

What, When, and Where Do You Mean?

Detecting Spatio-Temporal Concept Drift in Scientific Texts

Meilin Shi¹ 

Department of Geography and Regional Research, University of Vienna, Austria

Krzysztof Janowicz 

Department of Geography and Regional Research, University of Vienna, Austria

Zilong Liu 

Department of Geography and Regional Research, University of Vienna, Austria

Mina Karimi 

Department of Geography and Regional Research, University of Vienna, Austria

Ivan Majic 

Department of Geography and Regional Research, University of Vienna, Austria

Alexandra Fortacz 

Department of Geography and Regional Research, University of Vienna, Austria

Abstract

Inundated by the rapidly expanding AI research nowadays, the research community requires more effective research data management than ever. A key challenge lies in the evolving nature of concepts embedded in the growing body of research publications. As concepts evolve over time (e.g., keywords like *global warming* become more commonly referred to as *climate change*), past research may become harder to find and interpret in a modern context. This phenomenon, known as *concept drift*, affects how research topics and keywords are understood, categorized, and retrieved. Beyond temporal drift, such variations also occur across geographic space, reflecting differences in local policies, research priorities, and so forth. In this work, we introduce the notion of *spatio-temporal concept drift* to capture how concepts in scientific texts evolve across both space *and* time. Using a scientometric dataset in geographic information science, we detect how research keywords drifted across countries and years using word embeddings. By detecting spatio-temporal concept drift, we can better align archival research and bridge regional differences, ensuring scientific knowledge remains findable and interoperable within evolving research landscapes.

2012 ACM Subject Classification Information systems → Digital libraries and archives; Computing methodologies → Information extraction; Information systems → Ontologies

Keywords and phrases Concept Drift, Ontology, Large Language Models, Research Data Management

Digital Object Identifier 10.4230/LIPICs.GIScience.2025.16

Supplementary Material *Software (Source Code and Data):*

<https://github.com/meilinsi/Spatio-temporal-Concept-Drift-in-Scientific-Texts>

archived at `swb:1:dir:928946b53e1ebade1bc2cd2ddfe8a0b23a409111`

1 Introduction

The questions of *what*, *when*, and *where* permeate our daily conversations. When scheduling a group meeting, for instance, we agree on the topic of discussion (e.g., a proposal, *what*), the time (e.g., 10 a.m., *when*), and the location (e.g., a café, *where*). In communicating such

¹ Corresponding author



© Meilin Shi, Krzysztof Janowicz, Zilong Liu, Mina Karimi, Ivan Majic, and Alexandra Fortacz; licensed under Creative Commons License CC-BY 4.0

13th International Conference on Geographic Information Science (GIScience 2025).

Editors: Katarzyna Sila-Nowicka, Antoni Moore, David O'Sullivan, Benjamin Adams, and Mark Gahegan;

Article No. 16; pp. 16:1–16:18



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

information, we implicitly agree on a particular reference system. For time, we have temporal reference systems such as the Gregorian calendar, the yyyy-mm-dd date format, the 24-hour clock, and so on. For space, we have various geodetic datums, such as WGS 84 and NAD 83, as well as known place names we can refer to. For the *what* question, namely the thematic information, we also need a semantic reference system [23]. In this reference system, an ontology can help ensure that, by “proposal”, we are referring to a research proposal rather than a marriage proposal.

When it comes to ontology modeling and engineering, concepts are often represented as static entities [17]. For example, this is common in a foundational ontology (e.g., DOLCE [15]) to ensure a coherent view and interoperability across domains. In the real world, however, concepts are constantly evolving [19], and their meanings can vary across different social contexts and locations, as seen in the evolving sociocultural definitions of *gender* nowadays. Research in the Semantic Web and the broader knowledge representation and reasoning (KRR) communities has focused on *concept drift* to capture the dynamics of evolving concepts. In this respect, previous work in KRR [49, 14, 44, 8] focused mainly on the temporal aspect of a concept, i.e., the changing meaning of a concept over *time*, and overlooked the spatial perspective that often accompanies it.

In geographic information science (GIScience), constructing an ontology that maps geospatial concepts has always been challenging because of their unique spatio-temporal properties [12, 9]. Geospatial concepts, such as **Mountain** and **Forest**, are different from other general concepts because they do not have clearly defined boundaries nor can they be distinguished in bona fide fashion from neighboring concepts, e.g., **Hill** and **Woods** [42, 43]. For instance, the difference between **Lake** and **Pond** can be affected by seasonal water level variations [28]. This would make downstream tasks, e.g., question answering [33], more challenging. Furthermore, the conceptualization of such landscape concepts may also vary and evolve across languages, cultures, and regions [47, 11].

These challenges are not limited to modeling the aforementioned concepts that are vague geographic features. They also extend to research topics and keywords (e.g., *urban planning*, *climate change*), which we see as signifiers of concepts (i.e., mental representations that categorize areas of research [4]). Although many concepts in this regard exhibit concept drift, geospatial concepts are particularly susceptible due to their inherent dependence on physical, environmental, and sociopolitical contexts. To give a concrete example, the definition of *disasters* could vary significantly depending on local environmental conditions, infrastructure, and response systems. What qualifies as a natural *disaster* in one region (e.g., an earthquake with a magnitude of 6.0 in Haiti) may be labeled differently in another region (e.g., the United States) because of differences in local resilience. In comparison, concepts like *the speed of light* exhibit less spatial variability because of their underlying physical principles.

As Kuhn et al. [24] suggested, we should move space and time from merely being in application domains to becoming foundational aspects of ontologies. While the inherent vagueness in geographic features is not fully resolved, ontologies such as the GeoNames ontology² offer structured representations for these features. However, these efforts have yet to cover more abstract geospatial concepts embedded in research, such as those represented by keywords. Same as geographic features, these concepts are dynamic and spatially grounded, yet they are even more susceptible to societal changes and technological advancements (e.g., the coining of the concept *GeoAI* [21]). In this work, we look into these concepts in scientific research with explicit study areas. We propose an approach to quantify their fluidity and context-dependence across space *and* time via word embeddings. Our long-term goal is to

² <https://www.geonames.org/ontology>

develop an ontology that can address the dynamic nature of these concepts in scientific texts. This contributes to the FAIR principles (Findability, Accessibility, Interoperability, and Reusability) [50] by improving retrieval, reuse, and ensuring the long-term relevance of research [38].

The contributions of this work are as follows:

- We introduce spatio-temporal concept drift, which expands the previous notions of concept drift that focused mainly on temporal changes by incorporating both space *and* time.
- We propose a novel approach to detecting spatio-temporal concept drift in scientific texts via word embeddings.
- We demonstrate how accounting for spatio-temporal concept drift enhances the understanding of concepts in scientific texts, improves recall in FAIR-based research data management systems, and lays the groundwork for ontology learning with large language models (LLMs) in dynamic contexts.

The remainder of this paper is structured as follows. Section 2 introduces the theoretical background of our work. Section 3 reviews related work regarding word embeddings, with a focus on their ability to capture and quantify spatio-temporal variations in semantics. We describe our case study in Section 4, where we use a scientometric dataset in the field of GIScience to assess spatio-temporal concept drift. Section 5 presents the results. Section 6 discusses geographic bias within embeddings and future directions in ontology learning with LLMs. Finally, we conclude our work in Section 7.

2 Background

This section provides the theoretical background for the study on concept drift in knowledge representation. Here, we introduce the notion of spatio-temporal concept drift. This notion adds a spatial dimension to existing definitions of concept drift in the literature. In addition, we discuss the broader implications of spatio-temporal concept drift for research data management (RDM).

2.1 Concepts and Their Representation

The notion of *concept* has many different definitions across or even within domains, such as in linguistics, psychology, computer science, and cognitive science [35]. In this work, we adopt the definition by Stock [45] in information science, which defines a concept as a class containing objects that share certain properties.³

Concepts are fundamental units of meaning and serve as the building blocks of ontologies that help structure knowledge, enable reasoning, and facilitate interoperability. In terms of representation, previous work [49, 44] typically characterized a concept by its label (i.e., name), intension (i.e., defining properties), and extension (i.e., instances that fall under it), in the form $(label(C), int(C), ext(C))$ for a concept C . However, this would only apply to concepts that already existed in predefined ontologies. Verkijk et al. [48] proposed to use embedding techniques to derive vector representations of concepts. While their work focuses on knowledge-graph data, they showcased the ability of embedding techniques to capture flexible, context-aware representations of concepts for natural language data as well, in the form $(label(C), context(C))$.

³ It is worth noting that this definition of *concept* is closer to what cognitive scientists would call a *category*.

In this work, we treat keywords in research publications as representations of underlying concepts. Unlike established ontological categories, research keywords are rapidly evolving as science and society change. This makes them particularly relevant for studying spatio-temporal concept drift. Here, we focus on research keywords also with the aim of developing a structured ontology that can capture their changing meanings over time and space. Such an ontology could contribute to RDM by improving metadata organization, literature retrieval, question answering, and knowledge graph construction in scientific databases.

2.2 Spatio-Temporal Concept Drift

Adding a temporal dimension to concept representation accounts for changes in their meaning over time. The study of *concept drift*, as defined by Wang et al. [49], aims to capture these changes in concept meaning over time. For example, the keyword *global warming* was once the dominant term in research publications, referring to the rise in Earth’s temperature. Over time, *climate change* became more widely used to capture broader climate-induced impacts [26] and account for the fact that *warming* is not uniform. To model a concept C with a temporal component, it can be represented in the form $(label_t(C), int_t(C), ext_t(C))$ or $(label_t(C), context_t(C))$ at time t . Extending this notion to a spatio-temporal dimension means that a concept’s meaning may change both over time and space, e.g., at different rates. Here, we define this phenomenon as *spatio-temporal concept drift*. In this case, a concept can be represented as $(label_{t,s}(C), int_{t,s}(C), ext_{t,s}(C))$ or $(label_{t,s}(C), context_{t,s}(C))$ for a concept C at time t and region s .

Figure 1 illustrates how a concept moves in both time and geographic space. More abstractly, this can be thought of as the trajectory of a concept in a space-time prism [35]. The color gradient indicates (*semantic*) *concept similarity* [36, 39, 34, 22], thereby reflecting changes in its thematic dimension over time and space. Take the keywords *global warming* and *climate change* as an illustrative example. Region s_1 may have already adopted the use of *climate change* since time t_1 , while s_2 still has a mixed use of both terms. Note that, here at t_1 , the variation in the thematic dimension between s_1 and s_2 represents spatial variability, which differs from spatio-temporal concept drift, as it captures regional differences without a temporal dimension [25]. Later by t_3 , s_2 adopts the distinct use of *climate change* and *global warming*, aligning its concept representation with s_1 . Over time, as concepts evolve, their meanings may change gradually, showing a concept drift from t_2 to t_3 in s_1 , or change so much that they diverge into two, showing a concept split⁴ in region s_2 . Ultimately, the two concepts may converge into a shared understanding for both regions (indicated by the semantic similarity between C_1 and C'_1 at t_3).

Even with the advent of semantic search [18], which allows for more flexible query interpretation, concept drift remains a challenge, particularly in RDM and other archival systems. If the past and present keywords are not properly linked, search results may still be skewed toward more recent ones, simply because of their pertinence. This would lead to either incomplete retrieval results or misinterpretation of archival documents. Addressing spatio-temporal concept drift helps ensure that evolving knowledge remains accessible and meaningful across different time periods and regions. It could therefore enhance semantic interoperability overall and support the FAIR principles.

⁴ Definitions of *concept drift* and *concept split* are provided in our earlier work [40].




Figure 1 Representation of concepts drift and split over geographic space and time. The intensity of color indicates concept similarity.

3 Related Work

This section reviews existing work that provides means for measuring the latent semantics underlying words in their embeddings, with a spatio-temporal focus.

3.1 Spatial and Temporal Information in Word Embeddings

With the introduction of Word2Vec [29], word embeddings have revolutionized representation by converting words into dense vectors in a high-dimensional space,⁵ where semantically similar words are close to each other. Such representation enables a more flexible study of temporal and spatial variations in lexical semantics, offering an advantage over directly comparing different ontology versions. Early pre-trained word embeddings, such as GloVe [32], provide static representations, where each word is assigned a single vector, independent of context. Later, more advanced models like BERT [10] provide context-aware embeddings that capture more variations in meaning based on surrounding text using attention mechanisms.

Several studies [3, 20, 37] have explored the enrichment of word embeddings with temporal and spatial information. For example, Zhang et al. [54] focused on temporal counterpart search that detects semantically similar terms over time. The authors later also investigated the geographic variations in lexical semantics [55], e.g., showing that **typhoon** in Japan would be the most similar term to **hurricane** in the United States. Gong et al. [16] further extended this idea and proposed a model that conditions word embeddings on time *or* location (i.e., generating time- and location-specific embeddings). Their findings included word similarities over time (e.g., **bitcoin** in 2015 and **stocks** in 1992) and locations (e.g., **president** in the United States and **prime minister** in Canada). A few other studies in GIScience explored

⁵ Note that some literature uses the term “low-dimensional space” here when comparing the dimensionality to a one-hot encoding. We use the term here to signify that the resulting embeddings are in a, say, 300-dimensional vector space.

the representation learning of *places* via word embeddings, thereby using spatial information alone. Yan et al. [52] applied word-embedding techniques to learn embeddings of places based on their types and distances. Later, Zhai et al. [53] extended this approach to the representation learning of functional regions.

While these approaches captured variations in lexical semantics along one dimension effectively, they did not *jointly* consider spatial and temporal dimensions. As a result, they would fail to capture spatial-temporal lexical similarity, such as **chancellor** in Germany in 2010 and **prime minister** in the United Kingdom in 1980. Such similarity is centered in our study on spatio-temporal concept drift, which accounts for both dimensions at the same time.

3.2 Word Embedding Association Test

The Word Embedding Association Test (WEAT) [7], inspired by the Implicit Association Test in psychology, is a widely used method for quantifying semantic associations in word embeddings. It calculates association scores by comparing cosine similarities between two sets of target words (e.g., man and woman) and two sets of attribute words (e.g., doctor and gynecologist). For example, associations between $\overrightarrow{man}$ and $\overrightarrow{doctor}$ versus $\overrightarrow{woman}$ and $\overrightarrow{gynecologist}$ can be assessed using vector arithmetic [6, 30], expressed as $\overrightarrow{man} - \overrightarrow{woman} \approx \overrightarrow{doctor} - \overrightarrow{gynecologist}$. The resulting WEAT score indicates the degree of association between the two groups in the embedding space. WEAT provides a standardized measure and allows for statistical significance testing of observed changes. By adapting this method, we can compute the cosine similarities between different temporal snapshots and geographical regions, e.g., (*hurricane* to *Mexico*, 2005) and (*typhoon* to *China*, 2015). Note that 2005 and 2015 are not treated as vectors themselves but rather indicate the time periods associated with these concept-region pairs. If in a vector space, $\overrightarrow{hurricane} - \overrightarrow{typhoon} \approx \overrightarrow{Mexico} - \overrightarrow{China}$, this would allow us to quantify concept changes across space and time and reveal geographic prototypes underlying word embeddings. However, applying WEAT to spatio-temporal analysis also presents challenges, particularly in maintaining statistical power when data is sparse across certain regions or time periods.

4 Case Study: Geographical Information Science Publications

We employ a scientometric dataset from Wu et al. [51] to detect the spatio-temporal concept drift in a real-world dataset. This dataset includes research publications in the field of GIScience from 1991 to 2020, sourced from Scopus⁶. As the dataset focuses on international journals and conferences that publish exclusively in English, all included publications are in English. Here, we explicitly focus on papers that mention locations in their abstracts, including geopolitical entities (GPEs) – such as countries, states, and cities – as well as nationalities (NORP), using the spaCy transformer-based named entity recognition pipeline⁷. Table 1 presents the summary statistics of the dataset after filtering for these papers. For reference, we provide the full names of conference and journal abbreviations in Appendix A.

⁶ <https://www.scopus.com/>

⁷ <https://spacy.io/models>

■ **Table 1** Summary statistics of research publications with location mentions in abstracts.

Type	Name	Time Range	Number of Papers	Number of Keywords
Conference	COSIT	1993-2019	22	100
Conference	GIScience	2006-2020	18	78
Journal	CEUS	1999-2020	622	3177
Journal	CaGIS	1991-2020	196	1002
Journal	EPB	1998-2020	352	1744
Journal	GeoI	1997-2020	61	305
Journal	IJGIS	2005-2020	550	2640
Journal	JGS	1996-2020	196	922
Journal	JOSIS	2010-2020	21	133
Journal	SCC	2003-2020	19	92
Journal	TGIS	2007-2020	50	237
Total		1991-2020	2107	10430

4.1 Spatio-Temporal Dimensions of Concepts

As with the *what*, *when*, and *where* questions, we argue that each concept has thematic, temporal, and spatial dimensions. In this dataset, we treat each keyword as signifying an individual concept⁸ and represent these three dimensions accordingly.

To represent the (1) **thematic dimension**, we use the associated abstract, which provides contextual information of each concept. Since Scopus is an abstract and citation database without guaranteed full-text access, abstracts – being more consistently available across publications – are a practical choice for large-scale analysis. For the (2) **temporal dimension**, we use the publication year of the paper associated with each concept. Lastly, for the (3) **spatial dimension**, we use OpenStreetMap Nominatim⁹ to geocode the identified locations and extract the corresponding country for sub-national locations (e.g., cities). For location mentions like “East African”, we retain them at the continent level. If multiple countries are mentioned in an abstract, we document all of them. After retrieving 2,112 publications with location mentions, we manually reviewed 16 unidentified cases, assigning them to the country level or removing them where necessary. This resulted in a final dataset of 2,107 publications.

Table 2 includes examples of abstracts with location mentions and the extracted countries (or regions). The distribution of the 10 most mentioned countries in publications within our dataset is visualized in the heatmap in Figure 2. From this heatmap, we can observe that the United States and China lead in the number of publications, followed by other English-speaking countries (e.g., the United Kingdom, Canada, and Australia) and several European countries. Along the temporal axis, we also observe a notable increase in publications since 2005 in this scientometric dataset.

4.2 Spatio-Temporal Concept Drift in Embedding Space


With the defined spatial and temporal dimensions of each concept, we leverage word embeddings to capture their variations in the thematic dimension. We employ the pre-trained SciBERT model [5], which is designed for scientific texts, to compute embeddings. We use

⁸ The distinction between a symbol and a concept can be explained using the triangle of reference [31].

⁹ <https://nominatim.org>

■ **Table 2** Exemplar location mentions and extracted countries in abstracts.

Year	Abstract Excerpt	Location Mentions	Country
1999	“We illustrate...based on the street pattern of a small French town.”	[French]	[France]
2008	“A dataset describing...in New York City is analyzed to...the technique.”	[New York City]	[United States]
2017	“Using Austria and Slovenia as a study area,...modified IL.”	[Austria, Slovenia]	[Austria, Slovenia]



■ **Figure 2** Distribution of publications for the 10 most mentioned countries over time.

context-aware representations of concepts in natural language, i.e., $(label(C), context(C))$ for a concept C , as discussed in Section 2.1. Here, we compute these two types of embeddings for concepts (in this case, keywords): (1) **label embedding**, which is static and derived from the keywords themselves, and (2) **context embedding**, which is context-aware and based on their associated abstracts with location mentions.

Additionally, we investigate the sensitivity of context embeddings to location mentions by computing (3) **context embedding without locations** as well. This embedding is derived from associated abstracts of a concept, where each identified location mention is replaced with the placeholder “[Location]”. For instance, in the first example in Table 2, the sentence would become “We illustrate...based on the street pattern of a small [Location] town.” This helps reveal whether explicit geospatial references influence the context-aware representation of concepts.

For each keyword/concept C in a given year t and country (or region) s , we first average its context embeddings across all relevant abstracts, to ensure a single embedding for each unique keyword-year-country combination. We then integrate this with the label embedding through a convex combination to obtain the composite embedding $C_{t,s}$, formulated as:

$$C_{t,s} = \alpha \cdot label(C) + (1 - \alpha) \cdot \frac{1}{|D_{t,s}|} \sum_{d \in D_{t,s}} context(C, d) \quad (1)$$

where $label(C)$ is the embedding of the keyword itself through its label; $context(C, d)$ is the embedding of the abstract in document d containing the keyword; $D_{t,s}$ is the set of documents that contain the keyword from year t and country s ; and $|D_{t,s}|$ is the number of such documents. The parameter α determines the weight assigned to the label versus context embeddings.

To quantify how a concept C drifts across different space-time combinations, we use cosine similarity between their respective composite embeddings. Given two concept representations, e.g., C_{t_1,s_1} and C_{t_2,s_2} , their similarity is computed as:

$$sim(C_{t_1,s_1}, C_{t_2,s_2}) = \frac{C_{t_1,s_1} \cdot C_{t_2,s_2}}{\|C_{t_1,s_1}\| \cdot \|C_{t_2,s_2}\|} \quad (2)$$


5 Results

To start, we visualize these keywords in the embedding space. Figure 3 shows the distribution of keyword embeddings across countries and selected years (2000, 2005, 2010, 2015, and 2020), generated using t-Distributed Stochastic Neighbor Embedding (t-SNE) [46]. The label embedding weight α is empirically set to 0.3 to place greater emphasis on the context embedding while retaining sufficient label information. Higher values of α tend to produce overly label-driven clusters, whereas lower values may cause semantically related keywords to diverge (see Appendix B for examples).

From the figure, we can observe a cluster of keywords with location mentions of China over the years (represented by triangles of different colors in the middle left of the figure), and those with location mentions of European countries like Germany and the Netherlands appear closer to each other.

The t-SNE visualization provides an overview of keyword distributions; we then look into how individual keywords move along their semantic trajectories across different countries over the years. Note that all keywords are standardized to lowercase and American spelling. They are also lemmatized and expanded to their full forms (e.g., *DEM* to *digital elevation model*), with the exception of *GIS*, which we retain as an abbreviation due to its ambiguous reference to *GI Science* or *GI System*. For each keyword, we quantify its spatio-temporal coverage by multiplying its time span (in years) by the number of unique countries it is associated with, yielding a coverage score to reflect both its temporal persistence and geographic distribution. Table 3 presents the top 10 keywords ranked by their spatio-temporal coverage. From these, we plot the semantic trajectories for selected keywords – *GIS*, *urban planning*, *spatial analysis*, and *cellular automaton* – in Figure 4, using principal component analysis (PCA) [1] to reduce the dimensionality of their embeddings.

From these trajectories, we can notice that the extracted embeddings of *GIS* (Figure 4a) are quite consistent across Italy, Germany, Australia, and the UK in the early years of 1999 and 2000. Afterward, English-speaking countries, including the UK, the US, Canada, and Australia, along with China, have their embeddings clustered together. In contrast, European countries, e.g., Spain, France, and Italy, form a separate cluster between 1999 and 2015. This reflects that these countries might take different approaches to GIS theories and applications. Contrary to *GIS*, the semantic trajectories for *urban planning* (Figure 4b) and *spatial analysis* (Figure 4c) vary significantly across different countries. This indicates that these two keywords exhibit strong region-specific embeddings, reflecting that the research under these two keywords in our scientometric dataset is potentially more influenced by local policies, socioeconomic conditions, and so on. Their variations across countries also suggest that, even based on the same theoretical foundation, the practical applications of



■ **Figure 3** A t-SNE visualization of keyword embeddings with selected years and countries.

these concepts can vary and lead to country-dependent interpretations. Compared with *urban planning* and *spatial analysis*, *cellular automaton* (Figure 4d) shows similar semantic trajectories and clustered embeddings across countries. This country-wise consistency is likely attributed to the stronger mathematical and computational foundations of *cellular automaton*, which makes it potentially less influenced by local policies or conditions. This observation also indicates a more widely shared understanding and development of theories and applications in cellular automaton.

■ **Table 3** The top 10 keywords by spatio-temporal coverage.

Keyword	Time Span	Unique Years	Unique Countries	Total Count	Coverage Score
GIS	1993-2020 (27)	26	45	138	1215
Geographic Information System	1994-2020 (26)	20	36	72	936
Remote Sensing	1995-2020 (25)	17	29	48	725
Land Use	1995-2020 (25)	20	26	55	650
Model	1999-2020 (21)	12	28	41	588
Urban Planning	1998-2020 (22)	15	25	47	550
Spatial Analysis	1998-2020 (22)	18	25	54	550
Cellular Automaton	2000-2020 (20)	19	25	63	500
Visualization	1997-2019 (22)	14	20	39	440
Cadastre	2001-2020 (19)	8	20	24	380




Figure 4 PCA visualization of the semantic trajectory of keywords by country over time. Only the top 10 countries are included for visual clarity. We use the first and the second principle components for PCA visualization.

While semantic trajectories trace how a single keyword (used as a proxy for an underlying concept) evolves over time and across countries, they do not capture how it relates to other keywords in semantic space. Table 4 presents selected examples of keywords in different country-year combinations and how their meanings evolve. For each unique keyword-country-year combination, we identify its three most similar keywords from the other years, calculated based on Equation 2. For *GIS* in Austria (1999), its closest semantic matches are found in the UK and Australia to itself and *integrated model* in the Netherlands. The temporal distance between these matches is small, suggesting that the conceptualization of GIS remained relatively stable across these countries during this period. In contrast, *machine learning* in France (2003) follows a different pattern. It shares the strongest similarities with *ontology* and *temporal management* in Spain (2012) but also aligns with *machine learning* in Czechia (2020). This indicates that the early machine learning concept was probably integrated into various domains of GIScience over time. Lastly, *urban planning* in China (2011 and 2019) show a rather location-stable pattern that the strongest similarities are all found within China across different years and with related planning keywords. This suggests a more internally consistent evolution of urban planning concepts within the Chinese research community. In contrast, *urban planning* in Australia (2004 and 2018) shows similarities across many

countries (the US, the UK, Netherlands, and Finland). This indicates a more dynamic and globally connected evolution of the underlying concept. We can infer from this comparison that the concept of *urban planning* may drift slower and more localized in China, while at a faster rate and more international in Australia.

■ **Table 4** Selected cases of the top three similar keywords and their similarity scores across countries and years.

Query Keyword	Top 3 Similar Keywords	Sim.
GIS (Austria, 1999)	GIS (United Kingdom, 2000)	0.931
	Integrated Model (Netherlands, 2000)	0.894
	GIS (Australia, 2000)	0.892
Machine Learning (France, 2003)	Ontology (Spain, 2012)	0.916
	Temporal Management (Spain, 2012)	0.916
	Machine Learning (Czechia, 2020)	0.909
Urban Planning (China, 2011)	Urban Planning (China, 2019)	0.942
	Geographic Information System (China, 2016)	0.939
	Planning Support System (China, 2020)	0.938
Urban Planning (China, 2019)	Urban Spatial Dynamic (China, 2020)	0.953
	Scenario Planning (China, 2020)	0.948
	Urban Land Use (China, 2020)	0.948
Urban Planning (Australia, 2004)	Urban Planning (Finland, 2020)	0.947
	Urban Planning (United States, 1998)	0.937
	Urban Planning (Netherlands, 2016)	0.936
Urban Planning (Australia, 2018)	Urban Data (United States, 2019)	0.928
	Urban Land Use Change (United Kingdom, 2014)	0.927
	Urban Scaling Law (Europe, 2020)	0.924

Finally, we perform a sensitivity analysis on countries with at least 100 associated keywords to evaluate the impact of explicit location mentions on the extracted context embeddings. Using a two-tailed permutation test (with 1,000 permutations), we find that the average cosine distance ($1 - \text{cosine similarity}$) between embeddings with and without location mentions (0.010) is significantly smaller than what would be expected by chance (permutation mean = 0.242, std = 0.0005, $p < 0.001$). This indicates that explicit location mentions have minimal impact on the semantic representation of concepts in our case study. This is likely because geographical context is implicitly encoded in the text. We discuss this in more detail in the discussion section.

6 Discussion

This section discusses the challenges of defining spatio-temporal dimensions of concepts and the potential biases introduced in our case study. We also discuss findings from the sensitivity analysis, outline future research directions and the implications of spatio-temporal concept drift for ontology learning with large language models (LLMs).

6.1 Challenges in Defining Spatio-Temporal Dimensions of Concepts

Understanding spatio-temporal concept drift in scientific texts requires linking keywords (and their underlying concepts) to geographic locations and time, but this process inherently introduces biases.

In our case study, we use publication year as the temporal dimension of a concept. However, the publication year could be different from the actual study period (e.g., a paper published in 2010 on East Africa in the 1970s). We attribute spatial dimensions of concepts (keywords) based on location mentions in their associated abstracts. However, not all location mentions correspond to the actual study area; some may appear as examples, comparisons, or even counterexamples. We then aggregate these locations to use the country as the spatial unit of concept drift, overlooking regional variations within the country level. Take the concept of *urban planning* as an example. Its interpretation could differ significantly for New York City (e.g., a walkable, transit-oriented city) versus Los Angeles (e.g., a car-centric city) over the years. These regional disparities would become particularly pronounced in larger countries (e.g., the United States) with diverse geographic and socioeconomic conditions. Our spatial aggregation approach implicitly assumes concept homogeneity at the country level, which introduces biases into the learned embeddings.

Future work could explore improved spatio-temporal scoping techniques to capture study periods and areas more accurately; it should also include different spatial levels – cities, countries, and continents – to measure variations.

6.2 Sensitivity to Location Removal in Context Embeddings

Our sensitivity analysis reveals varying effects across countries when removing location mentions. For example, keywords associated with Japan show slightly larger differences in their embeddings, though the overall differences remained small. Here, several factors may complicate the interpretation of these results. First, the dataset has a substantial imbalance, with publications mentioning the United States far outnumbering those that mention other countries. When extracting unique country-year combinations, this imbalance leads to sparse samples for less represented countries, thus masking meaningful patterns. Second, we did not account for the proportion of location mentions within abstracts, e.g., some abstracts contain a list of study areas, while others mention one location briefly. The observed differences in embeddings with and without locations may be due to the removal of more contextual information rather than an inherent sensitivity to geographic reference. These factors need a more fine-grained analysis to quantify the impact of explicit location mentions on embedding representations in future work.

6.3 Implication for Ontology Learning with Large Language Models

Since large language models (LLMs) become more commonly used for ontology learning tasks [27, 2, 41], we need to ensure geographic and temporal variations in knowledge are accounted for to have more context-aware representations. Current LLMs are trained on vast corpora of text that usually lack explicit spatial and temporal structuring [13], which would likely overlook the spatio-temporal variations in concept representations unless specified in the prompt. Our findings show that concepts in scientific texts evolve differently across geographic space and over time. This suggests that ontologies derived from LLMs may inherit hidden geographic biases. For instance, when an LLM processes the concept of *smart city*, its interpretation might be overly influenced (and represented) by temporally and regionally dominant implementations (e.g., the recent decade in Singapore) if the training data is dominated by publications from this region and time period.

Our observations show that concepts in scientific texts can vary across geographic space and time, and suggest the need for a more context-aware mechanism when using LLMs for ontology learning. This could be achieved with region-specific knowledge validation

and/or the development of geographically aware prompting strategies. To capture the spatio-temporal dynamics in scientific concepts, future work could include the design of more few-shot learning approaches, where examples are carefully selected to represent diverse temporal and geographic interpretations of concepts.

7 Conclusions

Space and time are central to the study of geography and GIScience. These dimensions not only shape our daily physical interactions of *when* and *where*, but also influence the abstract representation of concepts in scientific knowledge. With the ever-increasing volume of research publications, we need methods to better structure concepts embedded in scientific research for organization and retrieval purposes. Encoded in an ontology, we could also account for the spatio-temporal dynamics of concepts, which are constantly evolving – often at varying rates across regions – due to technology and societal changes, for more effective research data management (RDM).

In this work, we introduce the notion of spatio-temporal concept drift. We complement previous work on concept drift by including the spatial dimension, and propose a novel approach using word embedding techniques to capture this drift over space and time. Using a scientometric dataset in the field of GIScience, we demonstrate that keywords (used as proxies for underlying concepts) show varying drift patterns over time and across countries. Spatially grounded concepts, such as *urban planning* (as compared to *cellular automaton*), can have substantial differences in meanings for different countries and over time.

The implications of this work extend beyond improving the understanding of concepts in scientific texts to enhancing FAIR-based RDM systems. This reminds us, for example, that concepts like *cellular automaton* may require less user intervention, while concepts like *urban planning* may need query enrichment to account for local and temporal variation to better match a user’s keyword. Given the observed spatio-temporal concept drift and the increasing use of ontology learning with large language models (LLMs), we also suggest that LLM-based ontology learning mechanisms should explicitly account for the spatial and temporal dimensions of concept representation. Making RDM systems and ontology learning approaches more sensitive to these variations will help improve retrieval and maintain the relevance of knowledge in scientific texts.

References

- 1 Hervé Abdi and Lynne J Williams. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459, 2010. doi:10.1002/wics.101.
- 2 Hamed Babaei Giglou, Jennifer D’Souza, and Sören Auer. LLMs4OL: Large language models for ontology learning. In *International Semantic Web Conference*, pages 408–427. Springer, 2023. doi:10.1007/978-3-031-47240-4_22.
- 3 David Bamman, Chris Dyer, and Noah A Smith. Distributed representations of geographically situated language. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 828–834, 2014. doi:10.3115/v1/p14-2134.
- 4 Lawrence Barsalou. Concepts and meaning. In L. Barsalou, W. Yeh, B. Luka, K. Olseth, K. Mix, and L. Wu, editors, *Chicago Linguistic Society 29: Papers From the Parasession on Conceptual Representations*, pages 23–61. University of Chicago, 1993.
- 5 Iz Beltagy, Kyle Lo, and Arman Cohan. SciBERT: A pretrained language model for scientific text. In *Conference on Empirical Methods in Natural Language Processing*, 2019. doi:10.18653/v1/D19-1371.

- 6 Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29, 2016. doi:10.48550/arXiv.1607.06520.
- 7 Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186, 2017. doi:10.1126/science.aal4230.
- 8 Giuseppe Capobianco, Danilo Cavaliere, Sabrina Senatore, et al. Ontodrift: a semantic drift gauge for ontology evolution monitoring. In *CEUR Workshop Proceedings*, volume 2821, pages 1–10. CEUR-WS, 2020. URL: <https://ceur-ws.org/Vol-2821/paper1.pdf>.
- 9 Christophe Claramunt. Ontologies for geospatial information: Progress and challenges ahead. *Journal of Spatial Information Science*, 20:35–41, 2020. doi:10.5311/JOSIS.2020.20.666.
- 10 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, 2019. doi:10.18653/v1/N19-1423.
- 11 Stephanie Duce and Krzysztof Janowicz. Microtheories for spatial data infrastructures-accounting for diversity of local conceptualizations at a global level. In *Geographic Information Science: 6th International Conference, GIScience 2010, Zurich, Switzerland, September 14-17, 2010. Proceedings 6*, pages 27–41. Springer, 2010. doi:10.1007/978-3-642-15300-6_3.
- 12 Max J Egenhofer and David M Mark. Naive geography. In *Spatial Information Theory A Theoretical Basis for GIS: International Conference COSIT'95 Semmering, Austria, September 21-23, 1995 Proceedings 2*, pages 1–15. Springer, 1995. doi:10.1007/3-540-60392-1_1.
- 13 Fahim Faisal and Antonios Anastasopoulos. Geographic and geopolitical biases of language models. In Duygu Ataman, editor, *Proceedings of the 3rd Workshop on Multi-lingual Representation Learning (MRL)*, pages 139–163, Singapore, December 2023. Association for Computational Linguistics. doi:10.18653/v1/2023.mrl-1.12.
- 14 Antske Fokkens, Serge Ter Braake, Isa Maks, Davide Ceolin, et al. On the semantics of concept drift: Towards formal definitions of semantic change. *Drift-a-LOD@ EKAW*, 2016. URL: https://ceur-ws.org/Vol-1799/Drift-a-LOD2016_paper_2.pdf.
- 15 Aldo Gangemi, Nicola Guarino, Claudio Masolo, Alessandro Oltramari, and Luc Schneider. Sweetening ontologies with DOLCE. In *International conference on knowledge engineering and knowledge management*, pages 166–181. Springer, 2002. doi:10.1007/3-540-45810-7_18.
- 16 Hongyu Gong, S. Bhat, and Pramod Viswanath. Enriching word embeddings with temporal and spatial information. In *Conference on Computational Natural Language Learning*, 2020. doi:10.18653/v1/2020.conll-1.1.
- 17 Nicola Guarino. Formal ontology, conceptual analysis and knowledge representation. *International journal of human-computer studies*, 43(5-6):625–640, 1995. doi:10.1006/ijhc.1995.1066.
- 18 Ramanathan Guha, Rob McCool, and Eric Miller. Semantic search. In *Proceedings of the 12th international conference on World Wide Web*, pages 700–709, 2003. doi:10.1145/775152.775250.
- 19 Prashant Gupta and Mark Gahegan. Categories are in flux, but their computational representations are fixed: That’s a problem. *Transactions in GIS*, 24(2):291–314, 2020. doi:10.1111/tgis.12602.
- 20 William L. Hamilton, Jure Leskovec, and Dan Jurafsky. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501, 2016. doi:10.18653/v1/P16-1141.
- 21 Krzysztof Janowicz, Song Gao, Grant McKenzie, Yingjie Hu, and Budhendra Bhaduri. GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond, 2020. doi:10.1080/13658816.2019.1684500.

- 22 Krzysztof Janowicz, Martin Raubal, and Werner Kuhn. The semantics of similarity in geographic information retrieval. *Journal of Spatial Information Science*, 2:29–57, 2011. doi:10.5311/JOSIS.2011.2.3.
- 23 Werner Kuhn. Semantic reference systems. *International Journal of Geographical Information Science*, 17(5):405–409, 2003. doi:10.1080/1365881031000114116.
- 24 Werner Kuhn, Martin Raubal, and Peter Gärdenfors. Cognitive semantics and spatio-temporal ontologies. *Spatial Cognition & Computation*, 7(1):3–12, 2007. doi:10.1080/13875860701337835.
- 25 Stephen C Levinson. Language and space. *Annual review of Anthropology*, 25(1):353–382, 1996. doi:10.1146/annurev.anthro.25.1.353.
- 26 Maurice Lineman, Yuno Do, Ji Yoon Kim, and Gea-Jae Joo. Talking about climate change and global warming. *PloS one*, 10(9):e0138996, 2015. doi:10.1371/journal.pone.0138996.
- 27 Huu Tan Mai, Cuong Xuan Chu, and Heiko Paulheim. Do LLMs really adapt to domains? an ontology learning perspective. In *International Semantic Web Conference*, pages 126–143. Springer, 2024. doi:10.1007/978-3-031-77844-5_7.
- 28 David M. Mark, Barry Smith, and Barbara Tversky. Ontology and geographic objects: An empirical study of cognitive categorization. In *Conference On Spatial Information Theory*, pages 283–298, 1999. doi:10.1007/3-540-48384-5_19.
- 29 Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations*, 2013. doi:10.48550/arXiv.1301.3781.
- 30 Malvina Nissim, Rik van Noord, and Rob van der Goot. Fair is better than sensational: Man is to doctor as woman is to doctor. *Computational Linguistics*, 46(2):487–497, 2020. doi:10.1162/coli_a_00379.
- 31 Charles Kay Ogden and Ivor Armstrong Richards. *The Meaning of Meaning: A Study of the Influence of Language upon Thought and of the Science of Symbolism*. Harcourt, Brace & World, Inc., 1923. doi:10.2307/2015195.
- 32 Jeffrey Pennington, Richard Socher, and Christopher D Manning. GloVe: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014. doi:10.3115/v1/d14-1162.
- 33 Dharmen Punjani, Kuldeep Singh, Andreas Both, Manolis Koubarakis, Iosif Angelidis, Konstantina Bereta, Themis Beris, Dimitris Bilidas, Theofilos Ioannidis, Nikolaos Karalis, et al. Template-based question answering over linked geospatial data. In *Proceedings of the 12th workshop on geographic information retrieval*, pages 1–10, 2018. doi:10.1145/3281354.3281362.
- 34 Martin Raubal. Formalizing conceptual spaces. In *Formal ontology in information systems, proceedings of the third international conference (FOIS 2004)*, volume 114, pages 153–164. Citeseer, 2004.
- 35 Martin Raubal. Representing concepts in time. In *Spatial Cognition VI. Learning, Reasoning, and Talking about Space: International Conference Spatial Cognition 2008, Freiburg, Germany, September 15-19, 2008. Proceedings 6*, pages 328–343. Springer, 2008. doi:10.1007/978-3-540-87601-4_24.
- 36 M Andrea Rodriguez and Max J. Egenhofer. Determining semantic similarity among entity classes from different ontologies. *IEEE transactions on knowledge and data engineering*, 15(2):442–456, 2003. doi:10.1109/TKDE.2003.1185844.
- 37 Maja Rudolph and David Blei. Dynamic embeddings for language evolution. In *Proceedings of the 2018 world wide web conference*, pages 1003–1011, 2018. doi:10.1145/3178876.3185999.
- 38 Christoph Schlieder. Digital heritage: Semantic challenges of long-term preservation. *Semantic Web*, 1(1-2):143–147, 2010. doi:10.3233/SW-2010-0013.
- 39 Angela Schwering. Approaches to semantic similarity measurement for geo-spatial data: a survey. *Transactions in GIS*, 12(1):5–29, 2008. doi:10.1111/j.1467-9671.2008.01084.x.

- 40 Meilin Shi, Krzysztof Janowicz, Zilong Liu, Mina Karimi, Ivan Majic, and Alexandra Fortacz. Defining concept drift and its variants in research data management: A scientometric case study on geographic information science. *Transactions in GIS*, 29(3):e70058, 2025. doi:10.1111/tgis.70058.
- 41 Cogan Shimizu and Pascal Hitzler. Accelerating knowledge graph and ontology engineering with large language models. *Journal of Web Semantics*, page 100862, 2025. doi:10.1016/j.websem.2025.100862.
- 42 Barry Smith and David M Mark. Geographical categories: an ontological investigation. *International journal of geographical information science*, 15(7):591–612, 2001. doi:10.1080/13658810110061199.
- 43 Barry Smith and David M Mark. Do mountains exist? towards an ontology of landforms. *Environment and Planning B: Planning and Design*, 30(3):411–427, 2003. doi:10.1068/b12821.
- 44 Thanos G Stavropoulos, Stelios Andreadis, Efstratios Kontopoulos, and Ioannis Kompatsiaris. SemaDrift: A hybrid method and visual tools to measure semantic drift in ontologies. *Journal of Web Semantics*, 54:87–106, 2019. doi:10.1016/j.websem.2018.05.001.
- 45 Wolfgang G Stock. Concepts and semantic relations in information science. *Journal of the American Society for Information Science and Technology*, 61(10):1951–1969, 2010. doi:10.1002/asi.21382.
- 46 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- 47 Saskia Van Putten, Carolyn O’Meara, Flurina Wartmann, Joanne Yager, Julia Villette, Claudia Mazzuca, Claudia Bieling, Niclas Burenhult, Ross Purves, and Asifa Majid. Conceptualisations of landscape differ across european languages. *Plos one*, 15(10):e0239858, 2020. doi:10.1371/journal.pone.0239858.
- 48 Stella Verkijk, Ritten Roothaert, Romana Pernisch, and Stefan Schlobach. Do you catch my drift? on the usage of embedding methods to measure concept shift in knowledge graphs. In *Proceedings of the 12th Knowledge Capture Conference 2023*, pages 70–74, 2023. doi:10.1145/3587259.3627555.
- 49 Shenghui Wang, Stefan Schlobach, and Michel Klein. Concept drift and how to identify it. *Journal of Web Semantics*, 9(3):247–265, 2011. doi:10.1016/j.websem.2011.05.003.
- 50 Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, et al. The FAIR guiding principles for scientific data management and stewardship. *Scientific data*, 3(1):1–9, 2016. doi:10.1038/sdata.2016.18.
- 51 Xiaohuan Wu, Weihua Dong, Lun Wu, and Yu Liu. Data and Code for "Research Themes of Geographical Information Science during 1991–2020: A Retrospective Bibliometric Analysis", 2022. doi:10.6084/m9.figshare.19242654.v1.
- 52 Bo Yan, Krzysztof Janowicz, Gengchen Mai, and Song Gao. From ITDL to Place2Vec: Reasoning about place type similarity and relatedness by learning embeddings from augmented spatial contexts. In *Proceedings of the 25th ACM SIGSPATIAL international conference on advances in geographic information systems*, pages 1–10, 2017. doi:10.1145/3139958.3140054.
- 53 Wei Zhai, Xueyin Bai, Yu Shi, Yu Han, Zhong-Ren Peng, and Chaolin Gu. Beyond Word2vec: An approach for urban functional region extraction and identification by combining Place2vec and POIs. *Computers, Environment and Urban Systems*, 74:1–12, 2019. doi:10.1016/j.compenvurbsys.2018.11.008.
- 54 Yating Zhang, Adam Jatowt, Sourav S Bhowmick, and Katsumi Tanaka. The past is not a foreign country: Detecting semantically similar terms across time. *IEEE Transactions on Knowledge and Data Engineering*, 28(10):2793–2807, 2016. doi:10.1109/TKDE.2016.2591008.
- 55 Yating Zhang, Adam Jatowt, and Katsumi Tanaka. Is tofu the cheese of asia?: Searching for corresponding objects across geographical areas. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 1033–1042, 2017. doi:10.1145/3041021.3055132.

A Conference and Journal Abbreviation Reference

■ **Table 5** Full names and abbreviations of selected conferences and journals.

Conference/Journal Name	Abbreviation
International Conference on Spatial Information Theory	COSIT
International Conference on Geographic Information Science	GIScience
Computers, Environment and Urban Systems	CEUS
Cartography and Geographic Information Science	CaGIS
Environment and Planning B: Urban Analytics and City Science	EPB
GeoInformatica	GeoI
International Journal of Geographical Information Science	IJGIS
Journal of Geographical Systems	JGS
Journal of Spatial Information Science	JOSIS
Spatial Cognition & Computation	SCC
Transactions in GIS	TGIS

B Sensitivity Analysis of the Label Embedding Weight

■ **Table 6** Top similar keywords retrieved under different values of the label embedding weight α , along with their similarity scores. The value of $\alpha = 0.3$ is used in the case study in this paper. Note that the examples are included post hoc to illustrate the qualitative effects of different α values.

Weight	Query Keyword	
	Urban Planning (Australia, 2018)	Climate Change (US, 2020)
$\alpha = 0.1$	urban scaling law (Europe, 2020): 0.925	climate change (US, 2014): 0.913
	zipf's law for city (Europe, 2020): 0.924	sea level rise (US, 2019): 0.910
	land use (Europe, 2020): 0.923	storm surge inundation (US, 2019): 0.909
	population density (Europe, 2020): 0.922	lidar (US, 2008): 0.909
	radial analysis (Europe, 2020): 0.922	greening scenario (US, 2018): 0.909
$\alpha = 0.3$	urban data (US, 2019): 0.928	climate change (US, 2014): 0.938
	urban land use change (UK, 2014): 0.927	climate change (UK, 2018): 0.919
	urban scaling law (Europe, 2020): 0.924	sea level rise (US, 2019): 0.912
	urban planning (US, 2019): 0.924	urban heat island (US, 2018): 0.911
	residential mobility (UK, 2014): 0.921	seasonal impact (US, 2018): 0.911
$\alpha = 0.5$	urban planning (US, 2019): 0.959	climate change (US, 2014): 0.967
	urban planning (Brazil, 2003): 0.948	climate change (UK, 2018): 0.956
	urban planning (Poland, 2017): 0.948	climate change (UK, 2012): 0.947
	urban planning (Spain, 2017): 0.947	climate change (US, 2013): 0.947
	urban planning (Netherlands, 2016): 0.943	climate change (US, 2015): 0.946