



# Topic space trajectories

## A case study on machine learning literature

Bastian Schaefermeier<sup>1</sup> · Gerd Stumme<sup>2</sup> · Tom Hanika<sup>2</sup>

Received: 26 October 2020 / Accepted: 26 February 2021  
© The Author(s) 2021

### Abstract

The annual number of publications at scientific venues, for example, conferences and journals, is growing quickly. Hence, even for researchers it becomes harder and harder to keep track of research topics and their progress. In this task, researchers can be supported by automated publication analysis. Yet, many such methods result in uninterpretable, purely numerical representations. As an attempt to support human analysts, we present *topic space trajectories*, a structure that allows for the comprehensible tracking of research topics. We demonstrate how these trajectories can be interpreted based on eight different analysis approaches. To obtain comprehensible results, we employ non-negative matrix factorization as well as suitable visualization techniques. We show the applicability of our approach on a publication corpus spanning 50 years of machine learning research from 32 publication venues. In addition to a thorough introduction of our method, our focus is on an extensive analysis of the results we achieved. Our novel analysis method may be employed for paper classification, for the prediction of future research topics, and for the recommendation of fitting conferences and journals for submitting unpublished work. An advantage in these applications over previous methods lies in the good interpretability of the results obtained through our methods.

**Keywords** Topic Models · Non-Negative Matrix Factorization · Multidimensional Scaling · Publication Dynamics · Interpretable Machine Learning

### Introduction

The number of publications published in scientific venues, such as journals and conferences, is vastly increasing. For instance, there were at least 5000 papers at major machine learning conferences and journals in 2018, more than twice as many as in 2008 (Ammar et al. 2018). Even though many venues emphasize a particular research field they do exhibit a plurality of topics. This is a natural consequence observed when research fields grow, which leads to new specializations and the emergence of new research topics. The related advance in knowledge is, however, overshadowed by the increasing inability to

---

Extended author information available on the last page of the article

maintain a comprehensive overview of a research field. Moreover, even for experts understanding and tracking the topic dynamics of contemporary research fields is an infeasible endeavor. Hence, the need for an automated approach to cope with the aforesaid vast amount of publication data is pressing.

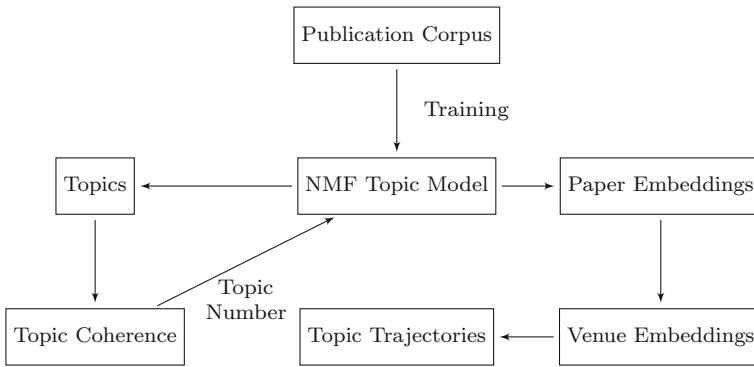
Answering to this we present a novel conceptualization for topic based analysis of scientific publication venues. Our approach is based on *topic models*, a common class of methods to analyze text corpora. By computing topic vectors for research papers we are able to position the associated publication venues in a low-dimensional topic space. This enables the comparison of conferences as well as journals and following their temporal (topical) dynamics. A governing constraint in our work is to compute human interpretable results. Thus, we employ non-negative matrix factorization (NMF), a model whose topics are comparatively well comprehensible. In detail, we aggregate document representations calculated by NMF over conferences/journals and years. This allows to capture the temporal topical development of scientific venues. Moreover, we semi-automatically select an appropriate number of topics based on a coherence measure. Additionally we take advantage of proper visualization techniques.

Altogether this results in the scientometric analysis of what we call *topic space trajectories* (TST). Despite the availability of almost all the building blocks, no one, to the best of our knowledge, has presented an akin notion for publication data analysis. We demonstrate the applicability of our approach on a publication data set of the top tier machine learning conferences covering the years from 1969 to 2018. In particular we track 32 publication venues in topic space and depict their topical drift. What is more, we introduce topic densities (i.e., distributions) for publication venues. Besides scientometric analysis we envision multiple applications of our work. First, we think that TST may be employed in conference or reviewer recommendations for new research papers in classification based approaches. Second, one may extrapolate the topical drift of conferences into the future. Third, our method may be transferred to the analysis of other text domains, such as news articles or patent applications.

Our work is organized as follows: First, we present our method for obtaining document and venue representations as well as their trajectories in topic space (“[Topic space trajectories](#)” section). In the subsequent (“[Case study on machine learning publications](#)” section), we introduce our case study on a machine learning publication corpus. We demonstrate for the resulting topic space a collection of interpretation approaches in (“[Interpreting topic space trajectories](#)” section). Before concluding our work in “[Conclusion](#)” section). we give an overview of related work in (“[Related work](#)” section).

## Topic space trajectories

In this section we describe our overall method, which is depicted in Fig. 1. The essential part is the publication corpus, from which we extract a subset of papers relevant for the analysis, for example, the subset of machine learning papers of certain venues. From this corpus we train a topic model using non-negative matrix factorization (NMF). Additionally to topics, we obtain embeddings (i.e., representations) of papers and venues, which will allow the definition of venue trajectories in topic space.



**Fig. 1** Depiction of our method. Each box contains some data or (intermediate) result. Arrows indicate where data or results are an input for the subsequent box

**Publication corpus**

We consider a corpus of publications  $\mathcal{D}$  as a set of papers, each of which appeared at a certain venue (conference or journal) in a certain year. At this point we omit a thorough discussion of preprocessing technicalities and defer the reader to (“[Case study on machine learning publications](#)” section). Equipped with this we may define what a publication corpus is in this work.

**Definition 2.1 (Publication Corpus)** A *publication corpus* is a ternary relation  $\mathcal{D} \subseteq P \times O \times Y$  where  $P$  is a set of papers,  $O$  is a set of publication venues (outlets) and  $Y \subset \mathbb{N}$  is a set of natural number indicating the publication years. Furthermore, it holds  $\forall (p, o, y) \in \mathcal{D} : |\{(o, y) \in O \times Y \mid (p, o, y) \in \mathcal{D}\}| = 1$ , i.e., there is a unique venue and year for each paper.

We may note that *papers*, *venues*, and *year* are just names. Hence, a publication corpus can be substituted by any structure bearing the same construction, e.g., news articles in certain newspapers on certain days.

**Document representation in topic space**

Say the number of papers (documents) in a publication corpus is  $d$  and the set of all words from all these documents is of size  $n$ . There are various methods for embedding these documents into a topic space (commonly referred to as *topic models*). Most of them employ a real-valued *word-document* matrix as starting point. Hence, we need to represent all elements  $p \in P$  as elements of  $\mathbb{R}_{\geq 0}^n$ . Therefore we consider in the following  $p$  to be an element of  $\mathbb{R}_{\geq 0}^n$ . In this representation each vector component denotes a word-weight. Such a weight for a particular word can be, for example, the word-frequency (*term-frequency*) or a more sophisticated approach like *tf-idf* (Ramos 2003). Having the vector representations of all publications we can construct the word-document matrix  $V$  simply by juxtaposing all those vectors. A *topic model* then finds a representation of that word-document matrix  $V$  as a product of two or more matrices. These factors represent particular structural elements of the resulting topic space, foremost the topics and some

representation for every document in this space, i.e., as a linear combination of *topic vectors*. An obvious goal here is that the number of topics  $t$  is substantially smaller than  $d$ .

A very prominent topic model, called *Latent Semantic Analysis* (LSA) (Deerwester et al. 1990), is based on the singular value decomposition of the word-document matrix. This method has been proven to work well for various natural language processing tasks, e.g., in Steinberger and Křišť'an (2007), in Pu and Yang (2006) or for recommendations of scientific articles in Chughtai et al. (2020). Nonetheless, we decided against its application in our work. The reason for this is that the results are difficult to comprehend or explain by humans. This problem arises from the fact that topics in LSA can contribute positively and negatively to the topical document representation. Another well-employed method is *latent Dirichlet allocation* (LDA) from Blei et al. (2003). It is based on a probabilistic model where documents are assumed to be generated from a distribution over topics, which themselves are distributions over words. While LDA generally achieves very good results, e.g., in text classification applications, it is known to fall short for small documents like paper abstracts. This has for example been shown empirically on Twitter posts in Hong and Davison (2010). Since we do prefer for our modeling a method that is able to cope with small documents we discard using LDA. The same holds for well-known investigated methods based on *doc2vec* (Le and Mikolov 2014). In doc2vec documents are embedded together with words in a real valued vector space of a chosen dimensionality. This is done such that related words and documents are mapped closely together and unrelated ones far from each other, with respect to some metric in that vector space. This method, however, does not directly result in a topic space representation of the documents. While nearby words of a document give some idea of its topic, these words may be different for each document. In contrast to this, we require that every document is represented by the same consistent set of topics. This allows for a meaningful comparison of documents in terms of their topics as well as the analysis of topic dynamics.

### Non-Negative Matrix Factorization

We now explain the used topic model NMF in more detail. Readers already familiar with or not concerned with the mathematical details of this method are advised to skip this section. We employ non-negative matrix factorization (NMF) (Lee and Seung 1999) for discovering paper topics and embedding papers into the topic space  $\mathbb{R}^t$  for a chosen dimensionality  $t \in \mathbb{N}_{>0}$ . We start with a word-document matrix  $V \in \mathbb{R}_{\geq 0}^{n \times d}$ , with  $n$  being the number of words and  $d$  the number of documents. In this matrix, each column is a document vector in which each component is a word weight. A word weight can be, for example, the term frequency in the document or the product of the term frequency and inverted document frequency (tf-idf). NMF finds an approximate factorization  $V \approx WH$ . The factorization is achieved through minimization of a *distance*  $d(V, WH)$ , typically achieved through multiplicative update rules. As a distance  $d$  we utilize (as commonly applied) the Frobenius norm of the matrix difference  $V - WH$ , i.e.,  $d(V, WH) := \|V - WH\|_F$ . The Frobenius norm  $\|A\|_F$  of an  $m \times n$  matrix  $A$  with entries  $a_{ij}$  is  $\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2$ . In the following we will address this distance also by the name *reconstruction error*. After the minimization the left factor  $W \in \mathbb{R}_{\geq 0}^{n \times t}$  contains in each column a topic vector of word weights. The factor  $H \in \mathbb{R}_{\geq 0}^{t \times d}$  contains in each column a representation of a document in topic space. Typically this representation is of much lower dimensionality than the original document representation, i.e.,  $t \ll n$ .

The reason for the superior interpretability of topic models obtained by NMF is that all elements in  $W$  and  $H$  are constrained to be positive. Hence, all components of a topic are additive. Additionally, in the NMF variant we use the topic vectors in  $W$  as well as the document vectors in  $H$  are constrained to have  $L_1$ -norm 1. Therefore, a document  $V_{*i}$  (denoting the  $i$ -th column of  $V$ ) can be (approximately) reconstructed as the weighted sum  $\sum_{j=1}^t H_{ji}W_{*j}$  of the topic vectors  $t_j := W_{*j}$  found through NMF. More precisely, any document is a convex combination of the topic vectors. The components  $H_{ji}$  are interpretable as the proportion of the topics  $j \in \{1 \dots t\}$  in the  $i$ -th document. The components of a topic vector  $t_j$  can be interpreted as the relative contribution of words to the topic. If the topic vectors  $t_j$  are linearly independent, they form a basis of the  $t$ -dimensional topic space. Hence, the vector  $d_i := H_{*i}$  contains the coordinates of the  $i$ -th document embedded into this topic space.

### Venue representation and trajectories in topic space

Since the ultimate goal of this work is to analyze the dynamics of publication venues in topic space, we need to find a (computable) representation of such venues, based on the so-far discussed document representations. Hence, we decided to represent publication venues through the centroid of their document embeddings. Employing centroids appeared to us as a natural modeling decision, since they represent the average of paper vectors, and thus the main topical research focus of a venue. Formally, for an index set  $I \subseteq \{1, \dots, d\}$  let the set  $D_I := \{d_i \mid i \in I\}$  be the representations of documents obtained through NMF, i.e., a set of the columns from matrix  $H$ . We then represent this set through their centroid  $r_I$ :

$$r_I := \frac{1}{|D_I|} \sum_{i \in I} d_i$$

A meaningful set  $I$  may index all documents from a specific venue from a specific year  $y \in Y$  or from a specific venue in a specific year. By abuse of notation we may refer to the index set of a specific venue  $o$  and year  $y$  by  $I(o, y)$ . Using this we may use the shorthand  $r_{I(o,y)}$  for the centroid of a specific venue  $o \in O$  in a specific year. We note that this centroid does exist if and only if the combination of  $y$  and  $o$  is present in the publication corpus.

**Definition 2.2 (Topic space trajectory (Venue))** For a publication corpus  $\mathcal{D} \subseteq P \times O \times Y$  and a venue  $o \in O$  we call

$$\tau(o) := \{(r_{I(o,y)}, y) \mid \exists p \in P : (p, o, y) \in \mathcal{D}\}$$

the *topic space trajectory (of venue  $o$ )*.

Hence, to obtain a trajectory of a venue, we calculate one centroid for each conference year. The set  $\tau(o)$  can easily be linearly ordered using the second element of all pairs  $(r_{I(o,y)}, y)$  and will therefore be considered a linearly ordered set in the following. This order in time justifies the name trajectory. The idea of venue trajectory is obviously adaptable to other entities, such as sets of venues or author trajectories in topic space.

## Calculating an optimal topic number

The utility of the just defined idea for topic space trajectories for analyzing a publication corpus depends on the properties of the concrete topic space. Most important is the number of topics. On the one hand this number has to be large enough to discriminate elements of the topic space. At the same time we require this number to be small in order to maintain a human interpretability. A general procedure to select an optimal topic number  $\hat{t}$ , is to calculate a topic model for different values of  $t$  as well as the value of a measure that rates the quality of the resulting topic model. For the final model, one then chooses the  $t$  with the maximum quality. In research on topic models and natural language processing in general, an established family of quality measures is called *coherence measures*. As the name suggests, the objective behind these is to evaluate the topical coherence of a set of words (which, for example, represent a topic), i.e., how strongly the words are interconnected semantically.

The concrete coherence measure we use is called *the  $C_V$  measure* and has been described in Rosner et al. (2014). The authors of that work have shown that amongst a variety of coherence measures,  $C_V$  has the highest correlation with topic interpretability ratings from human annotators. Hence, employing this measure coincides with our goal of selecting topics with good interpretability. We will now describe the measure in detail. Although we did not invent it, our contribution here is a much more concise explanation of how to calculate it. In the original work,  $C_V$  is embedded into an abstract framework for coherence measures. The explanation is therefore scattered in text and formulas across many pages. We also give an interpretation of the measure, which has not been done in the original work.

Generally speaking, the  $C_V$  measure for a topic is determined from co-occurrence statistics of the top- $n$  terms (ranked by their weights). The overall coherence of a model is calculated as the arithmetic mean of the topic coherences. The computation of the measure for one topic is achieved as follows: For each of the top  $n$  words  $w$  in the topic, a so-called *context vector*  $\mathbf{v}_w$  is created. The components of this context vector in the  $C_V$  measure are the normalized pointwise mutual information (*NPMI*) between  $w$  and all the top terms in the topic (including  $w$  itself). As an example, let  $n=4$  and let the top  $n$  terms of a topic be the set  $\{\text{search, query, engine, user}\}$ . Then the context vector for the term *search* is determined via

$$\mathbf{v}_{\text{search}} = \begin{pmatrix} \text{NPMI}(\text{search, search})^\gamma \\ \text{NPMI}(\text{search, query})^\gamma \\ \text{NPMI}(\text{search, engine})^\gamma \\ \text{NPMI}(\text{search, user})^\gamma \end{pmatrix}.$$

The parameter  $\gamma \in (0, \infty)$  is used to put more or less weight on higher values of the NPMI. The NPMI between two words  $w$  and  $v$  is calculated as follows:

$$\text{NPMI}(w, v) = \frac{\text{PMI}(w, v)}{-\log p(w, v) + \epsilon} = \frac{\log \frac{p(w, v) + \epsilon}{p(w)p(v)}}{-\log p(w, v) + \epsilon}$$

In the formula above,  $p(w)$  is the probability of the word  $w$  and  $p(w, v)$  is the probability of  $w$  and  $v$  occurring together. The value  $\epsilon$  is a small constant to avoid logarithms of zero. The probabilities are determined from the same corpus on which the topic model is trained. More specifically, a sliding window of size  $sw$  and step size 1 is put over each document.

Each step of the window over all documents defines what we will call here a *pseudo-document*. The probability of a word (or two words) is then calculated as the number of pseudo-documents in which the term occurs (or both terms occur) divided by the total number of pseudo-documents. More formally, let  $D$  be the set of pseudo-documents, and each  $d \in D$  a set of words  $w$ . Then  $p(w) := |\{d \in D \mid w \in d\}|/|D|$  and  $p(w, v) := |\{d \in D \mid \{w, v\} \subseteq d\}|/|D|$ . Once all context vectors have been determined, the coherence measure  $C_V$  can be calculated as follows:

$$C_V = \frac{1}{n} \sum_{i=1}^n s_{cos}(\mathbf{v}_{w_i}, \sum_{j=1}^n \mathbf{v}_{w_j})$$

The cosine similarity  $s_{cos}(w, v)$  between two vectors  $w$  and  $v$  is calculated by  $\frac{w \cdot v}{\|w\| \cdot \|v\|}$  and it equals the cosine of the angle  $\theta$  between the two vectors. Since the angle depends only on the direction of the two vectors, not their length, we can write  $s_{cos}(\mathbf{v}_{w_i}, \sum_{j=1}^n \mathbf{v}_{w_j}) = s_{cos}(\mathbf{v}_{w_i}, \frac{1}{n} \sum_{j=1}^n \mathbf{v}_{w_j})$ . The  $C_V$  measure can thus also be interpreted as the average cosine similarity between the context vectors and their centroid  $\frac{1}{n} \sum_{j=1}^n \mathbf{v}_{w_j}$ .

The  $C_V$  measure has three parameters: The number of top terms  $n$ , the size of the sliding window  $sw$  and the weight  $\gamma$ . We follow the recommendations and results from the paper Rosner et al. (2014) and the default values in gensim of  $n = 20$ ,  $sw = 110$  and  $\gamma = 1$ . Note that with our paper abstracts a window of size 110 often contains the whole document. Hence, it might be worth considering smaller window sizes in future research.

## Case study on machine learning publications

We conduct a case study for topic space trajectories applied to a machine learning paper corpus. The decision for using research articles from machine learning is motivated by the fact that we ourselves are researchers in that field and thus aware of its subfields and development. We claim that the choice of a more specialized topic does neither limit the applicability of our approach nor the usefulness of interpretation methods, which will be presented in Section 4.

### Dataset and filtering

We use the Semantic Scholar Open Research Corpus for our analyses (Ammar et al. 2018). Our version of this corpus is the dump from January 31st, 2019. We extract papers from the top 32 machine learning conferences and journals compiled in Kersting et al. (2019). We also added the *IJCAI* conference, which is another top-tier conference. For each paper, we have the paper title, the paper abstract as well as author and citation information. Hence, we emphasize here that our work is based on paper abstracts rather than full texts. Figure 2 shows the number of papers in our data set by venue and year. When all papers are included, these date back until 1969. However, in this figure we only depict years, in which more papers from different venues are included. Since our dump of the Open Research Corpus is from early 2019, for this year it only contains a small number of publications. Hence, we also leave out 2019 in the figure.

To ensure a good coverage of venues, we manually compared the number of papers of many randomly chosen years to the numbers counted in conference proceedings and journal volumes. We estimate that for the largest number of cases more than 80% of published papers are covered by our data. However, this is a very rough estimate and can





information, they either open up new problems on their own or, as in the case of keywords, can pose similar problems as paper titles. We hence restrict our work to incorporating only paper abstracts as a data source. Nonetheless, further data sources could open up interesting ways to the enhancements of our method in future research. Altogether, after this process, we are left with about 61 thousand machine learning papers prepared for our analyses, i.e., which are equipped with a venue name, a title, a year and an abstract text and, additionally, the title of the work.

## Preprocessing

For our analyses, we use the *gensim* library for the Python programming language (Řehůrek and Sojka 2010). This widely used library provides implementations of many different topic models together with various methods for natural language processing tasks, e.g., preprocessing. For each paper, we only use the text from the paper abstract. In Scholz et al. (2014) it has been shown that paper abstracts are sufficient to represent the content of publications and even can outperform models using the full paper text. This is a natural consequence of the fact that abstracts consist of a condensed summary of the research presented in the paper, but without additional, more “noisy” discourses that might be provided in the full text. We proceed by converting the texts to lower case, tokenizing and removing stop words with the method *remove\_stopwords* from *gensim*, which is based on a hard-coded list of stop words that apparently has been designed by the *gensim* developers. Stop words are words like “*the*”, “*such*” or “*and*”, i.e., words which frequently occur in text but do not provide much topical information. The concrete list can be found in the *gensim* source code<sup>1</sup>. In natural language processing tasks, attempts are often made to reduce different word forms to a common word stem, e.g., through the Porter stemming algorithm. Stemming reduces the size of the vocabulary and thus leads to increased performance (in terms of speed) of subsequent calculations. However, it is also known to lead to word representations that are often no “real words” anymore. It is also known to reduce unrelated words to the same stem (e.g. *university* and *universe* both being mapped to *univers*). Due to these disadvantages we do not use stemming. Finally, we represent each document in a vector space using tf-idf weights. In this, we remove all terms from the model that occur less than ten times in our data set. Our aim here was to speed up subsequent calculations considerably without losing the very important words. The cut-off value is of course dependent on the corpus size. We determined its order of magnitude by looking at the term distribution that we will address in more detail in Section (“[Abstract length and term distribution](#)”). The reasoning here is that many terms exist, which are very unlikely to occur in a document and thus do not contribute much to calculated topics. Through their removal, we considerably reduce the size of the vocabulary and hence of the matrices in NMF. This improves memory usage as well as calculation times. A similar approach has been used in van Eck and Waltman (2007a), where concepts (small sets of words) appearing in less than ten documents were removed from analysis. Through this process, we obtain the  $n \times d$  word-document matrix  $V = (w)_{ij}$  with entry  $w_{ij}$  containing the tf-idf-weight of word  $i$  in document  $j$ . The column vectors of  $V$  equate the papers  $P$ , which together with other data from the Open Research Corpus form a publication corpus  $\mathcal{D}$ , cf Definition 2.1.

<sup>1</sup> <https://github.com/RaRe-Technologies/gensim/blob/develop/gensim/parsing/preprocessing.py>.


## Matrix factorization

After having obtained a word-document-matrix  $V$  as described in the previous section, we apply to it the NMF-implementation from *gensim* using the default parameters. Since the minimization process in this implementation of NMF relies on sampling as well as randomly initialized matrices, it is nondeterministic. Hence, we obtain different results in each run of the algorithm. Consistent to the findings in Belford et al. (2018), our experience is that these results are nonetheless very stable. This especially holds in comparison to LDA, with which we also experimented. To determine the topical stability, we manually compared the rankings of the top terms of the learned topics. In almost all instances these ranked terms were nearly the same with slight variations in their order. Only some few topics were totally different in their content when running the learning algorithm repeatedly. We also briefly experimented with the parameters of NMF, namely the initialization method (e.g., to a non-negative double singular value decomposition instead of random values) and the minimized reconstruction error, which can be changed to the Kullback-Leibler divergence. Gensim does not have parameters to change the initialization method or the objective function of the minimization. We therefore experimented for this with the implementation from the python package *scikit-learn*. Subjectively we did not observe clear improvements or even strong differences and hence relied on the default parameters of the *gensim* implementation. To test for convergence of the *gensim* implementation of NMF, we plotted the error curves for repeated runs of the algorithm and with different parameters (such as a larger number of iterations over the data set). We found the error curve to very consistently converge to a low value after having seen about one third of our data set. Hence, no further parameter tweaking was necessary in our case.

## Optimal topic number

We calculated the coherence measure  $C_V$  for every topic number from 1 to 25. The results of these calculations are given in Fig. 3 (left). We also calculated  $C_V$  for some larger values, namely for each topic number  $t$  from the set  $\{30, 50, 100, 200, 300, 400, 500\}$ . The results are not depicted here, since the coherences did not improve. We achieved the maximum coherence with a topic number of  $t = 22$ . We therefore use 22 topics in our subsequent analysis. Note that for each topic number, we trained one single instance of the NMF model. We did this with a fixed random seed to ensure reproducibility. NMF is nondeterministic and the achieved model quality could actually vary from run to run. Our reasoning here was that NMF is still relatively stable. However, for future work we would recommend to train multiple models for each topic number and use the one with the maximum coherence. Using this method the full process of finding a model with a good topic number can be automated.

Before the  $C_V$  measure, we experimented with other heuristics to select a topic number. As an example, we tried to take advantage of the matrix reconstruction error mentioned in (“[Document representation in topic space](#)” section). This error optimally should be a monotonically decreasing function of the number of topics, since the convex combination of more topic vectors should allow for an equally or more accurate document reconstruction. Our intuition was that plotting the reconstruction error for an increasing number of topics should result in a curve that starts with a strong (negative) slope and slowly converges to a minimal error with a slope close to zero. Our idea was then to select the topic number through the elbow criterion, i.e., by finding a point where the further decrease



**Fig. 3** We present for different topic numbers **Left:** The coherence measure  $C_V$ . **Right:** The matrix reconstruction error


of the error with increasing  $t$  is neglectable. Figure 3 (right) shows the curve resulting from this approach. While it matches our intuition, the optimal number of topics would lie somewhere near 500. We find this number too large to give a comprehensible overview over research topics. We conjecture, however, that our approach would lead to success in selecting a good model for other downstream tasks such as classification.

**Abstract length and term distribution**

Since we work with relatively short paper abstracts instead of full paper texts, we deem it necessary to get an overview over some statistical properties of our data. Even more so, since our texts are shortened further through stop word removal and since real world data sets are sometimes erroneous, we need to ensure that mostly enough words are left with actual topical information. We also want to categorize our data in terms of a comparison to data sets from other text domains, to which our analysis methods could be applied.

First, we look at the distribution of the abstract lengths. In Fig. 4, this distribution is given before and after stop word removal. We see that the average abstract consists of about 152 words, while the median is 145 due to a skewed distribution and some exceptionally long outliers. Paper abstracts are therefore longer than Twitter posts, which originally had a character limit of 140 that has nowadays been increased to 280. Hence, more information about the topic is included in a paper abstract. The order of magnitude, however, is still similar if we assume an average word length of five characters (a common assumption made in typing speed tests). Other document types such as web sites typically are several orders of magnitude larger in their word number. Once we remove stop words from the abstracts, their average length decreases quite substantially to only 92 words and a median of 86. Hence, about one third of words in a paper abstract are stop words with little information.

Next, we analyze how frequently different words (terms) occur. Figure 5 shows the results of this analysis in a log-log-plot before (left) and after stop word removal (right). In the left plot we notice that before stop word removal, term frequencies follow Zipf’s law (Zipf 1949), as expected. Zipf’s law is the empirical observation that when terms are ranked by their frequency of occurrence, this frequency is distributed inversely proportional to their rank, i.e.,  $P_n \sim 1/r^a$  for some  $a > 0$ . Stop words are words which occur frequently while containing little information. Hence, after their removal, highly ranked



**Fig. 4** Abstract length distribution. The figure shows histograms of abstract lengths (counted as the number of words) in our data set before (left) and after stop word removal (right). For this plot, we removed outliers, i.e., papers with extremely long abstracts, which occur due to wrong data in the corpus. Vertical lines give the mean, median, and standard deviation

words are less frequent in the distribution as can be seen in Fig. 5 (right). This leads to the removal of words that are important for the matrix reconstruction error but have no use for topic discrimination.


### Interpreting topic space trajectories

After having described the circumstances of our case study, we are now presenting the actual topic space trajectories that we calculated from our paper corpus. What is more, we give interpretations and visualizations of our findings. While we conducted our study on a particular research field that we are familiar with, our approach here could be applied to other fields and text domains. Hence, our concrete analysis at this point simultaneously serves as an example of a generalizable approach.

### Found topics

Table 1 shows the topics identified through NMF. Each topic is represented by its top ten terms, determined from the word weights in its topic vector. We manually assigned a number and a name to each topic. The name is based on our own interpretation of the top words in a topic. For this, we considered up to 50 terms per topic. We may refer to some of these additional top words where we consider it useful for the interpretation of a topic.

Most topics found through NMF are well interpretable and often clearly correspond to one specific research area. To name some, we found topics related to Bayesian inference, neural networks, nonlinear control, optimization, social media, clustering, semantic web,



**Fig. 5** Term distribution. The x-axis in both plots shows the rank of each term when sorted in descending order of term frequency. The y-axis shows the term frequencies. Term frequencies are given before (left) and after stop word removal (right)

**Table 1** Topics identified through NMF. Each topic is represented here by the top ranked terms. Terms are given in order of their rank (i.e., from highest to lowest word weight). Topics are numbered arbitrarily for reference. We manually assigned names to topics based on our own interpretation. Sometimes interpretations reflect a tendency or focus of a topic

<b>1 Bayesian Inference</b>	<b>2 Search Engines</b>	<b>3 Neural Networks</b>	<b>4 Nonlinear Control</b>	<b>5 Optimization</b>
inference models bayesian model variational latent probabilistic variables posterior gaussian	search query queries engine engines user click ranking web users	network networks neural nodes layer deep training node recurrent hidden	control controller adaptive nonlinear systems tracking feedback nn robot loop	convex optimization algorithm function functions convergence gradient linear problems algorithms
<b>6 Neurons, Dynamic Networks</b>	<b>7 Classification, Pattern Mining</b>	<b>8 Information Retrieval</b>	<b>9 Social Media</b>	<b>10 Clustering</b>
neurons neural spike neuron time synaptic input activity spiking firing	tree classifiers classifier training mining decision data classification trees ensemble	retrieval topic topics information relevance ir models language model documents	social users media twitter content influence online services user people	clustering clusters cluster data algorithm means spectral algorithms similarity sets
<b>11 Learning &amp; Knowledge Bases</b>	<b>12 Semantic Web</b>	<b>13 Recommender Systems</b>	<b>14 Graphs</b>	<b>15 Reinforcement Learning</b>
learning task knowledge tasks entity entities text fuzzy online transfer	web pages page feature selection data semantic content information features	user recommendation users items item recommender recommendations collaborative filtering preferences	graph graphs nodes edges subgraph node structure mining vertices patterns	policy reinforcement agent agents reward action policies value learning decision
<b>16 Planning &amp; Reasoning, Association Rules, Logic</b>	<b>17 Document Retrieval, STS-Tags</b>	<b>18 Feature Extraction, Dimension Reduction</b>	<b>19 Support Vector Machines</b>	<b>20 Kernel Methods</b>
rules knowledge rule logic language reasoning planning semantic semantics domain	document documents inline retrieval query math tex formula term xml	feature multi features classification class dimensional sparse view supervised dimensionality	kernel svm kernels vector stability support matrix machines svms delays	density regression estimation kernel estimator gaussian distribution estimators probability distributions
<b>21 Matrix Methods</b>	<b>22 Image Recognition</b>			
matrix label rank domain labels learning target labeled data loss	image images visual object objects recognition segmentation features scene spatial			

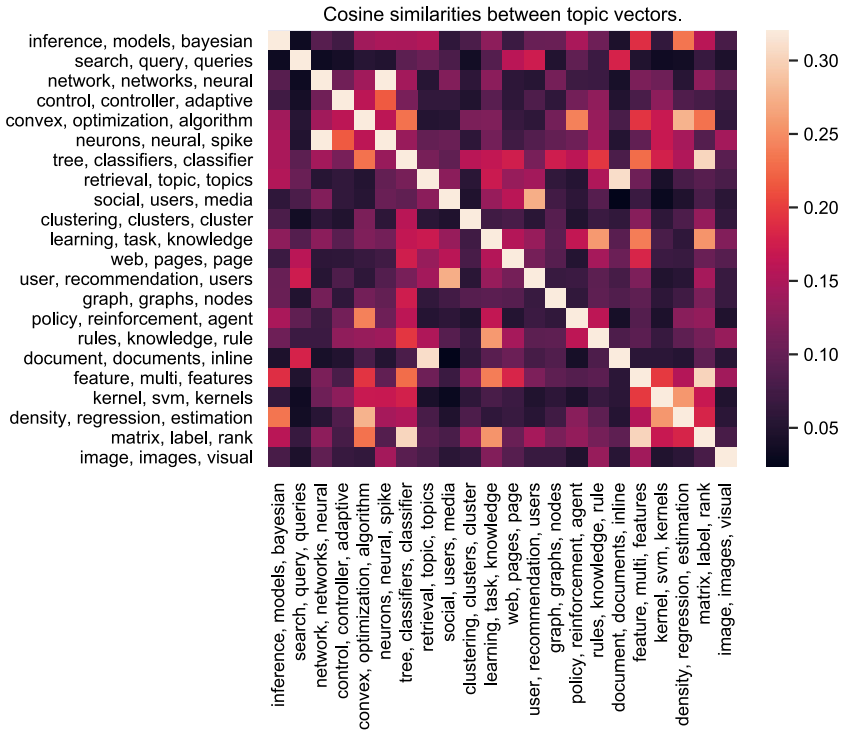
recommender systems, graphs, reinforcement learning and image recognition. However, we also found a few topics, which allow for ambiguous interpretations. As an example, topic 7 seems to be a mixture of pattern mining and classification methods. Topic 11 is a topic that appears to be related to different types of *learning*, e.g. (based on more than the top ten words), *transfer learning*, *online learning*, *representation learning*, *reinforcement learning* and *machine learning*. As such, it could be interpreted as a general topic on different concepts of machine learning, i.e., dealing with different approaches on *how* and *what* to learn. Additionally, this topic is related to knowledge bases. For some research areas, we found several topics with different focusses. As an example, we found two topics on neural networks (topic 3 and 6). Topic 3 here is more concerned with the architecture of networks (e.g., containing the words *layers*, *architecture*, *structure*, *convolutional*, *feed-forward*). Topic 6 puts more emphasis on neurons and the biological motivation of neural networks (e.g., *neuron*, *firing*, *stimulus*, *synapses*, *brain*, *signal*, *cortical*). Additionally, topic 6 puts more emphasis on advanced, dynamic neural networks (*spiking*, *temporal*, *recurrent*, *memory*).

Topic 17 is a mixture of document retrieval and XML-Tags from the Standards Tag Suite (STS), e.g., *tex*, *math*, *formula*. STS is an XML format used by publishers to exchange documents. In our data, STS was used for publications from one venue. Optimally, only the content from these tags should be parsed and added to the document representation. We spared this effort since few documents are concerned and a complicated checkup of the format of all documents with a subsequent parsing process would be required. Topic 18 is a mixture of feature extraction and methods related to dimensionality reduction (further top terms are *manifold*, *subspace*, *embedding*). We consider this a sensible mixture since these two topics are strongly related, i.e., dimensionality reduction methods are often used to extract features from data.

Altogether we found that NMF gives topics with good interpretability. We also found some limitations of the method. In some cases, two different topics are mixed together, although they are not strongly related semantically, e.g., in topic 7 *pattern mining* and *classification*. We surmise this behavior of NMF is encouraged when some third terms often co-occur with both topics, e.g., here *algorithm* and *data mining*. It could especially be encouraged through polysemic or homonymous terms (i.e., terms that, in a different context, have slightly or totally different meanings). Second, NMF sometimes learns two topics that could be one, e.g., two on neural networks. This behavior tends to occur for topics that are overrepresented in the training corpus. Where desired, sub- or supersampling based on (research) categories and paper numbers of venues could therefore mitigate such results. A third limitation of NMF is its flat structure. NMF hence fails to convey the taxonomy of topics, e.g., search engines being a subtopic of information retrieval. However, this lack of complexity is at the same time an advantage, since it improves the comprehensibility.

## Topic Similarities

In Fig. 6 we depict the cosine similarities between the calculated topic vectors  $t_j$  with  $j \in \{1 \dots t\}$ . Calculating the cosine similarity is a method often used in information retrieval to compare word weights of document vectors. The cosine similarity measures the cosine of the angle between two vectors. For vectors with non-negative components it lies between zero and one. Its maximum is reached for an angle of zero, i.e., when both vectors



**Fig. 6** Cosine similarities  $s_{cos}(x, y)$  between all pairs of topic vectors  $x$  and  $y$ . Note that similarities are symmetric, i.e., for all  $x, y$  we have  $s_{cos}(x, y) = s_{cos}(y, x)$ . Diagonal elements have a cosine similarity of 1. Best viewed in color

point into the same direction. Its minimum is reached here when the vectors are orthogonal. From the plot we can see how topics are related to each other. More specifically, brighter cells show that two topics have more similar word weights. Because the original topics have a dimension of more than 14000, it would be laborious to analyze this through a direct comparison of the word weights.

In the plot we see that most topic pairs have a low similarity. This indicates that most learned topics are a strong feature by their own and hence should enable us to represent well the variety of our data in topic space without redundancies. In the plot we notice that two topics on neural networks are most strongly related (“*network, networks*” and “*neurons, neural*”), but still have a low cosine similarity of about 0.3. We further see that two topics on information retrieval are closely related (“*retrieval, topics*” and “*document, documents*”). Such observations show that topics which come from the same research field (or supercategory) lead to more similar topic vectors. Turning this argument upside down, we can to some degree confirm or rebut our interpretations of the topic vectors. This is because similar topic vectors indicate similar research fields. It is interesting to notice that in some cases topics bear a comparatively high similarity to a variety of different topics. As an example, the topic “*tree, classifiers*” is similar to almost each of the other topics. The topic “*convex, optimization*” is similar to such topics as “*tree, classifiers*”, “*policy, reinforcement*”, “*feature, multi*”, “*density, regression*”, “*kernel, svm*”. Supposedly, such

cases occur due to a strong co-occurrence with other research topics and sometimes, these co-occurrences are also related semantically. As an example, optimization methods are often applied to a loss function for the training of classifiers such as some tree-based methods and support vector machines (SVMs), for fitting regression models, and in reinforcement learning. Hence, also the related topics occur together. The same holds for the topic on Bayesian inference, which is applied to a plethora of machine learning problems.


The lesson learned here is that research topics can be similar for at least three reasons: 1) They have the same supercategory. 2) One is a subcategory of the other. 3) One field is often applied to the other. Hence, calculating topic similarities enables us to interpret topics semantically to some degree. Note that the three cases can sometimes be hard to distinguish since learned topics sometimes do not reflect a pure, single research field. Methods for an automated analysis of these relations would open some interesting further research.

### Overall historical interest in topics

In Fig. 7 (top), we illustrate the topic trajectory for all papers in our data set from the years 1987 up to 2018. For this, we calculated one centroid per publication year. We depict the trajectory as chronologically sorted stacked bars, where each column depicts the topic space representation of one year. In Fig. 7 (bottom), we calculated the sum of the document vectors without dividing by the publication number, as done for the centroids. Through this method, the topic weights are given proportional to the total number of papers in a year instead of relative to a total sum of 1. Note that topic names and numbers correspond to the topics in Table 1. We sometimes used shortened names to improve the overview in our visualizations and will repeat on doing so in the remainder of this paper. However, we emphasize here that topics might comprise further notions than visible from the short names alone.

We now analyze some striking results. For this, we refer to topics by our manually assigned topic names. Observing the topics from 1987 up to 2000 and comparing these with the years from 2001 to 2018, it becomes apparent that three topics with the initially largest weights lose their relative importance over the years. These topics are topic 16 (*Planning & Reasoning*), topic 3 (*Neural Networks*) and topic 6 (*Neurons & Dynamic Neural Networks*). Note that the bars here are influenced by the number of venues, that published papers in a specific year. This is why biennially some bars become larger and smaller. The IJCAI conference here was held every second year in uneven years, i.e., those where aforementioned topics have more weight. We also have venues that did not yet exist at the beginning of our analysis. In 2001, for example, the Journal on Machine Learning Research (JMLR) was introduced (cf. Fig. 2). 2001 is also the first year in which a considerable amount of publications from the WWW conference appears in our data set (although founded earlier). Both these facts led to a larger number of publications in their specific research areas. Despite these influences, the overall tendencies of decline and raise in topic weights are visible. The mentioned decline of interest in neural networks took place at a time, when support vector machines (SVMs) became popular as a more efficient alternative (starting around mid-90s). SVMs stayed a widely used machine learning method. Neural networks, however, gained more interest again in more recent years. These facts become visible from our visualization and coincide with our personal background knowledge.





**Fig. 7** Topic importances by years. Columns are chronologically sorted by years. For each year, the relative importance of each topic is given by the proportion of the bar height. In the top image, the stacked bars are normalized to height 1. In the bottom plot, the total height of stacked bars depicts the number of papers in our data set. Best viewed in color

Further topics that have largely gained in interest include *Social Media*, *Recommender Systems*, *Optimization*, *Matrix Methods* and *Bayesian Inference*, among others. Social media became more popular through platforms like Facebook or Twitter. The interest in recommender systems has been promoted through the *Netflix prize*, a public contest on recommender systems that started in 2005, and through the introduction of the *RecSys* conference in 2007. Besides this, the growth of online platforms such as *Amazon* promoted the interest. Optimization, Bayesian inference and matrix methods have proven useful techniques that can be applied to a plethora of machine learning approaches. Topics that have recently lost weight in research are *Semantic Web*, *Search Engines* as well as *Clustering* and *Classification & Pattern Mining*. Again, this coincides with our personal

intuition. These latter two, very general topics have already been explored deeply in research and were often replaced by more specific problems, applications and methods.

Altogether, we notice that some machine learning topics gain, some lose relative importance. Sometimes, topics are almost invisible at the beginning and grow over time. This is a hint that these research topics just emerged or became more popular due to a certain event. Usually, popularity for machine learning methods is triggered by some milestone in their performance, i.e., when beating a benchmark on a data set by some large margin over previous methods. Sometimes other events trigger interest, such as public contests or the release of platforms.

We also aggregated topics over all papers in our data set without differentiating by year. This gives us the overall prevalence of topics in our data set. What we found here was, that the most prevalent topics in our data set are *Optimization*, *Classification & Pattern Mining* and the two topics on *Neural Networks*. Each of these topics makes out about 6-7% of our corpus. This is expected due to the many venues with large publication numbers on these topics in our data set. The prevalence of the least important topics, *Document Retrieval/STS-Tags* and *Clustering*, is only about half as high (about 3%). This is partly due to their loss in popularity in recent years, where the overall publication numbers were higher.

### Analyzing venue similarities through topical maps

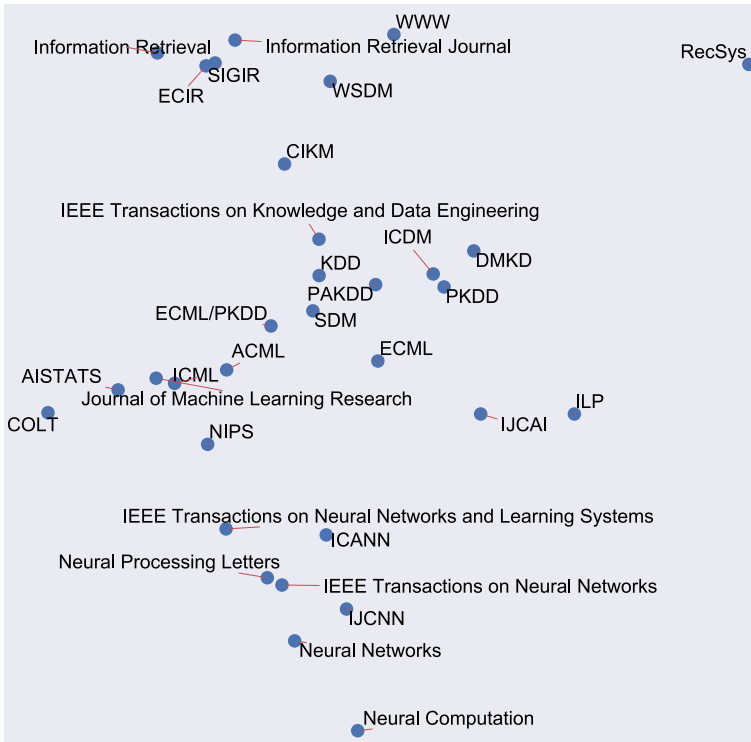
Topic space embeddings and trajectories are of a dimensionality that, in general, cannot directly be depicted in a coordinate system. On the other hand, it is often useful to visualize data in such a way. This is because it allows to find similarities and differences between entities (e.g., venues), indicated by nearness (or distance) in the plot. To be able to visualize high-dimensional data in such a way, we employ a well-known technique called multidimensional scaling (MDS), cf. Mead (1992). The idea behind MDS is to layout high-dimensional vectors in a low-dimensional space while preserving distances as best as possible. The low-dimensional representation of each input vector is found based on the squared differences between the pairwise distances of vectors in the input and the output space. More precisely, starting from random coordinates, points are aligned in the low-dimensional space s.t. an objective function is minimized as follows:

$$\min_{\hat{x}_1, \dots, \hat{x}_n} \sum_{i=1}^n \sum_{j=1}^n (d(x_i, x_j) - \|\hat{x}_i - \hat{x}_j\|_2)^2$$

In this,  $n$  is the number of vectors,  $x_i$  is one (high-dimensional) input vector and  $\hat{x}_i$  is the corresponding low-dimensional output vector to be determined through MDS. The function  $d$  is a measure of distance in the input space. Here, we use the Euclidean distance, i.e.,  $d(x, y) = \|x - y\|_2$ . Note that in theory a distance measure for compositional data such as the Aitchison distance (Martín-Fernández et al. 1998) or one between probability distributions, e.g., those presented in Schaefermeier et al. (2019), would be more suitable for our data. In our investigation, however, we found that the Euclidean distance leads to a better separation of venues into three different research fields, namely *neural networks*, *information retrieval* and *general machine learning*.

We use a dimensionality of two for the vectors  $\hat{x}_i$ , as this allows for good visualizations. In the resulting space, the two dimensions are not comprehensible as topics. Nonetheless, in this space we can analyze the topical similarity of venues based on their distance.

In Fig. 8 we depict topical representations of venues, calculated as the centroids of all their papers' topic vectors. We projected these centroids into a two-dimensional space



**Fig. 8** Topical map of venues. Dots show venue representations after dimensionality reduction through MDS. Closeness indicates topical similarity

using MDS. In this figure, the closer points are together, the more similar the research topics of their corresponding venues are. Therefore we call a visualization like in Fig. 8 a *topical map* of the venues from a publication corpus.

An interesting observation here is that venues cluster into three areas: In the top area we have venues that fall into the information retrieval category. In the middle part we have the more general machine learning venues. These two clusters are separated by *CIKM* (*Conference on Information and Knowledge Management*), which appears to fall somewhere in between both worlds. In the bottom area, we have conferences and journals specialized on artificial neural networks. *RecSys* (*Recommender Systems Conference*) in the top right of the plot is a conference with a strong topical focus on recommender systems. Hence, it appears to be a category by itself. However, *RecSys* lies most closely to the information retrieval world and is most dissimilar from neural networks.

Further interesting patterns emerge once we look at the specific venues. *NIPS* (*Neural Information Processing Systems*), as an example, was founded as a conference situated closer towards the neural networks topic. Over time, however, it developed into a more general machine learning conference, as can be seen by looking at current conference proceedings. Note that as a general tendency, the number of published papers per venue is growing from year to year as noticeable from Figs. 2 and 7 (bottom). Hence, more recent publication years have a stronger influence on the centroid of a venue. This explains why *NIPS* falls into the general machine learning category. However, it clearly is the one

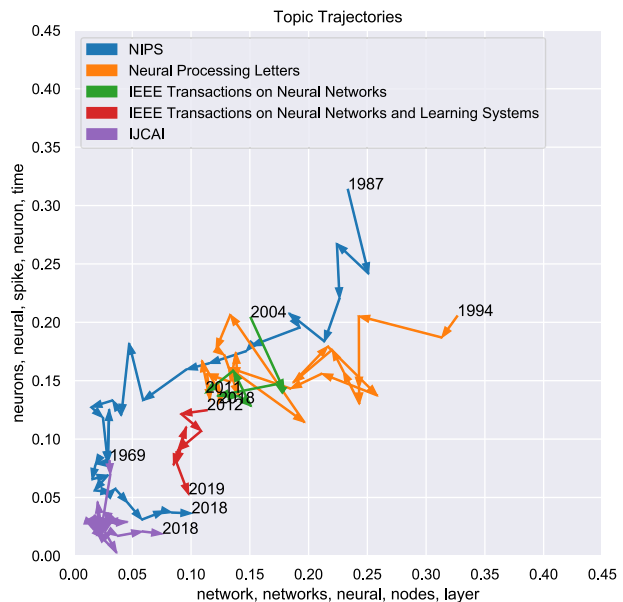
conference from this category that is closest to the neural networks cluster. *IJCAI* (*International Joint Conference on Artificial Intelligence*) and *ILP* (*Inductive Logic Programming*, right middle) are both conferences which lie closest to the general machine learning category. Nonetheless, they put more emphasis on specialized topics, such as knowledge representations and logic based systems. Hence, they lie at the border to general machine learning with a small but visible gap to the central part. Similarly, *COLT* (*Conference on Learning Theory*) lies at the left border of the same cluster.

In summary, the topical map shows that topics are captured well and as expected through our venue representations. Besides this, topical maps lead to interesting insights once we analyze (in this case visually perceived) clusters and edge cases, i.e., outliers and points which lie between several clusters. A natural enhancement of this method would be an analysis of *trajectories* in topical maps, i.e., how venues drift apart or together over time.

### Visualizing topic space trajectories through projection

Topic space trajectories exhibit too many dimensions (i.e., topics) for direct visualization in a coordinate system. To analyze trajectories, we hence project venue representations onto their most relevant two topics. We determine the relevance of a topic through its average weight in the trajectory, i.e., the topic weight averaged over all years. We demonstrate this method on the *NIPS* conference, which we selected due to its interesting trajectory. We selected some additional venues, most of which are related to neural networks, for comparison. Figure 9 depicts the trajectories created through this process. We marked the first and last year of each trajectory. Trajectories drift into the direction of the arrows. Through measuring the average weight in the trajectory, we identified the two topics on neural networks as the most relevant ones. The topic on the x-axis is the one we previously identified as being more concerned with the architecture of neural networks (we

**Fig. 9** Trajectories of selected Conferences. Projection on two selected topics on neural networks. Best viewed in color




simply called it *Neural Networks*). The topic on the y-axis is more concerned with the biologic motivation of neural networks and with dynamic neural networks (i.e., recurrent and spiking neural networks). We called this one *Neurons & Dynamic Neural Networks*.

We observe that initially all venues drift to the origin of the coordinate system, i.e., lose their relative interest in both topics. In recent years, however, research on neural network architecture has gained attraction again. By backtracking the trajectory of NIPS, we can see that this started around 2010. Around these years, the field of *deep learning* (Bengio 2009) gained much interest. This interest was motivated and accelerated through several discoveries and improvements in the field, e.g., the ReLU (rectified linear unit) activation function in Glorot et al. (2011), largely improved training times through GPU programming (Raina et al. 2009) and breakthroughs in performance on benchmark data sets, e.g., on the MNIST data set of handwritten digits in Ciresan et al. (2010) and on the ImageNet data set through convolutional neural networks in Krizhevsky et al. (2012). At this point we would like to note that other topics that gained interest, such as optimization, reinforcement learning and image recognition (cf. Fig. 7) are strongly connected to neural networks. Hence, the total interest in topics involving neural networks has increased even more than apparent from the trajectory in Fig. 9.

The conference IJCAI, which was held every two years starting from 1969 and yearly starting from 2015, is almost stationary most of the time. It exhibits only a small proportion of papers on neural networks. In recent years it shows a movement to the right, i.e., an increasing relevance of research on neural network architecture. The NIPS conference has a comparatively smooth, easy to follow trajectory, which again ends close to the origin. This result is supported by the fact that NIPS has become a more general conference on machine learning. *Neural Processing Letters* starts close to NIPS but ends at a different location, with more relevance on both neural network topics. This is a reasonable result, since it is a journal focussed specifically on this research area. An interesting case is *IEEE Transactions on Neural Networks*, which has been renamed to *IEEE Transactions on Neural Networks and Learning Systems* in the year 2012. In our data set these two are handled as separate venues. We noticed the name change through the behavior of the trajectories. More specifically, the endpoint of the trajectory under the first venue name lies close to the starting point of the trajectory under the second name.

Altogether, we see that topic space trajectories are an effective method for a human interpretable analysis of topic drift. A drawback resulting from high-dimensional data, like topic vectors, is that we can only visualize the trajectory for up to three topics. This problem, however, can be mitigated through the selection of relevant topics. One possibility here is a manual topic selection through the user. Likewise, an automated solution can be established through a measure of topic relevance. For our example in Fig. 9, this measure is the average weight of the topic. We can imagine further measures for different applications. Using the maximal topic weight over all years would yield trajectories for topics which were strongly relevant, even when only for a short time. Measures based on the absolute difference between topic weights for different years open another promising direction. Such measures would return trajectories with strong movement in topic space. This can be fine-tuned based on which years are considered (e.g., the difference between the first and the last year, between all consecutive years or between all pairs of years) and how these are aggregated (e.g., using the average or maximum of the differences).




**Fig. 10** Conference topics for 2018. In this visualization, each row represents a conference. The conference is represented through their mean topic vector. Each column represents a topic. Colors indicate topic weights. For each topic, the most important terms are given

### Visualizing topic space trajectories as heat maps

We found that heat maps as depicted in Fig. 10 are an effective method for interpretable topic visualizations. In such plots, topic weights can effortlessly be compared across venues. Each row here visualizes the topic space representation of one venue. In this specific instance, we calculated the centroid of papers from 2018. The topic weights are visualized through different color shades in the columns, with brighter colors indicating stronger topic weights.

Building upon this idea, in Fig. 11 (a) and (b) we visualize all topic space trajectories of our data set through one heat map per venue. In these heat maps, each row represents the topic space representation of a venue for a specific year. By following the rows from top to bottom, we see how interest in specific topics evolves over time. We only calculated and displayed centroids for years, in which at least ten papers were published at a venue. We do this for two reasons: First, occasionally we have instances of papers which seem to have a wrong year or venue. Second, sometimes our data set contains very few papers for a year. Both these lead to venue representations, that do not reflect reality well. We hence only calculate trajectories over years with more samples, i.e., papers. This leads to the effect, that we have no trajectory at all for the venue *DMKD*.

In the resulting heat maps we make two particularly interesting general observations: First, the topics evolve smoothly from year to year, despite the fact that each row was calculated from completely different papers. Second, we see that venues with similar research exhibit similar patterns in their heat maps. The heat maps thus can be regarded as distinguishable fingerprints that research areas leave. As an example, the heat maps of different venues on Information Retrieval (e.g., *Information Retrieval* and *SIGIR*) exhibit a visually similar appearance. This is due to their topic weights being stronger in the same columns as well as being similar in their development over time (i.e., over the rows).



**Fig. 11 (a)** Topic space trajectories as heat maps. Each heat map visualizes the topic space trajectory of a venue. In each heat map, rows represent years and columns represent topics. Cell colors indicate topic weights. Topic numbers are given on the x-axis in each subplot and corresponding topic names in the legend at the bottom. DMKD is missing since we require a minimum number of ten papers per year. Trajectories are continued on the next page.

**(b)** Topic space trajectories as heat maps (cont'd)

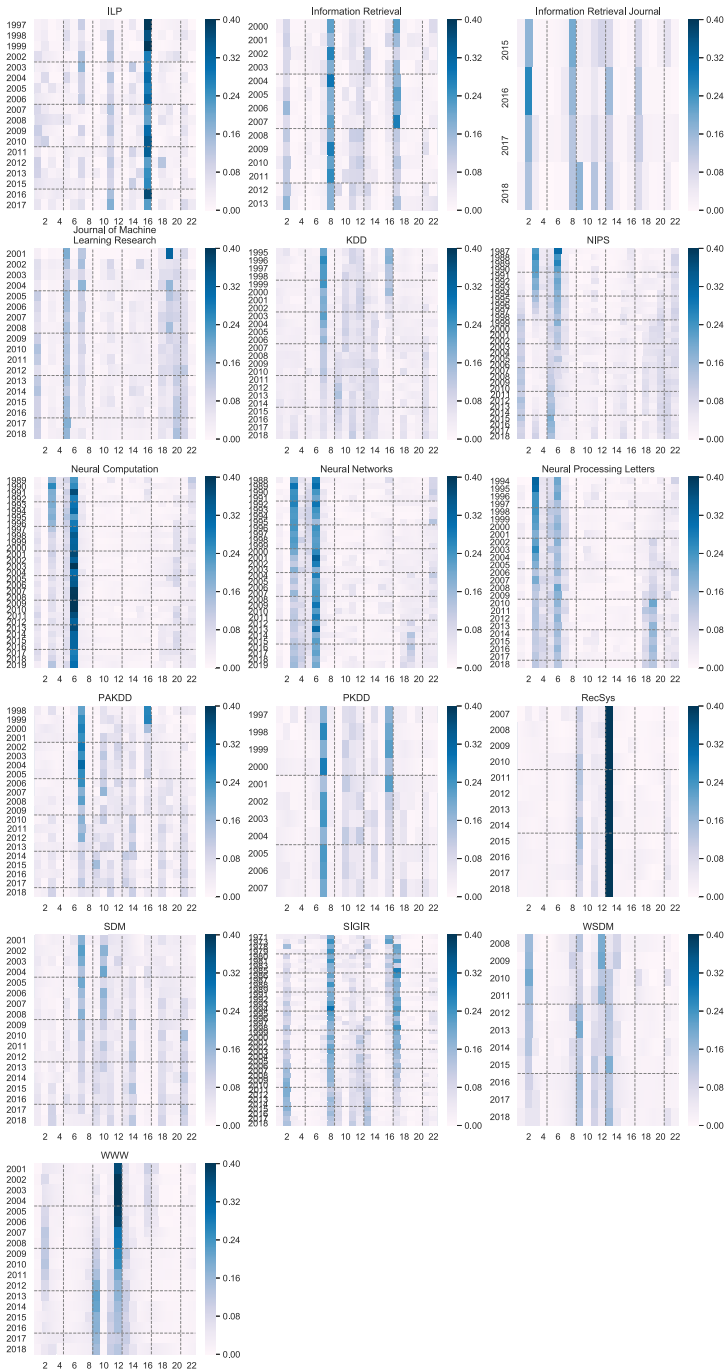


Fig. 11 continued



Likewise, such a similarity is strongly visible for conferences and journals on artificial neural networks (e.g., *NIPS*, *Neural Computation* and *Neural Networks*). In particular, we have two important topics on neural networks with a similar development over time.

Altogether the strongly visible patterns show that topic space trajectories capture topical specifics about venues as well as about their historical development. Our visualizations as heat maps make these specifics visually perceivable. We argue that heat maps are among the best possible visualization method for these kinds of trajectories. The reason is that heat maps capture all dimensions, i.e., topics, while providing a good and comparable overview.

We now analyze the trajectories of some particular venues. One important machine learning conference is the *International Joint Conference on Artificial Intelligence* (IJCAI). What strikes most for this venue, is that it has a strong focus on one particular topic, which is *Planning & Reasoning, Logic & Association Rules*. This focus, however, has decreased since 1995 and the conference has become more diverse. In particular, *Reinforcement Learning* has gained importance since then, which is a different approach to solve similar tasks.

Another interesting case are the *European Conference on Machine Learning* (ECML) and *Principles and Practice of Knowledge Discovery in Databases* (PKDD). Both conferences, similarly to IJCAI, started with a high interest in planning & reasoning (etc.) that declined since 1995 (for PKDD some years later). Both additionally exhibit a strong focus on *Classification & Pattern Mining* throughout their existence until 2007, as well as some interest in *Clustering*. What distinguishes these conferences is how weights are distributed across other topics. PKDD here is concerned with web pages, graphs and knowledge bases. ECML is concerned with matrix and kernel methods as well as SVMs and feature extraction. In 2008 these two conferences were merged and since then called *ECML/PKDD*. The resulting heat map of this conference thus exhibits even more diversely distributed topics than the two alone.

## Topic diversity

In this part, we analyze the topical diversity of venues based on their topic space representations. For this, we calculate a measure of diversity from each of these vectors. As the components of the vectors are all positive and sum up to 1, we can interpret each topic space representation as a probability distribution over topics. A higher diversity should be obtained, the more evenly distributed these topics are. This can be achieved through the Shannon-Entropy. The Shannon-Entropy of a probability distribution  $p$  over a discrete random variable  $x$  with outcomes  $X$  is calculated as follows:

$$H(p) = - \sum_{x \in X} p(x) \cdot \ln p(x)$$


While this measure becomes larger, the more evenly distributed topics are, its concrete value is not well interpretable. In Jost (2006), the entropy is therefore converted to a more interpretable measure through taking the exponential, i.e. by calculating  $\exp(H(p))$ . This measure is often used in biology to calculate *the effective number of species*, i.e., the number of evenly distributed species that would be necessary to obtain the same calculated entropy. Hence, the maximum possible value of this measure for a distribution over  $n$  outcomes is exactly  $n$ . This is reached when all outcomes have the probability  $1/n$ . In

**Table 2** Ranking of venues by topical diversity. Topical diversity is measured as the effective number of species, which is the exponential of the Shannon entropy. Journals are printed in italic. There is a tendency of journals to be more focussed

Rank	Venue (Conference or <i>Journal</i> )	Diversity
1	ECML/PKDD	19.30
2	<i>IEEE Transactions on Knowledge and Data Engineering</i>	19.24
3	KDD	19.08
4	CIKM	18.21
5	ACML	18.02
6	SDM	17.99
7	PAKDD	17.54
8	NIPS	17.13
9	ICML	16.64
10	IJCAI	16.51
11	ECML	16.36
12	ICANN	16.25
13	<i>Journal of Machine Learning Research</i>	16.18
14	<i>IEEE Transactions on Neural Networks and Learning Systems</i>	16.01
15	ICDM	15.81
16	PKDD	15.44
17	WSDM	15.42
18	AISTATS	15.25
19	SIGIR	14.58
20	ECIR	14.42
21	<i>Neural Processing Letters</i>	14.35
22	WWW	14.15
23	<i>IEEE Transactions on Neural Networks</i>	14.10
24	<i>Information Retrieval Journal</i>	13.96
25	DMKD	13.92
26	<i>Neural Networks</i>	13.74
27	IJCNN	13.73
28	<i>Information Retrieval</i>	13.49
29	COLT	12.95
30	ILP	12.82
31	<i>Neural Computation</i>	11.73
32	RecSys	7.87

information theory, the measure is also sometimes referred to as the *perplexity* of a distribution, although with a different connotation and interpretation.

We calculate the effective number of species for the topic space representation of every venue. We then rank the conferences in decreasing order of diversity. The results are given in Table 2. The results agree with our background knowledge about these venues. More general conferences or journals on knowledge discovery, which contain research from almost every machine learning field are ranked high in the list. First rank is ECML/



**Fig. 12** Topic diversities for venues over years. The diversity of a venue in a specific year is indicated by the color (or shade) in the plot. Black boxes indicate that not enough data is available. In most cases this means, that the venue did not take place or exist during that year. The last row contains the average diversity over all venues in that year

PKDD with a diversity of 19.30. This is closely followed by other more general conferences and journals on machine learning and knowledge discovery, such as KDD and CIKM. In contrast to this, in the lower part of the table we see venues which are more specialized. As an extreme example, RecSys, with its strong focus on recommender systems, is ranked lowest with a diversity of 7.87. Second lowest rank is Neural Computation with a diversity of 11.73. COLT and ILP are comparatively specialized conferences as well. Almost all other conferences and journals on neural networks or information retrieval.

In the middle part of the table we have the cases which are in between. Interesting examples are ECML and PKDD, which starting from 2008 merged to the single conference ECML/PKDD. While both already exhibited a strong topical diversity, it increased even more through the merge. Another interesting case is NIPS (Conference on Neural Information Processing Systems). Although this originally was a conference on neural networks as its name indicates, it later evolved into a more general conference on machine learning and artificial intelligence. Hence, its topical diversity is considerably higher than the diversity of all other conferences and journals focussing on neural networks.

In a second analysis, we look at the development of diversity over the years. This is depicted as a heat map in Fig. 12. In this figure, each row depicts the topic diversity of a venue for all years our data set spans. The diversity is indicated by the color (or shade) of a cell in a row. The last row contains the average diversity over all conferences. Diversities were only calculated, where at least ten papers were available for a venue and year. Note that in some cases, we still have erroneous data. As an example, SIGIR did not take place before 1978. Analyzing the results, it is interesting to note that many venues start out with an increase in diversity (e.g., NIPS, WWW, KDD). In later years, however, and especially the last few years, the diversity often is lower than previously (e.g., NIPS, ICML, IEEE Transactions on Knowledge and Data Engineering, AISTATS). The tendency of an early

increase is also noticeable in the last row containing the average diversity. Here, however, we see that later on the diversity is close to constant. It seems that during the last years, some conferences put more focus on specific topics again, while on average the diversity does not change. One reason for this could lie in the growing number of different conferences and journals. Venues here sometimes might want to distinguish themselves more from others. This leads to more focussed single venues, while the average diversity remains nearly constant. Venues could also have become more selective in their review process, choosing only papers that fit the conference well. This could be a consequence of the increasing popularity of machine learning in the recent years, which has also led to a larger number of submitted papers that venues can choose from.

Another interesting case is the WWW conference. From 2001 to 2006, this conference has a diversity that ranges between 9.7 and 10.7. In 2007, there is a sudden increase to a diversity of 12.9. Our topic trajectory heat map in Fig. 11 indicates that in this year the previously strongest topic on web pages has lost popularity. Instead, interest in search engines and social media has grown. Additionally, the topic on recommender systems starts to gain relevance, although already starting one year before in 2006. We assume that two events played a big role for this result: In 2006, the social media platform *Twitter* was released. Twitter quickly became a popular resource for research due to its large, global user base and public API. The second event was a competition on recommender systems, called the *Netflix prize*<sup>2</sup>. In this competition, which started in 2006, participants were invited to develop a recommender system that predicted user ratings for films. The prize money of one million dollars led to an increasing interest in recommender systems with more than 5000 teams actively participating in the competition.


## Topic densities

In this part, we analyze the topical density of some venues. For this, we use the topic space representations of the documents in our data set. For one venue, we project the topic space representations of papers into two-dimensional space using MDS. If we have more than 1000 papers for a venue, we only utilize a random sample of 1000 papers. We do this to speed up our calculations, since MDS has a complexity of  $O(n^2)$  for calculating the distances between the sample pairs. After MDS, we do a kernel density estimation (KDE) to estimate a probability density of papers in the two-dimensional space. The KDE is performed with a gaussian kernel and a grid search for the optimal bandwidth, a hyper-parameter of this method. Finally, we plot the density as a heat map together with the locations of the projected papers.

While the dimensions in the projected space are not directly interpretable as topics, we can interpret distance in this space as topical distance. An advantage of representing conferences by their distribution over only their centroid is, that we can see the whole distribution of topics instead of only an aggregated mean value. This distribution captures additional information such as various topical hot spots, i.e., dense areas in the projected space.

Figure 13 shows the results of this process for some selected conferences. We notice here that venues with a broader focus (ECML, KDD) tend to have several “blobs” at the margin of the distribution. We suspect that these blobs are all from different topical focusses. Conferences with a stronger single topical focus (COLT, RecSys) do not exhibit

<sup>2</sup> <https://www.netflixprize.com>.



**Fig. 13** Topic densities of selected venues. Dots represent papers after dimensionality reduction of topic vectors through MDS. Proximity in this space hence indicates topical similarity. Colors indicate the probability density of papers estimated through KDE

this behavior or only slightly. For RecSys, we have a considerably dense area at the right part, while to the left of this papers drift apart from each other. This indicates one single very strong topical focus. Papers deviating from this focus are distributed evenly across other topics, i.e., without any further strongly visible cluster.

The lower four plots all show venues with research on neural networks. In our heat maps in Fig. 11 we showed that their topic space trajectories exhibit similar patterns. In the density plots, however, we see that nonetheless their distributions are different. Topic densities hence reveal additional information to the topic space representations of venues. This is because the venue representations only give an average of the paper vectors in topic space instead of the full distribution. For Neural Processing Letters, two distinct topical focusses are visible, one being broader (i.e., with more variance) than the other one. Neural Computation has a clear focus, similar to RecSys. This is backed by the fact, that it has the second lowest topic diversity directly after RecSys (cf. Table 2). NIPS has a broad distribution and papers are focussed on different topics, visible through blobs distributed around the margin. Neural Networks, similar to Neural Processing Letters, seems to have two main focusses. However, there is an area, where these two clusters almost merge. This indicates that there is a smooth thematic transition between both focusses. Hence, we suspect that these blobs are the two topics on neural network.

We saw in this section that topic densities give supplementary information to topic space embeddings. However, in the resulting space, we do not know which area of a density belongs to which topic. At this point, we leave the development of a method for a more specific topical interpretation of areas open for further research. One possibility here would be to create a “pseudo-paper” vector for each topic, with full weight on the topic. One could then project these pseudo-papers into the two-dimensional space together with the real papers and mark their positions. It might also be beneficial to apply MDS to the papers of several venues together before estimating the density of each. Finally, we envision that topic space trajectories are extended to such densities, i.e., the chronological development of densities is investigated.

## Related work

We discussed topic models in more detail in (“[Document representation in topic space](#)” section). Hence, we only give a brief outline here. Topic models generally take a document-word matrix as input, in which documents are represented in a vector space model. The classic topic model is Latent Semantic Analysis (LSA) (Deerwester et al. 1990), which performs a singular value decomposition of this matrix. It has been used for various text mining tasks, e.g., recently in recommendations of scientific articles (Chughtai et al. 2020). Similarly to LSA, non-negative matrix factorization (NMF) (Lee and Seung 1999) decomposes a matrix into two factors, however with the constraint that all parts which contribute to the reconstruction of the input matrix must be positive. The decomposition is therefore better understandable for humans than through LSA, where topics and words can also contribute negatively. Latent Dirichlet allocation (LDA) (Blei et al. 2003) is a probabilistic approach on topic modelling that has gained much interest and is probably the de-facto standard in topic modelling. Many extensions exist, for example dynamic topic models, where topics change over time (Blei and Lafferty 2006, 2007; Wang and McCallum 2006) and which incorporate word embeddings (Dieng et al. 2019). For LDA, some research has reported that it cannot handle short documents well (Hong and Davison 2010). In our work, however, we look at (relatively short) paper abstracts. Additionally, in

Belford et al. (2018) it was shown that LDA produces considerably less stable results than NMF in repeated runs of the algorithm, hence leading to impoverished reproducibility. Therefore we refrain from LDA and related methods in our research. Incorporating dynamic topics as well as word embeddings complicates interpretability (since topics change over time) and leads to even worse reproducibility than LDA alone. The reason for the latter is that these models introduce additional elements of randomness. Hence, we also refrain from methods from these realms. Instead, we utilize the well-investigated NMF, which gives interpretable and reproducible results.

The work Berry et al. (2007) gives an overview over research on NMF. As the authors point out, research on NMF algorithms is focussed on first, applications to different kinds of data, second, enhancing or improving aspects of the algorithm (such as convergence and space and runtime complexity) and third, analyzing algorithmic properties (e.g., concerning the training convergence). In Yang et al. (2011), it was found that normalizing input data and minimizing Kullback-Leibler (KL) divergence instead of the Frobenius norm may lead to faster convergence and more better approximations. Following this line, Hien and Gillis (2020) recently proposed new algorithms for KL based NMF with guaranteed non-increasingness of the objective function. In Ding et al. (2008) it has been shown that with KL divergence, NMF minimize the same objective function as the topic model probabilistic latent semantic indexing (PLSA), which is regarded as a predecessor of LDA. Different optimization methods have been used for NMF, namely based on multiplicative update rules, gradient descent or alternating least squares (Berry et al. 2007). The minimization of NMF is often done starting from randomly initialized matrices, which can greatly affect the results (Berry et al. 2007). Other initialization methods of the matrices have thus been proposed, e.g., based on a nonnegative singular value decomposition (Boutsidis and Gallopoulos 2008). The work Zhao and Tan (2016) proposes an online algorithm for NMF, which is able to learn and update topics without having to store the documents. It is also shown to have good convergence properties as well as being able to handle outliers. The gensim implementation we use in this work is based on the same paper. By the same authors in Zhao et al. (2017), NMF is extended to optimize a broader class of divergences as their target function. In this work, the performance, stability and convergence speed in the application as a topic model is demonstrated in a document clustering task, in which documents are assigned to the topic with the strongest weight. Many further attempts have been made to improve certain properties of NMF or to deal with certain implementation problems. These are often very specific, e.g. to the used optimization algorithm or target function.

Different attempts have been made to deal with dynamics over time in topic analysis. Generally, there are two basic approaches: 1. Dynamic topic models, which explicitly incorporate time. 2. Using static topic models or word frequencies and to comparing results over time in a post-processing step. Examples of the first approach include the aforementioned dynamic topic model (D-LDA), an extension of LDA in which topic distributions exist for time steps and are conditioned on their previous step (Blei and Lafferty 2006). In *Topics over Time* (TOT) by Wang and McCallum (2006), each topic is modelled as a distribution, which generates words as well as continuous timestamps. This is achieved by extending LDA with an observable timestamp. More recently, the dynamic embedded topic model (D-ETM) has been proposed (Dieng et al. 2019). In this, D-LDA is extended with the recently popular word embeddings. Despite the interestingness of these approaches, they also come with some drawbacks in their practical application: The additional parameters lead to more complicated training with often unstable results, which has already been shown for traditional LDA in comparison to NMF (Belford et al. 2018).

Furthermore, interpretability is often reduced. This is especially the case when topics (as word distributions) may change over time and need to be reinterpreted. A different approach is thus to employ static topic models or word frequencies and to compare results over time in a follow-up step. This typically involves sorting results for time steps or slices, sometimes followed by some form of aggregation (as done in our approach) and, finally, an analysis of the results over time. Often, only two slices are employed to provide a before/after comparison in topics and maintaining overview. This, however, may often be too coarse-grained for a deeper analysis. In Scientometrics, topics are often analyzed by frequent terms in articles or through keywords instead of using explicit topic models. Keywords can be added by authors to their papers to facilitate automatic indexing, e.g., for search engines. A field sometimes using this approach and related to our work is the one of *burst detection*, i.e., the detection of sudden increases or decreases of interest in a topic (Kleinberg 2002). As a more recent application Tattershall et al. (2020) applies burst detection methods from stock market trend analysis to detect bursts in research topics based on normalized term frequencies. A disadvantage over our approach here is the term-based analysis, which requires additional manual work and does not give an overall summary. Another popular tool incorporating burst detection as well as bibliometric techniques is CiteSpace (Chen 2006). Emerging trend detection here is performed through terms which rapidly grow in frequency over time. Visualizations of scientific research is enriched based on citation and co-citation data. The topical analyses are, however, very coarse-grained. VOSviewer is a more recent tool for the visualization of scientific research (van Eck and Waltman 2010). VOSViewer is based on *Visualization of Similarities* (VOS), a method inspired by and improving on multidimensional scaling when applied to e.g., bibliometric data (van Eck and Waltman 2007b). VOSViewer allows the discovery of topics in terms of clustered co-occurring terms or keywords. Comparisons over time can be achieved by comparing visualizations derived from different time slices. VOSViewer is still popular and has been used, e.g., for the analysis of machine learning research in van Eck and Waltman (2007a) and more recently in Bhattacharya (2019). Scholia (Nielsen et al. 2017) is a more recent tool based on Wikidata. It allows to visualize topics of scientific publications based on keywords. All these approaches do not allow for a fine-grained, topical analysis over time. In contrast to our approach, they do not give an overall summary and comparison over all topics on the venue level and are sometimes difficult to interpret.

There exists much research where textual analysis of scientific articles has been performed, often involving topic models. In Mimno (2012), 24 journals from philology and archaeology were analyzed regarding paper locations in a vector space over time and regarding topical variation. In Griffiths and Steyvers (2004) article abstracts from the *PNAS* journal were analyzed for hot topics by performing a linear trend analysis on the weights in a topic model. In Wang and McCallum (2006), articles from NIPS conference were analyzed with a dynamic, continuous extension of LDA. The authors showed that through this method, the increasing and decreasing popularity of recurrent neural networks during the 1990s could be recovered, similar to the findings in our work. In this work full paper texts were used, in contrast to our work where abstracts are employed. In Sipos et al. (2012), based on *PNAS* articles, a temporal summary consisting of landmark documents, authors and topics was generated. Additionally, in this work papers were laid out in a two-dimensional space based on paper similarities, bibliographic coupling and force-directed layouting. The method for this, called *Vxinsight*, was presented in Davidson et al. (1998). In Blei and Lafferty (2007) and Dieng et al. (2019) the authors demonstrated the applicability of extensions of LDA on articles from the *Science* magazine and the *ACL*



conference (Annual Conference of the Association for Computational Linguistics). In Skupin (2004) paper abstracts from geographic research were visualized based on a vector space model, a self-organizing map and hierarchical clustering. What all of the above attempts have in common, is that they do not present a thorough, full analysis and comparison of different venues. Additionally, these previous works operate on the paper level, i.e., do not find representations of conferences or journals in vector space. The only exception from this is Mimno (2012), which, however, solely operates in a dimensionality-reduced space obtained from a bag-of-words representation instead of an interpretable topic space. In contrast to this, we envision methods to analyze topic trajectories of conferences and journals in an interpretable and reproducible manner.

Comparing our results to existing scientometric analyses of machine literature, we often found confirmation and sometimes differences. In van Eck and Waltman (2007a), articles from the computational intelligence field, which strongly overlaps with machine learning, were analyzed using VOS. Some of the clusters they found are similar to our findings, namely *control problems*, *classification*, *regression* and *optimization*. A surprising result of this work different from ours, was that clustering and classification were determined as a single topical cluster. Tattershall et al. (2020) analyze computer science research through burst detection and found some results similar to ours, e.g., the rise of popularity of Social Networks and related terms between 2004 and 2014, the lowered interest in *web 2.0* in recent years. In Bhattacharya (2019), an analysis of machine learning literature has been performed based on measures from social network analysis, VOSviewer and keywords from the web of science. A frequency ranking of topics for 2018 was performed, which was very different from the popularity we found based on our topics. As an example, support vector machines were found to be very popular, which is a contrast to our results as well as our personal background knowledge. An attempt of clustering keywords was made, which, in contrast to our topics, produced clusters of often very different keywords. Finally we want to emphasize that, to the best of our knowledge, we are the first to provide such a deep topical analysis of a large corpus of machine learning literature over time in a long period and on the venue level.

## Conclusion

In this work we introduced *topic space trajectories*, a novel approach to analyze conferences and journals. Loosely following the notion of *interpretable AI*, our focus here lay on interpretable results and reproducible methods. Overall, additionally to introducing our method, our aim was to present a deep analysis of the results achievable through it. We therefore demonstrated our approach on a set of machine learning conferences and journals and came to various interesting insights. Additionally to being reasonable and fascinating by themselves, the presented results support the applicability of our approach. Nonetheless, we also found some limitations. Most of these limitations stem from the nature of high-dimensional data. The reduction of high-dimensional documents in a *tf-idf* representation to a low-dimensional topic space representation mitigates such problems. However, the topic space still has a dimension that cannot be visualized in a coordinate system. As a consequence, we further reduced our data to a two-dimensional space through multidimensional scaling. In this representation we can analyze topical (dis)similarities between objects (i.e., documents or venues) in what we called *topical maps* and *topic densities*. However, we obtain dimensions that are not interpretable as topics anymore and lose information about *which* topics are (dis)similar. We therefore found salvation in two

solutions: First, plotting trajectories for only the two most interesting topics, and second, visualizations of trajectories as heat maps. Based on topic space trajectories, we also analyzed topical diversities of venues and their development over time. Last but not least, we made our discoveries plausible throughout our analysis through a comparison with historical events in research.

Many of our approaches open room for further research. An interesting direction would be automated name suggestions for the topics found through NMF. We imagine that this could be achieved through frequent bigrams in the most relevant documents for a topic. Topical maps and topic densities could be extended by their development in time, i.e., by calculating them for trajectories. Both would benefit from methods for improved interpretability.

As already indicated in our introduction, we envision further applications of our method. First, our venue embeddings could be used in a recommendation scenario, e.g., similar as in Chughtai et al. (2020). Second, it would be interesting to explore whether topic space trajectories can be extrapolated into the future. This could then, as an example, lead to even better recommendations. Possible applications lie in researchers looking for a suitable conference or journal for an unpublished paper or for recommending reviewers for a submitted paper. Last, we are eager to see our method applied to other kinds of data. As one very similar example, instead of venues we could analyze the trajectories of authors in topic space. Our method can, however, also be applied to totally different text domains.

**Acknowledgements** This work was funded by the German Federal Ministry of Education and Research (BMBF) in its program “Quantitative Wissenschaftsforschung” as part of the REGIO project under Grant 01PU17012 as well as the German Research Foundation (DFG) priority programme (SPP) 1894 project topikos.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Ammar W, Groeneveld D, Bhagavatula C, Beltagy I, Crawford M, Downey D, Dunkelberger J, Elgohary A, Feldman S, Ha V, Kinney R, Kohlmeier S, Lo K, Murray T, Ooi H, Peters ME, Power J, Skjonsberg S, Wang LL, Wilhelm C, Yuan Z, van Zuylen M, Etzioni O (2018) Construction of the literature graph in semantic scholar. In: Bangalore S, Chu-Carroll J, Li Y (eds) Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 3 (Industry Papers), (pp. 84–91). Association for Computational Linguistics.
- Belford, M., Mac Namee, B., & Greene, D. (2018). Stability of topic modeling via matrix factorization. *Expert Syst Appl*, 91(C), 159–169.
- Bengio, Y. (2009). Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1), 1–127.
- Berry, M. W., Browne, M., Langville, A. N., Pauca, V. P., & Plemmons, R. J. (2007). Algorithms and applications for approximate nonnegative matrix factorization. *Computational Statistics & Data Analysis*, 52(1), 155–173.

- Bhattacharya, S. (2019). Some salient aspects of machine learning research: A bibliometric analysis. *Journal of Scientometric Research*, 8, s85–s92.
- Blei, D., & Lafferty, J. (2007). A correlated topic model of science. *Annals of Applied Statistics*, 1, 17–35.
- Blei DM, Lafferty JD (2006) Dynamic topic models. In: Cohen WW, Moore AW (eds) ICML, ACM, ACM International Conference Proceeding Series, vol 148, pp 113–120
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *J Mach Learn Res*, 3, 993–1022.
- Boutsidis, C., & Gallopoulos, E. (2008). Svd based initialization: A head start for nonnegative matrix factorization. *Pattern Recognition*, 41(4), 1350–1362.
- Chen, C. (2006). CiteSpace II: detecting and visualizing emerging trends and transient patterns in scientific literature. *Journal of the American Society for Information Science and Technology*, 57(3), 359–377. <https://doi.org/10.1002/asi.20317>.
- Chughtai, G. R., Lee, J., Shahzadi, M., Kabir, A., & Hassan, M. A. S. (2020). An efficient ontology-based topic-specific article recommendation model for best-fit reviewers. *Scientometrics*, 122(1), 249–265.
- Ciresan, D. C., Meier, U., Gambardella, L. M., & Schmidhuber, J. (2010). Deep, big, simple neural nets for handwritten digit recognition. *Neural Computation*, 22(12), 3207–3220.
- Davidson, G. S., Hendrickson, B., Johnson, D. K., Meyers, C. E., & Wylie, B. N. (1998). Knowledge mining with vxinsight: Discovery through interaction. *Journal of Intelligent Systems*, 11, 259–285.
- Deerwester, S. C., Dumais, S. T., Landauer, T. K., Furnas, G. W., & Harshman, R. A. (1990). Indexing by latent semantic analysis. *Journal of the American Society of Information Science*, 41(6), 391–407.
- Dieng AB, Ruiz FJR, Blei DM (2019) The dynamic embedded topic model. CoRR Preprint abs/1907.05545, <http://arxiv.org/abs/1907.05545>
- Ding, C. H. Q., Li, T., & Peng, W. (2008). On the equivalence between non-negative matrix factorization and probabilistic latent semantic indexing. *Comput Stat Data Anal*, 52(8), 3913–3927.
- Glorot, X., Bordes, A., Bengio, Y. (2011). Deep sparse rectifier neural networks. In: Gordon GJ, Dunson DB, Dudík M (eds) AISTATS, JMLR.org, JMLR Proceedings, vol 15, (pp. 315–323).
- Griffiths, T. L., & Steyvers, M. (2004). Finding scientific topics. *Proceedings of the National Academy of Science*, 101, 5228–5235.
- Hien, L. T. K., Gillis, N. (2020). Algorithms for nonnegative matrix factorization with the kullback-leibler divergence. Preprint at <https://arxiv.org/abs/2010.01935>
- Hong, L., Davison, B. D. (2010). Empirical study of topic modeling in twitter. In: Proceedings of the First Workshop on Social Media Analytics, ACM, New York, NY, USA, SOMA '10, (pp. 80–88).
- Jost, L. (2006). Entropy and diversity. *Oikos*, 113(2), 363–375.
- Kersting, K., Peters, J., Rothkopf, C. (2019). Was ist eine Professur fuer Kuenstliche Intelligenz? CoRR Preprint abs/1903.09516, <http://arxiv.org/abs/1903.09516>
- Kleinberg, J. (2002). Bursty and hierarchical structure in streams. In: KDD '02: Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, New York, NY, USA, pp 91–101
- Krizhevsky, A., Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In: Bartlett PL, Pereira FCN, Burges CJC, Bottou L, Weinberger KQ (eds) Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012., pp 1106–1114
- Le QV, Mikolov T (2014) Distributed representations of sentences and documents. In: Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21–26 June 2014, JMLR.org, JMLR Workshop and Conference Proceedings, vol 32, pp 1188–1196
- Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755), 788–791.
- Martín-Fernández, J., Barceló-Vidal, C., Pawlowsky-Glahn, V., Buccianti, A., Nardi, G., & Potenza, R. (1998). Measures of difference for compositional data and hierarchical clustering methods. *Proceedings of IAMG*, 98, 526–531.
- Mead, A. (1992). Review of the development of multidimensional scaling methods. *Journal of the Royal Statistical Society Series D (The Statistician)*, 41(1), 27–39.
- Mimno, D. (2012). Computational historiography: Data mining in a century of classics journals. *J Comput Cult Herit*, 5(1), 3:1-3:19.
- Nielsen FA, Mietchen D, Willighagen EL (2017) Scholia, scientometrics and wikidata. In: Blomqvist E, Hose K, Paulheim H, Lawrynowicz A, Ciravegna F, Hartig O (eds) ESWC (Satellite Events), Springer, Lecture Notes in Computer Science, vol 10577, pp 237–259
- Pu Q, Yang GW (2006) Short-text classification based on ica and lsa. In: International Symposium on Neural Networks, Springer, pp 265–270

- Raina R, Madhavan A, Ng AY (2009) Large-scale deep unsupervised learning using graphics processors. In: Danyluk AP, Bottou L, Littman ML (eds) ICML, ACM, ACM International Conference Proceeding Series, vol 382, pp 873–880
- Ramos, J., et al. (2003). Using tf-idf to determine word relevance in document queries. *Proceedings of the first instructional conference on machine learning, New Jersey, USA, 242*, 133–142.
- Řehůřek R, Sojka P (2010) Software Framework for Topic Modelling with Large Corpora. In: Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks, ELRA, Valletta, Malta, pp 45–50
- Rosner F, Hinneburg A, Röder M, Nettling M, Both A (2014) Evaluating topic coherence measures. CoRR abs/1403.6397
- Schaefermeier B, Hanika T, Stumme G (2019) Distances for wifi based topological indoor mapping. In: Proceedings of the 16th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, Association for Computing Machinery, New York, NY, USA, MobiQuitous '19, p 308-317
- Scholz C, Illig J, Atzmueller M, Stumme G (2014) On the predictability of talk attendance at academic conferences. In: Proceedings of the 25th ACM Conference on Hypertext and Social Media, ACM, Santiago, Chile, HT '14, pp 279–284
- Sipos R, Swaminathan A, Shivaswamy P, Joachims T (2012) Temporal corpus summarization using sub-modular word coverage. In: Chen X, Lebanon G, Wang H, Zaki MJ (eds) 21st ACM International Conference on Information and Knowledge Management, CIKM'12, Maui, HI, USA, October 29 - November 02, 2012, ACM, pp 754–763
- Skupin A (2004) The world of geography: Visualizing a knowledge domain with cartographic means. In: Proc. Nat'l Academy Sciences vol. 101, Suppl. 1, pp 5274–5278
- Steinberger J, Křišt'án M (2007) Lsa-based multi-document summarization. In: Proceedings of 8th International Workshop on Systems and Control, vol 7
- Tattershall, E., Nenadic, G., & Stevens, R. D. (2020). Detecting bursty terms in computer science research. *Scientometrics*, 122(1), 681–699.
- van Eck, N. J., & Waltman, L. (2007a). Bibliometric mapping of the computational intelligence field. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 15(5), 625–645.
- van Eck, N. J., & Waltman, L. (2007b). Vos: A new method for visualizing similarities between objects. In R. Decker & H. J. Lenz (Eds.), *Advances in Data Analysis* (pp. 299–306). Heidelberg: Springer, Berlin Heidelberg, Berlin.
- van Eck, N. J., & Waltman, L. (2010). Software survey: Vosviewer, a computer program for bibliometric mapping. *Scientometrics*, 84(2), 523–538.
- Wang, X., McCallum, A. (2006). Topics over time: a non-markov continuous-time model of topical trends. In: Eliassi-Rad T, Ungar LH, Craven M, Gunopulos D (eds) Proceedings of the Twelfth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Philadelphia, PA, USA, August 20-23, 2006, ACM, (pp. 424–433).
- Yang, Z., Zhang, H., Yuan, Z., Oja, E. (2011). Kullback-leibler divergence for nonnegative matrix factorization. In: Honkela, T., Duch, W., Girolami, M. A., Kaski, S. (eds) ICANN (1), Springer, Lecture Notes in Computer Science, vol 6791, pp 250–257
- Zhao, R., Tan, V. Y. F. (2016). Online nonnegative matrix factorization with outliers. Preprint at <http://arxiv.org/abs/1604.02634>
- Zhao, R., Tan, V. Y. F., Xu, H. (2017). Online nonnegative matrix factorization with general divergences. In: Singh, A., Zhu, X. J. (eds) AISTATS, PMLR, Proceedings of Machine Learning Research, vol 54, (pp. 37–45).
- Zipf, G. K. (1949). *Human Behavior and the Principle of Least Effort*. Addison-Wesley.

## Authors and Affiliations

**Bastian Schaefermeier<sup>1</sup> · Gerd Stumme<sup>2</sup> · Tom Hanika<sup>2</sup>**

✉ Bastian Schaefermeier  
bsc@cs.uni-kassel.de

Gerd Stumme  
stumme@cs.uni-kassel.de

Tom Hanika  
tom.hanika@cs.uni-kassel.de

<sup>1</sup> L3S Research Center, Appelstraße 9a, Hannover, Germany

<sup>2</sup> Knowledge and Data Engineering Group, University of Kassel, Wilhelmshoeher Allee 73, Kassel, Germany