

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From - To)	
		New Reprint		-	
4. TITLE AND SUBTITLE			5a. CONTRACT NUMBER		
Computational Fact Checking from Knowledge Networks			W911NF-12-1-0037		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHORS			5d. PROJECT NUMBER		
Giovanni Luca Ciampaglia , Prashant Shiralkar , Luis M. Rocha , Johan Bollen , Filippo Menczer , Alessandro Flammini					
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES			8. PERFORMING ORGANIZATION REPORT NUMBER		
Indiana University at Bloomington 509 E 3RD ST Bloomington, IN 47401 -3654					
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 61766-NS-DRP.57		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT Traditional fact checking by expert journalists cannot keep up with the enormous volume of information that is now generated online. Computational fact checking may significantly enhance our ability to evaluate the veracity of dubious information. Here we show that the complexities of human fact checking can be approximated quite well by finding the shortest path between concept nodes under properly defined semantic proximity metrics on knowledge graphs. Framed as a network problem this approach is feasible with efficient computa-					
15. SUBJECT TERMS fact-checking; knowledge graph; semantic proximity; transitive closure					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU		Alessandro Flammini
UU	UU	UU			19b. TELEPHONE NUMBER
					812-856-1830

Report Title

Computational Fact Checking from Knowledge Networks

ABSTRACT

Traditional fact checking by expert journalists cannot keep up with the enormous volume of information that is now generated online. Computational fact checking may significantly enhance our ability to evaluate the veracity of dubious information. Here we show that the complexities of human fact checking can be approximated quite well by finding the shortest path between concept nodes under properly defined semantic proximity metrics on knowledge graphs. Framed as a network problem this approach is feasible with efficient computational techniques. We evaluate this approach by examining tens of thousands of claims related to history, entertainment, geography, and biographical information using a public knowledge graph extracted from Wikipedia. Statements independently known to be true consistently receive higher support via our method than do false ones. These findings represent a significant step toward scalable computational fact-checking methods that may one day mitigate the spread of harmful misinformation.

REPORT DOCUMENTATION PAGE (SF298)
(Continuation Sheet)

Continuation for Block 13

ARO Report Number 61766.57-NS-DRP
Computational Fact Checking from Knowledge I...

Block 13: Supplementary Note

© 2015 . Published in PLoS ONE, Vol. Ed. 0 10, (6) (2015), (, (6). DoD Components reserve a royalty-free, nonexclusive and irrevocable right to reproduce, publish, or otherwise use the work for Federal purposes, and to authorize others to do so (DODGARS §32.36). The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.

Approved for public release; distribution is unlimited.

RESEARCH ARTICLE

Computational Fact Checking from Knowledge Networks

Giovanni Luca Ciampaglia^{1*}, Prashant Shiralkar¹, Luis M. Rocha^{1,2}, Johan Bollen¹, Filippo Menczer¹, Alessandro Flammini¹

1 Center for Complex Networks and Systems Research, Indiana University, Bloomington, Indiana, United States of America, **2** Instituto Gulbenkian de Ciencia, Oeiras, Portugal

* gciampag@indiana.edu



OPEN ACCESS

Citation: Ciampaglia GL, Shiralkar P, Rocha LM, Bollen J, Menczer F, Flammini A (2015) Computational Fact Checking from Knowledge Networks. PLoS ONE 10(6): e0128193. doi:10.1371/journal.pone.0128193

Academic Editor: Alain Barrat, Centre de Physique Théorique, FRANCE

Received: January 28, 2015

Accepted: April 24, 2015

Published: June 17, 2015

Copyright: © 2015 Ciampaglia et al. 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: All relevant data are within the paper and its Supporting Information files with the exception of the DBpedia files, which are available from the DBpedia website (<http://dbpedia.org>).

Funding: This work was supported in part by the Swiss National Science Foundation (fellowship 142353, <http://snf.ch>), the Lilly Endowment (<http://www.lillyendowment.org/>), the James S. McDonnell Foundation (<https://www.jsmf.org/>), the National Science Foundation (grant CCF-1101743, <http://www.nsf.gov/>), and Department of Defense (grant W911NF-12-1-0037, <http://www.darpa.mil/>). The

Abstract

Traditional fact checking by expert journalists cannot keep up with the enormous volume of information that is now generated online. Computational fact checking may significantly enhance our ability to evaluate the veracity of dubious information. Here we show that the complexities of human fact checking can be approximated quite well by finding the shortest path between concept nodes under properly defined semantic proximity metrics on knowledge graphs. Framed as a network problem this approach is feasible with efficient computational techniques. We evaluate this approach by examining tens of thousands of claims related to history, entertainment, geography, and biographical information using a public knowledge graph extracted from Wikipedia. Statements independently known to be true consistently receive higher support via our method than do false ones. These findings represent a significant step toward scalable computational fact-checking methods that may one day mitigate the spread of harmful misinformation.

Introduction

Online communication platforms, in particular social media, have created a situation in which the proverbial lie “can travel the world before the truth can get its boots on.” Misinformation [1], astroturf [2], spam [3], and outright fraud [4] have become widespread. They are now seemingly unavoidable components of our online information ecology [5] that jeopardize our ability as a society to make rapid and informed decisions [6–10].

While attempts to partially automate the detection of various forms of misinformation are burgeoning [11–15], automated reasoning methods are hampered by the inherent ambiguity of language and by deliberate deception. However, under certain conditions, reliable knowledge transmission can take place online [16]. For example, Wikipedia, the crowd-sourced online encyclopedia, has been shown to be nearly as reliable as traditional encyclopedias, even though it covers many more topics [17]. It now serves as a large-scale knowledge repository for millions of individuals, who can also contribute to its content in an open way. Vandalism, bias, distortions, and outright lies are frequently repaired in a matter of minutes [18]. Its continuous editing process even indicates signs of collective human intelligence [19].

funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

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

Here we show that we can leverage any collection of factual human knowledge, such as Wikipedia, for automatic fact checking [20]. Loosely inspired by the principle of epistemic closure [21], we computationally gauge the support for statements by mining the connectivity patterns on a knowledge graph. Our initial focus is on computing the support of simple statements of fact using a large-scale knowledge graph obtained from Wikipedia. More in general, fact checking can be seen as a special case of link prediction in knowledge graphs [22].

Knowledge Graphs

Let a *statement of fact* be represented by a subject-predicate-object triple, e.g., (“Socrates,” “is a,” “person”). A set of such triples can be combined to produce a *knowledge graph* (KG), where nodes denote *entities* (i.e., subjects or objects of statements), and edges denote predicates. Given a set of statements that has been extracted from a knowledge repository—such as the aforementioned Wikipedia—the resulting KG network represents all factual relations among entities mentioned in those statements. Given a new statement, we expect it to be true if it exists as an edge of the KG, or if there is a short path linking its subject to its object within the KG. If, however, the statement is untrue, there should be neither edges nor short paths that connect subject and object.

In a KG distinct paths between the same subject and object typically provide different factual support for the statement those nodes represent, even if the paths contain the same number of intermediate nodes. For example, paths that contain generic entities, such as “United States” or “Male,” provide weaker support because these nodes link to many entities and thus yield little specific information. Conversely, paths comprised of very specific entities, such as “positronic flux capacitor” or “terminal deoxynucleotidyl transferase,” provide stronger support. A fundamental insight that underpins our approach is that the definition of path length used for fact checking should account for such information-theoretic considerations.

To test our method we use the DBpedia database [23], which consists of all factual statements extracted from Wikipedia’s “infoboxes” (see Fig 1(a)). From this data we build the large-scale *Wikipedia Knowledge Graph* (WKG), with 3 million entity nodes linked by approximately 23 million edges (see [Materials and Methods](#)). Since we use only facts within infoboxes, the WKG contains the most uncontroversial information available on Wikipedia. This conservative approach is employed to ensure that our process relies as much as possible on a human-annotated, collectively-vetted factual basis. The WKG could be augmented with automatic methods to infer facts from text and other unstructured sources available online. Indeed, other teams have proposed methods to infer knowledge from text [24] to be employed in large and sophisticated rule-based inference models [24–26]. Here we focus on the feasibility of automatic fact checking using *simple* network models that leverage DBpedia. For this initial goal, we do not need to enhance the WKG, but such improvements can later be incorporated.

Semantic Proximity from Transitive Closure

Let the WKG be an undirected graph $G = (V, E)$ where V is a set of concept nodes and E is a set of predicate edges (see [Materials and Methods](#)). Two nodes $v, w \in V$ are said to be *adjacent* if there is an edge between them $(v, w) \in E$. They are said to be *connected* if there a sequence of $n \geq 2$ nodes $v = v_1, v_2, \dots, v_n = w$, such that, for $i = 1, \dots, n-1$ the nodes v_i and v_{i+1} are adjacent. The *transitive closure* of G is $G^* = (V, E^*)$ where the set of edges is closed under adjacency, that is, two nodes are adjacent in G^* iff they are connected in G via at least one path. This standard notion of closure has been extended to weighted graphs, allowing adjacency to be generalized by measures of path length [27], such as the semantic proximity for the WKG we introduce next.

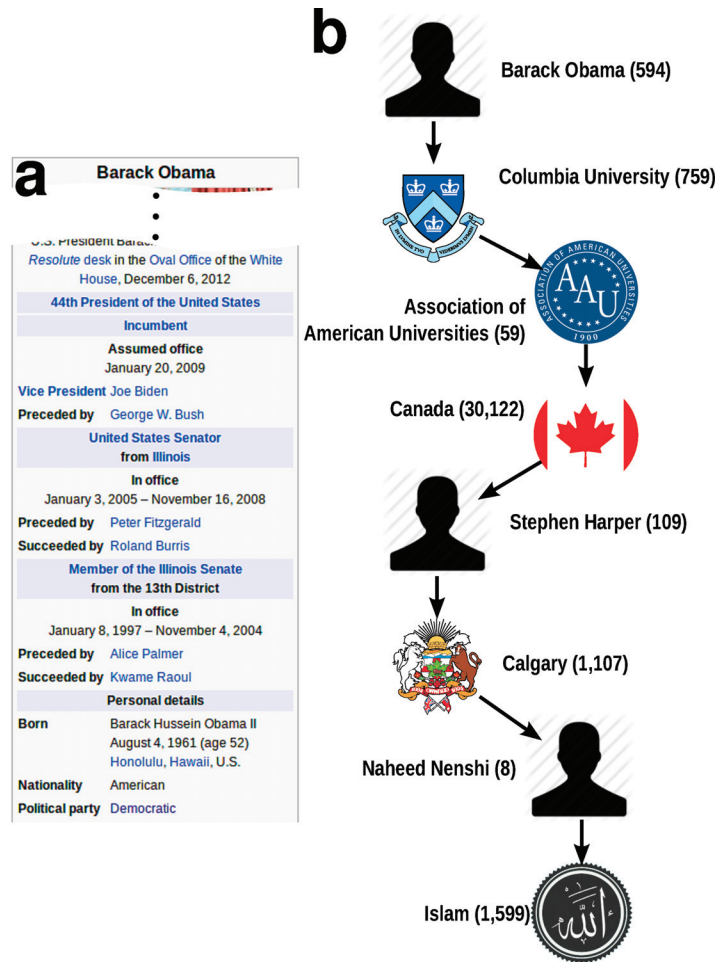


Fig 1. Using Wikipedia to fact-check statements. (a) To populate the knowledge graph with facts we use structured information contained in the ‘infoboxes’ of Wikipedia articles (in the figure, the infobox of the article about *Barack Obama*). (b) Using the Wikipedia Knowledge Graph, computing the truth value of a subject-predicate-object statement amounts to finding a path between subject and object entities. In the diagram we plot the shortest path returned by our method for the statement “*Barack Obama is a muslim.*” Numbers in parentheses indicate the degree of the nodes. The path traverses high-degree nodes representing generic entities, such as *Canada*, and is assigned a low truth value.

doi:10.1371/journal.pone.0128193.g001

The truth value $\tau(e) \in [0, 1]$ of a new statement $e = (s, p, o)$ is derived from a transitive closure of the WKG. More specifically, the truth value is obtained via a path evaluation function: $\tau(e) = \max \mathcal{W}(P_{s,o})$. This function maps the set of possible paths connecting s and o to a truth value τ . A path has the form $P_{s,o} = v_1 v_2 \dots v_n$, where v_i is an entity node, (v_i, v_{i+1}) is an edge, n is the path length measured by the number of its constituent nodes, $v_1 = s$, and $v_n = o$. Various characteristics of a path can be taken as evidence in support of the truth value of e . Here we use the *generality* of the entities along a path as a measure of its length, which is in turn aggregated to define a *semantic proximity*:

$$\mathcal{W}(P_{s,o}) = \mathcal{W}(v_1 \dots v_n) = \left[1 + \sum_{i=2}^{n-1} \log k(v_i) \right]^{-1} \tag{1}$$

where $k(v)$ is the degree of entity v , i.e., the number of WKG statements in which it participates; it therefore measures the generality of an entity. If e is already present in the WKG (i.e., there is an edge between s and o), it should obviously be assigned maximum truth. In fact $\mathcal{W} = 1$ when $n = 2$ because there are no intermediate nodes. Otherwise an indirect path of length $n > 2$ may be found via other nodes. The truth value $\tau(e)$ maximizes the semantic proximity defined by Eq 1, which is equivalent to finding the shortest path between s and o [27], or the one that provides the maximum information content [28, 29] in the WKG. The transitive closure of weighted graphs equivalent to finding the shortest paths between every pair of nodes is also known as the *metric closure* [27]. This approach is also related to the Path Ranking Algorithm [30], except that here we use the shortest path (equivalent to maximum probability) rather than combining a sample of bounded-length paths in a learning framework.

Fig 1(b) depicts an example of a shortest path on the WKG for a statement that yields a low truth value. Note that in this specific formulation we disregard the semantics of the predicate, therefore we are only able to check statements with the simplest predicates, such as “is a”; negation, for instance, would require a more sophisticated definition of path length.

Alternative definitions of $\tau(e)$ are of course possible. Instead of shortest paths, one could use a different optimization principle, such as widest bottleneck, also known as the *ultra-metric closure* [27], which corresponds to maximizing the path evaluation function \mathcal{W}_u :

$$\mathcal{W}_u(P_{s,o}) = \mathcal{W}_u(v_1 \dots v_n) = \begin{cases} 1 & n = 2 \\ [1 + \max_{i=2}^{n-1} \{\log k(v_i)\}]^{-1} & n > 2. \end{cases} \quad (2)$$

Or it could be possible to retain the original directionality of edges and have a directed WKG instead of an undirected one. As described next, we evaluated alternative definitions of $\tau(e)$ and found Eq 1 to perform best.

Results

Calibration

Our fact-checking method requires that we define a measure of path semantic proximity by selecting a transitive closure algorithm (the shortest paths of Eq 1 or the widest bottleneck paths of Eq 2) and a directed or undirected WKG representation. To evaluate these four combinations empirically, let us attempt to infer the party affiliation of US Congress members. In other words, we want to compute the support of statements like “ x is a member of y ” where x is a member of Congress and y is a political party. We consider all members of the 112th US Congress that are affiliated with either the Democratic or Republican party (Senate: $N = 100$; House: $N = 445$). We characterize each member of Congress with its semantic proximity to all nodes in the WKG that represent ideologies. This yields an $N \times M$ feature matrix \mathcal{F}_{tc} for each of the four transitive closure methods. Panel (a) of Fig 2 illustrates the proximity network obtained from \mathcal{F}_{tc} that connects members of the 112th Congress and their closest ideologies, as computed using Eq 1. A high degree of ideological polarization can be observed in the WKG, consistent with blogs [31] and social media [32].

We feed \mathcal{F}_{tc} into off-the-shelf classifiers (see Materials and Methods). As shown in Table 1, the metric closure on the undirected graph gives the most accurate results. Therefore, we continue to use this combination in our semantic proximity computations when performing the validation tasks described below.

To evaluate the overall performance of the calibrated model, we also compared it against DW-NOMINATE, the state of the art in political classification [34]. This model is not based on data from a knowledge graph, but on explicit information about roll-call voting patterns.

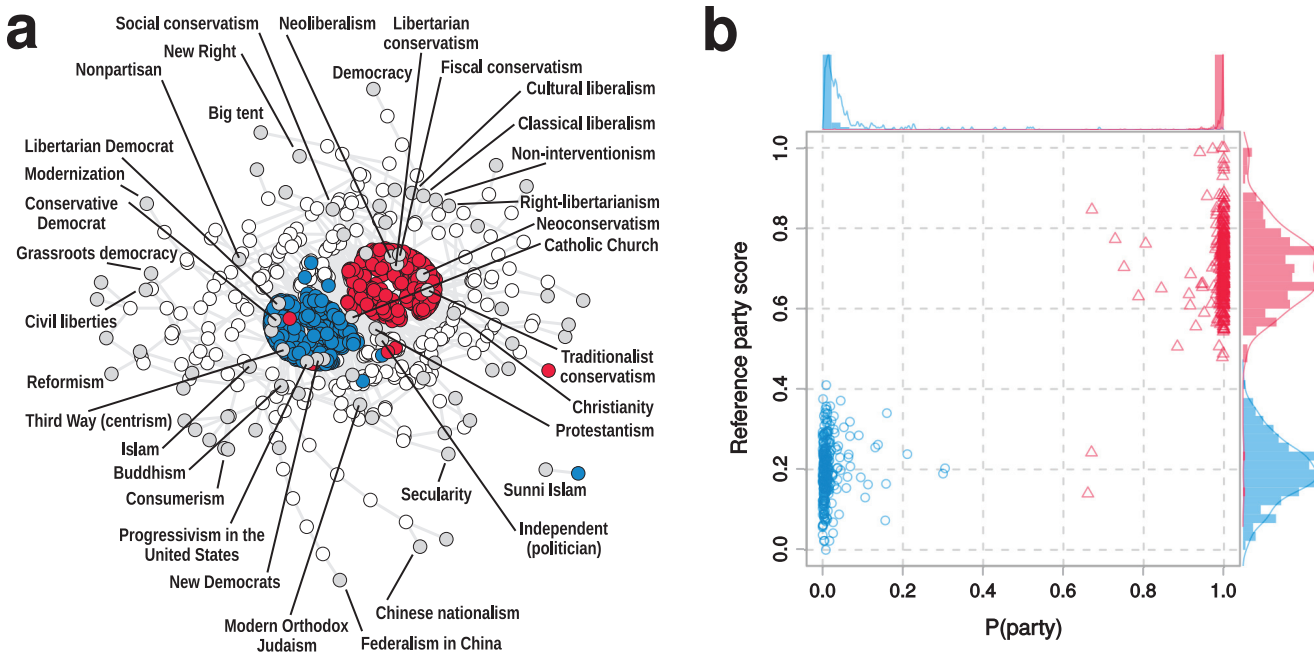


Fig 2. Ideological classification of the US Congress based on truth values. (a) Ideological network of the 112th US Congress. The plot shows a subset of the WKG constituted by paths between Democratic or Republican members of the 112th US Congress and various ideologies. Red and blue nodes correspond to members of Congress, gray nodes to ideologies, and white nodes to vertices of any other type. The position of the nodes is computed using a force-directed layout [33], which minimizes the distance between nodes connected by an edge weighted by a higher truth value. For clarity only the most significant paths, whose values rank in the top 1% of truth values, are shown. (b) Ideological classification of members of the 112th US Congress. The plot shows on the x axis the party label probability given by a Random Forest classification model trained on the truth values computed on the WKG, and on the y axis the reference score provided by DW-NOMINATE. Red triangles are members of Congress affiliated to the Republican party and blue circles to the Democratic party. Histograms and density estimates of the two marginal distributions, color-coded by actual affiliation, are shown on the top and right axes.

doi:10.1371/journal.pone.0128193.g002

Table 1. Transitive closure calibration.

		Directed		Undirected	
		k-NN	RF	k-NN	RF
House	Metric	96	99	97	99
House	Ultra-metric	56	57	53	57
Senate	Metric	70	100	96	100
Senate	Ultra-metric	49	39	70	61

Area under Receiver Operating Characteristic (ROC) curve of two classifiers, random forests (RF) and *k*-nearest neighbors (*k*-NN) on the ideological classification task.

doi:10.1371/journal.pone.0128193.t001

Table 2. Ideological classification results.

Dataset	RF		k-NN	
	F-score	AUROC	F-score	AUROC
TRANSITIVE CLOSURE \mathcal{F}_{tc}				
Senate	0.99	1.00	0.91	0.96
House	0.99	1.00	0.90	0.97
INFOBOXES \mathcal{F}_b				
Senate	0.66	0.46	0.62	0.54
House	0.54	0.66	0.68	0.54

Out-of-sample F-score and Area Under Receiver Operating Characteristic (AUROC) of random forest (RF) and k -nearest neighbors (k -NN) classifiers trained on truth scores computed by the fact checker, using either the transitive closure or solely information from infoboxes.

doi:10.1371/journal.pone.0128193.t002

Comparing our classification results with such a baseline is also useful to gauge the quality of the latent information contained in the WKG for the task of political classification. As shown in panel (b) of Fig 2, a Random Forests classifier trained on our truth values matches the performance of DW-NOMINATE.

Value of indirect connections

Most of the WKG information that our fact checker exploits is provided by indirect paths (i.e., comprising $n > 2$ nodes). To demonstrate this, we compare the calibrated model of Eq 1 to the fact checker’s performance with only the information in the infoboxes.

In practice, we compute an additional feature matrix \mathcal{F}_b , using the same sequence of steps outlined in the calibration phase, but additionally constraining the shortest path algorithm to use only paths (if any) with exactly $n = 2$ nodes, i.e., direct edges. Thus \mathcal{F}_b encodes only the information of the infoboxes of the politicians. The results from 10-fold cross validation using \mathcal{F}_{tc} and \mathcal{F}_b are shown in Table 2. The same off-the-shelf classifiers, this time trained on \mathcal{F}_b , perform only slightly better than random, thus confirming that the truth signal is yielded by the structure of indirect connections in the WKG.

Validation on factual statements

We test our fact-checking method on tasks of increasing difficulty, and begin by considering simple factual statements in four subject areas related to entertainment, history, and geography. We evaluate statements of the form “ d_i directed m_j ,” “ p_i was married to s_j ,” and “ c_i is the capital of r_j ,” where d_i is a director, m_j is a movie, p_i is a US president, s_j is the spouse of a US president, c_i is a city, and r_j is a country or US state. By considering all combinations of subjects and objects in these classes, we obtain matrices of statements (see Materials and Methods). Many of them, such as “Rome is the capital of India,” are false. Others, such as “Rome is the capital of Italy,” are true. To prevent the task from being trivially easy, we remove any edges that represent true statements in our test set from the graph. Fig 3 shows the matrices obtained by running the fact checker on the factual statements. Let e and e' be a true and false statement, respectively, from any of the four subject areas. To show that our fact checker is able to correctly discriminate between true and false statements with high accuracy, we estimate the probability that $\tau(e) > \tau(e')$. To do so we plot the ROC curve of the classifier (see Fig 4) since the area

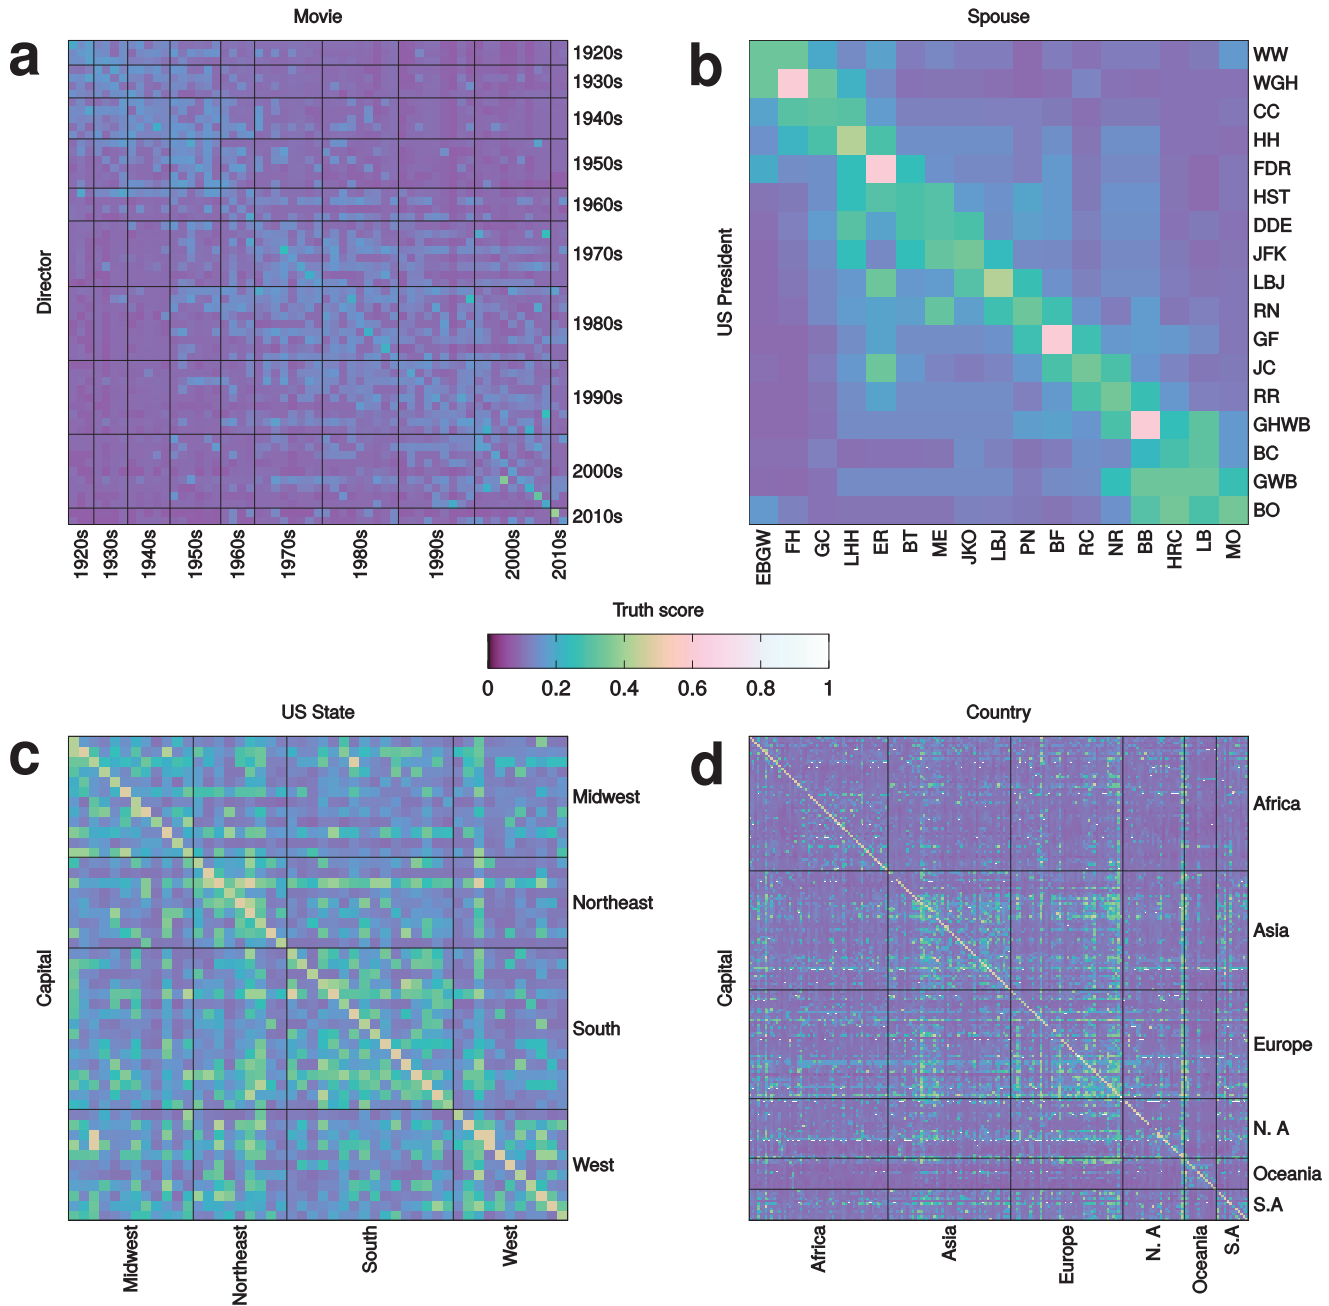


Fig 3. Automatic truth assessments for simple factual statements. In each confusion matrix, rows represent subjects and columns represent objects. The diagonals represent true statements. Higher truth values are mapped to colors of increasing intensity. **(a)** Films winning the Oscar for Best Movie and their directors, grouped by decade of award (see the complete list in the [S1 Text](#)). **(b)** US presidents and their spouses, denoted by initials. **(c)** US states and their capitals, grouped by US Census Bureau-designated regions. **(d)** World countries and their capitals, grouped by continent.

doi:10.1371/journal.pone.0128193.g003

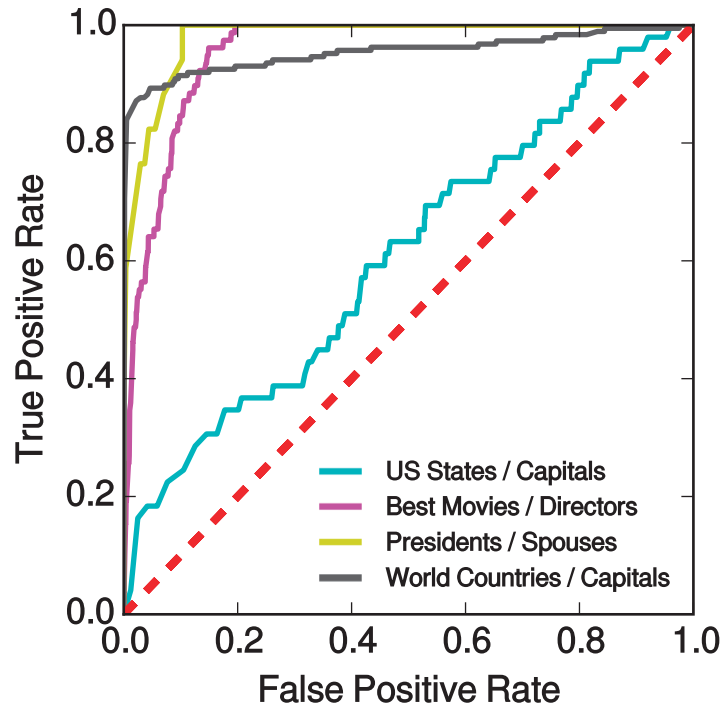


Fig 4. Receiver Operating Characteristic for the multiple questions task. For each confusion matrix depicted in Fig 3 we compute ROC curves where true statements correspond to the diagonal and false statements to off-diagonal elements. The red dashed line represents the performance of a random classifier.

doi:10.1371/journal.pone.0128193.g004

under the ROC curve is equivalent to this probability [35]. With this method we estimate that, in the four subject areas, true statements are assigned higher truth values than false ones with probability 95%, 98%, 61%, and 95%, respectively.

Validation on annotated corpus

In a second task, we consider an independent corpus of novel statements extracted from the free text of Wikipedia and annotated as true or false by human raters [36] (see [Materials and Methods](#)). We compare the human ratings with the truth values provided by our automatic fact checker (Fig 5). Although the statements under examination originate from Wikipedia, they are not usually represented in the WKG, which is derived from the infoboxes only. When a statement is present in the WKG, the link is removed. The information available in the WKG about the entities involved in these particular statements is very sparse, therefore this task is more difficult than the previous case.

We find that the truth values computed by the fact checker are positively correlated to the average ratings given by the human evaluators. Table 3 shows the positive correlation between GREC human annotations and our computational truth scores.

As shown in Fig 5, our fact checker yields consistently higher support for true statements than false ones. Using only information in the infoboxes however yields worse performance, closer to random choice: AUROC = 0.47 and 0.52 for the ‘degree’ and ‘institution’ predicates, respectively. We conclude that the fact checker is able to integrate the strength of indirect paths in the WKG, which pertain to factual information not originally included in the infoboxes.

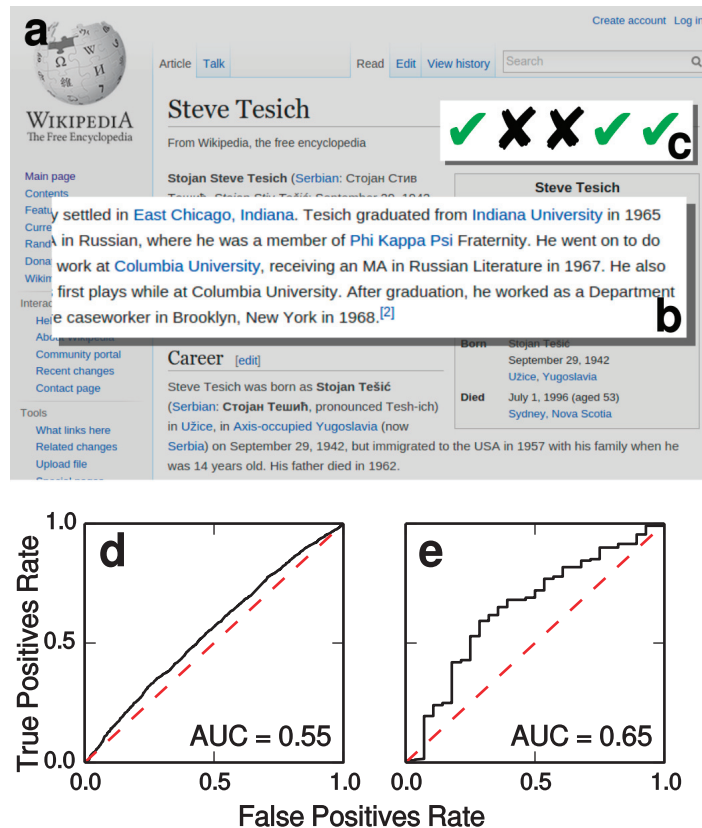


Fig 5. Real-world fact-checking scenario. (a) A document from the ground truth corpus. (b) Statement to fact-check: *Did Steve Tesich graduate from Indiana University, Bloomington?* This information is not present in the infobox, and thus it is not part of the WKG. (c) Annotations from five human raters. In this case, the majority of raters believe that the statement is true, and thus we consider it as such for classification purposes. (d) Receiver operating characteristic (ROC) curve of the classification for subject-predicate-object statements in which the predicate is “institution” (e.g., “Albert Einstein,” “institution,” “Institute for Advanced Studies”). A true positive rate above the false positive rate (dashed line), and correspondingly an area under the curve (AUC) above 0.5, indicate better than random performance. (e) ROC curve for statements with “degree” predicate (e.g., “Albert Einstein,” “degree,” “University Diploma”).

doi:10.1371/journal.pone.0128193.g005

Table 3. Agreement between fact checker and human raters.

Relation	ρ	p -value	τ	p -value
Degree	0.17	2×10^{-5}	0.13	10×10^{-6}
Institution	0.09	4×10^{-19}	0.07	1×10^{-24}

We use rank-order correlation coefficients (Kendall’s τ and Spearman’s ρ) to assess whether the scores are correlated to the ratings. Significance tests rule out the null hypothesis that the correlation coefficients are zero.

doi:10.1371/journal.pone.0128193.t003

Discussion

These results are both encouraging and exciting: a simple shortest path computation maximizing information content can leverage an existing body of collective human knowledge to assess the truth of new statements. In other words, the important and complex human task of fact checking can be effectively reduced to a simple network analysis problem, which is easy to solve computationally. Our approach exploits implicit information from the topology of the WKG, which is different from the statements explicitly contained in the infoboxes. Indeed, if we base our assessment only on direct edges in the WKG, performance decreases significantly. This demonstrates that much of the correct measurement of the truthfulness of statements relies on indirect paths. Because there are many ways to compute shortest paths in distance graphs, or transitive closures in weighted graphs [27], there is ample room for improvement on this method.

Our WKG is built from statement of facts, which are represented as subject-predicate-object triples, i.e. information with an inherent directionality. Our results show that an undirected KG yielded the best outcomes [37, 38]. This is somehow surprising, given that in transforming a directed graph into undirected we are destroying potentially useful information. However, while some semantic relations are inherently one-way, it can be argued that some relations can be navigated in both directions (e.g. “Barack Obama” “Married-To” “Michelle Obama”) [39]. Thus we conjecture that the loss of information from disposing of the direction of edges is balanced by the possibility of finding paths, and hope that future research will elucidate this conjecture.

We live in an age of overabundant and ever-growing information, but much of it is of questionable veracity [10, 40]. Establishing the reliability of information in such circumstances is a daunting but critical challenge. Our results show that network analytics methods, in conjunction with large-scale knowledge repositories, offer an exciting new opportunity towards automatic fact-checking methods. As the importance of the Internet in our everyday lives grows, misinformation such as panic-inducing rumors, urban legends, and conspiracy theories can efficiently spread online in variety of new ways [5, 8]. Scalable computational methods, such as the one we demonstrate here, may hold the key to mitigate the societal effects of these novel forms of misinformation.

Materials and Methods

Wikipedia Knowledge Graph

To obtain the WKG we downloaded and parsed RDF triples data from the DBpedia project (dbpedia.org). We used three datasets of triples to build the WKG: the “Types” dataset, which contains subsumption triples of the form (subject, “is-a,” Class), where Class is a category of the DBpedia ontology; the “Properties” dataset, which contains triples extracted from infoboxes; and the DBpedia ontology, from which we used all triples with predicate “subClassOf.” This last data was used to reconstruct the full ontological hierarchy of the graph. We then discarded the predicate part of each triple and conflated all triples having the same subject and object, obtaining an edge list. In this process, we discarded all triples whose subject or object belonged to external namespaces (e.g., FOAF and schema.org). We also discarded all triples from the “Properties” dataset whose object was a date or any other kind of measurement (e.g., “Aristotle,” “birthYear,” “384 B.C.”), because by definition they never appear as subjects in other triples.

Ideological classification of the US Congress

To get a list of ideologies we consider the “Ideology” category in the DBpedia ontology and look up in the WKG all nodes Y connected to it by means of a statement (Y , “is-a,” “Ideology”). We found $M = 819$ such nodes (see [S1 Text](#) for the complete list). Given a politician X and an ideology Y we then compute the truth value of the statement “ X endorses ideology Y .” To perform the classification, we use two standard classifier algorithms: k -Nearest Neighbors [41] and Random Forests [42]. To assess the classification accuracy we computed F-score and area under Receiver Operating Characteristic (ROC) curve using 10-fold cross-validation.

Simple factual statements

We formed simple statements by combining each of N subject entities with each of N object entities. We performed this procedure in four subject areas: (1) Academy Awards for Best Movie ($N = 59$), (2) US presidential couples ($N = 17$), (3) US state capitals ($N = 48$), and (4) world capitals ($N = 187$). For directors with more than one award, only the first award was used. All data were taken from Wikipedia (see [S1 Text](#) for data tables). To make the test fair, if a triple indicating a true statement was already present in the WKG, we removed it from the graph before computing the truth value. This step of the evaluation procedure is typical of link prediction algorithms [43].

Independent corpus of statements

The second ground truth dataset is based on the Google Relation Extraction Corpus (GREC) [36]. For simplicity we focus on two types of statements, about education degrees ($N = 602$) and institutional affiliations ($N = 10,726$) of people, respectively. Each triple in the GREC comes with truth ratings by five human raters ([Fig 5\(c\)](#)), so we map the ratings into an ordinal scale between -5 (all raters replied ‘No’) and $+5$ (all raters replied ‘Yes’), and compare them to the truth values computed by the fact checker. The subject entities of several triples in the GREC appear in only a handful of links in the WKG, limiting the chances that our method can find more than one path. Therefore we select from the two datasets only triples having a subject with degree $k > 3$. Similarly to the previous task, if the statement is already present in the WKG, we remove the corresponding triple before computing the truth value.

Supporting Information

S1 Text. Data tables and list of ideologies. data tables for [Fig 3](#) and list of ideologies used in the ideological classification of the US Congress. (PDF)

Acknowledgments

The authors would like to thank Karissa McKelvey for original exploratory work and Deborah Rocha for editing the manuscript. We acknowledge the Wikimedia Foundation, the DBpedia project, and Google for making their respective data freely available.

Author Contributions

Conceived and designed the experiments: GLC PS LMR FM JB AF. Performed the experiments: GLC PS. Analyzed the data: GLC PS. Contributed reagents/materials/analysis tools: GLC PS. Wrote the paper: GLC LMR FM JB AF.

References

1. Mendoza M, Poblete B, Castillo C. Twitter Under Crisis: Can We Trust What We RT? In: Proceedings of the First Workshop on Social Media Analytics. SOMA'10. New York, NY, USA: ACM; 2010. p. 71–79.
2. Ratkiewicz J, Conover M, Meiss M, Gonçaves B, Flammini A, Menczer F. Detecting and Tracking Political Abuse in Social Media. In: Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media. Barcelona, Spain: AAAI; 2011.
3. Cranor LF, LaMacchia BA. Spam! Commun ACM. 1998 Aug; 41(8):74–83. doi: [10.1145/280324.280336](https://doi.org/10.1145/280324.280336)
4. Jagatic TN, Johnson NA, Jakobsson M, Menczer F. Social Phishing. Commun ACM. 2007 Oct; 50(10):94–100. doi: [10.1145/1290958.1290968](https://doi.org/10.1145/1290958.1290968)
5. Friggeri A, Adamic LA, Eckles D, Cheng J. Rumor Cascades. In: Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media. Ann Arbor, MI: AAAI; 2014.
6. Flanagin AJ, Metzger MJ. Perceptions of Internet Information Credibility. Journalism & Mass Communication Quarterly. 2000; 77(3):515–540. doi: [10.1177/107769900007700304](https://doi.org/10.1177/107769900007700304)
7. Rieh SY, Danielson DR. Credibility: A multidisciplinary framework. Annual Review of Information Science and Technology. 2007; 41(1):307–364. doi: [10.1002/aris.2007.1440410114](https://doi.org/10.1002/aris.2007.1440410114)
8. Kata A. A postmodern Pandora's box: Anti-vaccination misinformation on the Internet. Vaccine. 2010; 28(7):1709–1716. doi: [10.1016/j.vaccine.2009.12.022](https://doi.org/10.1016/j.vaccine.2009.12.022) PMID: [20045099](https://pubmed.ncbi.nlm.nih.gov/20045099/)
9. Castillo C, Mendoza M, Poblete B. Information credibility on Twitter. In: Proceedings of the 20th international conference on World Wide Web. Hyderabad, India: ACM; 2011. p. 675.
10. Lewandowsky S, Ecker UKH, Seifert CM, Schwarz N, Cook J. Misinformation and Its Correction: Continued Influence and Successful Debiasing. Psychological Science in the Public Interest. 2012; 13(3):106–131. doi: [10.1177/1529100612451018](https://doi.org/10.1177/1529100612451018)
11. Wilner T. Meet the robots that factcheck. Columbia Journalism Review. 2014 September/October;.
12. Gupta A, Kumaraguru P, Castillo C, Meier P. TweetCred: A Real-time Web-based System for Assessing Credibility of Content on Twitter. In: Proc. 6th International Conference on Social Informatics (SocInfo). Barcelona, Spain; 2014.
13. Resnick P, Carton S, Park S, Shen Y, Zeffer N. RumorLens: A System for Analyzing the Impact of Rumors and Corrections in Social Media. In: Proc. Computational Journalism Conference; 2014.
14. Wu Y, Agarwal PK, Li C, Yang J, Yu C. Toward Computational Fact-Checking. In: Proceedings of the VLDB Endowment. vol. 7. Hangzhou, China; 2014.
15. Finn S, Metaxas PT, Mustafaraj E, O'Keefe M, Tang L, Tang S, et al. TRAILS: A System for Monitoring the Propagation of Rumors On Twitter. In: Proc. Computational Journalism Conference; 2014.
16. Berners-Lee T, Hendler J, Lassila O. The Semantic Web. Scientific American. 2001; 284(5):28–37. doi: [10.1038/scientificamerican0501-34](https://doi.org/10.1038/scientificamerican0501-34)
17. Giles J. Internet encyclopaedias go head to head. Nature. 2005; 438(7070):900–901. doi: [10.1038/438900a](https://doi.org/10.1038/438900a) PMID: [16355180](https://pubmed.ncbi.nlm.nih.gov/16355180/)
18. Priedhorsky R, Chen J, Lam STK, Panciera K, Terveen L, Riedl J. Creating, Destroying, and Restoring Value in Wikipedia. In: Proceedings of the 2007 International ACM Conference on Supporting Group Work. GROUP'07. New York, NY, USA: ACM; 2007. p. 259–268.
19. DeDeo S. Collective Phenomena and Non-Finite State Computation in a Human Social System. PLoS ONE. 2013; 8(10):e75818. doi: [10.1371/journal.pone.0075818](https://doi.org/10.1371/journal.pone.0075818) PMID: [24130745](https://pubmed.ncbi.nlm.nih.gov/24130745/)
20. Cohen S, Hamilton JT, Turner F. Computational Journalism. Commun ACM. 2011 Oct; 54(10):66–71. doi: [10.1145/2001269.2001288](https://doi.org/10.1145/2001269.2001288)
21. Luper S. The Epistemic Closure Principle. In: Zalta EN, editor. The Stanford Encyclopedia of Philosophy. Fall 2012 ed. The Metaphysics Research Lab, Stanford University; 2012.
22. Nickel M, Murphy K, Tresp V, Gabrilovich E. A Review of Relational Machine Learning for Knowledge Graphs: From Multi-Relational Link Prediction to Automated Knowledge Graph Construction. arXiv preprint arXiv:150300759. 2015;.
23. Auer S, Bizer C, Kobilarov G, Lehmann J, Cyganiak R, Ives Z. DBpedia: A Nucleus for a Web of Open Data. In: Aberer K, Cudré-Mauroux P, Choi KS, Noy N, Allemang D, Lee KI, et al., editors. The Semantic Web. vol. 4825 of Lecture Notes in Computer Science. Springer Berlin Heidelberg; 2007. p. 722–735.

24. Dong X, Gabrilovich E, Heitz G, Horn W, Lao N, Murphy K, et al. Knowledge Vault: A Web-scale Approach to Probabilistic Knowledge Fusion. In: Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD'14. New York, NY, USA: ACM; 2014. p. 601–610.
25. Etzioni O, Banko M, Soderland S, Weld DS. Open Information Extraction from the Web. *Commun ACM*. 2008 Dec; 51(12):68–74. doi: [10.1145/1409360.1409378](https://doi.org/10.1145/1409360.1409378)
26. Niu F, Zhang C, Ré C, Shavlik J. Elementary: Large-scale knowledge-base construction via machine learning and statistical inference. *International Journal on Semantic Web and Information Systems*. 2012; 8(3):42–73. doi: [10.4018/jswis.2012070103](https://doi.org/10.4018/jswis.2012070103)
27. Simas T, Rocha LM. Distance Closures on Complex Networks. *Network Science*. 2014;p. In Press. (arXiv:1312.2459).
28. Markines B, Menczer F. A Scalable, Collaborative Similarity Measure for Social Annotation Systems. In: Proceedings of the 20th ACM Conference on Hypertext and Hypermedia. HT'09. New York, NY, USA: ACM; 2009. p. 347–348.
29. Aiello L, Barrat A, Schifanella R, Cattuto C, Markines B, Menczer F. Friendship prediction and homophily in social media. *ACM Trans WEB*. 2012; 6(2):9. doi: [10.1145/2180861.2180866](https://doi.org/10.1145/2180861.2180866)
30. Lao N, Mitchell T, Cohen WW. Random Walk Inference and Learning in a Large Scale Knowledge Base. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing. EMNLP'11. Stroudsburg, PA, USA: Association for Computational Linguistics; 2011. p. 529–539.
31. Adamic LA, Glance N. The Political Blogosphere and the 2004 U.S. Election: Divided They Blog. In: Proceedings of the 3rd International Workshop on Link Discovery. ACM, Chicago, Illinois; 2005. p. 36–43.
32. Conover M, Ratkiewicz J, Francisco M, Gonçaves B, Flammini A, Menczer F. Political Polarization on Twitter. In: Proc. 5th International AAAI Conference on Weblogs and Social Media (ICWSM). Barcelona, Spain: AAAI; 2011.
33. Kamada T, Kawai S. An algorithm for drawing general undirected graphs. *Information Processing Letters*. 1989; 31(1):7–15. doi: [10.1016/0020-0190\(89\)90102-6](https://doi.org/10.1016/0020-0190(89)90102-6)
34. Poole KT, Rosenthal H. *Ideology and Congress: A Political Economic History of Roll Call Voting*. Transaction Publishers; 2007.
35. Fawcett T. An Introduction to ROC Analysis. *Pattern Recogn Lett*. 2006 Jun; 27(8):861–874. doi: [10.1016/j.patrec.2005.10.010](https://doi.org/10.1016/j.patrec.2005.10.010)
36. Orr D. 50,000 Lessons on How to Read: a Relation Extraction Corpus; 2013. (Date of access: 01/04/2014). Available at: <http://goo.gl/FBuP20>.
37. Masucci AP, Kalampokis A, Eguíluz VM, Hernández-García E. Wikipedia information flow analysis reveals the scale-free architecture of the semantic space. *PLoS ONE*. 2011; 6(2):e17333. doi: [10.1371/journal.pone.0017333](https://doi.org/10.1371/journal.pone.0017333) PMID: [21407801](https://pubmed.ncbi.nlm.nih.gov/21407801/)
38. Capitán JA, Borge-Holthoefer J, Gómez S, Martínez-Romo J, Araujo L, Cuesta JA, et al. Local-based semantic navigation on a networked representation of information. *PLoS ONE*. 2012; 7(8):e43694. doi: [10.1371/journal.pone.0043694](https://doi.org/10.1371/journal.pone.0043694) PMID: [22937081](https://pubmed.ncbi.nlm.nih.gov/22937081/)
39. Steyvers M, Tenenbaum JB. The Large-Scale Structure of Semantic Networks: Statistical Analyses and a Model of Semantic Growth. *Cognitive Science*. 2005; 29(1):41–78. doi: [10.1207/s15516709cog2901_3](https://doi.org/10.1207/s15516709cog2901_3) PMID: [21702767](https://pubmed.ncbi.nlm.nih.gov/21702767/)
40. Nyhan B, Reifler J, Ubel PA. The Hazards of Correcting Myths About Health Care Reform. *Medical Care*. 2013; 51(2):127–132. doi: [10.1097/MLR.0b013e318279486b](https://doi.org/10.1097/MLR.0b013e318279486b) PMID: [23211778](https://pubmed.ncbi.nlm.nih.gov/23211778/)
41. Bishop CM. *Pattern recognition and machine learning*. Springer-Verlag; 2006.
42. Breiman L. Random forests. *Machine learning*. 2001; 45(1):5–32. doi: [10.1023/A:1010933404324](https://doi.org/10.1023/A:1010933404324)
43. Liben-Nowell D, Kleinberg J. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology*. 2007 May; 58(7):1019–1031. doi: [10.1002/asi.20591](https://doi.org/10.1002/asi.20591)