

See discussions, stats, and author profiles for this publication at: <http://www.researchgate.net/publication/282200688>

Interlinking Documents Based on Semantic Graphs with an Application

CHAPTER · JANUARY 2015

DOI: 10.1007/978-3-319-13545-8_9

READS

92

6 AUTHORS, INCLUDING:



Bernardo Pereira Nunes

Pontifícia Universidade Católica do Rio de Ja...

53 PUBLICATIONS 97 CITATIONS

SEE PROFILE



Ricardo Kawase

Forschungszentrum L3S

60 PUBLICATIONS 118 CITATIONS

SEE PROFILE



Stefan Dietze

Leibniz Universität Hannover

100 PUBLICATIONS 416 CITATIONS

SEE PROFILE



Marco A. Casanova

Pontifícia Universidade Católica do Rio de Ja...

251 PUBLICATIONS 1,794 CITATIONS

SEE PROFILE

K14is – 76

Interlinking documents based on semantic graphs with an application

Bernardo Pereira Nunes, Besnik Fetahu, Ricardo Kawase, Stefan Dietze, Marco Antonio Casanova, and Diana Maynard

Abstract Connectivity and relatedness of Web resources are two concepts that define to what extent different parts are connected or related to one another. Measuring connectivity and relatedness between Web resources is a growing field of research, often the starting point of recommender systems. Although relatedness is liable to subjective interpretations, connectivity is not. Given the Semantic Web's ability of linking Web resources, connectivity can be measured by exploiting the links between entities. Further, these connections can be exploited to uncover relationships between Web resources. This chapter describes the application and expansion of a relationship assessment methodology from social network theory to measure the connectivity between documents. The connectivity measures are used to identify connected and related Web resources. The approach is able to expose relations that traditional text-based approaches fail to identify. The proposed approaches are validated and assessed through an evaluation on a real-world dataset, where results show that the proposed techniques outperform state of the art approaches. Finally, a Web-based application called *Cite4Me* that uses the proposed approach is presented.

Key words: Document connectivity, semantic connections, semantic graphs, document recommendation

Bernardo Pereira Nunes and Marco Antonio Casanova
Department of Informatics, Pontifical Catholic University of Rio de Janeiro/RJ – Brazil
e-mail: {bnunes, casanova}@inf.puc-rio.br

Ricardo Kawase, Besnik Fetahu and Stefan Dietze
L3S Research Center, Leibniz University Hannover, Appelstr. 9a, 30167 Hannover, Germany
e-mail: {kawase, fetahu, dietze}@l3s.de

Diana Maynard
Department of Computer Science, University of Sheffield, Sheffield, UK
e-mail: diana@dcs.shef.ac.uk

76.1 Introduction

User-generated content is characterized by a high degree of diversity and heavily varying quality. Given the ever increasing pace at which this form of Web content is evolving, adequate preservation and detection of correlations has become a cultural necessity. Extraction of entities from Web content, in particular social media, is a crucial challenge in order to enable the interlinking of related Web content, semantic search and navigation within Web archives, and to assess the relevance of a given set of Web objects for a particular query or crawl. Earlier work by Dietze et al. [3] developed a processing chain dealing with entity extraction and enrichment, consisting of a set of dedicated components which handle Named Entity Recognition (NER) and consolidation (enrichment, clustering, disambiguation) as part of one coherent workflow.

Traditional approaches to finding related Web resources, such as documents, are often addressed using a combination of Information Retrieval (IR) and Natural Language Processing (NLP) techniques. These techniques compute the similarities between a set of terms from specific resources based on their overlap, or through latent semantic analysis [4] measuring relatedness of individual terms and resources. Nonetheless, most of these techniques require large corpora and a partially common vocabulary/terminology between the resources. Thus, in such cases, they fail to detect latent semantic relationships between resources.

On the other hand, semantic approaches exploit knowledge defined in a data graph to compute notions of similarity and connectivity [18]. The approach taken in this work explicitly targets *connectivity* as a measure of the *relationship* between two Web resources, as opposed to their *similarity*.

An example is derived from datasets specific to the ARCOMEM project¹, which primarily consist of extracted information about events and entities (see Dietze et al. [3]). ARCOMEM follows a use case-driven approach based on scenarios aimed at creating focused Web archives, particularly of social media, by adopting novel entity extraction and interlinking mechanisms. These archives deploy a document repository of crawled Web content and a structured Resource Description Framework (RDF) knowledge base containing metadata about entities and events detected in the archived content.

For instance, Figure 76.1 shows three sets of Web resources (depicted at the top), each associated with one particular entity/event, where the entity (“Jean Claude Trichet”) and event (“Trichet warns of systemic debt crisis”) are both enriched with the same DBpedia entity². This enables clustering of the respective entity and event and their connected Web resources, as an example of direct connection (solid red line in the diagram). However, the third set of Web resources is connected with a third entity “ECB” (which refers to the *European Central Bank*), enriched with the corresponding DBpedia resource³. While NLP and standard IR approaches would

¹ Archiving Community Memories <http://www.arcomem.eu>

² <http://dbpedia.org/resource/Jean-Claude-Trichet>

³ <http://dbpedia.org/resource/ECB>

fail to detect a connection between them, analysing the DBpedia graph uncovers a close connection between ECB and *Jean Claude Trichet* (being a former ECB president), and hence allows us to establish a connection (dashed line) between all involved entities/events and their connected Web resources. Analysis of the reference data graph thereby allows us to identify implicit connections between entities and documents.




Fig. 76.1: Example: connections between Web documents, extracted entities and DBpedia enrichments within ARCOMEM dataset.

This chapter presents a general-purpose approach to detect and measure semantic connectivity between entities within reference datasets as a means to compute connectivity between Web resources (documents) in disparate datasets and document corpora. Our semantic connectivity score is based on the Katz index [13], a score for measuring relatedness of actors in social networks, which has been adopted and expanded to take into account the semantics of data graphs.

In previous work [16, 15], the authors have introduced the semantic connectivity score between entities and an in-depth analysis of how semantic graphs can be exploited to uncover latent connections between entities. In this chapter, the authors extend the previous approach based on entity connectivity to find latent connections across documents. Moreover, they present a Web-based application named *Cite4Me* as a use case scenario for their document interlinking approach.

The remainder of the chapter is structured as follows. Section 76.2 presents an overview of related research. Section 76.3 introduces the semantic connectivity score between documents. Sections 76.4 and 76.5 show the evaluation method and the outcomes of the method. Section 76.6 introduces *Cite4Me* a Web-based application as a use case scenario of the approach. Finally, Section 76.7 presents the conclusion and future work.

76.2 Related Work

Kaldoudi et al. [12] discusses how to apply the overall approach of actor/network theory to data graphs. Graph summarization is an interesting approach to exploit semantic knowledge in annotated graphs. Thor et al. [20] exploited this technique for link prediction between genes in the area of Life Sciences. Their approach relies on the fact that summarisation techniques can create compact representations of the original graph, by adopting a set of criteria for creation, correction and deletion of edges and grouping of nodes. Thus, a prediction function ranks the edges with the most potential, and then suggests possible links between two given genes.

Potamias et al. [17] presents another approach based on Dijkstra's shortest path along with random walks in probabilistic graphs to define distance functions that identify the k closest nodes from a given source node. Lehmann et al. [11] introduces RelFinder, which shows semantic associations between two different entities from RDF datasets. This is based on a breadth-first search algorithm responsible for finding all related entities in the tripliset. The work described here uses the RelFinder approach to exploit the connectivity between entities.

In the field of Social Networks, Hasan and Zake [10] present a survey of link prediction techniques, where they classify the approaches into the following categories: *feature based link prediction*, *bayesian probabilistic models*, *probabilistic relational models* and *linear algebraic methods*. According to this classification, the approach described here can be classified as a *feature based link prediction* method. Work from Leskovec et al. [14] presents a technique suggesting positive and negative relationships between people in a social network. This notion is also addressed in our method, but takes into account the path length, as mentioned previously.

Finding semantic relationships between two given entities is also discussed in the context of ontology matching [9, 21, 22]. In the case described here, hub ontologies could also be used to infer missing relationships into another ontology. From the approaches outlined, we combine different techniques to uncover connections between disparate entities, which enables the exploitation of the relationships between entities to identify connections between Web resources.

76.3 Document Connectivity

This section presents the main steps of the process chain of the approach. The whole process is composed of four steps, described as follows:

Entity Extraction: Pre-processing of documents for finding and extracting term references and named entities;

Entity Enrichment: Matching of references in external knowledge bases such as DBpedia and Freebase⁴;

Entity Connectivity: Uncovering of latent relationships between entities and induction of connections amongst entities.

Document Connectivity: Uncovering latent relationships between documents through entity connections and inducing connections amongst documents.

Steps 1-3 have been introduced in the authors' previous work [16, 15] and therefore, in this section, the focus is on Step 4, the contribution of this work, in which latent connections are discovered between documents. However, Step 3 (defined in previous work) is of paramount importance in order to fully appreciate Step 4.

76.3.1 A novel approach to document connectivity

In this section, a document connectivity score is defined, which relies on connections between entities based on reference graphs. Before introducing the document connectivity approach, it should be recalled how Step 3 (described in [15]) uncovers latent connections between entities which the approach described here builds upon.

76.3.1.1 Entity Connectivity

As the main goal of this work is to uncover latent information between documents, the content of the documents is first exploited to find connections between terms and entities that occur in the documents that would in turn induce connections between the documents themselves. For this, the documents are processed to find and extract term references and named entities, and then enrich these mentions using reference datasets (DBpedia).

Assuming that this process is already solved by previous approaches, the remaining problem is that of finding latent connections between entities. For this, the Semantic Connectivity Score (SCS_e), responsible for discovering latent connections between entity pairs, is used. SCS_e is based on the Katz index [13] which is used to estimate the relatedness of actors in social networks. To adapt the Katz index for finding latent connections between entity pairs in large graphs, three main adaptations have been applied: maximum path length, undirected graphs, and transversal paths.

1. *Maximum path length:* Traversing large graphs is computationally expensive and the computation of all paths between entity pairs is computationally intractable. Thus, to make our approach feasible, we restrict the computation of

⁴ <http://www.freebase.com>

paths between entity pairs with a maximum path length of four intermediate edges (links) between them.⁵

2. *Undirected graphs*: Reference graphs like DBpedia and Freebase have object properties that are often found in their inverse form. For instance, as described in [8], the property *fatherOf* is the inverse property of *sonOf*. Thus, connectivity is explored between entity pairs without taking into account the edge direction. Hence, the semantic connectivity scores between entities are the same for both directions.
3. *Transversal paths*: As described in [2], relation types found in reference graphs are distinguished as *hierarchical* and *transversal*. Concisely, *hierarchical* relations indicate similarity through typical hierarchical relations between entity pairs. Examples of hierarchical relations are: `rdfs:subclassOf`, `dcterms:subject` and `skos:broader`. Unlike *hierarchical* relations, *transversal* relations indicate entity connectivity independent of their similarity, i.e. non-hierarchical relations. Thus, to compute the semantic connectivity score between entity pairs only transversal relations are considered. An example of transversal relation is given by the entity pairs “Jean Claude Trichet” and “European Central Bank” introduced in Section 76.1, where the “European Central Bank” is linked to the entity “President of the European Central Bank” through the transversal RDF property⁶ that, for its part, links to “Jean Claude Trichet” through another transversal RDF property⁷.

Having introduced and defined the scenario in which the SCS_e is computed between an entity pair (e_1, e_2) , Equation 1 is now presented:

$$SCS_e(e_1, e_2) = 1 - \frac{1}{1 + (\sum_{l=1}^{\tau} \beta^l \cdot |paths_{(e_1, e_2)}^{<l>}|)} \quad (76.1)$$

where $|paths_{(e_1, e_2)}^{<l>}|$ is the number of *transversal* paths of length l between entities e_1 and e_2 , τ is the maximum length of paths considered (in our case $\tau = 4$), and $0 < \beta \leq 1$ is a positive damping factor. The damping factor β^l is responsible for exponentially penalizing longer paths. The smaller this factor, the smaller the contribution of longer paths is to the final score. Obviously, if the damping factor is 1, all paths will have the same weight independent of the length. In previous experiments, we observed that $\beta = 0.5$ achieved better results in terms of precision [16]. Equation 76.1 is normalised to range between $[0, 1]$.

Returning to the example presented in Section 76.1, the SCS_e is computed for the entities *Jean Claude Trichet* (JCT) and *European Central Bank* (ECB), using DBpedia as the reference triples set. Omitting the details, assume that 8 paths of length 2 are obtained, and 14 paths of length 3, resulting in the following score:

⁵ Note that the maximum path length exploited was previously determined after comprehensive tests presented in [15], and also adopted in [5].

⁶ <http://dbpedia.org/property/leaderTitle>

⁷ <http://dbpedia.org/property/title>

$$\begin{aligned}
SCS_e(JCT, ECB) &= 1 - \frac{1}{1 + (0.5^2 \cdot 8 + 0.5^3 \cdot 14)} \\
&= 1 - \frac{1}{1 + (2 + 1.75)} \\
&= 0.79
\end{aligned} \tag{76.2}$$

Note that even for a small number (in this case 8) of short paths (of length 2), the contribution to the overall score is larger than for longer paths (of length 3). Evidently, the score obtained by a longer path can overcome a shorter path depending on the number of paths found and the damping factor assigned.

76.3.1.2 Document Connectivity

Based on the semantic connectivity score between entity pairs (SCS_e), the Semantic Connectivity Score (SCS_d) between two Web resources W_1 and W_2 is defined as follows:

$$SCS_d(W_1, W_2) = \begin{cases} 0, \text{ iff } |E_1| = 0 \text{ or } |E_2| = 0 \\ \left(\sum_{\substack{e_1 \in E_1 \\ e_2 \in E_2 \\ e_1 \neq e_2}} SCS_e(e_1, e_2) + |E_1 \cap E_2| \right) \cdot \frac{1}{|E_1| \cdot |E_2|}, \text{ otherwise} \end{cases} \tag{76.3}$$

where E_i is the set of entities found in W_i , for $i = 1, 2$. Note that the score is normalised between $[0, 1]$. The score $SCS_d(W_1, W_2)$ is 0 when no connection between entity pairs across documents exists or iff $|E_1| = 0$ or $|E_2| = 0$. Otherwise, the score is represented by the sum of semantic connectivity scores between entities, normalised over the total number of entity pair comparisons.

To illustrate the semantic connectivity score between document pairs, we present two descriptions of documents extracted from the USAToday⁸ corpus. It can be observed that the underlined terms are entities previously recognised through the entity recognition and enrichment process (Steps 1 and 2).

- (i) The Charlotte Bobcats could go from the NBA's worst team to its best bargain.
- (ii) The New York Knicks got the big-game performances they desperately needed from Carmelo Anthony and Amar'e Stoudemire to beat the Miami Heat.

Thus, for each entity in document (i) and document (ii), the semantic connectivity score (SCS_e) between entities can be computed. Table 76.1 summarises the scores between entity pairs between documents (i) and (ii). The final score between the documents (i) and (ii) is:

⁸ <http://www.usatoday.com>

$$\begin{aligned}
 SCS_d(W_1, W_2) &= \frac{(0.87 + 0.63 + 0.60 + 0.89) + (0.85 + 0.60 + 0.63 + 0.87)}{2 \cdot 4} \\
 &= \frac{5.96}{8} = 0.74
 \end{aligned}
 \tag{76.4}$$

Table 76.1: SCS scores between entity pairs in document (i) and (ii).

Entities from document (i)	Entities from document (ii)	SCS_e
<i>Charlotte Bobcats</i>	<i>New York Knicks</i>	0.87
<i>Charlotte Bobcats</i>	<i>Carmelo Anthony</i>	0.63
<i>Charlotte Bobcats</i>	<i>Amar'e Stoudemire</i>	0.60
<i>Charlotte Bobcats</i>	<i>Miami Heat</i>	0.89
<i>NBA</i>	<i>New York Knicks</i>	0.85
<i>NBA</i>	<i>Carmelo Anthony</i>	0.60
<i>NBA</i>	<i>Amar'e Stoudemire</i>	0.63
<i>NBA</i>	<i>Miami Heat</i>	0.87

76.4 Evaluation Method

This section describes in detail the evaluation methodology and experiment setup used to validate the hypothesis of uncovering latent relationships between Web resources (*entities* and *documents*) using the semantic connectivity score SCS_d .

76.4.1 Dataset

The dataset used to evaluate our approach consists of a subset of randomly selected news articles (documents) from the USA Today news Website. In total, document connectivity is considered for 40,000 document pairs. Each document contains a title and a summary, where the latter is 200 characters long on average. The *entity extraction* step was performed using DBpedia Spotlight⁹. The resulting set of annotations consists of approximately 80,000 entity pairs.

⁹ <http://spotlight.dbpedia.org/>

76.4.2 Gold standard

In order to validate the results of our evaluation, the first step is to obtain a ground truth of relationships between documents. Given the lack of such benchmarks, a user evaluation was conducted to collect user judgements with the aim of creating a gold standard. The user evaluation was set up in CrowdFlower¹⁰, a crowdsourcing platform. In order to construct the gold standard, 600 document pairs were randomly selected to be evaluated. The evaluation process consisted of a questionnaire on a 5-point Likert scale model where participants were asked to rate their agreement of the suggested semantic connection between a given document pair.

Additionally, participants' expectations regarding declared connected document pairs were investigated. In this case, presenting two documents deemed to be connected, participants were asked if such connections were expected (from *extremely unexpected* to *extremely expected*, also on a 5-point Likert scale). The judgements collected provide a gold standard for the analysis of the techniques. Note that in this work, additional challenges are posed with respect to the gold standard, because our semantic connectivity score is aimed at detecting possibly unexpected relationships which are not always obvious to the user. To this end, a gold standard created by humans provides an indication of the performance of the approach with respect to precision and recall, but it may lack appreciation of some of the relationships found.

76.4.3 Evaluation Methods

To emphasise the benefits of measuring connectivity between documents using our approach, it was compared against competing methods which measure connectivity via co-occurrence-based metrics to detect entity and document connectivity. In the first evaluation, the performance of SCS_d was compared against the co-occurrence-based method and the Explicit Semantic Analysis (ESA) method.

Co-occurrence-based method (CBM) is a score between entities that relies on an approximation of the number of existing Web pages that contain these entities. For example, Nunes et al. [16] estimates the co-occurrence score of entity pairs by issuing queries (such as “*Jean Claude Trichet*” + “*European Central Bank*”) to a search engine and retrieving the total number of search results that contain the entity labels in their text body. We interpret a large number of pages as an indicator of high connectivity, and a small number of pages as an indicator of low connectivity between the queried terms (which represent entities in our case). Besides CBM, there are other similar approaches to quantify the relatedness between entities, such as Pointwise Mutual Information (PMI) [1] and Normalised Google Distance (NGD) [7]. However, they take into account the joint distribution and the probability of their individual distributions, which

¹⁰ <https://www.crowdfunder.com/>

requires knowing a priori the total number of Web pages searched by a search engine. Thus, in this case, the document connectivity score is given by a small adjustment in Equation (76.3) where, instead of SCS_d , CBM is used.

Explicit Semantic Analysis (ESA) proposed by Gabrilovich and Markovitch [6] measures the relatedness between Wikipedia¹¹ concepts by using a vector space model representation, where each vector entry is assigned using the *tf-idf* weight between the entities and its occurrence in the corresponding Wikipedia article. The final score is given by the cosine similarity between the weighted vectors.

In order to evaluate the document connectivity, we compared our method with the traditional statistical *tf-idf* method, in addition to ESA and CBM. As mentioned, the latter method was slightly modified to measure the connectivity between documents, where in Equation (76.3) the semantic score was replaced by the co-occurrence-based score.

76.4.4 Evaluation Metrics

For measuring the performance of the document connectivity approaches, standard evaluation metrics of precision (P), recall (R) and $F1$ measure were used. Note that in these metrics, as relevant pairs, only those marked in the gold standard (gs) as connected according to the 5-point Likert Scale (*Strongly Agree & Agree*) were considered.

For the document connectivity, the precision measure (P_d) is the ratio of the set of all retrieved document pairs deemed as relevant over the set of document pairs that are connected. Thus, the relevant documents are those that were marked as *Strongly Agree & Agree*, while the set of document pairs that are connected consists of those that have a semantic connectivity score greater than a given threshold (see Equation (76.5)).

$$P_d = \begin{cases} 0, & \text{iff } |\Phi_{retrieved}^\tau| = 0 \\ \frac{|\Phi_{retrieved}^\tau \cap \Phi_{relevant}|}{|\Phi_{retrieved}^\tau|}, & \text{otherwise} \end{cases} \quad (76.5)$$

where $\Phi_{relevant}$ is the set of retrieved document pairs that are relevant and $\Phi_{retrieved}^\tau$ is the set of all connected document pairs greater than a given threshold (τ).

The recall (R_d) is the ratio of the set of retrieved documents that are relevant over the set of all relevant document pairs according to the gold standard (see Equation (76.6)).

¹¹ <http://www.wikipedia.org>

$$R_d = \begin{cases} 0, & \text{iff } |\Phi_{relevant(gs)}| = 0 \\ \frac{|\Phi_{retrieved}^r \cap \Phi_{relevant}|}{|\Phi_{relevant(gs)}|}, & \text{otherwise} \end{cases} \quad (76.6)$$

where $\Phi_{relevant(gs)}$ is the set of all relevant document pairs.

Finally, $F1$ measure shows the balance between precision and recall, and is computed as in Equation (76.7).

$$F1_w = \begin{cases} 0, & \text{iff } (P_d + R_d) = 0 \\ 2 \cdot \frac{P_d \cdot R_d}{P_d + R_d}, & \text{otherwise} \end{cases} \quad (76.7)$$

76.5 Results

In this section, the authors report evaluation results for the document connectivity approaches. For each method, the authors present the results for their ability to discover latent connections between pairs of resources. Furthermore, the authors also present an in-depth analysis of their shortcomings and advantages for discovering connections between documents.

76.5.1 Document connectivity results

Table 76.2 shows the results according to the gold standard presented in the Likert scale, where users evaluated if a given entity pair could be connected in a document. Compared with the gold standard, 368 entity pairs out of 812 could have some connection.

From the set of entities that could co-occur in a document, 51% of those entities were also connected based on our gold standard, while 34% were *Undecided*. Analysis of the results for the *Undecided* category will be provided in Section 76.5.2, since these results are of particular interest in establishing latent relationships between Web resources.

Table 76.2: Total number of results for the GS in Likert-scale.

Strongly Agree	Agree	Undecided	Disagree	Strongly Disagree
96	272	139	165	140

The performance of each method is shown in Figure 76.2. As in the task of entity connectivity, SCS_d performs slightly better than CBM in terms of precision, while

CBM is better in terms of recall. $F1$ measure is similar, with 60.0% and 59.6% for SCS_d and CBM, respectively. In both cases, ESA has the lowest performance.




Fig. 76.2: Precision, recall and $F1$ measure amongst methods.

The positive correlation of entity connectedness and their co-occurrence in the same document was 79.6%, 78.0% and 23.5% for SCS_d , CBM and ESA respectively, considering only the *Strongly Agree* and *Agree* relevance judgement results.

As already indicated in the introduction in Section 76.1, the proposed semantic approach can be exploited to measure document connectivity by taking into account the connectedness of entities that describe a document and their semantic connections. Indeed, as shown by the positive correlation of entity connectivity and entity co-occurrence in a document, the authors claim that their approach can be used as a method for inferring document “relatedness” where other statistical models would fail.

To validate the usefulness of our approach, the authors compared the results against the well established document relatedness measure *tf-idf*. The approach was able to find 500 unique connections between documents, whereas *tf-idf* found only 25. As described in Section 76.4.1, the corpus is composed of small descriptions of the news articles, which severely limits the ability of *tf-idf* to identify connections between them.

The authors also conducted an experiment to evaluate the uncovered connections by both methods. They found that 16% of the connections found by their approach were relevant, compared with 12% using *tf-idf*. They took into consideration that the recall achieved by *tf-idf* is only 3.6%, whereas for SCS_d , it is close to 86%.

76.5.2 Analysis of the Results

Table 76.3 shows the results for the task of document connectivity. The mixed approach $CBM+SCS_d$ performs best on finding the co-occurrence of entity pairs in a

Table 76.3: Precision, recall and F1 measure amongst methods.

	CBM	SCS_d	ESA	CBM+ SCS_d
Precision	0.47	0.49	0.21	0.51
Recall (GS)	0.80	0.77	0.25	0.89
Recall	0.49	0.48	0.15	0.54
F1 (GS)	0.59	0.60	0.23	0.64
F1	0.48	0.48	0.18	0.52

document. It is worth noting as well that the co-occurrence of entity pairs for documents can be retrieved with high recall (90%) when using the proposed combination of CBM+ SCS_d .

A positive correlation of entity connectivity and co-occurrence in a document is of high importance for the proposed approach, allowing to establish newly constructed knowledge that can be represented as an aggregate of the entity connections.

The authors would also like to point out the challenges posed by our approach to creating a gold standard. As previously mentioned, while this work aims at detecting semantic connectivity of entities beyond traditional co-occurrence based approaches, this results in connections that might be to some extent unexpected yet correct, according to background knowledge (such as DBpedia in this case). Hence, using a manually created gold standard, though being the only viable option, might impact the precision values for this work in a negative way, as correct connections might have been missed by the evaluators. This has been partially confirmed by the large number of detected co-occurrences which were marked as *undecided* by the users, where manual inspection of samples in fact confirmed meaningful connections between entity pairs. This confirms that in a number of cases connections were not necessarily incorrect, but simply presented information that was unknown to the users. Thus, the authors believe that a more thorough evaluation providing the evaluators with information on how a connection emerged, where showing all properties and entities that are part of a path greater than one would give more reliable judgements.

76.6 Use Case Scenarios

This section describes a practical application scenario that makes use of the four-step processing chain introduced in Section 76.3 to interlink concepts and documents based on semantic graphs. The authors have developed *Cite4Me*, a Web-based application for searching academic papers and recommending related work in a user-friendly interface that assists users in exploring relationships between authors, institutions, papers and query terms. The main features of the *Cite4Me* application are explained in the following subsections.

76.6.1 Exploratory Search

The exploratory search feature assists users to discover related work, people and institutions that are working on a specific topic. To provide this type of search, *entity extraction* is used (see Section 76.3), and the publications' content is annotated with structured information from DBpedia. For instance, the categories of the extracted concepts are used to interlink publications through the topics they cover. In cases where two publications share the same category (`dcterms:subject` property), a link between both publications is created. Figure 76.3 shows an example of topically related publications.

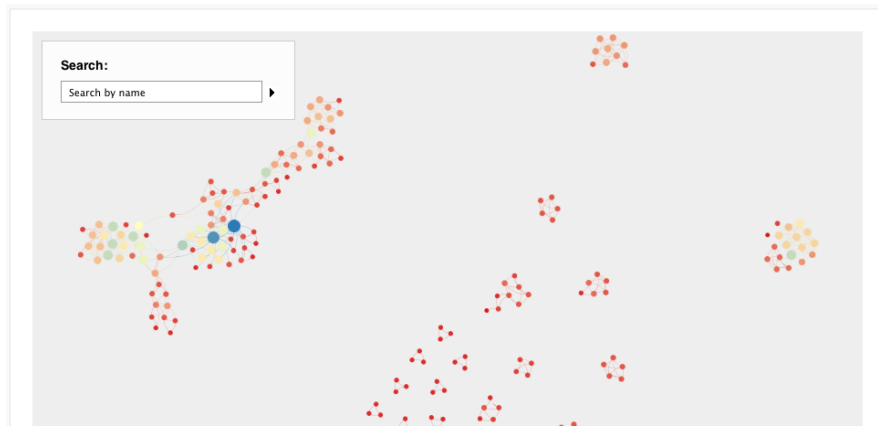


Fig. 76.3: Preview of the exploratory search functionality.

76.6.2 Semantic Search

The semantic search component of *Cite4Me* is similar to the Explicit Semantic Analysis (ESA) technique [6]. After running the annotation process (see Section 76.3), the relatedness score between the enriched concepts (DBpedia entities) found in the user query terms and the publications' content are computed and ranked. The relatedness score is computed based on the tf-idf score for the entities found in the publications' content. The ranking of the retrieved documents is based on the sum of the tf-idf scores of the matching concepts. Figure 76.4 illustrates the semantic search functionality. Alongside the results of the semantic search a tag cloud shows the most prominent terms for a given user query. The tag cloud is updated while browsing through the list of results. The tags are selected based on the *tf-idf* score for the entities found in the abstract of the retrieved papers.

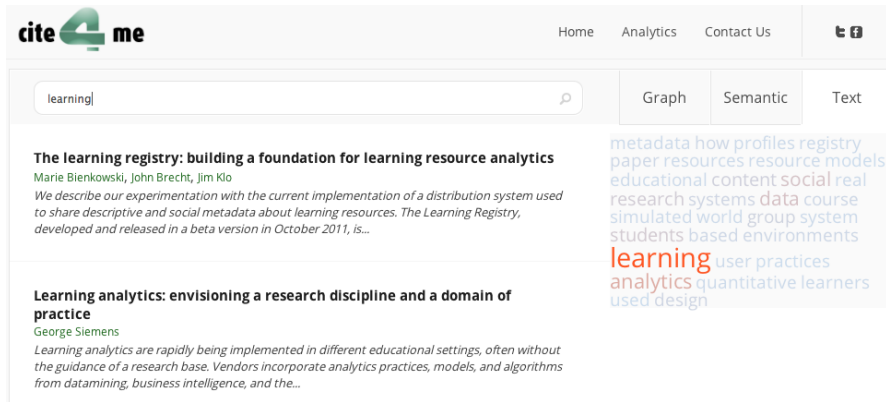


Fig. 76.4: Preview of the semantic search functionality.

76.6.3 Paper recommendation

An important feature of *Cite4Me* and which differentiates it from similar tools is the semantic paper recommendation. It is based on the step of *Document Connectivity*. For a retrieved scientific publication based on some user query, the tool recommends a related paper based on a score calculated according to the connectivity of the visited publication with other publications in a given corpus. To measure such connectivity the abstract and keywords of the paper are analysed (see Section 76.3). The top ranked publication is suggested to the user as a related material as shown in Figure 76.5.

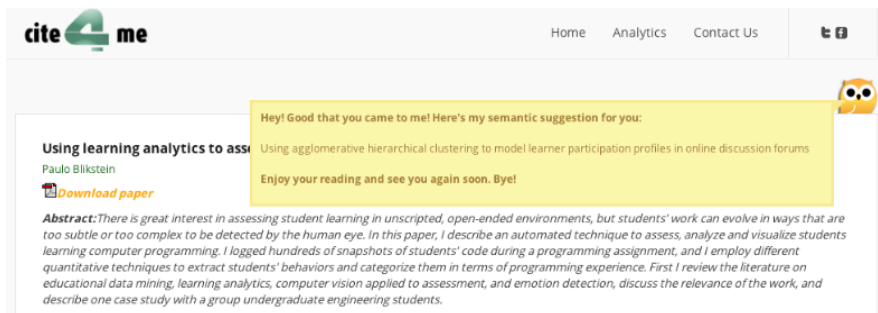


Fig. 76.5: An example of paper recommendation based on SCS_d .

76.6.4 Data

Currently, *Cite4Me* is linked to a dataset (LAK Dataset¹²) which contains semistructured research publications from the ACM Digital Library (under a special license) and other public datasets (see also Taibi et al. [19] for details). The dataset contains 315 full papers along with their descriptive metadata while new publications are added continuously. Metadata as well as the full text body are freely available in a variety of formats, including RDF accessible via a public SPARQL¹³ endpoint.

76.7 Conclusion and future work

This chapter has presented a general-purpose approach to discover and quantify document connectivity. To compute document connectivity, a semantic-based entity connectivity approach (SCS_e) was introduced, which adapts a measure from social network theory (Katz) to data graphs, in particular Linked Data, and extended to interlink documents (SCS_d). SCS_d was able to uncover 16% of unique inferred document connections based on entity co-occurrence, not found by the state of the art method CBM. Additionally, while using a combination of CBM+ SCS_d an $F1$ measure of 52% was achieved.

The experiments presented show that SCS_d enables the detection and establishment of document connectivity that a priori linguistic and co-occurrence approaches would not reveal. Contrary to the latter, the former approach relies on semantic relations between entities as represented in structured background knowledge, available via reference datasets. A combination of the semantic approach and traditional co-occurrence-based measures provided very promising results for detecting connected documents. While both approaches (CBM and SCS_d) produce fairly good indicators for document connectivity, an evaluation based on Kendall's τ rank correlation showed that the approaches differ in the relationships they uncover [16]. A comparison of agreement and disagreement between different methods revealed that both approaches are complementary and produce particularly good results when combined: the semantic approach is able to find connections between entities that do not necessarily co-occur in documents (found on the Web), while the CBM tends to emphasize entity connections between entities that are not necessarily strongly connected in reference datasets.

As for future work, the aim is to apply weights to different edge/property types according to their inherent semantics in order to provide a more refined score and to investigate means to combine our complementary approaches.

¹² <http://www.solaresearch.org/resources/lak-dataset/>

¹³ SPARQL Protocol and RDF Query Language

76.8 Acknowledgements

This work has been partially supported by CAPES (Process n° 9404-11-2), the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement No 270239 (ARCOMEM) and No 317620 (LinkedUp) as well as CNPQ under grants 301497/2006-0 and 475717/2011-2 and by FAPERJ under grants E-26/103.070/ 2011.

References

1. Church, K.W., Hanks, P.: Word association norms, mutual information, and lexicography. *Comput. Linguist.* **16**(1), pp. 22–29 (1990)
2. Damljanovic, D., Stankovic, M., Laublet, P.: Linked data-based concept recommendation: Comparison of different methods in open innovation scenario. In: E. Simperl, P. Cimiano, A. Polleres, O. Corcho, V. Presutti (eds.) *The Semantic Web: Research and Applications, Lecture Notes in Computer Science*, vol. 7295, pp. 24–38. Springer Berlin Heidelberg (2012)
3. Dietze, S., Maynard, D., Demidova, E., Risse, T., Peters, W., Doka, K., Stavrakas, Y.: Entity extraction and consolidation for social web content preservation. In: A. Mitschick, F. Loizides, L. Predoiu, A. Nurnberger, S. Ross (eds.) *SDA, CEUR Workshop Proceedings*, vol. 912, pp. 18–29. CEUR-WS (2012)
4. Dumais, S.T.: Latent semantic analysis. *Annual Review of Information Science and Technology* **38**(1), pp. 188–230 (2004)
5. Fang, L., Sarma, A.D., Yu, C., Bohannon, P.: Rex: explaining relationships between entity pairs. *Proc. VLDB Endow.* **5**(3), pp. 241–252 (2011)
6. Gabrilovich, E., Markovitch, S.: Computing semantic relatedness using wikipedia-based explicit semantic analysis. In: *Proceedings of the 20th international joint conference on Artificial intelligence, IJCAI'07*, pp. 1606–1611. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2007)
7. Gligorov, R., ten Kate, W., Aleksovski, Z., van Harmelen, F.: Using google distance to weight approximate ontology matches. In: *Proceedings of the 16th international conference on World Wide Web, WWW '07*, pp. 767–776. ACM, New York, NY, USA (2007)
8. Graves, A., Adali, S., Hendler, J.: A method to rank nodes in an rdf graph. In: C. Bizer, A. Joshi (eds.) *Proceedings of the Poster and Demonstration Session at the 7th International Semantic Web Conference (ISWC2008)*, Karlsruhe, Germany, October 28, 2008, *CEUR Workshop Proceedings*, vol. 401. CEUR-WS (2008)
9. Groß, A., Hartung, M., Kirsten, T., Rahm, E.: Mapping Composition for Matching Large Life Science Ontologies. In: *Proceedings of the 2nd International Conference on Biomedical Ontology, ICBO*, pp. 109–116 (2011)
10. Hasan, M., Zaki, M.: A survey of link prediction in social networks. In: C.C. Aggarwal (ed.) *Social Network Data Analytics*, pp. 243–275. Springer US (2011)
11. Heim, P., Hellmann, S., Lehmann, J., Lohmann, S., Stegemann, T.: RelFinder: Revealing relationships in RDF knowledge bases. In: *Proceedings of the 3rd International Conference on Semantic and Media Technologies (SAMT), Lecture Notes in Computer Science*, vol. 5887, pp. 182–187. Springer (2009)
12. Kaldoudi, E., Dovrolis, N., Dietze, S.: Information organization on the internet based on heterogeneous social networks. In: *Proceedings of the 29th ACM international conference on Design of communication, SIGDOC '11*, pp. 107–114. ACM, New York, NY, USA (2011)
13. Katz, L.: A new status index derived from sociometric analysis. *Psychometrika* **18**(1), pp. 39–43 (1953)

14. Leskovec, J., Huttenlocher, D., Kleinberg, J.: Predicting positive and negative links in online social networks. In: Proceedings of the 19th international conference on World wide web, WWW '10, pp. 641–650. ACM, New York, NY, USA (2010)
15. Pereira Nunes, B., Dietze, S., Casanova, M., Kawase, R., Fetahu, B., Nejdl, W.: Combining a co-occurrence-based and a semantic measure for entity linking. In: P. Cimiano, O. Corcho, V. Presutti, L. Hollink, S. Rudolph (eds.) The Semantic Web: Semantics and Big Data, *Lecture Notes in Computer Science*, vol. 7882, pp. 548–562. Springer Berlin Heidelberg (2013)
16. Pereira Nunes, B., Kawase, R., Dietze, S., Taibi, D., Casanova, M.A., Nejdl, W.: Can entities be friends? In: G. Rizzo, P. Mendes, E. Charton, S. Hellmann, A. Kalyanpur (eds.) Proceedings of the Web of Linked Entities Workshop in conjunction with the 11th International Semantic Web Conference, *CEUR-WS*, vol. 906, pp. 45–57 (2012)
17. Potamias, M., Bonchi, F., Gionis, A., Kollis, G.: k-nearest neighbors in uncertain graphs. *Proc. VLDB Endow.* **3**(1-2), pp. 997–1008 (2010)
18. Sheth, A., Aleman-Meza, B., Arpinar, F.S., Sheth, A., Ramakrishnan, C., Bertram, C., Warke, Y., Anyanwu, K., Aleman-meza, B., Arpinar, I.B., Kochut, K., Halaschek, C., Ramakrishnan, C., Warke, Y., Avant, D., Arpinar, F.S., Anyanwu, K., Kochut, K.: Semantic association identification and knowledge discovery for national security applications. *Journal of Database Management* **16**, pp. 33–53 (2005)
19. Taibi, D., Dietze, S.: Fostering analytics on learning analytics research: the lak dataset. In: M. d'Aquin, S. Dietze, H. Drachslar, E. Herder, D. Taibi (eds.) Proceedings of the LAK Data Challenge, Leuven, Belgium, April 9, 2013, *CEUR Workshop Proceedings*, vol. 974. CEUR-WS (2013)
20. Thor, A., Anderson, P., Raschid, L., Navlakha, S., Saha, B., Khuller, S., Zhang, X.N.: Link prediction for annotation graphs using graph summarization. In: 10th International Conference on The Semantic Web, Vol. Part I, ISWC'11, pp. 714–729. Berlin, Heidelberg (2011)
21. Vidal, V.M.P., de Macedo, J.A.F., Pinheiro, J.C., Casanova, M.A., Porto, F.: Query processing in a mediator based framework for linked data integration. *IJBDCN* **7**(2), pp. 29–47 (2011)
22. Xu, L., Embley, D.W.: Discovering direct and indirect matches for schema elements. In: Proceedings of the Eighth International Conference on Database Systems for Advanced Applications, DASFAA '03, p. 39. IEEE Computer Society, Washington, DC, USA (2003)