did not have such a diverse set of temporal expressions (mostly the year 2016 *now*, *tonight*). Such differences capture the nature of contributions where the temporal aspect of the pandemic is indeed of much more significance due to the interest in how fast the rate of cases change and speculation about when things would improve. The resulting topics were processed to see when various issues emerged. Fig 10 shows when the *about* topic elements (*topico:isAbout*) were observed daily. Upon observing the references to drugs, we checked if other drugs were also referenced simply by querying DBpedia if the element type is *dbo:Drug*. This identified the other drugs referenced in tweets as: BCG vaccine, Cocaine, Doxycycline, Favipiravir, Generic drug, Pharmaceutical drug, Polio vaccine, Ibuprofen, Paracetamol, Chloroquine, Azithromycin, Antiviral drug, Hydroxychloroquine. The most frequently referenced one was Hydroxychloroquine—a drug mentioned by the president of the United States several times. Obviously, tweets from such small intervals are insufficient to inspect such a vast issue. A more comprehensive examination with collections that covers all time zones would be

required. Nevertheless, even with this small set, it is evident that this approach produces relevant topics.

The results we obtained are encouraging leaving us with many future directions to pursue, which we elaborate in the remainder of this section.

### Topic element detection

Since semantic topics consist of topic elements, correctly identifying them is important. Here, the main challenges are the inability to link to anything at all and incorrect linking. Obviously, entity linking fails when suitable entities are not represented on LOD, such as when new subjects emerge. In recent years the significance given to the creation and accessibility to open data resources has led to a rapid increase in the data represented on LOD [74] (see [Background](#) section for information about LOD). In this work, we eliminate unlinked spots, which could be particularly problematic for spots with high frequencies since that indicates common interest. To alleviate this matter, such spots could be linked to an instance of *owl:Thing* indicating that there is some *thing* of significance whose type is unknown.

For entities that exist but not successfully linked, better approaches are required. Named entity recognition and linking are active research areas that are improving across all domains and languages. Another approach to determining the correct entity type ranks all relevant types using taxonomies and ontologies such as YAGO, and FREEBASE [107, 108]. Also, additional pre-processing steps can be taken prior to entity linking, such as tweet normalization [109] and hashtag segmentation [110].

As discussed in the Experiments and results section, identifying locations is challenging since many entities can be considered a location in some context. We imposed some rules to determine if such elements qualified as a location in the context it was used. Our evaluation revealed that although all elements we deemed to be locations were correct. Unfortunately, we missed identifying some of them since they did not match our rules—mostly due to how tweets were articulated. Location prediction on Twitter is known to be challenging and is of significant interest since there are many areas of application [111, 112]. It is of interest for many purposes, such as disaster tracking and mitigation and with the emergence of the COVID-19 pandemic crisis, this work has been intensified. Many studies appear as preprints, which are not yet vetted, however with the immense motivation improved location detection is expected. Our rules for detecting locations must be revisited. Also, indirect location indicators such as those found in profiles and geotagged content [113], and in co-occurrence patterns [114] could offer hints that improve the detection of locations. For ethical reasons, we do not (and do not intend to) use any profile information, but could consider utilizing other indirect signals.

A more troublesome issue stems from ambiguous terms, which is most prevalent in person names. For example, the spot *Clinton* was inaccurately linked to the 42<sup>nd</sup> U.S. president Bill Clinton instead of the 2016 U.S. presidential candidate Hillary Clinton in tweets regarding lying about Obamacare. In this case, both persons are politicians, one was a U.S. president and the other a U.S. presidential candidate, and they are spouses. Although this example is particularly challenging, the ambiguity of person names is generally challenging. A similar issue arises since the titles of songs, movies, albums, and books are terms of ordinary conversation, such as time (**Time** magazine), cure (music band **The Cure**), and WHO (television character **Dr. Who**). We encountered such cases in our experiments, albeit not frequently, since the entity linker typically assigns low confidence rates for such links. Furthermore, our approach eliminates the links to entities that occur infrequently (see the *Identifying topics* section) in a collaborative filtering manner.

Recently, word-embedding techniques that capture semantic similarity among terms are being applied to named entity recognition (NER) and disambiguation [115, 116], and entity



linking [117, 118]. These techniques represent terms as vectors in a high dimensional vector space and obtain them via machine learning given a corpus. The semantics of terms are captured from the context of the terms. The vectors of semantically similar terms are close in the vector space. Since emerging entities are expected to be included in the knowledge bases, the topic and/or topic elements could be periodically revisited for opportunities for improvement. These advances in named entity detection and linking are very promising and are expected to positively improve the detection of topic elements.

In our experiments, we focused on English tweets and named entity recognition so that we could interpret the results. Several tools work for other languages (including TagMe that we used in our prototype). Furthermore, the natural language processing community strongly emphasizes work on low resource languages, which is resulting in additional knowledge resources and tools. The goal is to work with multiple languages, which link to the same conceptual entity. Thereby, being able to glean information regarding content that is globally produced. This is important for many tasks of global interest, such as pandemic diseases, disasters, news, entertainment, and learning material.

### Semantic topic identification

In this work, we chose maximal cliques to identify the topics so as to assure that all the elements are related by virtue of having been posted together. The co-occurrence graphs from which topics are extracted have relatively few nodes with high degree centrality (i.e., Hillary Clinton and Donald Trump in the debate sets) with the remaining node being relatively weak (see Fig 7). Thus, several topics extracted from such graphs tend to share the dominant nodes, which reflect the narratives related to the dominant nodes. On the other hand, the nodes that are connected to the dominant nodes tend to fall into different topics since they are usually not connected to each other. This fairly accurately reflects micropost content (i.e., many different topics involve Donald Trump). However, this results in some topics seeming very similar or repetitive. It is worth investigating more relaxed graph algorithms to increase the elements of topics while preserving the context. For example,  $k$ -cliques (the maximal sub-graphs where the largest geodesic distance between any two vertices is  $k$ ) constrained by type rules could yield richer topics. However, caution must be exercised, since the volume of microposts and their limited context is likely to yield many potential yet unrelated candidates for  $k > 1$ , which is computationally challenging and costly. Note that the goal is not to simply increase the number of topics or their elements since we are aiming to reduce large sets of tweets to higher-level topics. Rather, the aim is to increase the quality of topics by associating related elements.

The size of the post collections we used was limited by the rate limits of the Twitter streaming API and our computational resources. For the debates, this corresponded to the tweets posted within 2-minute intervals. During heavy posting conditions the subjects change frequently and short windows are suitable since the topics change and the number of tweets to detect collective interest is sufficient. During slower posting conditions the subjects don't change as fast, thus sets collected over longer durations are appropriate. Dynamically varying collection durations based on how frequently the subjects of topics change over time would be valuable. *Topico* is capable of representing such intervals, however, they must be determined through time series analysis.

### Semantic topic representation

*Topico* specifies an elementary set of topic element types, namely person, location, temporal expressions, and other entities (those related by *topico:isAbout*). It encompasses basic classes, object properties, and data properties to represent commonly occurring elements. Inferred



relations and classes support convenient processing. This ontology could be extended to cover additional types (such as events, art, currency, character, products, drug, and natural objects) as well as refine existing types (such as facility, address, astral body, organization, and market) [119]. While covering a wider type of cross-domain entities is of interest, we expect that customized for a specific domain will be quite interesting. For this purpose, ontologies relevant to the domain of interest and associated data resources are required. There are many useful ontologies and resources, especially in the life-sciences domain. Naturally, domain-specific tasks would also be defined. The tasks shown in the *Experiments and results* illustrate the kinds of topic-related tasks that could be of interest to campaign managers, journalists, and political enthusiasts.

### Semantic topics utilization

The purpose of focusing on semantic topics is for their semantic processing potentials. We demonstrated how semantic topics can be utilized through various semantic tasks in the *Experiments and results* section. The vision is to deliver this power to an end-user who is following the rapidly flowing distributed microposts. Towards this end, higher-level tasks should be defined such as similarity, sentiment analysis, and recommendations. Tracking topics, such as when they emerge, if they persist, if and when they spike, and if they exhibit some pattern is useful information. Reports that provide statistical information will enable those who are interested in the topics to take action. Other interesting tasks are tracking the evolution of and predicting topics. One of the future directions is an explorer for S-BOUN-TI topics which requires the generation of human-interpretable topics. Using such an explorer users could search and browse topics, view ranked topics, graphs, and charts that provide relational and temporal information (trends), view social network analysis. They must be able to view the results of the processing of semantic topics, which may be predefined domain-specific. Such a system should recommend topics and topic related observations such as trends, newly emerging issues. Furthermore, multimedia presentations that depict the lifecycles of topics persist over a given time can be generated with dynamic summarization techniques [120, 121].

Eventually, a tool that is customizable with domain-specific knowledge resources for detecting and processing topics. The specific nature of the subject and desired processing will vary depending on the context. A domain-specific topic detection system customized for diseases with knowledge bases like the International Classification of Diseases (ICD) [122] for diseases and SNOMED-CT (Systematized Nomenclature of Medicine—Clinical Terms) [123] would be useful in tracking the for pandemics related topics of public interest. Such topic explorers would enable users to glean domain-specific insights that are very difficult to obtain by direct experience with a vast number of microposts.

One of the most interesting potentials of semantic resources is revealed via *federated queries* that search across distributed resources. Unfortunately, the performance of federated queries can be quite inefficient. The order in which queries are executed must be carefully designed to achieve reasonable response times. Generally, electing to execute the more restrictive queries prior to others which restricts the search space is considered a good approach. Finally, generating streams of semantic topics could be facilitated with stream reasoning [124] and queried with a stream query language such as C-SPARQL [125] and C-SPRITE [126].

### Conclusions

This work investigates the viability of extracting semantic topics from collections of microposts via processing their corresponding linked entities that are LOD resources. To this end, an

ontology (*Topico*) to represent topics is designed, an approach to extracting topics from sets of microposts is proposed, a prototype of this approach is implemented, and topics are generated from large sets of posts from Twitter.

The main inquiry of this work was to examine whether an approach based on linking microposts to LOD resources could be utilized in generating semantically represented and machine-interpretable topics. The proposed approach extracts a significantly rich set of information from the posts in terms of relating them to web-resources which themselves are related to other resources via data and object properties. We demonstrate the benefits of using LOD and ontologies while identifying as well as utilizing the topics. During the identification phase, we were able to identify candidate elements and resolve their types. The ontologically represented topics consisting of entities enabled processing opportunities that revealed information about collections of microposts that are not readily observable even if each post were to be manually inspected. Also, we notice an increase in the quality of the generated topics over time thanks to the efforts related to the continued expansion and correction of LOD resources.

Our main goal of producing machine-interpretable topics was for their utilization in further processing. We demonstrated such utilization with several examples of various levels of complexity, where information that is not readily available in the original posts is revealed. A user evaluation (with 81.0% precision and  $F_1$  of 93.3%) and regularly performed manual inspections show that the identified topics are relevant. In summary, we are encouraged by the results we obtained and list many research opportunities to improve the topic identification approach and process topics in general and in domain-specific manners.

## Supporting information

### **S1 Appendix. The main object properties of *topico:Topic*.**

(PDF)

**S1 Fig. The evaluation tool showing *BOUN-TI* and *S-BOUN-TI* topics from [VP] [68-70].** The *See tweets* and *See word cloud* links show the related tweets and a word cloud generated from them. The *Entity frequencies* link shows the list of linked entities and their frequencies. All resources are reachable through web-links for inspection.

(PNG)

### **S2 Fig. A sample co-occurrence graph ( $G'$ ) of topic elements from the first presidential debate.**

(PNG)

### **S1 Table. The namespace prefixes that are utilized in *Topico* and referred to in this paper.**

(PNG)

**S2 Table. Percentage of tweets in the post sets that produce the vertices (topic elements), edges (co-occurring elements), and topics.** This table shows the percentage of tweets in the post sets that produce the vertices (topic elements), edges (co-occurring elements), and topics. The columns labeled *Before* and *Pruned* show the impact of pruning the graph. The columns labeled *Topic* show how many were retained in the topic.

(PNG)

### **S3 Table. The intervals within the datasets that were used for evaluating topics.**

(PNG)

## Acknowledgments

We thank Dr. T. B. Dinesh and Dr. Jayant Venkatanatha for valuable contributions during the preparation of this work. We are grateful for the feedback received from the members of SosLab (Department of Computer Engineering, Boğaziçi University) during the development of this work.

## Author Contributions

**Conceptualization:** Ahmet Yıldırım, Suzan Uskudarli.

**Data curation:** Ahmet Yıldırım.

**Formal analysis:** Ahmet Yıldırım, Suzan Uskudarli.

**Investigation:** Ahmet Yıldırım.

**Methodology:** Ahmet Yıldırım, Suzan Uskudarli.

**Resources:** Suzan Uskudarli.

**Software:** Ahmet Yıldırım.

**Supervision:** Suzan Uskudarli.

**Validation:** Ahmet Yıldırım, Suzan Uskudarli.

**Visualization:** Ahmet Yıldırım.

**Writing – original draft:** Ahmet Yıldırım, Suzan Uskudarli.

**Writing – review & editing:** Ahmet Yıldırım, Suzan Uskudarli.

## References

1. Twitter. Twitter. 2020 June 15 [cited 15 Jun 2020]. In: [Internet]—. [about 1 screen]. Available from: <https://twitter.com>.
2. Internet Live Stats. Twitter statistics. 2020 Jun 15 [cited 15 Jun 2020]. In: [Internet]—. [about 1 screen]. Available from: <http://www.internetlivelstats.com/twitter-statistics/>.
3. Eisenstein J. What to do about bad language on the internet. In: Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Atlanta, Georgia: ACL; 2013. p. 359–369. Available from: <http://www.aclweb.org/anthology/N13-1037>.
4. Yan X, Guo J, Lan Y, Cheng X. A Biterm Topic Model for Short Texts. In: Proceedings of the 22Nd International Conference on World Wide Web. WWW'13. New York, NY, USA: ACM; 2013. p. 1445–1456.
5. Li C, Duan Y, Wang H, Zhang Z, Sun A, Ma Z. Enhancing Topic Modeling for Short Texts with Auxiliary Word Embeddings. *ACM Trans Inf Syst.* 2017; 36(2):11:1–11:30. <https://doi.org/10.1145/3091108>
6. Fang A. Analysing political events on Twitter: topic modelling and user community classification. Doctoral dissertation, University of Glasgow. 2019. Available from: <https://theses.gla.ac.uk/41135/>.
7. Perrier A. Segmentation of Twitter Timelines via Topic Modeling; 2015. Available from: [https://alexisperrier.com/nlp/2015/09/16/segmentation\\_twitter\\_timelines\\_lda\\_vs\\_lsa.html](https://alexisperrier.com/nlp/2015/09/16/segmentation_twitter_timelines_lda_vs_lsa.html).
8. Chodhary S, Jagdale RS, Deshmukh SN. Semantic Analysis of Tweets using LSA and SVD. *International Journal of Emerging Trends and Technology in Computer Science.* 2016; 5(4).
9. Ozer M, Kim N, Davulcu H. Community detection in political Twitter networks using Nonnegative Matrix Factorization methods. In: Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM. Institute of Electrical and Electronics Engineers Inc. 2016. p. 81–88.
10. Chen Y, Zhang H, Liu R, Ye Z, Lin J. Experimental explorations on short text topic mining between LDA and NMF based Schemes. *Knowledge-Based Systems.* 2019; 163:1–13. <https://doi.org/10.1016/j.knosys.2018.08.011>

11. Alvanaki F, Michel S, Ramamritham K, Weikum G. See What's enBlogue: Real-time Emergent Topic Identification in Social Media. In: Proceedings of the 15th International Conference on Extending Database Technology. EDBT'12. New York, NY, USA: ACM; 2012. p. 336–347.
12. Cataldi M, Di Caro L, Schifanella C. Emerging Topic Detection on Twitter Based on Temporal and Social Terms Evaluation. In: Proceedings of the Tenth International Workshop on Multimedia Data Mining. MDMKDD'10. New York, NY, USA: ACM; 2010. p. 4:1–4:10.
13. Kasiviswanathan SP, Melville P, Banerjee A, Sindhvani V. Emerging Topic Detection Using Dictionary Learning. In: Proceedings of the 20th ACM International Conference on Information and Knowledge Management. CIKM'11. New York, NY, USA: ACM; 2011. p. 745–754.
14. Marcus A, Bernstein MS, Badar O, Karger DR, Madden S, Miller RC. Twitinfo: Aggregating and Visualizing Microblogs for Event Exploration. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. CHI'11. New York, NY, USA: ACM; 2011. p. 227–236.
15. Mathioudakis M, Koudas N. TwitterMonitor: Trend Detection over the Twitter Stream. In: Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data. SIGMOD'10. New York, NY, USA: ACM; 2010. p. 1155–1158.
16. Sayyadi H, Raschid L. A Graph Analytical Approach for Topic Detection. *ACM Trans Internet Technol.* 2013; 13(2):4:1–4:23. <https://doi.org/10.1145/2542214.2542215>
17. Bicalho P, Pita M, Pedrosa G, Lacerda A, Pappa GL. A General Framework to Expand Short Text for Topic Modeling. *Inf Sci.* 2017; 393(C):66–81. <https://doi.org/10.1016/j.ins.2017.02.007>
18. Celebi HB, Uskudarli S. Content Based Microblogger Recommendation. In: 2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing. IEEE; 2012. p. 605–610.
19. Degirmencioglu EA, Uskudarli S. Exploring area-specific microblogging social networks. In: *WebSci10: Extending the Frontiers of Society On-Line*. Raleigh, North Carolina, USA: Web Science; 2010.
20. Sharifi BP, Inouye DI, Kalita JK. Summarization of Twitter Microblogs. *The Computer Journal.* 2014; 57(3):378–402. <https://doi.org/10.1093/comjnl/bxt109>
21. Yildirim A, Uskudarli S, Özgür A. Identifying Topics in Microblogs Using Wikipedia. *PLoS ONE.* 2016; 11(3):1–20.
22. Han H, Viriyothai P, Lim S, Lameter D, Mussell B. Yet Another Framework for Tweet Entity Linking (YAFTEL). In: 2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR); 2019. p. 258–263.
23. Sakor A, Onando Mulangí, Singh K, Shekarpour S, Esther Vidal M, Lehmann J, et al. Old is Gold: Linguistic Driven Approach for Entity and Relation Linking of Short Text. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). Minneapolis, Minnesota: ACL; 2019. p. 2336–2346. Available from: <https://www.aclweb.org/anthology/N19-1243>.
24. Ferragina P, Scaiella U. Fast and Accurate Annotation of Short Texts with Wikipedia Pages. *IEEE Software.* 2012; 29(1):70–75. <https://doi.org/10.1109/MS.2011.122>
25. Gattani A, Lamba DS, Garera N, Tiwari M, Chai X, Das S, et al. Entity Extraction, Linking, Classification, and Tagging for Social Media: A Wikipedia-based Approach. *Proc VLDB Endow.* 2013; 6(11):1126–1137. <https://doi.org/10.14778/2536222.2536237>
26. Bizer C, Heath T, Berners-Lee T. Linked data: The story so far. In: *Semantic services, interoperability and web applications: emerging concepts*. IGI Global; 2011. p. 205–227.
27. Shadbolt N, Berners-Lee T, Hall W. The Semantic Web Revisited. *IEEE Intelligent Systems.* 2006; 21(3):96–101. <https://doi.org/10.1109/MIS.2006.62>
28. W3C. Semantic Web [cited 15 Jun 2020]. In: W3C web site [Internet]—. [about 3 screens]; 2015. Available from: <https://www.w3.org/standards/semanticweb/>.
29. Yildirim A. S-BounTI: Semantic Topic Identification approach from Microblog post sets. An application. [cited 15 Jun 2020]. Database: figshare [Internet]. Available from: <https://doi.org/10.6084/m9.figshare.5943211>.
30. SoSLab. Explore semantic topics. 2020 Jun 15 [cited 15 Jun 2020]. In: SoSLab Web Site [Internet]—. [about 1 screen]. Available from: <http://soslab.cmpe.boun.edu.tr/sbounti/>.
31. Yildirim A, Uskudarli S. S-BounTI: Semantic Topic Identification approach from Microblog post sets using Linked Open Data, published datasets. [cited 15 Jun 2020]. Database: figshare [Internet]. 2018. Available from: <https://doi.org/10.6084/m9.figshare.7527476>.
32. Blei DM, Ng AY, Jordan MI. Latent Dirichlet Allocation. *Journal of machine Learning research.* 2003; 3(Jan):993–1022.

33. Qiang J, Qian Z, Li Y, Yuan Y, Wu X. Short Text Topic Modeling Techniques, Applications, and Performance: A Survey. *CoRR*. 2019;abs/1904.07695.
34. Cheng X, Yan X, Lan Y, Guo J. BTM: Topic Modeling over Short Texts. *IEEE Transactions on Knowledge and Data Engineering*. 2014; 26(12):2928–2941. <https://doi.org/10.1109/TKDE.2014.2313872>
35. Lin T, Tian W, Mei Q, Cheng H. The Dual-sparse Topic Model: Mining Focused Topics and Focused Terms in Short Text. In: *Proceedings of the 23rd International Conference on World Wide Web. WWW'14*. New York, NY, USA: ACM; 2014. p. 539–550. Available from: <http://doi.acm.org/10.1145/2566486.2567980>.
36. Qiang J, Li Y, Yuan Y, Wu X. Short text clustering based on Pitman-Yor process mixture model. *Applied Intelligence*. 2018; 48(7):1802–1812. <https://doi.org/10.1007/s10489-017-1055-4>
37. Yin J, Chao D, Liu Z, Zhang W, Yu X, Wang J. Model-based Clustering of Short Text Streams. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD'18*. New York, NY, USA: ACM; 2018. p. 2634–2642. Available from: <http://doi.acm.org/10.1145/3219819.3220094>.
38. Weng J, Lim EP, Jiang J, He Q. TwitterRank: Finding Topic-sensitive Influential Twitterers. In: *Proceedings of the Third ACM International Conference on Web Search and Data Mining. WSDM'10*. New York, NY, USA: ACM; 2010. p. 261–270. Available from: <http://doi.acm.org/10.1145/1718487.1718520>.
39. Bauer S, Noulas A, Séaghdha DO, Clark S, Mascolo C. Talking Places: Modelling and Analysing Linguistic Content in Foursquare. In: *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*; 2012. p. 348–357.
40. Mehrotra R, Sanner S, Buntine W, Xie L. Improving LDA Topic Models for Microblogs via Tweet Pooling and Automatic Labeling. In: *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR'13*. New York, NY, USA: ACM; 2013. p. 889–892.
41. Qiu M. Latent Dirichlet Allocation (LDA) Model for Microblogs (Twitter, Weibo etc.); 2017. Available from: <https://github.com/minghui/Twitter-LDA>.
42. Chen W, Wang J, Zhang Y, Yan H, Li X. User Based Aggregation for Biterm Topic Model. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Beijing, China: ACL; 2015. p. 489–494. Available from: <https://www.aclweb.org/anthology/P15-2080>.
43. Wang W, Zhou H, He K, Hopcroft JE. Learning Latent Topics from the Word Co-occurrence Network. In: Du D, Li L, Zhu E, He K, editors. *Theoretical Computer Science*. Singapore: Springer Singapore; 2017. p. 18–30.
44. Yi F, Jiang B, Wu J. Topic Modeling for Short Texts via Word Embedding and Document Correlation. *IEEE Access*. 2020; 8:30692–30705. <https://doi.org/10.1109/ACCESS.2020.2973207>
45. Qiang J, Chen P, Wang T, Wu X. Topic Modeling over Short Texts by Incorporating Word Embeddings. In: Kim J, Shim K, Cao L, Lee JG, Lin X, Moon YS, editors. *Advances in Knowledge Discovery and Data Mining*. Cham: Springer International Publishing; 2017. p. 363–374.
46. Li C, Wang H, Zhang Z, Sun A, Ma Z. Topic Modeling for Short Texts with Auxiliary Word Embeddings. In: *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR'16*. New York, NY, USA: ACM; 2016. p. 165–174. Available from: <http://doi.acm.org/10.1145/2911451.2911499>.
47. Petrović S, Osborne M, Lavrenko V. Streaming First Story Detection with Application to Twitter. In: *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics. ACL*; 2010. p. 181–189.
48. Genc Y, Sakamoto Y, Nickerson JV. Discovering Context: Classifying Tweets through a Semantic Transform based on Wikipedia. In: *Proceedings of the 6th international conference on Foundations of augmented cognition: directing the future of adaptive systems. FAC'11*. Springer-Verlag; 2011. p. 484–492. Available from: <http://dl.acm.org/citation.cfm?id=2021773.2021833>.
49. Prieto VM, Matos S, Álvarez M, Cacheda F, Oliveira JL. Twitter: A Good Place to Detect Health Conditions. *PLoS ONE*. 2014; 9(1):1–11. <https://doi.org/10.1371/journal.pone.0086191>
50. Parker J, Wei Y, Yates A, Frieder O, Goharian N. A Framework for Detecting Public Health Trends with Twitter. In: *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining. ASONAM'13*. New York, NY, USA: ACM; 2013. p. 556–563.
51. Eissa AHB, El-Sharkawi ME, Mokhtar HMO. Towards Recommendation Using Interest-Based Communities in Attributed Social Networks. In: *Companion Proceedings of the The Web Conference 2018. WWW'18*. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee; 2018. p. 1235–1242.



52. Gruetze T, Kasneci G, Zuo Z, Naumann F. CohEEL: Coherent and Efficient Named Entity Linking through Random Walks. *Web Semantics: Science, Services and Agents on the World Wide Web*. 2016; 37(0).
53. Kumar VK. Knowledge Representation Technologies Using Semantic Web. In: *Web Services: Concepts, Methodologies, Tools, and Applications*. IGI Global; 2019. p. 1068–1076.
54. Liao X, Zhao Z. Unsupervised Approaches for Textual Semantic Annotation, A Survey. *ACM Comput Surv*. 2019; 52(4):66:1–66:45. <https://doi.org/10.1145/3324473>
55. Martinez-Rodriguez JL, Lopez-Arevalo I, Rios-Alvarado AB, Hernandez J, Aldana-Bobadilla E. Extraction of RDF Statements from Text. In: Villazón-Terrazas B, Hidalgo-Delgado Y, editors. *Knowledge Graphs and Semantic Web*. Cham: Springer International Publishing; 2019. p. 87–101.
56. Matentzoglou N, Malone J, Mungall C, Stevens R. MIRO: guidelines for minimum information for the reporting of an ontology. *Journal of Biomedical Semantics*. 2018; 9(1):6. <https://doi.org/10.1186/s13326-017-0172-7> PMID: 29347969
57. Rospocher M, Corcoglioniti F, Dragoni M. Boosting Document Retrieval with Knowledge Extraction and Linked Data. *Semantic Web*. 2019; 10(4):753–778. <https://doi.org/10.3233/SW-180325>
58. Gottschalk S, Demidova E. EventKG—the hub of event knowledge on the web—and biographical timeline generation. *Semantic Web*. 2019;(Pre-press):1–32.
59. Abebe MA, Tekli J, Getahun F, Chbeir R, Tekli G. Generic metadata representation framework for social-based event detection, description, and linkage. *Knowledge-Based Systems*. 2019; <https://doi.org/10.1016/j.knsys.2019.06.025>.
60. van Aggelen A, Hollink L, Kemman M, Kleppe M, Beunders H. The debates of the European Parliament as Linked Open Data. *Semantic Web Journal*. 2017; 8(2):271–281. <https://doi.org/10.3233/SW-160227>
61. Nanni F, Menini S, Tonelli S, Ponzetto SP. Semantifying the UK Hansard (1918–2018). In: *Proceedings of the 19th ACM/IEEE Joint Conference on Digital Libraries: JCDL'19, June 2019, Urbana-Champaign, Illinois*. New York, NY: ACM; 2019. p. 1–2. Available from: <https://madoc.bib.uni-mannheim.de/49597/>.
62. Nédellec C. OntoBiotope. [cited 15 June 2020]. Database: Inra [Internet]. 2018.
63. Karadeniz I, Özgür A. Linking entities through an ontology using word embeddings and syntactic re-ranking. *BMC Bioinformatics*. 2019; 20(1):156. <https://doi.org/10.1186/s12859-019-2678-8> PMID: 30917789
64. W3C. Web Ontology Language (OWL) [cited 15 Jun 2020]. In: [Internet]—. [about 3 screens]. Available from: <https://www.w3.org/OWL/>.
65. Antoniou G, Groth P, Harmelen Fv, Hoekstra R. *A Semantic Web Primer*. The MIT Press; 2012.
66. Brickley D, Miller L. FOAF Vocabulary Specification 0.99. 2014 Jan 14 [cited 15 Jun 2020]. In: [Internet]—. [about 61 screens]. Available from: <http://xmlns.com/foaf/spec/>.
67. W3C. WGS84 Geo Positioning: an RDF vocabulary. 2009 Apr 20 [cited 15 Jun 2020]. In: W3C Web site [Internet]—. [about 5 screens]. Available from: [http://www.w3.org/2003/01/geo/wgs84\\_pos](http://www.w3.org/2003/01/geo/wgs84_pos).
68. W3C Semantic Web Interest Group. Basic Geo (WGS84 lat/long) Vocabulary. 2004 Feb 06 [cited 15 Jun 2020]. In: W3C Web site [Internet]—. [about 5 screens]. Available from: <https://www.w3.org/2003/01/geo/>.
69. GeoNames. GeoNames Ontology—Geo Semantic Web. 2010 Oct 5 [cited 15 Jun 2020]. In: [Internet]—. [about 3 screens]. Available from: <http://www.geonames.org/ontology/documentation.html>.
70. Guha RV, Brickley D, Macbeth S. Schema.Org: Evolution of Structured Data on the Web. *Commun ACM*. 2016; 59(2):44–51. <https://doi.org/10.1145/2844544>
71. Schema.org. Home—schema.org. [cited 15 Jun 2020]. In: [Internet]—. [about 1 screen]. Available from: <http://schema.org/>.
72. Cox S, Little C, Hobbs JR, Pan F. Time Ontology in OWL. W3C; 2020. Available from: <https://www.w3.org/TR/2020/CR-owl-time-20200326/>.
73. TagMe. TagMe API Documentation. [cited 15 Jun 2020]. In: d4science services [Internet]—. [about 4 screens]. Available from: <https://sobigdata.d4science.org/web/tagme/tagme-help>.
74. Linked Data community. Linked Data | Linked Data—Connect Distributed Data across the Web. [cited 25 May 2020]. In: [Internet]—. [about 1 screen]; 2020. Available from: <http://linkeddata.org/>.
75. Auer S. Introduction to LOD2. In: Auer S, Bryl V, Tramp S, editors. *Linked Open Data—Creating Knowledge Out of Interlinked Data: Results of the LOD2 Project*. Springer International Publishing; 2014. p. 1–17.
76. Schmachtenberg M, Bizer C, Paulheim H. The Semantic Web—ISWC 2014. In: *Adoption of the Linked Data Best Practices in Different Topical Domains*. Cham: Springer International Publishing; 2014. p. 245–260.



77. McCrae JP, Abele A, Buitelaar P, Cyganiak R, Jentzsch A, Andryushechkin V, et al. The Linked Open Data Cloud Diagram. 2020 May [cited 15 Jun 2020]. In: [Internet]—. [about 6 screens]. Available from: <http://lod-cloud.net>.
78. Bizer C, Lehmann J, Kobilarov G, Auer S, Becker C, Cyganiak R, et al. DBpedia—A Crystallization Point for the Web of Data. *Web Semantics: Science, Services and Agents on the World Wide Web*. 2009; 7(3):154–165. <https://doi.org/10.1016/j.websem.2009.07.002>
79. DBpedia. The Release Circle—A Glimpse behind the Scenes. [cited 15 Jun 2020]. In: DBpedia Blog [Internet]—. [about 4 screens]. Available from: <https://blog.dbpedia.org/2018/10/17/the-release-circle-a-glimpse-behind-the-scenes/>.
80. Vrandečić D, Krötzsch M. Wikidata: A Free Collaborative Knowledgebase. *Commun ACM*. 2014; 57(10):78–85. <https://doi.org/10.1145/2629489>
81. Wikimedia Foundation. Wikidata. 2019 Dec 30 [cited 15 Jun 2020]. In: Wikidata [Internet]—. [about 3 screens]. Available from: [https://www.wikidata.org/wiki/Wikidata:Main\\_Page](https://www.wikidata.org/wiki/Wikidata:Main_Page).
82. Wikidata. Wikidata:Statistics. [cited 15 Jun 2020]. In: Wikidata Web site [Internet]—. [about 3 screens]. Available from: <https://www.wikidata.org/wiki/Wikidata:Statistics>.
83. Wikimedia Foundation. Wikidata query service. [cited 15 Jun 2020]. In: Wikidata [Internet]—. [about 1 screen]. Available from: <https://query.wikidata.org/>.
84. DBpedia. Virtuoso SPARQL Query Editor. [cited 15 Jun 2020]. In: DBpedia [Internet]—. [about 1 screen]. Available from: <http://dbpedia.org/sparql>.
85. Seaborne A, Harris S. SPARQL 1.1 Query Language. W3C; 2013. Available from: <http://www.w3.org/TR/2013/REC-sparql11-query-20130321/>.
86. Pérez J, Arenas M, Gutierrez C. Semantics and Complexity of SPARQL. *ACM Trans Database Syst*. 2009; 34(3).
87. Stanford Center for Biomedical Informatics Research. Protégé; 2018. Available from: <https://protege.stanford.edu/>.
88. Noy NF, McGuinness DL. *Ontology Development 101: A Guide to Creating Your First Ontology*; 2001.
89. Yıldırım A, Uskudarli S. About Topico ontology. 2019 Aug 25 [cited 15 Jun 2029]. In: SoSLab Web Site [Internet]—. [4 screens]. Available from: <http://soslab.cmppe.boun.edu.tr/sbounti/AboutTopico.pdf>.
90. Eppstein D, Löffler M, Strash D. Listing All Maximal Cliques in Sparse Graphs in Near-optimal Time. *CoRR*. 2010;abs/1006.5440.
91. Bailey F. Phirehose. 2018 Mar 15 [cited 15 Jun 2020]. In: GitHub [Internet]—. [about 3 screens]; 2018. Available from: <https://github.com/fennb/phirehose>.
92. Twitter. Filter realtime Tweets [cited 15 Jun 2020]. In: Twitter Developer Documentation [Internet]—. [about 2 screens]. Available from: <https://developer.twitter.com/en/docs/tweets/filter-realtime/api-reference/post-statuses-filter>.
93. R Core Team. R: A Language and Environment for Statistical Computing; 2013. Available from: <http://www.R-project.org/>.
94. Apache Jena. Fuseki. [cited 15 Jun 2020]. In: Apache Web site [Internet]—. [about 2 screens]. Available from: <https://jena.apache.org/documentation/fuseki2/index.html>.
95. Politico Staff. Full transcript: First 2016 presidential debate. 2016 Sep 27 [cited 15 Jun 2020]. In: Politico Web site [Internet]—. [about 67 screens]. Available from: <https://www.politico.com/story/2016/09/full-transcript-first-2016-presidential-debate-228761>.
96. Politico Staff. Full transcript: Second 2016 presidential debate. 2016 Oct 10 [cited 15 Jun 2020]. In: Politico Web site [Internet]—. [about 52 screens]. Available from: <https://www.politico.com/story/2016/10/2016-presidential-debate-transcript-229519>.
97. Politico Staff. Full transcript: Third 2016 presidential debate. 2016 Oct 20 [cited 15 Jun 2020]. In: Politico Web site [Internet]—. [about 56 screens]. Available from: <https://www.politico.com/story/2016/10/full-transcript-third-2016-presidential-debate-230063>.
98. Politico Staff. Full transcript: 2016 vice presidential debate. 2016 Oct 5 [cited 15 Jun 2020]. In: Politico Web site [Internet]—. [about 66 screens]. Available from: <https://www.politico.com/story/2016/10/full-transcript-2016-vice-presidential-debate-229185>.
99. Yıldırım A, Uskudarli S, Ozgur A. Tf values, word frequency values for gathering idf values, and the evaluation data submitted to PLoS One, titled Identifying Topics in Microblogs Using Wikipedia. [cited 15 Jun 2020]. Database: figshare [Internet]. Available from: [https://figshare.com/articles/data\\_tar\\_gz/2068665](https://figshare.com/articles/data_tar_gz/2068665).
100. McKinney J, Iannella R. vCard Ontology—for describing People and Organizations. W3C; 2014. Available from: <http://www.w3.org/TR/2014/NOTE-vcards-20140522/>.

101. Horrocks I, Patel-Schneider PF, Boley H, Tabet S, Grosz B, Dean M. SWRL: A Semantic Web Rule Language Combining OWL and RuleML. W3C; 2004. Available from: <https://www.w3.org/Submission/2004/SUBM-SWRL-20040521/>.
102. Wang J, Bansal M, Gimpel K, Ziebart BD, Yu CT. A Sense-Topic Model for Word Sense Induction with Unsupervised Data Enrichment. *Transactions of the Association for Computational Linguistics*. 2015; 3:59–71. [https://doi.org/10.1162/tacl\\_a\\_00122](https://doi.org/10.1162/tacl_a_00122)
103. Newman D, Hagedorn K, Chemudugunta C, Smyth P. Subject Metadata Enrichment Using Statistical Topic Models. In: *Proceedings of the 7th ACM/IEEE-CS Joint Conference on Digital Libraries*. JCDL'07. New York, NY, USA: ACM; 2007. p. 366–375.
104. Tuarob S, Pouchard LC, Giles CL. Automatic Tag Recommendation for Metadata Annotation Using Probabilistic Topic Modeling. In: *Proceedings of the 13th ACM/IEEE-CS Joint Conference on Digital Libraries*. JCDL'13. New York, NY, USA: ACM; 2013. p. 239–248.
105. Amazon Inc. Amazon Mechanical Turk. [cited 15 Jun 2020]. In: Amazon Web site [Internet]—. [about 4 screens]. Available from: <https://www.mturk.com>.
106. Hripcsak G, Rothschild AS. Agreement, the f-measure, and Reliability in Information Retrieval. *Journal of the American Medical Informatics Association*. 2005; 12(3):296–298. <https://doi.org/10.1197/jamia.M1733> PMID: 15684123
107. Garigliotti D, Hasibi F, Balog K. Identifying and exploiting target entity type information for ad hoc entity retrieval. *Information Retrieval Journal*. 2019; 22(3):285–323. <https://doi.org/10.1007/s10791-018-9346-x>
108. Tonon A, Catasta M, Prokofyev R, Demartini G, Aberer K, Cudrè-Mauroux P. Contextualized Ranking of Entity Types Based on Knowledge Graphs. *Web Semantics: Science, Services and Agents on the World Wide Web*. 2016; 37-38:170–183. <https://doi.org/10.1016/j.websem.2015.12.005>
109. Sönmez Ç, Özgür A. A Graph-based Approach for Contextual Text Normalization. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL; 2014. p. 313–324. Available from: <http://aclweb.org/anthology/D14-1037>.
110. Çelebi A, Özgür A. Segmenting hashtags and analyzing their grammatical structure. *Journal of the Association for Information Science and Technology*. 2018; 69(5):675–686. <https://doi.org/10.1002/asi.23989>
111. Zheng X, Han J, Sun A. A Survey of Location Prediction on Twitter. *IEEE Transactions on Knowledge and Data Engineering*. 2018; 30(9):1652–1671. <https://doi.org/10.1109/TKDE.2018.2807840>
112. Kumar A, Singh JP. Location reference identification from tweets during emergencies: A deep learning approach. *International Journal of Disaster Risk Reduction*. 2019; 33:365–375. <https://doi.org/10.1016/j.ijdr.2018.10.021>
113. Chi L, Lim KH, Alam N, Butler CJ. Geolocation Prediction in Twitter Using Location Indicative Words and Textual Features. In: *NUT@COLING, 2nd Workshop on Noisy User-generated Text*; 2016. p. 227–234.
114. Ozdakis O, Ramampiaro H, Nørnvåg K. Locality-adapted kernel densities of term co-occurrences for location prediction of tweets. *Information Processing & Management*. 2019; 56(4):1280–1299. <https://doi.org/10.1016/j.ipm.2019.02.013>
115. Chiu JPC, Nichols E. Named Entity Recognition with Bidirectional LSTM-CNNs. *Transactions of the Association for Computational Linguistics*. 2016; 4:357–370. [https://doi.org/10.1162/tacl\\_a\\_00104](https://doi.org/10.1162/tacl_a_00104)
116. Yamada I, Shindo H, Takeda H, Takefuji Y. Joint Learning of the Embedding of Words and Entities for Named Entity Disambiguation. *CoRR*. 2016;abs/1601.01343.
117. Francis-Landau M, Durrett G, Klein D. Capturing Semantic Similarity for Entity Linking with Convolutional Neural Networks. *CoRR*. 2016;abs/1604.00734.
118. Yin X, Huang Y, Zhou B, Li A, Lan L, Jia Y. Deep Entity Linking via Eliminating Semantic Ambiguity With BERT. *IEEE Access*. 2019; 7:169434–169445. <https://doi.org/10.1109/ACCESS.2019.2955498>
119. Sekine S, Sudo K, Nobata C. Extended Named Entity Hierarchy. In: *Proceedings of the Third International Conference on Language Resources and Evaluation (LREC'02)*. Las Palmas, Canary Islands—Spain: European Language Resources Association (ELRA); 2002. p. 1818–1824. Available from: <http://www.lrec-conf.org/proceedings/lrec2002/pdf/120.pdf>.
120. Ragunath R, Sivarajani N. Ontology based text document summarization system using concept terms. *ARPJ Eng Appl Sci*. 2015; 10:2638–2642.
121. Kalender M, Eren MT, Wu Z, Cirakman O, Kutluk S, Gultekin G, et al. Videolization: knowledge graph based automated video generation from web content. *Multimedia Tools Appl*. 2018; 77(1):567–595. <https://doi.org/10.1007/s11042-016-4275-4>
122. World Health Organization. WHO International Classification of Diseases. [cited 28 May 2020]. In WHO Web site [Internet]—. [about 2 screens].

123. SNOMED International. SNOMED CT website. [cited 28 May 2020]. In [Internet] -. [about 6 screens]. Available from: <http://www.snomed.org/snomed-ct/>.
124. Pham TL, Ali MI, Mileo A. Enhancing the scalability of expressive stream reasoning via input-driven parallelization. *Semantic Web Journal*. 2019; 10(3):457–474. <https://doi.org/10.3233/SW-180330>
125. Barbieri DF, Braga D, Ceri S, Della Valle E, Grossniklaus M. C-SPARQL: SPARQL for Continuous Querying. In: *Proceedings of the 18th International Conference on World Wide Web. WWW'09*. New York, NY, USA: ACM; 2009. p. 1061–1062.
126. Bonte P, Tommasini R, De Turck F, Ongenaë F, Valle ED. C-Sprite: Efficient Hierarchical Reasoning for Rapid RDF Stream Processing. In: *Proceedings of the 13th ACM International Conference on Distributed and Event-based Systems. DEBS'19*. New York, NY, USA: ACM; 2019. p. 103–114. Available from: <http://doi.acm.org/10.1145/3328905.3329502>.