

Large-Scale Semantic Exploration of Scientific Literature using Topic-based Hashing Algorithms

Editor(s): Tomi Kauppinen, Aalto University, Finland; Daniel Garijo, University of Southern California, USA; Natalia Villanueva, The University of Texas at El Paso, USA

Solicited review(s): Anita de Waard, Elsevier Labs; Daniel Garijo, University of Southern California; One anonymous reviewer

Carlos Badenes-Olmedo^{a,*}, José Luis Redondo-García^b and Oscar Corcho^a

^a *Ontology Engineering Group, Universidad Politécnica de Madrid, Boadilla del Monte, Spain*

E-mails: cbadenes@fi.upm.es, ocorcho@fi.upm.es

ORCID: 0000-0002-2753-9917, 0000-0002-9260-0753

^b *Amazon Research, Cambridge, UK*

E-mail: jluisred@amazon.com

ORCID: 0000-0002-7413-447X

Abstract. Searching for similar documents and exploring major themes covered across groups of documents are common activities when browsing collections of scientific papers. This manual knowledge-intensive task can become less tedious and even lead to unexpected relevant findings if unsupervised algorithms are applied to help researchers. Most text mining algorithms represent documents in a common feature space that abstract them away from the specific sequence of words used in them. Probabilistic Topic Models reduce that feature space by annotating documents with thematic information. Over this low-dimensional latent space some locality-sensitive hashing algorithms have been proposed to perform document similarity search. However, thematic information gets hidden behind hash codes, preventing thematic exploration and limiting the explanatory capability of topics to justify content-based similarities. This paper presents a novel hashing algorithm based on approximate nearest-neighbor techniques that uses hierarchical sets of topics as hash codes. It not only performs efficient similarity searches, but also allows extending those queries with thematic restrictions explaining the similarity score from the most relevant topics. Extensive evaluations on both scientific and industrial text datasets validate the proposed algorithm in terms of accuracy and efficiency.

Keywords: Document Similarity, Information Search and Retrieval, Clustering, Topic Models, Hashing

1. Introduction

Huge amounts of documents are publicly available on the Web offering the possibility of extracting knowledge from them (e.g. scientific papers in digital journals). Document similarity comparisons in many information retrieval (IR) and natural language processing (NLP) areas are too costly to be performed in such huge collections of data and require more efficient ap-

proaches than having to calculate all pairwise similarities.

In this paper we address the problem of programmatically generating annotations for each of the items inside big collections of textual documents, in a way that is computationally affordable and enables a semantic-aware exploration of the knowledge inside it that state-of-the-art methods relying on topic models are not able to materialize.

Most text mining algorithms represent documents in a common feature space that abstracts the specific

*Corresponding author. E-mail: cbadenes@fi.upm.es.

sequence of words used in each document and, with appropriate representations, facilitate the analysis of relationships between documents even when written using different vocabularies. Although a sparse word or n-gram vectors are popular representational choices, some researchers have explored other representations to manage these vast amounts of information. Latent Semantic Indexing (LSI) [15], Probabilistic Latent Semantic Indexing (PLSI) [23] and more recently, Latent Dirichlet Allocation (LDA) [10], which is the simplest probabilistic topic model (PTM) [9], are algorithms focused on reducing feature space by annotating documents with thematic information. PLSI and PTM also allow a better understanding of the corpus through the topics discovered, since they use probability distributions over the complete vocabulary to describe them. However, only PTM's are able to identify topics in previously unseen texts.

One of the greatest advantages of using PTM in large document collections is the ability to represent documents as probability distributions over a small number of topics, thereby mapping documents into a low-dimensional latent space (the K-dimensional probability simplex, where K is the number of topics). A document, represented as a point in this simplex, is said to have a particular topic distribution. This brings a lot of potential when applied over different IR tasks, as evidenced by recent works in different domains such as scholarly [21][17], health [42] [35] [47], legal [40][18], news [22] and social networks [44][13]. This low-dimensional feature space could also be suitable for document similarity tasks, especially on big real-world data sets, since topic distributions are continuous and not as sparse as discrete-term feature vectors.

Exact similarity computations for most topic distributions require to have complexity $O(n^2)$ for neighbours detection tasks or $O(kn)$ computations when k queries are compared against a dataset of n documents. Computation can be an approximate nearest neighbor (ANN) search problem. ANN search is an optimization problem that finds nearest neighbors of a given query q in a metric space of n points. Due to the low storage cost and fast retrieval speed, hashing is one of the most popular solutions for ANN search [33] [4] [58]. This technique transforms data points from the original feature space into a binary-code space, so that similar data points have larger probability of collision (i.e. having the same hash code). This type of formulation for the document similarity comparison problem has proven to yield good results in the metric space due to the fact that ANN search has been designed to handle

distance metrics (e.g. cosine, Euclidean, Manhattan) [46][43][27], even in high-dimensional simplex spaces handling information-theoretically motivated metrics (e.g. Hellinger, Kullback-Leibler divergence, Jensen-Shannon divergence) as demonstrated by [36].

However, the smaller space created by existing hashing methods loses the exploratory capabilities of topics to support document similarity. The notion of topics is lost and therefore the ability to make thematic explorations of documents. Moreover, metrics in simplex space are difficult to interpret and the ability to explain the similarity score on the basis of the topics involved in the exploration can be helpful. While other models based on vector representations of documents are simply agnostic to the human concept of themes, topic models can help finding the reasons why two documents are similar.

Semantic knowledge can be thought of as knowledge about relations among several types of elements, including words, concepts, and percepts [20]. Since topic models create latent themes from word co-occurrence statistics in corpus, a topic (i.e latent theme) reflects the knowledge about the word-word relations it contains. This abstraction can be extended to cover the knowledge derived from sets of topics. The topics obtained via state-of-the art methods (LDA) are hierarchically divided into groups with different degrees of semantic specificity in a document. Documents can then be annotated with the semantic inferred from the topics detected, and from their relation between topics inside each hierarchy level. Let's look at a practical example to clarify this idea. A topic model is created from texts labeled with Eurovoc ¹ categories. This model² annotates texts with categories inferred from their topic distributions. For the document "*Commission Decision of 23 December 2003.. on seeds and propagating material of gramineae, Triticum aestivum..*"³, the top5 categories are: (1) *research*, (2) *sugar*, (3) *fats*, (4) *textile_industry* and (5) *marketing*. In contrast to these categories that standard topic modelling methods are able to offer, a 3-level hierarchical set of topics would be: (1) *research*, (2) *sugar* and *fats*, and (3) *textile_industry* and *marketing*. The knowledge provided by each of these annotations is derived from the relations between the topics that compose it. Based on these semantic annotations, the content-based similarity among documents is calcu-

¹<http://publications.europa.eu/resource/dataset/eurovoc>

²<http://library.linkeddata.es/jrc-en-model/>

³[https://eur-lex.europa.eu/eli/dec/2004/57\(1\)/oj](https://eur-lex.europa.eu/eli/dec/2004/57(1)/oj)

lated and the exploration of large document collections is performed following an ANN search.

Thus, in this paper, we propose a hashing algorithm that (1) groups similar documents, (2) preserves their topic distributions, and (3) works over unseen documents. Therefore our contributions are:

- a **novel hashing algorithm** based on topic models that not only performs efficient searches, but also introduces semantic in the hierarchy of concepts as a way to restrict those queries and provide explanatory information.
- an optimized and easily customizable open-source **implementation of the algorithm** [8]
- **data-sets** and **pre-trained models** to facilitate other researchers to replicate our experiments and validate and test their own ideas [8]

2. Document Similarity

In the probability simplex space created from topic models, documents are represented as vectors containing topic distributions. Distance metrics based on vector-type data such as Euclidean distance (l_2), Manhattan distance (l_1), and angular metric (θ) are not optimal in this space [36]. Information-theoretically motivated metrics such as Kullback-Leibler (KL) divergence (Eq.1) (also known as relative entropy), Jensen-Shannon (JS) divergence (Eq.2) (as its symmetric version) and Hellinger (He) distance (Eq.3) are often more reasonable [36]:

$$KL(P, Q) = \sum_{i=1}^K p(x_i) \log \frac{p(x_i)}{q(x_i)} \quad (1)$$

$$JS(P, Q) = \frac{1}{2} KL\left(p, \frac{p+q}{2}\right) + \frac{1}{2} KL\left(q, \frac{p+q}{2}\right) \quad (2)$$

$$He(P, Q) = \sum_{i=1}^K \left(\sqrt{p(x_i)} - \sqrt{q(x_i)} \right)^2 \quad (3)$$

where P and Q are two known distributions, K is the dimensionality of P and Q, and p_i and q_i are the values of the i_{th} component of P and Q, respectively.

He distance is also symmetric and, along with JS divergence, are usually used in various fields where a comparison between two probability distributions is re-




Fig. 1. Distance values based on KL-divergence between 10 pair of documents from topic models with 100-to-2000 dimensions.

quired. However, all these metrics are not well-defined distance metrics, that is, they do not satisfy triangle inequality [12]. This inequality considers $d(x, z) \leq d(x, y) + d(y, z)$ for a metric d [20]. It places strong constraints on distance measures and on the locations of points in a space given a set of distances. As a metric axiom the triangle inequality must be satisfied in order to take advantage of the inferences that can be deduced from it. Thus, if similarity is assumed to be a monotonically decreasing function of distance, this inequality avoids the calculation of all pairs of similarities by considering that if x is similar to y and y is similar to z , then x must be similar to z .

S2JSD was introduced by [16] to satisfy the triangle inequality. It is the square root of two times the JS divergence:

$$S2JSD(P, Q) = \sqrt{2 * JS(P, Q)} \quad (4)$$

However, making sense out of the similarity score is not easy. As shown in figures 1 to 4, given a set of pairs of documents, their similarity scores vary according to the number of topics. So the distances between those pairs fluctuate from being more to less distant when changing the number of topics.

Distances between documents generally increase as the number of dimensions of the space increases. This is due to the fact that as the number of topics describing the model increases, the more specific the topics will be. Topics shared by a pair of documents can be broken down into more specific topics that are not shared by those documents. Thus, similarity between pairs of documents is dependent on the model used to represent them when considering this type of metrics. We know that absolute distances between documents vary when we tune hyperparameters differently, but in this study




Fig. 2. Distance values based on *JS*-divergence between 10 pair of documents from topic models with 100-to-2000 dimensions.




Fig. 3. Distance values based on *He*-divergence between 10 pair of documents from topic models with 100-to-2000 dimensions.




Fig. 4. Distance values based on *S2JSD* between 10 pair of documents from topic models with 100-to-2000 dimensions.

we also see that "relative distances" also change: e.g. for model M1, A is closer to B than C, but according to a M2 trained in the same corpora with different parameters, A is closer to C than B (cross-lines in figs 1-4). This behaviour highlights the difficulty of establishing absolute similarity thresholds and the complexity to measure distances taking into account all dimensions. Distance thresholds should be model-dependent

rather than general and metrics flexible enough to handle dimensional changes. These challenges are tackled through the proposed hashing algorithms by means of clusters of topics to measure similarity, instead of directly using their weights.

3. Hashing Topic Distributions

Hashing methods transform the data points from the original feature space into a binary-code Hamming space, where the similarities in the original space are preserved. They can learn hash functions (data-dependent) or use projections (data-independent) from the training data [54]. Data-independent methods unlike data-dependent ones do not need to be re-calculated when data changes, i.e. adding or removing documents to the collection. Taking large-scale scenarios into account (e.g. Document clustering, Content-based Recommendation, Duplicate Detection), this is a key feature along with the ability to infer hash codes individually (for each document) rather than on a set of documents.

Data-independent hashing methods depend on two key elements: (1) data type and (2) distance metric. For vector-type data, as introduced in section 2, based on l_p distance with $p \in [0, 2)$ lots of hashing methods have been proposed, such as p -stable Locality-Sensitive Hashing (LSH) [14], Leech lattice LSH [3], Spherical LSH [48], and Beyond LSH [4]. Based on the θ distance many methods have been developed such as Kernel LSH [29] and Hyperplane hashing [51]. But only few methods handle density metrics in a simplex space. A first approach transformed the *He* divergence into an Euclidean distance so that existing ANN techniques, such as LSH and k-d tree, could be applied [28]. But this solution does not consider the special attributions of probability distributions, such as Non-negative and Sum-equal-one. Recently, a hashing schema [36] taking into account the symmetry has been proposed, non-negativity and triangle inequality features of the S2JSD metric for probability distributions. For set-type data, Jaccard Coefficient is the main metric used. Some examples are K-min Sketch [31], Min-max hash [26], B-bit minwise hashing [30] and Sim-min-hash [57].

All of them have demonstrated efficiency in the search for similar documents, but none of them allows the search for documents (1) by thematic areas or (2) by similarity levels, nor they offer (3) an explanation about the similarity obtained beyond the vectors used to calculate it. Binary-hash codes drop a very precious information: the topic relevance.




Fig. 5. Hash method based on hierarchical set of topics from a given topic distribution

A new hierarchical set-type data is proposed. Each level of the hierarchy indicates the importance of the topic according to its distribution. Level 0 contains the topics with the highest score. Level 1 contains the topics with highest score once the first ones have been eliminated, and so on. From a vector of components, where each of the components is the score of topic t , a vector containing set of topics is proposed, where each of the dimensions means a topic relevance. Thus, for the topic distribution $q = [0.3, 0.15, 0.4, 0.15]$, a hierarchical set of topics may be $h = \{(t2), (t0), (t1, t3)\}$. It means that topic $t2$ (0.4) is the most relevant, then topic $t0$ (0.3) and, finally, topics $t1$ (0.15) and $t3$ (0.15). This is just an example about the data structure that will support the different hashing strategies. In section 3.3 some approaches to create hash codes based on this data structure are described.

3.1. Data Type

A traditional approach to text representation usually requires encoding of documents into numerical vectors. Words are extracted from a corpus as feature candidates and based on a certain criterion they are assigned values to describe the documents: term-frequency, TF-IDF, information gain, and chi-square are typical measures. But this causes two main problems: huge number of

dimensions and sparse distribution. The use of topics as feature space has been extended to mapping documents into low-dimensional vectors. However, as shown in Figures 1 to 4, the distance metrics based on probability densities vary according to the dimensions of the model and reveal the difficulty of calculating the similarity values using the vectors with the topic distributions.

Since hashing techniques can transform both vector and set-based data [36] [26] into a new space where the similarity (i.e. closeness of points) in the original feature space is preserved, a new set-based data structure is proposed in this paper. It is created from clusters of topics organized by relevance levels and it aims to extend the ability of building queries with topic-based restrictions over the searching space while maintaining high level of accuracy.

The new hierarchical set-type data describes each document as a sequence of sets of topics sorted by relevance. Each level of the hierarchy expresses how important those topics are in that document. In the first level (i.e level 0) are the topics with the highest score. In the second level (i.e level 1) are the topics with the highest score once the first ones have been removed, and so on. In this work, several clustering approaches have been considered to assign topics to each level.

In a feature space created from a PTM with eight topics, for example, each data point p is described

by a eight-dimensional vector with the topic distributions: $vp = [t0, t1, t2, t3, t4, t5, t6, t7]$. Then, given a point $q1 = [0.18, 0.15, 0.2, 0.05, 0.14, 0.11, 0.09, 0.08]$, the three-level hierarchical set of topics may be $h = [\{t2\}, \{t0\}, \{t1, t4\}]$. It means that $t2$ is the most relevant topic, then topic $t0$ and finally topics $t1$ and $t4$. This is just an example about the data structure that will support the hashing strategies. In section 3.3 some approaches to create hash codes based on this data structure are described.

Domain-specific features such as vocabulary, writing style, or speech type, have a major influence on the topic models, but not in the hashing algorithms described in this article. The methods for creating hash codes are agnostic of these particularities since they are only based on the topic distributions generated by the models.

3.2. Distance Metric

Since documents are described by set-type data, the proposed distance metric is based on the Jaccard coefficient. This metric computes the similarity of sets by looking at the relative size of their intersection as follows:

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|} \quad (5)$$

where A and B are set of topics.

More specifically, d_J is based on the Jaccard distance, which is obtained by subtracting the Jaccard coefficient J from 1:

$$d_J(A, B) = 1 - J(A, B) \quad (6)$$

The proposed distance measure d_H used to compare hash codes created from set of topics is the sum of the Jaccard distances d_j for each hierarchy level, i.e. for each set of topics:

$$d_H(H_1, H_2) = \sum_{l=1}^L \left(d_J(H_1(x_l), H_2(x_l)) \right) \quad (7)$$

where H_1 and H_2 are hash codes, $H_1(x_l)$ and $H_2(x_l)$ are the set of topics up to level l for each hash code H and L is the maximum hierarchy level. A corner case is $L = T$, where T is the number of topics in the model.

Threshold-based Hashing




Fig. 6. Threshold-based Hierarchical Hash ($L=3$)

3.3. Hash Function

The hash function clusters topics based on relevance levels. Three approaches are proposed depending on the criteria used to group topics: threshold-based, centroid-based and density-based.

3.3.1. Threshold-based Hierarchical Hashing Method

This approach is just an initial and naive way of grouping topics by threshold values into each relevance level. They can be manually defined or automatically generated by thresholds dividing the topic distributions as follows:

$$th_{inc} = \frac{1}{(L + 1) \cdot T} \quad (8)$$

where L is the number of hierarchy levels, and T the number of topics.

If $L = 3$ and $T = 10$ for a topic distribution td defined as follows:

$$td = [0.017, 0.141, 0.010, 0.172, 0.030, 0.090, 0.199, 0.133, 0.031, 0.171] \quad (9)$$

Then, a threshold-based hierarchical hash H_T , with an automatically created threshold defined by equation 8, is equals to $H_T = \{(t1, t3, t5, t6, t7, t9), (), (t4, t8)\}$ with $th_{inc} = 0.025$ (Fig 6).

3.3.2. Centroid-based Hierarchical Hashing Method

This approach assumes topic distributions can be partitioned into k clusters where each topic belongs to the cluster with the nearest mean score. It is based on


Centroid-based Hashing

Fig. 7. Centroid-based Hierarchical Hash (L=3)

the k-Means clustering algorithm, where k is obtained by adding 1 to the number of hierarchy levels. Unlike the previous method, threshold values used to define the hierarchy levels may vary between documents, i.e. for each topic distribution, since they are calculated for each distribution separately.

Following the previous example, if $L = 3$ and $T = 10$ for a topic distribution td defined in equation 9, then a centroid-based hierarchical hash H_C equals to $H_C = \{(t6), (t9, t7, t3, t1), (t5)\}$ (Fig 7).

3.3.3. Density-based Hierarchical Hashing Method

This approach also considers relative hierarchical thresholds for each relevance level. Now, a topic distribution is described by points in a single dimension. In this space, topics closely packed together are grouped together. This approach does not require a fixed number of groups. It only requires a maximum distance (eps) to consider two points close and grouped together. This value can be estimated from the own distribution of topics (e.g. variance).

Following the above example, if $L = 3$ and td is the topic distribution defined in equation 9, then a density-based hierarchical hash H_D is equals to $H_D = \{(t6), (t9, t3), (t1)\}$ when eps equals to the variance of the topic distribution (Fig 8).

3.4. Online-mode Hashing

Hashing methods are batch-mode learning models that require huge data for learning an optimal model and cannot handle unseen data. Recent work address on-line mode by learning algorithms [24] that get hashing model accommodate to each new pair of data. But these


Density-based Hashing

Fig. 8. Density-based Hierarchical Hash (L=3)

approaches require the hashing model to be updated during each round based on the new pairs of data.

Our methods rely on topic models to build hash codes. These models do not require to be updated to make inferences about data not seen during training. In this way, the proposed hashing algorithms can work on large-scale and real-time data, as the size and the novelty of the collection does not influence the annotation process.

4. Experiments

As mentioned above (Section 2), it is difficult to interpret the similarity score calculated by metrics in a probability space. Since all of them are based on adding the distance between each dimension of the model (eq. 1, 2 and 3), distributions that share a fair amount of the less representative topics may still get higher similarity values than those that share the most representative ones specially if the model has a high number of dimensions.

Figures 9 and 10 show overlapped topic distributions of two pairs of documents. In the first case (fig 9), none of the most representative topics of each document is shared between them. However, the similarity score calculated from divergence-based metrics (eq 2) is higher than in the second case (fig 10), where the most representative topic is shared (topic 26). This behavior is due to the sum of the distances between the less representative topics (i.e. topics with a low weight value) being greater than the sum of the distances between the most representative ones (i.e. topic with a high weight value). In high-dimensional models, that sum may be more representative than the one obtained with the most




Fig. 9. Topic Distribution of two documents. Similarity score, based on JSD, is equals to 0.74




Fig. 10. Topic Distribution of two documents. Similarity score, based on JSD, is equals to 0.71

relevant topics, which are fewer in number than the less relevant ones.

The following experiments aim to validate that **hash codes based on hierarchical set of topics not only make it possible to search for similar documents with high accuracy, but also to extend queries with new restrictions and to offer information that helps explaining why two documents are similar.**

4.1. Datasets and Evaluation Metrics

Three datasets [8] are used to validate the proposed approach. The OPEN-RESEARCH⁴ dataset consist of 500k research papers in Computer Science, Neuroscience, and Biomedical randomly selected from the Open Research Corpus [52]. The CORDIS⁵ dataset contains 100k documents describing research and innovation projects funded by the European Union under a framework programme since 1990. The PATENTS

dataset consists of 1M patents randomly selected from the USPTO⁶ collection. For each dataset, documents are mapped to two latent topic spaces with different dimensions using LDA. We perform parameter estimation using collapsed Gibbs sampling for LDA [19] from the open-source librAIry [6] software. It is a framework that combines natural language processing (NLP) techniques with machine learning algorithms on top of the Mallet toolkit [37], an open-source machine learning package. The number of topics varies to study their influence on the performance of the algorithm (i.e. CORDIS-70 indicates a latent space created with 70 topics).

Experiments use JS divergence as an information-theoretically motivated metric in the probabilistic space created by topic models. Since it is a smoothed and symmetric alternative to the KL divergence, which is a standard measure for comparing distributions [50], it has been extensively used as state-of-the-art metric over topic distributions in literature [49][1][36]. Our upper bound is created from the brute-force comparison of the reference documents with all documents in the collection to obtain the list of similar documents.

In this scenario the goal is to minimize the accuracy loss introduced by hashing algorithms. Since this is a large-scale problem and an accuracy-oriented task, recall is not a good measure to be considered and precision is only relevant for sets much smaller than the total size of data (between 3-5 candidates).

All the experimental results are averaged over random training/set partitions. For each topic space, 100 documents are selected as references, and the remaining documents as search space. As noted above, only p@5 will be used to report the results of the experiments.

4.2. Retrieving Similar Documents

It is challenging to create an exhaustive gold standard, given the significant amount of human labour that is required to get a comprehensive view of the subjects being covered in it. In order to overcome this problem, the list of similar documents to a given one is obtained after comparing the document with all the documents of the repository and sorting the result. We have observed that different distance functions perform similarly in this scenario (figs 1 to 4), so we have decided to use only the JS divergence (eq. 2) in our experiments.

⁴<https://labs.semanticscholar.org/corpus/>

⁵<https://data.europa.eu/euodp/data/dataset/cordisref-data>

⁶<https://www.uspto.gov/learning-and-resources/ip-policy/economic-research/research-datasets>

OPEN-RES-100 (p@5)

LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	0.22	0.20	0.86	1.00	0.66	0.80
3	0.23	0.20	0.87	1.00	0.81	1.00
4	0.27	0.20	0.89	1.00	0.86	1.00
5	0.27	0.20	0.92	1.00	0.89	1.00
6	0.27	0.20	0.94	1.00	0.92	1.00

Table 1

Precision at 5 (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on Open Research dataset using a model with 100 topics. LEVEL column indicates the number of hierarchies used.

OPEN-RES-500 (p@5)

LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	0.23	0.20	0.76	0.80	0.67	0.80
3	0.24	0.20	0.80	1.00	0.71	0.80
4	0.25	0.20	0.83	1.00	0.74	0.80
5	0.25	0.20	0.86	1.00	0.81	1.00
6	0.24	0.20	0.89	1.00	0.86	1.00

Table 2

Precision at 5 (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on Open Research dataset using a model with 500 topics. LEVEL column indicates the number of hierarchies used.

Only the top N documents obtained from this method are used as reference set to measure the performance of the algorithms proposed in this paper. The value of N is equals to 0.5% of the corpus size (i.e. if the corpus size is equal to 1000 elements, only the top 5 most similar documents are considered relevant for a given document). This value has been considered after reviewing datasets used in similar experiments [28][36]. In those experiments, the reference data is obtained from existing categories, and the minimum average between corpus size and categorized documents is around 0.5%.

Once the reference list of documents similar to a given one is defined, the most similar documents through the proposed methods (i.e. threshold-based hierarchical hashing method (thhm), centroid-based hierarchical hashing method (chhm) and density-based hierarchical hashing method (dhhm)) are also obtained. An inverted index has been implemented by using Apache Lucene⁷ as document repository. The source code of both the algorithms and tests is publicly available [8].

Let's look at an example to better understand the procedure. We want to measure the accuracy and data size ratio used to identify the top5 similar documents to

CORDIS-70 (p@5)

LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	0.18	0.20	0.92	1.00	0.66	0.70
3	0.20	0.20	0.92	1.00	0.80	0.80
4	0.22	0.20	0.94	1.00	0.86	1.00
5	0.23	0.20	0.91	1.00	0.89	1.00
6	0.19	0.20	0.92	1.00	0.91	1.00

Table 3

Precision at 5 (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on CORDIS dataset using a model with 70 topics. LEVEL column indicates the number of hierarchies used.

CORDIS-150 (p@5)

LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	0.19	0.20	0.88	1.00	0.78	0.80
3	0.19	0.20	0.92	1.00	0.80	1.00
4	0.25	0.20	0.91	1.00	0.82	1.00
5	0.25	0.20	0.91	1.00	0.83	1.00
6	0.27	0.20	0.91	1.00	0.86	1.00

Table 4

Precision at 5 (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on CORDIS dataset using a model with 150 topics. LEVEL column indicates the number of hierarchies used.

a new document d_1 from a corpus of 1000 documents. The similarity between d_1 and all the documents in the corpus is calculated based on JS divergence. The top50 (0.5%) documents with the highest values will be the set of documents considered as similar to d_1 . As we are going to use an ANN-based approach, we need the hash expressions of all documents to measure similarity. The data structure proposed in this work is a hierarchy of sets of topics, so that the most similar documents are those that share most of the topics at the highest levels of the hierarchy.

The representational model for this example only considers 8 topics, that is, a document is described by a vector with 8 dimensions where each dimension corresponds to a topic (i.e. $[t_0, t_1, t_2, t_3, t_4, t_5, t_6, t_7]$) and its value will be the weight of that topic in the document, for example $d_1 = [0.18, 0.15, 0.2, 0.05, 0.14, 0.11, 0.09, 0.08]$. The hierarchy level (L) will be equal to 2, i.e. the hash expression has two hierarchical sets of topics: $h = \{h_0, h_1\}$.

According to methods described at Section 3.3, there are 3 ways to create the hierarchical hash codes for documents:

1. threshold-based (thhm): 2 thresholds are defined as described in section 3.3.1, for example 0.15 and

⁷<http://lucene.apache.org>

0.1. h_0 includes the topics with a weight greater than 0.15, and h_1 the remaining topics with a weight greater than 0.1. Then $h_0 = \{t_0, t_1, t_2\}$ and $h_1 = \{t_4, t_5\}$. Based on the hash expression $h = \{(t_0, t_1, t_2), (t_4, t_5)\}$, the documents that share more topics in those levels (i.e $h_0 = (t_0 \text{ OR } t_1 \text{ OR } t_2)$, $h_1 = (t_4 \text{ OR } t_5)$) or in other levels but with less relevance are ordered. Since there are many topics in the expression, potentially many documents are similar when sharing at least one of them. This increases the data ratio. Accuracy is also affected, as the algorithm is not able to bring under the same bucket similar documents. In short, the hash expression is not representative of the document, for the given exploratory task.

2. **centroid-based (chhm)**: sets of topics are created using a clustering algorithm based on centroids as described in section 3.3.2. The cardinalities of the hierarchical groups are generally more uniform with this method. Since $k = L + 1 = 3$ in this example, $h_0 = \{t_0, t_2\}$ and $h_1 = \{t_1, t_4\}$. The number of representative topics at each level of the hierarchy is usually lower, and this causes the data ratio used to discover similar documents to decrease as well. This approach increases the precision because now the hierarchy is more selective to distinguish similar documents. However, the size of region of similar candidates is still high.
3. **density-based (dhhm)**: now the clustering algorithm is based on how dense certain regions in the topic relevance dimensions are as described in section 3.3.3. It can group topics that have unbalanced distributions and, therefore, generates more discriminating hash expressions than with the previous algorithm. In the example, we would have a hash expression like this: $h_0 = \{t_2\}$ and $h_1 = \{t_0\}$. This significantly reduces the data ratio used to discover similar documents and does not excessively penalize accuracy. Obviously, increasing L (i.e. number of hierarchies) increases precision, but with $L > 3$ that gain is not so significant.

As it can be seen in tables 1 to 6, the mean and median of precision are calculated to compare the performance of the methods. In this assessment environment, the variance is not robust-enough because score values don't follow a normal distribution. We consider the result obtained as significant, based on the fact that mean and median values are fairly close. The centroid-based method (chhm) and the density-based method (dhhm)

PATENTS-250 (p@5)						
LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	0.03	0.00	0.71	0.80	0.67	0.80
3	0.08	0.00	0.91	1.00	0.90	1.00
4	0.11	0.00	0.95	1.00	0.95	1.00
5	0.12	0.00	0.95	1.00	0.96	1.00
6	0.11	0.00	0.97	1.00	0.97	1.00

Table 5

Precision at 5 (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on Patents dataset using a model with 250 topics. LEVEL column indicates the number of hierarchies used.

PATENTS-750 (p@5)						
LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	0.02	0.00	0.77	0.80	0.76	0.80
3	0.04	0.00	0.94	1.00	0.95	1.00
4	0.06	0.00	0.97	1.00	0.97	1.00
5	0.08	0.00	0.97	1.00	0.97	1.00
6	0.06	0.00	0.97	1.00	0.97	1.00

Table 6

Precision at 5 (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on Patents dataset using a model with 750 topics. LEVEL column indicates the number of hierarchies used.

OPEN-RES-100 (data-ratio)						
LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	99.8	99.9	45.2	45.9	4.9	2.5
3	99.9	99.9	74.4	77.6	13.4	10.7
4	99.9	99.9	87.4	90.2	27.2	22.8
5	99.9	99.9	95.4	96.3	49.9	42.6
6	99.9	99.9	97.9	98.7	72.2	65.8

Table 7

Data size ratio used (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on Open Research dataset and 100 topics.

OPEN-RES-500 (data-ratio)						
LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	95.9	96.3	22.2	22.1	1.4	0.3
3	99.1	99.2	43.9	43.7	5.1	4.1
4	99.6	99.6	57.1	57.3	11.7	10.3
5	99.6	99.6	70.7	70.7	28.8	22.0
6	99.9	99.9	81.5	80.6	50.3	40.1

Table 8

Data size ratio used (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on Open Research dataset and 500 topics.

CORDIS-70 (data-ratio)

LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	99.9	99.9	51.3	56.3	5.1	5.0
3	99.9	99.9	84.8	89.5	10.5	10.6
4	99.9	99.9	96.1	97.6	20.8	19.5
5	99.9	99.9	98.9	99.4	35.0	32.7
6	99.9	99.9	99.7	99.8	53.1	51.2

Table 9

Data size ratio used (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on CORDIS dataset and 70 topics.

CORDIS-150 (data-ratio)

LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	99.9	99.9	40.9	41.2	3.1	2.9
3	99.9	99.9	75.3	76.7	6.2	6.1
4	99.9	99.9	90.0	92.1	12.1	11.8
5	99.9	99.9	96.4	96.9	21.6	20.6
6	99.9	99.9	98.1	98.9	36.5	33.9

Table 10

Data size ratio used (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on CORDIS dataset and 150 topics.

show a similar behaviour to the one offered by the use of brute force by means of JS divergence.

In terms of efficiency, we consider the times to compare pairs of topic distributions constant, and we focus on the number of comparisons needed. Thus, algorithms with larger candidate spaces will be less efficient than others when the accuracy in both is the same. Tables 7-12 show the percentage of the corpus used by each of the algorithms to discover similar documents. Tables 1-6 show the accuracy of each algorithm for each of these scenarios. Density-based algorithm (dhhm) shows better balance between accuracy and volume of information (efficiency). It uses smaller samples (i.e lower ratio size) than others in all tests and even when it only uses a subset that is a 6.2% (Table 10) of the entire corpus, it obtains an accuracy of 0.808 (Table 4).

The precision achieved by the algorithm based on density (dhhm), which is much more restrictive than the others, suggests that few topics are required to represent a document in order to obtain similar ones. In addition, the number of topics does not seem to influence the performance of the algorithms, since their precision values are similar among the datasets of the same corpus. This shows that hashing methods based on hierarchical set of topics are robust to models with different dimensions.

The behavior of the algorithms have also been analyzed when the number of topics in the model varies.

PATENTS-250 (data-ratio)

LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	99.9	99.9	43.2	32.7	35.1	23.0
3	99.9	100.0	82.4	100.0	78.2	100.0
4	99.9	100.0	96.5	100.0	95.1	100.0
5	99.9	99.9	99.2	100.0	98.9	100.0
6	100.0	100.0	99.8	100.0	99.7	100.0

Table 11

Data size ratio used (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on Patents dataset and 250 topics.

PATENTS-750 (data-ratio)

LEVEL	THHM		CHHM		DHHM	
	mean	median	mean	median	mean	median
2	99.9	100.0	35.2	23.6	31.8	19.9
3	99.9	99.9	81.4	99.8	79.6	98.8
4	99.9	99.9	96.5	99.9	95.5	99.5
5	97.7	96.6	99.0	99.9	98.6	99.7
6	99.1	98.6	99.7	99.9	99.5	99.8

Table 12

Data size ratio used (*mean* and *median*) of threshold-based (THHM), centroid-based (CHHM) and density-based (DHHM) hierarchical hashing methods on Patents dataset and 750 topics.

Models with 100, 200, 300, 400, 500, 600, 700, 800, 900 and 1000 topics were created from the CORDIS corpus. For each model, the p@5 of the hashing methods is calculated taking into account the hierarchy levels: 2, 3, 4, 5 and 6. Figures 11 to 13 show the results obtained for each algorithm. It can be seen how the performance, i.e precision, of each of the algorithms is not influenced by the dimensions of the model.

4.3. Exploration

In a certain domain, we may want to retrieve similar documents to one given. For example, searching for articles in the Biomedical domain that are similar to an article about Semantic Web. In terms of topics this kind of search requires to narrow down the initial search space to a subset with only documents that contain the topics that better describe the queried domain.

Existing hashing techniques based on a binary-code Hamming space do not allow to customize the search query beyond the reference document itself. However, the algorithms proposed in this work allow adding new restrictions to the initial query based on the reference document, since they use a hierarchy of set of topics as hash codes.

Through the following example we describe the workflow to enable such retrieval operations. For sim-




Fig. 11. Precision at 5 (*mean*) of threshold-based hashing method when number of topics varies in CORDIS dataset.




Fig. 12. Precision at 5 (*mean*) of centroid-based hashing method when number of topics varies in CORDIS dataset.




Fig. 13. Precision at 5 (*mean*) of density-based hashing method when number of topics varies in CORDIS dataset.

licity we consider hash expressions with only two hierarchy levels. The reference document d_1 has the following hash expression: $h = \{h_0, h_1\} = \{(t10), (t18)\}$.

The first query, Q_1 , searches for documents similar to the reference document d_1 among all documents in the corpus. One of the ways to formalise this query

OPEN-RESEARCH-100			
hash	q1	q2	ratio
thhm	499,755	160,660	67.8
chhm	356,111	1,976	99.44
dhhm	49,068	766	98.43

Table 13

Number of documents similar to a given one (q1) and also in a specific domain (q2) for threshold-based (thhm), centroid-based (chhm) and density-based (dhhm) hierarchical hashing methods.

looks like this: $Q_1 = h_0 : t10^{100}$ or $h_0 : t18^{50}$ or $h_1 : t10^{50}$ or $h_1 : t18^{100}$. It sets a maximum boost (100) when the same restrictions as the reference document ($t10$ in h_0 and $t18$ in h_1) are fulfilled, and a lower boost (50) for the others ($t18$ in h_0 and $t10$ in h_1). In the specific case of applying this query to the CORDIS dataset, we observed that most of the retrieved documents included topic $t18$ (fig 14).

But if we were only interested in similar documents to d_1 that have topic $t10$, we could restrict the previous query Q_1 to express this condition in the following way: $Q_2 = (h_0 : t10^{100}$ or $h_1 : t10^{50})$ and $(h_1 : t10^{50}$ or $h_1 : t18^{100})$. The result obtained by Q_2 (fig 14) shows that the condition has been considered since there is a balance between topics $t10$ and $t18$ among the documents similar to d_1 .

This type of restrictions based on the semantics offered by topics in the hash expression get enabled thanks to the methods proposed in this work.

5. Conclusions

The usefulness of topics created by probabilistic models when exploring collections of scientific articles on large-scale has been widely studied in the literature. Each document in the corpus is described by probability distributions that measure the presence of those topics in their content. These vectors can also be used to measure the similarity between documents by using metrics such as Jensen-Shannon divergence. But with large amounts of items in the collection, discovering the entire set of nearest neighbors to a given document would be infeasible. Due to the low storage cost and fast retrieval speed, hashing is one of the popular solutions for approximate nearest neighbors. However, existing hashing methods for probability distributions only focus on the efficiency of searches from a given document, without handling complex queries or offering hints about why one document is considered more similar than another. A new data structure is proposed to represent




Fig. 14. Most relevant topics in similar documents from using a document as query (Q1) and setting topic t10 as mandatory (Q2).

hash codes, based on topic hierarchies created from the topic distributions. This approach has proven to obtain high-precision results and can accommodate additional query restriction. This way of encoding documents can also help to understand why two documents are similar, based on the intersection of topics at hierarchies of relevance.

In this paper we have focused on (1) comparing the performance of topic-based hashing methods with respect to the distance metrics based on probability distributions (e.g. JS divergence), (2) their ability to support more complex queries based on topic-based filters and (3) the expressiveness of their annotations (topics hierarchically divided into groups with different degrees of semantic specificity) to justify the relations obtained.

A manually annotated corpus with content similarity relations would further confirm the ability of the metrics proposed in this paper to reflect similarity as humans perceive it. Ongoing work on this line includes the creation of questionnaires⁸ to more accurately capture how similar two documents are from the perspective of human evaluators who read them both. This is an ambitious task that need to deals with the evaluators' own interpretation of similarity. What an expert perceives as different (since his knowledge in the domain allows him to identify discrepancies between the two texts), may be considered as similar by an inexperienced user that might not be able to capture those fine grained differences.

The next steps in our research are to extend the metric proposed in this paper from the point of view of the perception of similarity that a human makes, and to perform a more in-depth investigation about the meaning of the topics grouped by levels of relevance.

6. Acknowledgments

This research was supported by the Spanish Ministerio de Economía, Industria y Competitividad and EU FEDER funds under the DATOS 4.0: RETOS Y SOLUCIONES - UPM Spanish national project (TIN2016-78011-C4-4-R)

References

- [1] N. Aletras, T. Baldwin, J. H. Lau, and M. Stevenson. Evaluating topic representations for exploring document collections. *Journal of the Association for Information Science and Technology*, 68(1):154–167, 2017. doi:10.1002/asi.23574.
- [2] N. Aletras and M. Stevenson. Measuring the similarity between automatically generated topics. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, volume 2, pages 22–27, Gothenburg, Sweden, 2014. ACM Press. doi:10.3115/v1/E14-4005.
- [3] A. Andoni and P. Indyk. Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions. In *Proceedings of 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 459–468. IEEE Press, 2006. doi:10.1109/FOCS.2006.49.
- [4] A. Andoni, P. Indyk, H. L. Nguyễn, and I. Razenshteyn. Beyond locality-sensitive hashing. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1018–1028, Portland, Oregon, USA, 2014. SIAM Press. doi:10.1137/1.9781611973402.76.
- [5] A. Andoni and I. Razenshteyn. Optimal Data-Dependent Hashing for Approximate Near Neighbors. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*, pages 793–801, Portland, Oregon, USA, 2015. ACM Press. doi:10.1145/2746539.2746553.
- [6] C. Badenes-Olmedo, J. L. Redondo-García, and O. Corcho. Distributing Text Mining tasks with librAIry. In *Proceedings of the 2017 ACM Symposium on Document Engineering*, pages 63–66, Valletta, Malta, 2017. ACM Press. doi:10.1145/3103010.3121040.
- [7] C. Badenes-Olmedo, J. L. Redondo-García, and O. Corcho. Efficient Clustering from Distributions over Topics. In *Proceedings of the Knowledge Capture Conference*, pages 1–8, Austin, Texas, USA, 2017. ACM Press. doi:10.1145/3148011.3148019.
- [8] C. Badenes-Olmedo, J. L. Redondo-García, and O. Corcho. Large-scale Topic-based Search Java Algorithm, 2019. doi:10.5281/zenodo.3465855.
- [9] D. Blei, L. Carin, and D. Dunson. Probabilistic Topic Models. *IEEE Signal Processing Magazine*, 55(4):77–84, 2010. doi:10.1109/MSP.2010.938079.
- [10] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3(4-5):993–1022, 2003. doi:10.1162/jmlr.2003.3.4-5.993.
- [11] A. Z. Broder, S. C. Glassman, M. S. Manasse, and G. Zweig. Syntactic clustering of the Web. *Computer Networks and ISDN Systems*, 29(8-13):1157–1166, 1997. doi:10.1016/S0169-7552(97)00031-7.
- [12] M. S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the 34th annual ACM symposium*

⁸<http://library.linkeddata.es/survey>

- on *Theory of computing*, pages 380–388, Montreal, Quebec, Canada, 2002. ACM Press. doi:10.1145/509907.509965.
- [13] X. Cheng, X. Yan, Y. Lan, and J. Guo. BTM: Topic modeling over short texts. *IEEE Transactions on Knowledge & Data Engineering*, 26(12):2928–2941, 2014. doi:10.1109/TKDE.2014.2313872.
- [14] M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the Twentieth Annual Symposium on Computational Geometry*, pages 253–262, Brooklyn, New York, USA, 2004. ACM Press. doi:10.1145/997817.997857.
- [15] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407, 1990. doi:10.1002/(SICI)1097-4571(199009)41:6<391::AID-ASII>3.0.CO;2-9.
- [16] D. M. Endres and J. E. Schindelin. A new metric for probability distributions. *IEEE Transactions on Information Theory*, 49(7):1858–1860, 2006. doi:10.1109/TIT.2003.813506.
- [17] C. J. Gatti, J. D. Brooks, and S. G. Nurre. A Historical Analysis of the Field of OR/MS using Topic Models. *Computing Research Repository*, abs/1510.0, 2015. arxivId:1510.05154.
- [18] D. Greene and J. P. Cross. Exploring the Political Agenda of the European Parliament Using a Dynamic Topic Modeling Approach. *Political Analysis*, 25(1):77–94, 2017. doi:10.1017/pan.2016.7.
- [19] T. L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(Supplement 1):5228–5235, 4 2004. doi:10.1073/pnas.0307752101.
- [20] T. L. Griffiths, M. Steyvers, and J. B. Tenenbaum. Topics in semantic representation. *Psychological Review*, 114(2):211–244, 2007. doi:10.1037/0033-295X.114.2.211.
- [21] D. Hall, D. Jurafsky, and C. D. Manning. Studying the history of ideas using topic models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, page 363, Morristown, NJ, USA, 2008. ACM Press. doi:10.3115/1613715.1613763.
- [22] J. He, L. Li, and X. Wu. A self-adaptive sliding window based topic model for non-uniform texts. In V. Raghavan, S. Aluru, G. Karypis, L. Miele, and W. Xindong, editors, *IEEE International Conference on Data Mining*, volume 2017-Novem, pages 147–156, New Orleans, LA, USA, 2017. IEEE Computer Society. doi:10.1109/ICDM.2017.24.
- [23] T. Hofmann. Probabilistic Latent Semantic Indexing. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, volume 51, pages 211–218. ACM Press, 2017. doi:10.1145/3130348.3130370.
- [24] L.-K. Huang, Q. Yang, and W.-S. Zheng. Online Hashing. *IEEE Transactions on Neural Networks and Learning Systems*, 29(6):2309–2322, 2018. doi:10.1109/TNNLS.2017.2689242.
- [25] P. Indyk and R. Motwani. Approximate nearest neighbors. In *Proceedings of the 30th annual ACM symposium on Theory of computing*, pages 604–613, Dallas, Texas, USA, 1998. ACM Press. doi:10.1145/276698.276876.
- [26] J. Ji, J. Li, S. Yan, Q. Tian, and B. Zhang. Min-Max Hash for Jaccard Similarity. In *Proceedings of 13th International Conference on Data Mining*, pages 301–309. IEEE, 12 2013.
- [27] K. Krstovski and D. A. Smith. A Minimally Supervised Approach for Detecting and Ranking Document Translation Pairs. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 207–216, Edinburgh, Scotland, UK., 2011. ACM Press. isbn:9781937284121.
- [28] K. Krstovski, D. A. Smith, H. M. Wallach, and A. McGregor. Efficient Nearest-Neighbor Search in the Probability Simplex. In *Proceedings of the 2013 Conference on the Theory of Information Retrieval*, pages 101–108, Copenhagen, Denmark, 2013. ACM Press. doi:10.1145/2499178.2499189.
- [29] B. Kulis and K. Grauman. Kernelized Locality-Sensitive Hashing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(6):1092–1104, 2012. doi:10.1109/TPAMI.2011.219.
- [30] P. Li and C. König. b-bit minwise hashing. In *Proceedings of the 19th International Conference on World Wide Web*, pages 671–680. ACM Press, 2010. doi:10.1145/1772690.1772759.
- [31] P. Li, A. B. Owen, and C. H. Zhang. One permutation hashing. In *Advances in Neural Information Processing Systems*, volume 4, pages 3113–3121. Curran Associates, Inc., 2012. isbn:9781627480031.
- [32] W. Liu, C. Mu, S. Kumar, and S.-F. Chang. Discrete Graph Hashing. In *Proceedings of the 27th International Conference on Neural Information Processing Systems*, volume 2, pages 3419–3427, Montreal, Canada, 2014. MIT Press.
- [33] Y. Liu, J. Cui, Z. Huang, H. Li, and H. T. Shen. K-LSH: An Efficient Index Structure for Approximate Nearest Neighbor Search. *Proceedings of the VLDB Endowment*, 7(9):745–756, 2014. doi:10.14778/2732939.2732947.
- [34] N. Ljubesic, D. Boras, N. Bakaric, and J. Njavro. Comparing measures of semantic similarity. In *Proceedings of the 30th International Conference on Information Technology Interfaces*, pages 675–682, Dubrovnik, Croatia, 2008. IEEE. doi:10.1109/ITI.2008.4588492.
- [35] H.-M. Lu, C.-P. Wei, and F.-Y. Hsiao. Modeling healthcare data using multiple-channel latent Dirichlet allocation. *Journal of Biomedical Informatics*, 60:210–223, 2016. doi:10.1016/j.jbi.2016.02.003.
- [36] X. Mao, B.-S. Feng, Y.-J. Hao, L. Nie, H. Huang, and G. Wen. S2JSD-LSH: A Locality-Sensitive Hashing Schema for Probability Distributions. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 3244–3251, San Francisco, California, USA, 2017. AAAI Press. issn:2374-3468.
- [37] A. McCallum. MALLETT: A Machine Learning for Language Toolkit, 2002.
- [38] A. Niekler and P. Jähnichen. Matching results of latent dirichlet allocation for text. In *Proceedings of the 11th International Conference on Cognitive Modeling (ICCM)*, pages 317–322, Berlin, Germany, 2012. isbn:9783798324084.
- [39] R. O’Donnell, Y. Wu, and Y. Zhou. Optimal Lower Bounds for Locality-Sensitive Hashing (Except When q is Tiny). *ACM Transactions on Computation Theory*, 6(1):1–13, 2014. doi:10.1145/2578221.
- [40] J. O’Neill, C. Robin, L. O’Brien, and P. Buitelaar. An analysis of topic modelling for legislative texts. In *Proceedings of the 2nd Workshop on Automated Semantic Analysis of Information in Legal Texts, co-located with the 16th International Conference on Artificial Intelligence and Law*, London, UK, 2016. CEUR-WS.org.
- [41] M. Paul and R. Girju. Topic Modeling of Research Fields: An Interdisciplinary Perspective. In *Proceedings of the International Conference Recent Advances in Natural Language Processing*, pages 337–342, Borovets, Bulgaria, 2009. ACL Press.
- [42] M. J. Paul and M. Dredze. Discovering Health Topics in Social Media Using Topic Models. *PLoS ONE*, 9(8):1–11, 2014.

- doi:10.1371/journal.pone.0103408.
- [43] S. Petrovic, M. Osborne, and V. Lavrenko. Streaming First Story Detection with application to Twitter. In *Proceedings of the 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 181–189, Los Angeles, California, USA, 2010. ACL Press. isbn:978-1-932432-65-7.
- [44] D. Ramage, S. Dumais, and D. Liebling. Characterizing Microblogs with Topic Models. In *International AAAI Conference on Weblogs and Social Media*, pages 1–8, Washington, DC, USA, 2010. American Association for Artificial Intelligence. isbn:9781577354451.
- [45] D. Ramage, E. Rosen, J. Chuang, C. D. Manning, and D. A. McFarland. Topic Modeling for the Social Sciences. 1, 2009.
- [46] D. Ravichandran, P. Pantel, and E. Hovy. Randomized algorithms and NLP. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 622–629, Morristown, NJ, USA, 2005. ACM Press. doi:10.3115/1219840.1219917.
- [47] M. D. Tapi Nzali, S. Bringay, C. Lavergne, C. Mollevi, and T. Opitz. What Patients Can Tell Us: Topic Analysis for Social Media on Breast Cancer. *JMIR Medical Informatics*, 5(3):e23, 2017. doi:10.2196/medinform.7779.
- [48] K. Terasawa and Y. Tanaka. Spherical LSH for Approximate Nearest Neighbor Search on Unit Hypersphere. In *Algorithms and Data Structures*, pages 27–38. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. doi:10.1007/978-3-540-73951-7_4.
- [49] W. B. Towne, C. P. Rosé, and J. D. Herbsleb. Measuring Similarity Similarly. *ACM Transactions on Intelligent Systems and Technology*, 8(1):1–28, 2016. doi:10.1145/2890510.
- [50] P. D. Trobisch, M. Bauman, K. Weise, F. Stuby, and D. J. Hak. Histologic analysis of ruptured quadriceps tendons. *Knee Surgery, Sports Traumatology, Arthroscopy*, 18(1):85–88, 2010. doi:10.1007/s00167-009-0884-z.
- [51] S. Vijayanarasimhan, P. Jain, and K. Grauman. Hashing Hyperplane Queries to Near Points with Applications to Large-Scale Active Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(2):276–288, 2014. doi:10.1109/TPAMI.2013.121.
- [52] A. Waleed, G. Dirk, B. Chandra, B. Iz, C. Miles, D. Doug, D. Jason, E. Ahmed, F. Sergey, H. Vu, K. Rodney, K. Sebastian, L. Kyle, M. Tyler, O. Hsu-Han, P. Matthew, P. Joanna, S. Sam, W. Lucy, Lu, W. Chris, Y. Zheng, v. Z. Madeleine, and E. Oren. Construction of the Literature Graph in Semantic Scholar. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 84–91, New Orleans, Louisiana, USA, 2018. ACM Press. doi:10.18653/v1/N18-3011.
- [53] H. M. Wallach, D. Mimno, and A. McCallum. Rethinking LDA: Why priors matter. In *Proceedings of the 22nd International Conference on Neural Information Processing Systems*, pages 1973–1981, Vancouver, British Columbia, Canada, 2009. Curran Associates Inc. isbn:9781615679119.
- [54] J. Wang, W. Liu, S. Kumar, and S.-F. Chang. Learning to Hash for Indexing Big Data-A Survey. *Proceedings of the IEEE*, 104(1):34–57, 2016. doi:10.1109/JPROC.2015.2487976.
- [55] L. Xing and M. J. Paul. Diagnosing and Improving Topic Models by Analyzing Posterior Variability. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 6005–6012, New Orleans, Louisiana USA, 2018. AAAI Press. isbn:978-1-57735-800-8.
- [56] T. Zhang, G. J. Qi, J. Tang, and J. Wang. Sparse composite quantization. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 07-12-June, pages 4548–4556, Boston, MA, USA, 2015. IEEE. doi:10.1109/CVPR.2015.7299085.
- [57] W.-L. Zhao, H. Jégou, and G. Gravier. Sim-min-hash. In *Proceedings of the 21st ACM international conference on Multimedia*, pages 577–580, Barcelona, Spain, 2013. ACM Press. doi:10.1145/2502081.2502152.
- [58] Y. Zhen, Y. Gao, D.-Y. Yeung, H. Zha, and X. Li. Spectral Multimodal Hashing and Its Application to Multimedia Retrieval. *IEEE Transactions on Cybernetics*, 46(1):27–38, 2016. doi:10.1109/TCYB.2015.2392052.