



Multiple Texts as a Limiting Factor in Online Learning: Quantifying (Dis-)similarities of Knowledge Networks

Alexander Mehler^{1*}, Wahed Hemati¹, Pascal Welke², Maxim Konca¹ and Tolga Uslu¹

¹ Text-Technology Lab, Institute of Computer Science, Faculty of Computer Science and Mathematics, Goethe University Frankfurt, Frankfurt, Germany, ² Machine Learning and Artificial Intelligence Lab, Institute for Computer Science, University of Bonn, Bonn, Germany

We test the hypothesis that the extent to which one obtains information on a given topic through Wikipedia depends on the language in which it is consulted. Controlling the size factor, we investigate this hypothesis for a number of 25 subject areas. Since Wikipedia is a central part of the web-based information landscape, this indicates a language-related, linguistic bias. The article therefore deals with the question of whether Wikipedia exhibits this kind of linguistic relativity or not. From the perspective of educational science, the article develops a computational model of the information landscape from which multiple texts are drawn as typical input of web-based reading. For this purpose, it develops a hybrid model of intra- and intertextual similarity of different parts of the information landscape and tests this model on the example of 35 languages and corresponding Wikipedias. In the way it measures the similarities of hypertexts, the article goes beyond existing approaches by examining their structural and semantic aspects intra- and intertextually. In this way it builds a bridge between reading research, educational science, Wikipedia research and computational linguistics.

Keywords: multiple texts, information landscape, knowledge graphs, intratextual similarity, intertextual similarity, three-level topic model, network similarity measurement, linguistic relativity

OPEN ACCESS

Edited by:

Olga Zlatkin-Troitschanskaia,
Johannes Gutenberg University
Mainz, Germany

Reviewed by:

Guillermo Solano-Flores,
Stanford University, United States
Mary Frances Rice,
University of New Mexico,
United States

*Correspondence:

Alexander Mehler
mehler@em.uni-frankfurt.de

Specialty section:

This article was submitted to
Digital Education,
a section of the journal
Frontiers in Education

Received: 15 May 2020

Accepted: 23 September 2020

Published: 03 November 2020

Citation:

Mehler A, Hemati W, Welke P,
Konca M and Uslu T (2020) Multiple
Texts as a Limiting Factor in Online
Learning: Quantifying (Dis-)similarities
of Knowledge Networks.
Front. Educ. 5:562670.
doi: 10.3389/feduc.2020.562670

1. INTRODUCTION

Reading is increasingly carried out by means of online multiple texts, which can simultaneously consist of (segments of) texts of diverse genres, registers, authorships, credibilities etc. (Barzilai and Zohar, 2012; Goldman et al., 2012; Britt et al., 2018). That is, learning takes place, so to speak, on the basis of “document collages” whose components are gathered from a constantly growing, nowadays mostly web-based information landscape (Zlatkin-Troitschanskaia et al., 2018) or space (Hartman et al., 2018)¹. The multiplicity of the texts involved and the diversity of their genres and register

¹The term information landscape (Zlatkin-Troitschanskaia et al., 2019a,b) includes but is not limited to the web as both a huge and highly diverse text data repository or source of online reading (Cho and Afflerbach, 2015; Wolf, 2018), which is substructured along countless web genres (Mehler et al., 2010), registers and thematic domains. Online reading resembles a traversal of this landscape, each of which involves numerous decisions about what to read, in what sequence and in what depth. The term “landscape” manifests this dual character: on the one hand, the information landscape is a repository that offers innumerable decision possibilities (intertext-as-product perspective), which on the other hand are to be decided by the reader (Cho and Afflerbach, 2015; Britt et al., 2018) in such a way that for each reading process (usually different) multiple text is delimited in this landscape (intertext-as-process perspective).

(Halliday and Hasan, 1989) are text-linguistic characteristics of online reading (Britt et al., 2018). A third, so to speak macroscopic aspect of this process is the starting point of this article. It is about the *Information Landscape* (IL) from which innumerable readers in countless reading processes delineate ever new multiple texts and thus manifest a distributed process through which this landscape is opened up. To introduce our research agenda regarding this IL, we start from the *Documents Model* (DM) of Perfetti et al. (1999) and Britt et al. (2012). While multiple texts are studied by a wide range of approaches², the reason for choosing the DM as a starting point is due to its text-linguistic heritage—based on the *Construction-Integration Model* (CIM) of Kintsch (1998)—and its context model, which facilitates modular extensions. As far as the text-linguistic orientation of the DM is concerned, its notion of the so-called intertext model is of particular interest for our study of the IL.

Generally speaking, the DM distinguishes two outcomes of multiple text comprehension: the *Intertext Model* (IM), which comprises representations of the constituents of multiple texts and their links, and the *Mental Model* (MM)³ as a result of comprehension processes that operate within and beyond the boundaries of these constituents. This includes the process of integration, a term borrowed from the CIM, which in the DM also concerns information from different texts. In contrast to text linguistics, which predicts that cohesion and coherence relations should be resolvable within the boundaries of a text to facilitate its understanding (Kintsch, 1998), this condition does not usually apply to multiple texts: they induce additional intrinsic cognitive loads (Sweller, 1994) as a result of interacting elements of separate texts (such as conflicting, contradicting, or otherwise incoherent statements about the same event; Barzilai and Zohar, 2012) and increased efforts in decision making as a result of hyperlinkage (DeStefano and LeFevre, 2007)⁴. Consequently, Goldman et al. (2012, p. 357) speak of (as we may add: online) reading as an *intertextual process*.

Britt et al. (2012) assume that intertext models represent selected constituents of multiple texts as “document entities” together with entity-related information (e.g., on authorship). This is supplemented by three types of links: IM-related source-to-source (e.g., *x supports or contradicts y*), MM-related content-to-content and source-to-content links (see **Figure 1**). A prediction of the DM, which is crucial for our work, is that the probability of generating an intertext model as a result of reading a multiple text is a function of the number of the texts involved, their authors, the perspectives they provide on the corresponding described situation (Britt et al., 2012, p. 171), the tasks to be accomplished and other contextual factors (Britt

et al., 2018). This suggests to speak of the intertext model as a kind of *cognitive map* (Downs and Stea, 1977) of the underlying multiple text, where the MM abstracts from this textbase (e.g., by applying macro operations; van Dijk, 1980; van Dijk and Kintsch, 1983): that is, readers produce intertext models as cognitive maps of multiple texts as parts of the underlying IL, while groups or communities of readers produce distributed cognitive maps (Mehler et al., 2019) of larger sections of the IL or the IL as a whole. This duality of small- and large-scale reading processes leads to the object of this article. That is, we ask how the IL looks like from the perspective of these distributed cognitive maps or vice versa, how it presents itself to different reader communities.

The latter question will be in the focus of this article, which is organized as follows: section 2 outlines a model of distributed reading of multiple texts by multiple readers that overcomes the single-reader perspective of the DM. Section 3 explains the relevance of Wikipedia for the educational science pursued here and gives an overview of related research. Section 4 explains our research questions and describes in detail the methods we have developed to answer them. In section 5, we describe our experiments and discuss their results. More specifically, in section 5.2 we explain how our approach relates to research on linguistic relativity and in section 5.3 we discuss two implications of our findings for research on information processing and online reasoning in the context of higher education, which is the focus of the special issue to which this article contributes. This includes two implications: The first one concerns a reformulation of the closed world assumption which describes a problematic attitudinal basis for reading learning resources like Wikipedia; the second concerns the problem of blurred domain boundaries regarding a microstructural, learner-related and a macrostructural perspective in terms of distributed reading processes. Finally, in section 6 we give a conclusion and an outlook on future work.

2. TOWARD A MODEL OF DISTRIBUTED READING OF MULTIPLE TEXTS

Although the DM takes the necessary step of generalizing the CIM toward modeling multiple texts, it is largely single reader-oriented. To broaden this focus, we generalize the DM conceptually in two steps:

1. Step 1: *Microscopic alignment*: In a first step, which is still within the boundaries of the DM, we consider intertext models as reader-centered approximations of parts of the IL with varying degrees of explicitness (reflecting the probability mentioned above). This predicts that different readers can approximate different parts (i.e., multiple texts) of the IL just as they can align (Pickering and Garrod, 2004) their intertext models of overlapping parts depending on their interaction, which according to Salmerón et al. (2018, p. 286) is a characteristic of non-academic online reading and also regards online collaborative learning (Coiro et al., 2018, p. 487). Starting from the context model of Britt et al. (2018, p. 53), **Figure 2** illustrates this alignment scenario in terms of a situation semantic adaptation (Mehler and

²See Primor and Katzir (2018) for a current overview and List and Alexander (2019) for an integrated view of this model landscape.

³Instead of situation models, Britt et al. (2018) speak of integrated models as mental models of the semantic content of multiple texts, to emphasize that this content is not limited to descriptions of situations (as provided, e.g., by narrative texts). In this way, depending on the task, texts of different types get into the focus of their reading model. We prefer to speak of mental models to emphasize the openness of our approach to the DM.

⁴For a review of aspects of cognitive load of Internet-based online reading see Loh and Kanai (2016).

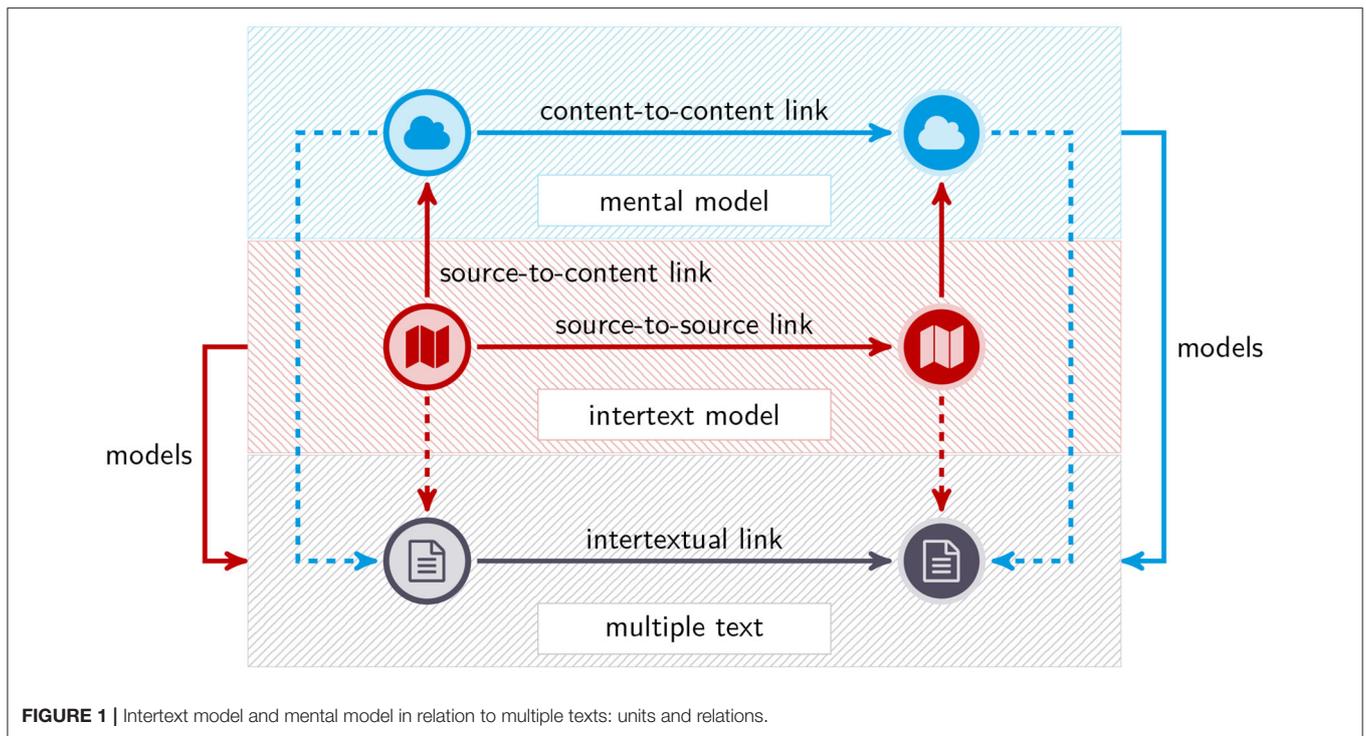


FIGURE 1 | Intertext model and mental model in relation to multiple texts: units and relations.

Ramesh, 2019): two readers *A* and *B* read not necessarily different multiple texts in related contexts (related to space, time, etc.) to solve the same or related tasks (manifested by task instructions etc.) in order to achieve the same or related goals. We assume that the texts originate from an IL in whose context they describe a situation *S* and that *A* and *B* refer to the same resource situation *R* to understand which situation their multiple texts actually describe, where all references to contextual units are indirect: they are mediated through mental representations (cf. Britt et al., 2018, p. 45) of multiple texts (IM), described situations (MM), reading contexts (context model), task contexts (task model) and resource situations [long-term memory (LTM) or knowledge space]. As a result of collaborative, cooperative or simply parallel reading processes, the representations of the readers may align with each other, in the short-term (concerning, e.g., their MMs) or long-term (regarding their LTMs). That is, as illustrated in **Figure 2**, *A* and *B* have the possibility to align their mental representations so that they understand the same or similar multiple texts as descriptions of the same or similar situations⁵. Evidently, such an alignment requires many things, but at least the chance that both readers have access to the same or semantically sufficiently similar texts from which they can extract the same or sufficiently similar

multiple texts. *But do they?* This question brings us to the second step of generalizing the DM:

2. Step 2: *Macroscopic alignment*: From the perspective of reader communities, reading is a distributed process that approximates a multifaceted IL, so that both the IL and its distributed representation by innumerable intertext models jointly develop. Obviously, parallel to the diversity of the IL, communities of readers are also diverse as a result of a wide range of factors (Hsieh, 2012; Hargittai and Dobransky, 2017; Braasch et al., 2018b): for example, membership in different language communities (as focused by this article), ethnicities, cultures, age groups, social groups (e.g., families), residency in different places (or geographies; Graham et al., 2015) or practice of different social roles. In any event, in analogy to Step 1, we may expect that different communities dealing with the same or semantically similar parts of the IL should be able to align their corresponding intertext models among each other. In shorter terms: different groups should be able to represent similar parts of the IL in a similar way. *But do they actually have access to the same or at least sufficiently similar parts of the IL—especially under the condition that they deal with the same topic?*
3. Step 1 concerns the microscopic alignment of intertext models as a result of reading situations that are paired together via shared, overlapping, or otherwise related multiple texts, tasks, etc. Step 2 refers to the macroscopic alignment of reader communities as a result of countless such pairings, whereby these communities and consequently their alignments are subdivided according to their social structure. In the range of these extremes there are meso-level alignment processes manifested by smaller groups of agents (such as learning

⁵Adding the notion of a described situation to the context model of Britt et al. (2018) allows for distinguishing communication scenarios, such as misunderstandings (*A* and *B* represent the same situation *S* in mutually incompatible ways), disinformation (*A* pretends to describe a situation he or she knows to be unreal, whereas *B* considers it to be factual) or misinformation (*A* and *B* represent the same unreal situation *S*, which both consider real) (Kendeou et al., 2019) and related phenomena.

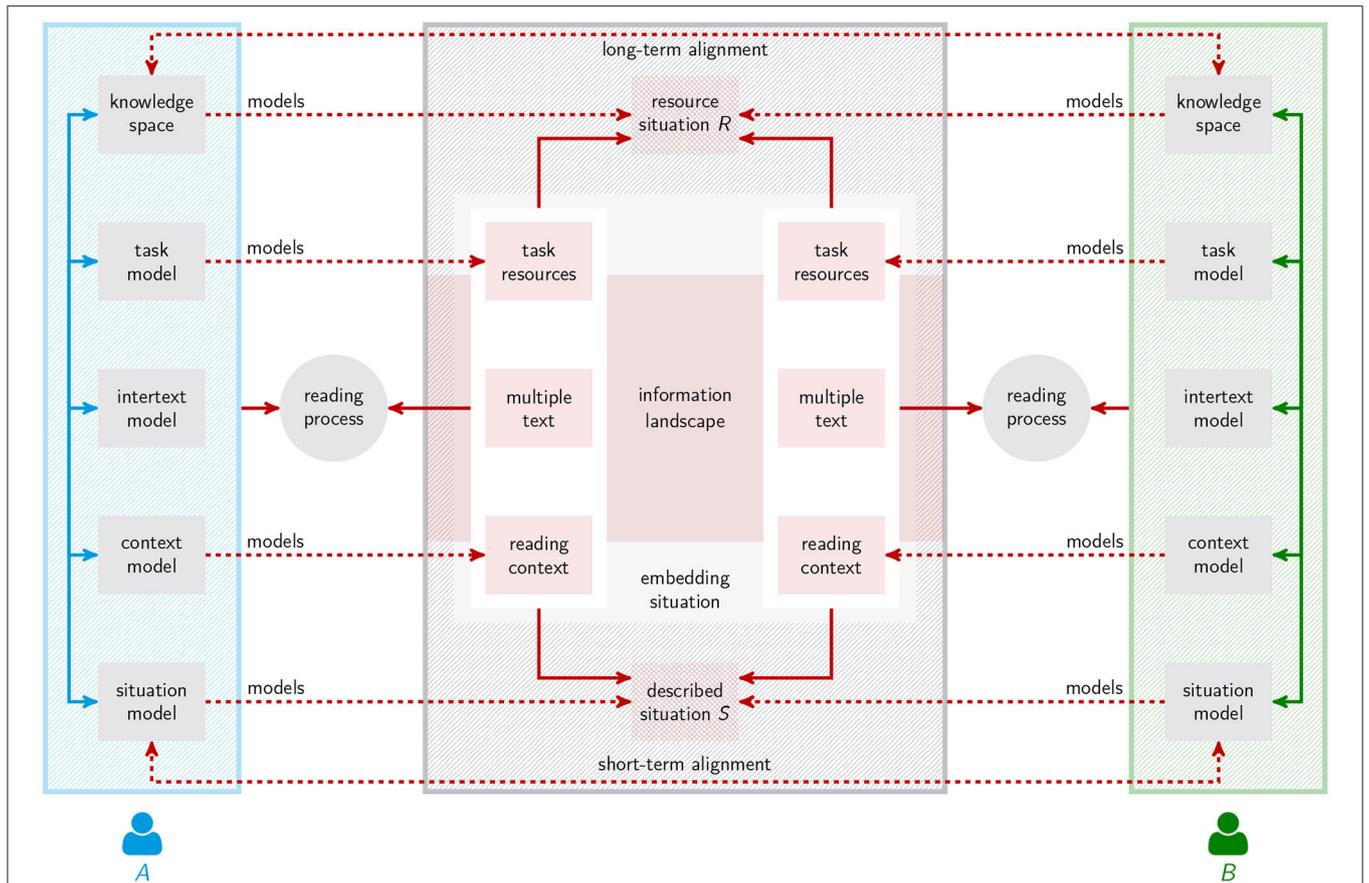


FIGURE 2 | Extending the DM in terms of alignment research.

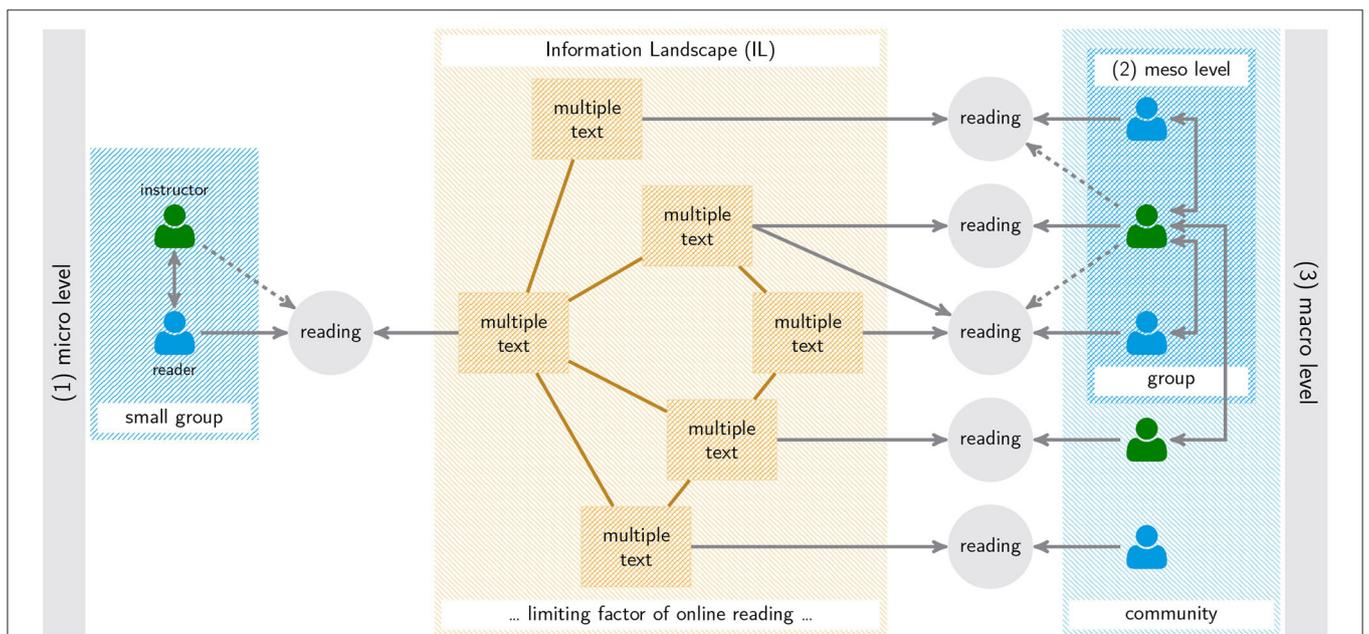


FIGURE 3 | Three research perspectives on multiple texts.

groups) (see **Figure 3**). The three levels have in common that the underlying IL is temporally on a different scale: it is subject to a lower dynamic than the multiple texts that are extracted from it as a result of reading processes—even in the case of algorithmically (generated and) linked documents (where the underlying algorithms may reflect user profiles). But how uniformly does the IL present itself to its (communities or groups of) readers? Obviously this question is currently outside the scope of the framework of the DM and its relatives.

Step 2 concerns precisely the viewpoint of this article. That is, we are concerned with a central prerequisite for alignable intertext models among readers as members of large communities. This refers to the intertextual shape of the IL from the perspective of different communities who may have different accesses to it or “see” different landscapes, even in situations where the opposite would be assumed. The DM and related approaches do not model what the multiple texts are extracted from and what countless intertext models in their distributed totality ultimately represent, that is, an underlying multifaceted, highly dynamic IL, its numerous document nodes and their relational, intertextual embeddings.

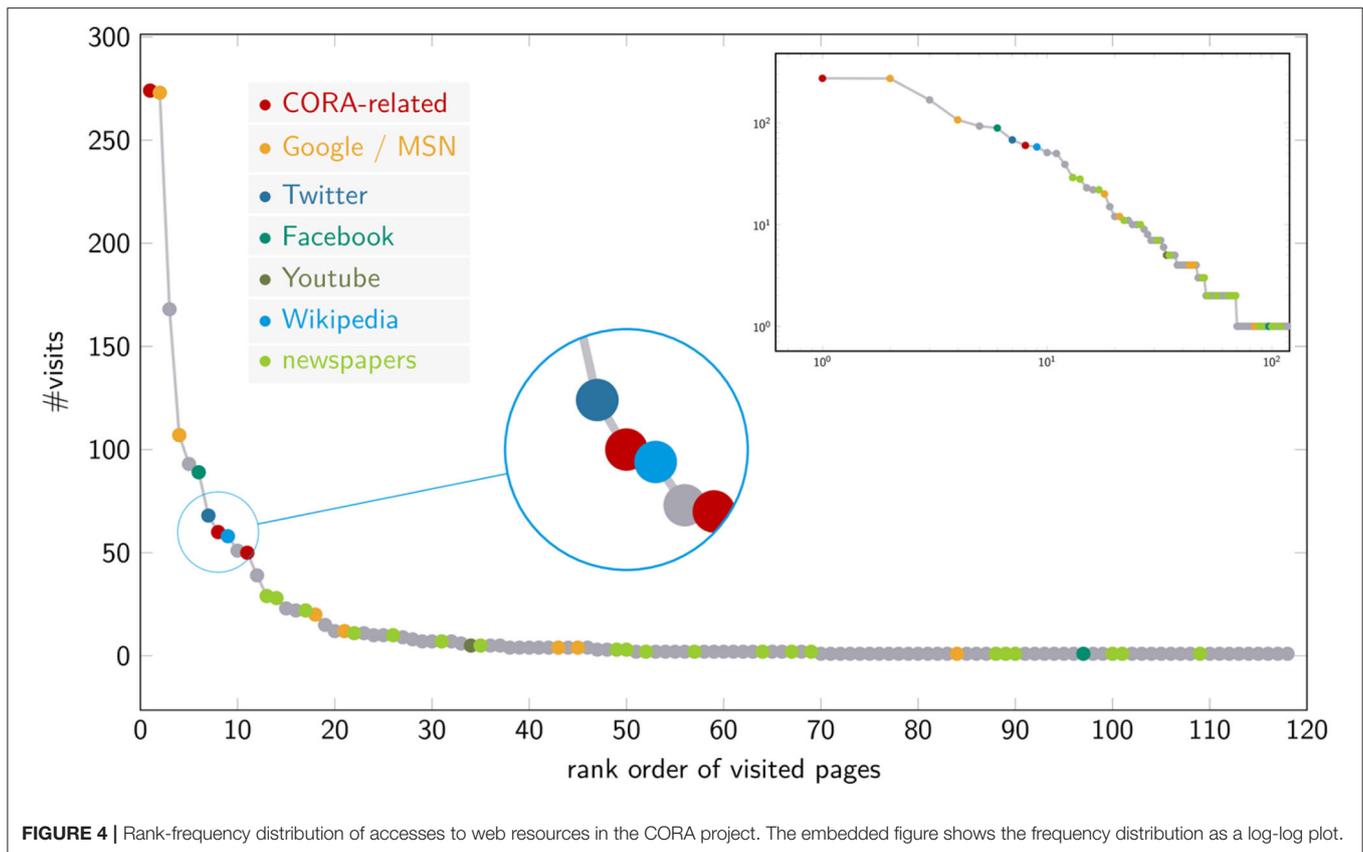
According to Hartman et al. (2018, p. 56), reading research mostly considers small amounts of offline texts pre-selected by the experimenter rather than open ILs in which users decide what to read. But if reading is a kind of problem solving that involves multiple search and decision processes (Britt et al., 2018, p. 43) (e.g., about what to search for and where to find it), then the question arises as to the limits of these processes as imposed by the IL and how they differ for which reader communities. Apparently, approaches to multiple texts focus on micro-models that leave the corresponding macro-models, which inform about the shape of the IL and its organizational laws, under-specified. The present paper takes a step in the direction of filling this gap: it develops a macroscopic model of the IL and examines how its shape appears from the perspective of certain large-scale reader communities. Our aim is, so to speak, to impart knowledge about the “wild” in which the sort of reading takes place which according to Braasch et al. (2018b, p. 535) is to become the subject of reading research. Thus, our approach is complementary to current research on the intertext model: we study the IL underlying the construction of intertext models from a macroscopic perspective, in contrast to reading research, which starts from a microscopic perspective of small groups or individual readers (see **Figure 3**). In terms of the integrated framework of multiple texts (List and Alexander, 2019) we are concerned with the intertextuality of those information units to which the cognitive strategies and behavioral skills of readers are related. That is, in modification of the fourth goal of future research on the use of multiple sources according to Braasch et al. (2018b), we deal with the phenomenon that different communities are offered different information, especially in the context of the same topic. The extent to which this phenomenon applies to different language communities will be examined using the example of the most frequently used knowledge resource on the Web, that is Wikipedia (RRID:SCR_004897).

3. WIKIPEDIA: EDUCATIONAL RELEVANCE AND BIAS

Wikipedia is a primary multilingual source for web-based knowledge acquisition and online learning (Lucassen and Schraagen, 2010; Okoli et al., 2012, 2014; Head, 2013; Mesgari et al., 2015; Konieczny, 2016; Singer et al., 2017; Lemmerich et al., 2019). It not only offers releases for numerous languages, but is changing and growing to a degree that makes it a first reference for an open topic universe, a resource which—besides the web as a whole—reflects the breadth, depth, and dynamics of human knowledge generation to an outstanding degree: whenever a new topic comes up, little time passes before Wikipedia provides information about how it is classified and related to already established topics⁶. Its dynamics and multilingualism even make it a source for the automatic extraction of learning content (Pentzold et al., 2017; Conde et al., 2020). The educational relevance refers to Wikipedia as a whole as well as to specific domains, such as law or health (Okoli et al., 2014; Smith, 2020). Singer et al. (2017) show that 16% of English Wikipedia uses are work/school-related tending to involve long page views in thematically coherent sessions. Lemmerich et al. (2019) additionally show that these uses vary considerably across languages, especially in the case of work/school-related uses. Independent of academic reservations (cf. Mesgari et al., 2015; Konieczny, 2016), Wikipedia establishes itself as a source for reading by students (Okoli et al., 2014), as an additional learning resource (Konieczny, 2016), partly with advantages over textbooks (Scaffidi et al., 2017), a resource to which students themselves contribute (Okoli et al., 2012). **Figure 4** illustrates the relevance of Wikipedia for online academic reading: it shows the Zipfian distribution of online sources frequented by students in the CORA study (Nagel et al., accepted). In this context, Wikipedia appears as an outstanding reference. Even more so: taking into account the “mutual beneficial relationship” of Google and Wikipedia (McMahon et al., 2017) (according to which traffic on Wikipedia originates mainly from Google), both resources largely dominate this usage scenario.

Given the exceptional position of Wikipedia, the research community has investigated numerous biases regarding its content and use. Denning et al. (2005) speak of accuracy, motive, expertise, volatility, coverage, and source as risks of Wikipedia, which are potential reference points for biases. In terms of coverage, this concerns, for example, the “self-focus bias” (Hecht and Gergle, 2009) or the tendency that topic selections reflect author preferences (cf. Holloway et al., 2007; Halavais and Lackaff, 2008). In this context, Hecht and Gergle (2010a) show for 25 language editions that they differ enormously in the coverage of topics, that this diversity is not explained by their size and that English Wikipedia does not cover its sister editions. A similar approach is taken by Warncke-Wang et al. (2012), who investigate a variant of Tobler’s first law (Tobler, 1970) according to which geographically nearby Wikipedias are more similar

⁶Given its relevance, it is astonishing how rarely Wikipedia is mentioned as a source for reading in Braasch et al. (2018a).



in terms of *Inter-Language Links* (ILL). What distinguishes these studies from ours is that, with few exceptions, conceptual alignments of Wikipedias concern only paired articles identified by ILLs, so that hypertext structure, which is crucial for online learning in terms of the DM, is ignored. Furthermore, similarities of articles are quantified by degree statistics, so that one can hardly speak of a content-based comparison. Comparable to the latter studies, however, we also refer to aligned articles to map shared topics (cf. Bao et al., 2012), but do so via Wikidata (RRID:SCR_018492) to identify commonly referred concepts and retain their network structure as manifested by Wikipedia's article graph.

Biased topic coverage is related to what Massa and Scrinzi (2012) call the linguistic point of view (contradicting Wikipedia's NPOV), which predicts that different (e.g., cultural) communities tend to present the same topics differently. In line with this view, thematic biases reflect cultural differences. This perspective is further developed by Miquel-Ribé and Laniado (2016), who speak of cultural identities, according to which editors tend to write about topics related to their culture. Using geo-referenced data and keyword-related heuristics, they identify cultural identity related articles (on average 25% of articles in 40 languages) to diagnose language-specific thematic preferences (regarding 15 languages) and translation-related associations of editions that are dominated by certain topics (e.g., *geography*). Thematic preferences are modeled using 18 topics derived with the help of Wikipedia's category system. Based on this analysis,

they distinguish types of culture-related articles: language-specific articles shared by a few editors and articles appearing in many languages (in terms of ILLs). Another example is Miz et al. (2020), who examine English, French and Russian Wikipedia by exploring clusters of trending articles using topic modeling based on eight topics. In contrast to these studies, we not only consider a larger number of languages and many more topics, but especially semantically coherent article subnetworks, which are examined for their differences along intra- and intertextual dimensions.

Biases of Wikipedia were also analyzed regarding selected areas: Lorini et al. (2020) observe a variant of Tobler's Law according to which authors tend to write about geographically close events. Similarly, Samoilenko et al. (2017) describe a preference for recency. Oeberst et al. (2019) investigate a bias of groups who present their views more positively (cf. Álvarez et al., 2020). A related example regarding biographical articles that includes linguistic analyses is given by Callahan and Herring (2011). Wagner et al. (2016) also present a multidimensional content analysis, now of a gender bias. Given the importance of Wikipedia as a knowledge repository and taking into account its various biases, the question of its influence on knowledge formation on the part of readers comes up (Oeberst et al., 2018).

The research considered so far shares the observation of a biased topic coverage, which relativizes Wikipedia's domain independence (Jiang et al., 2017), since certain topics (Kittur et al., 2009) or views dominate, be it due to cultural preferences

(Massa and Scrinzi, 2012; Laufer et al., 2015; Miquel-Ribé and Laniado, 2016), **language differences** (Hecht and Gergle, 2010a; Massa and Scrinzi, 2012; Warncke-Wang et al., 2012; Samoilenko et al., 2016), **geographical factors** (Hecht and Gergle, 2010b; Karimi et al., 2015; Laufer et al., 2015; Samoilenko et al., 2016; Lorini et al., 2020), or the fact that group membership influences POV (Oeberst et al., 2019). However, though these observations should be based on content-related analyses of large amounts of data, they often concern rather non-content related features (e.g., degree statistics) taking into account a maximum of 20 topics, so that topic resolution is kept low while hypertext structure is underrepresented. On the other hand, concentrating on selected areas allows accurate linguistic analyses to be carried out, but these are difficult to automate and thus difficult to apply across languages. What is needed, therefore, is a procedure that allows for more precise thematic and article network-related analyses and which can be automatically applied to many languages. Exactly such a procedure is presented here: it uses Wikidata to identify subjects of articles of whatever subject areas, and about 100 topic categories to model their diversity. This will allow us to investigate biased topic coverage intra- and intertextually.

4. RATIONALE AND METHOD

To investigate the topic coverage of Wikipedia in educationally relevant areas, we investigate how the descriptions of the same entities or knowledge objects (from the fields of economics, physics, chemistry, biology, etc.) are distributed across its language editions. That is, we investigate how Wikipedia presents itself to its readers as part of the IL in the area of education-related reading. We test the hypothesis that the extent to which one obtains information on a given topic depends on the language in which Wikipedia is consulted. Given the skewed size distribution of Wikipedia's releases for different languages, this may sound obvious at first. But we will control the size factor and examine this hypothesis for individual subject areas and topics. Since Wikipedia is a highly frequented part of the IL, this would indicate a *language-related bias*, that is, a sort of *linguistic relativity*. Thus, our article is ultimately concerned with the question of whether Wikipedia exhibits this kind of relativity or not. We test this hypothesis on the example of 35 Wikipedias. To this end, we focus on three research questions:

Q1 How do languages resemble each other in terms of the knowledge networks that manifest themselves in the associated Wikipedias?

Obviously, high dissimilarities in the latter sense mean that students who consult the associated Wikipedias are informed very differently about the same field of knowledge. Differences in knowledge between learners of different languages may then be consolidated or even expanded as a result of such a bias. An example of such a scenario is shown in **Figure 5**, which contrasts networks of articles about paintings from German and Dutch Wikipedia. The extracted networks are obviously very different; they show very different parts of the information landscape, although on the same subject area.

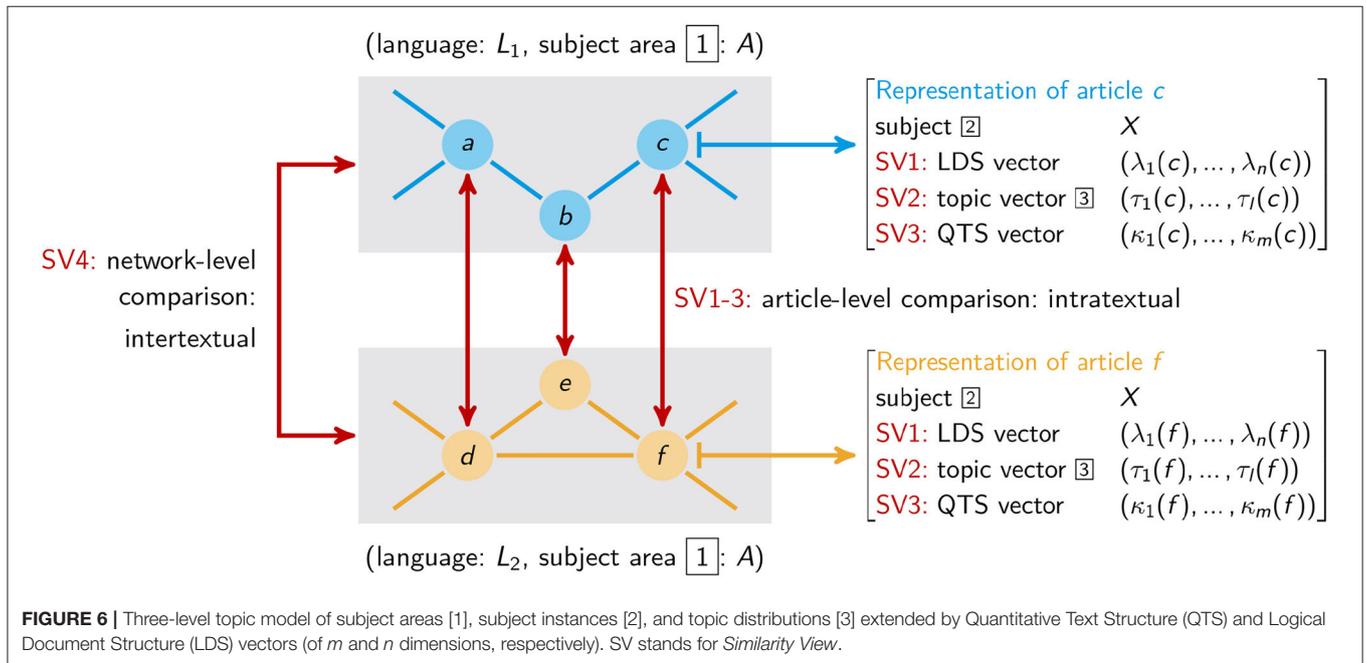
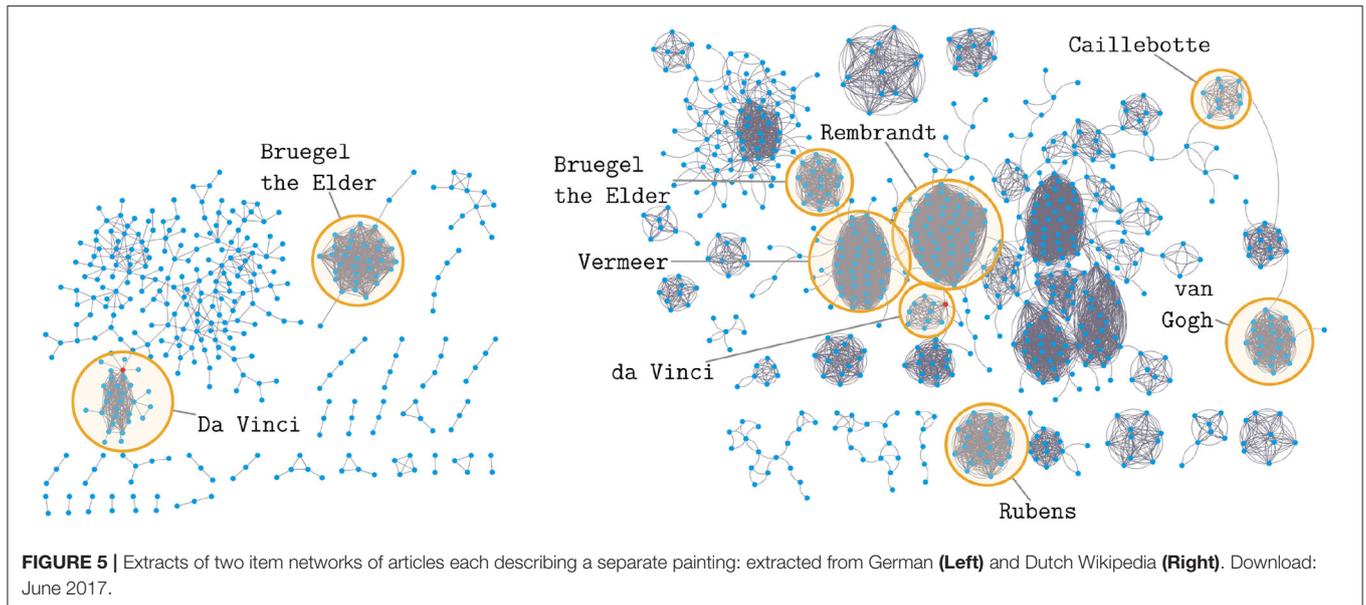
Q2 How do the similarity ratings after Q1 differ depending on the underlying knowledge domain?

The various fields of knowledge and scientific disciplines that contribute to their development have been developed in different ways. Domain-specific learning can therefore benefit considerably from discipline-specific strategies of knowledge acquisition and processing (List and Alexander, 2019). If this is true, then we might expect a shadowing of these differences in Wikipedia: it is then likely that different fields of knowledge manifest themselves differently in Wikipedia, while the similarities of languages as manifested in Wikipedia are strongly conditioned by the reference to these fields. By answering Q2, we inform educational research about the manifestation density of certain knowledge domains in certain language editions. This research can help to avoid wrong conclusions from generalizations for knowledge domains or languages: different outcomes of learners of different languages, for example, may be the result of differences in such resources and not necessarily the result of different linguistic structures of the underlying task descriptions (Mehler et al., 2018).

Q3 Regardless of such differences, is there a knowledge-related “lingua franca” which, by its Wikipedia, makes the dissemination of knowledge in other languages predictable and thus serves as a reference for knowledge dissemination?

English Wikipedia could play the role of such a reference due to its size and status as the primary source of translation between Wikipedias (cf. Warncke-Wang et al., 2012). However, several studies question this role (Hecht and Gergle, 2010a; Samoilenko et al., 2016). Thus, the question arises whether these results also apply to our combined intra- and intertextual model.

To answer Q1–3, we develop a method to extract and compare Wikipedia article graphs on the same topic, henceforth called *item networks*. Since we have to apply this method across languages, it must be both easy to implement and systematically reproducible, in such a way as to ensure that we are dealing with the same subject(s) regardless of the languages under consideration. To realize this measurement operation with the help of Wikidata we compare item networks intra- and intertextually by means of a three-level topic model as depicted in **Figure 6**: starting from any subject area (e.g., *painting*) (1st level), we identify all its instances in Wikidata, which we use to extract all the articles within Wikipedia's language editions that address these instances. In this way, we identify the 2nd level of our topic model, whose elements we refer to here as subjects or subject instances (see **Figure 6**). That is, a Wikipedia article is assigned a unique subject (e.g., *Mona Lisa*) based on its Wikidata mapping, which specifies the entity (i.e., a *Wikidata item* in the role of the *definiendum*) that the article describes. At this stage we get two topic assignments for each article: the corresponding subject area (e.g., *painting*) and its subject instance (e.g., *Mona Lisa*). Using Wikipedia's article graph, where connections between articles are given by hyperlinks, we then obtain one article network per subject area and language, with the semantic coherence of these item networks resulting from the reference to the underlying subject area common to their articles. In the third



step we characterize each node of these item networks by three vectors (see **Figure 6**). This applies in particular to the thematic perspectives under which subjects are treated. In this way we reach the 3rd level of our topic model: articles as manifestations of the same subject area are characterized according to the thematic perspectives (contributing to the *definiens*) under which they describe their subject where these perspectives are modeled as topic vectors.

Suppose the subject of a given article in a certain language is the painting *Bal du moulin de la Galette* (by Pierre-Auguste Renoir) which instantiates the subject area *painting*. Further,

assume that our algorithm detects that this article deals with the topic economy (since it reports that this painting is one of the most expensive ever auctioned). Then we get three topic-related assignments (see **Figure 6**): regarding the article's (abstract) subject area, its (concrete) subject and the thematic perspective of their descriptions. Having done this for whole article networks of different languages, we can then compare these networks at the subject level by asking whether articles about the same subject thematize it in similar ways—this concerns what we call *intratextual similarity*—and at the subject area level by asking whether the article networks are structurally similar in terms of

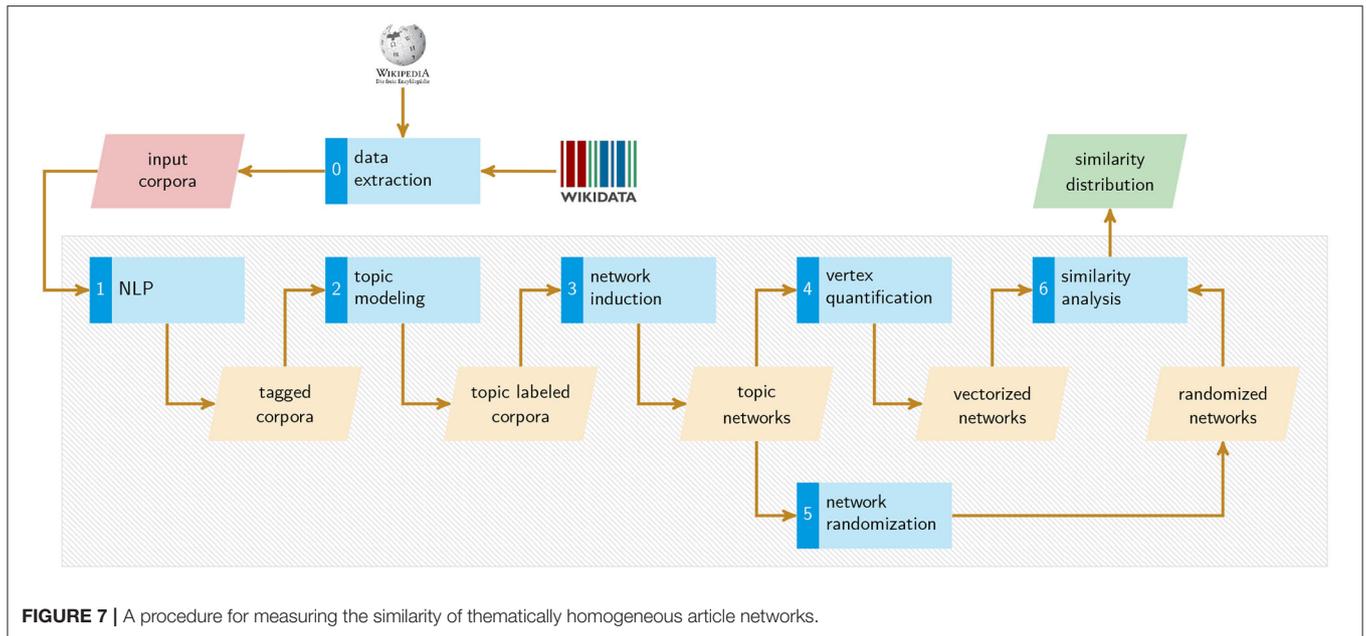


FIGURE 7 | A procedure for measuring the similarity of thematically homogeneous article networks.

TABLE 1 | Variants of using Wikidata items to identify Wikipedia articles as vertices of INs.

		B-hierarchy	
		Unexpanded	Expanded
A-hierarchy	Unexpanded ($A_{\hat{x}}$)	$U_{x \in \mathbb{O}} \{ \hat{B}_{\hat{x}} \}, U_{x \in \mathbb{O}} \{ \hat{\mathbb{B}}_{\hat{x}} \}$ 1,333,421	$U_{x \in \mathbb{O}} \{ \hat{B}_{\hat{x}} \}, U_{x \in \mathbb{O}} \{ \hat{\mathbb{B}}_{\hat{x}} \}$ 24410338
	Expanded ($A_{\check{x}}$)	$U_{x \in \mathbb{O}} \{ \check{B}_{\check{x}} \}, U_{x \in \mathbb{O}} \{ \check{\mathbb{B}}_{\check{x}} \}$ 3,190,562	$U_{x \in \mathbb{O}} \{ \hat{B}_{\hat{x}} \}, U_{x \in \mathbb{O}} \{ \hat{\mathbb{B}}_{\hat{x}} \}$ >32,949,992

what we call *intertextual similarity*. In this way, we implement the procedure to answer Q1–3, as shown in Figure 7. Under the null hypothesis, different language article networks for the same subject area are very similar, both in terms of the articles’ descriptions of the same subjects (intratextual similarity) and in terms of their hypertext structure (intertextual similarity). We will essentially falsify this hypothesis and answer Q1–3 accordingly. We now turn to explaining Step 0–6 of Figure 7 in detail.

4.1. Data Extraction

In order to extract article networks on the same subject area relevant to educational science, that is, in order to perform Step 0 and 3 of Figure 7, we use the classification of the fields of science and technology of OECD (2007) and its correspondents in Wikidata. The resulting networks, induced separately for each Wikipedia, are called *Item Networks* (IN) to emphasize the way they are induced by means of Wikidata items. In this way we guarantee three things: (i) the reference to subject areas, such as those addressed by the PISA studies, (ii) a thematic breadth of the selected topics, and (iii) their transferability between Wikipedias. INs are generated as follows: for each of the OECD categories that

can be assigned to Wikidata, we consult its *studies*-statements to determine all Wikidata classes that are related to the respective OECD category in this way. This is necessary because we need to move from OECD categories (e.g., *art*) (which we refer to for ensuring thematic diversity of general topics) to subject areas (e.g., *painting*) as likely destinations of searches in the context of the former: while the OECD categories and their Wikidata subclasses induce *subclass of*-hierarchies of fields of science and technology (henceforth called *A-hierarchies*), the latter induce *subclass of*-hierarchies of subject areas (cf. Figure 6) that are studied in these fields (called *B-hierarchies*). At this point we need an extension to ensure a higher coverage rate of OECD categories. The reason for this is that some of them do not have *studies*-statements in Wikidata, so that they would fall out of the selection process to the detriment of the targeted thematic breadth. Therefore, we additionally examine the descendants of OECD categories in A-hierarchies to determine additional classes from B-hierarchies by means of these descendants’ *studies* statements. This leads to a number of alternatives for sampling INs (see Table 1), of which only a subset is feasible, as is now formally explained: let $\mathcal{D} = (V_{\mathcal{D}}, A_{\mathcal{D}}, \lambda)$ be a representation of Wikidata as a directed graph with the set $V_{\mathcal{D}}$ of vertices (so-called Wikidata items), the set $A_{\mathcal{D}}$ of arcs (links) between these items and the arc labeling function λ . Further, let $\mathbb{O} = \{x, y, \dots\}$ be the set of OECD categories. For a given OECD category $x \in \mathbb{O}$, we generate the set $A_{\hat{x}}$ of all items belonging to the A-hierarchy dominated by $x \in V_{\mathcal{D}}$:

$$A_{\hat{x}} = \{v \in V_{\mathcal{D}} \mid 0 < ged_{sco}(x, v) < |V_{\mathcal{D}}|\} \cup \{x\} \quad (1)$$

where $ged_{sco}(x, v)$ is the length of the shortest directed path from x to v in \mathcal{D} crossing only *subclass of*-links. Alternatively, we dispense with this expansion and get the set $A_{\check{x}} = \{x\}$. In the case of $A_{\check{x}}$, we then explore *studies*-links of x and of its subclasses,

while in the case of $A_{\hat{x}}$, only those of x are explored. The resulting sets of B -level items that are “studied” in this sense are denoted by

$$B_{\hat{x}} = \{w \in V_{\mathcal{D}} \mid \exists v \in A_{\hat{x}}: (v, w) \in A_{\mathcal{D}} \wedge \lambda((v, w)) = studies\} \tag{2}$$

$$B_{\check{x}} = \{w \in V_{\mathcal{D}} \mid (x, w) \in A_{\mathcal{D}} \wedge \lambda((x, w)) = studies\} \tag{3}$$

Obviously, $A_{\check{x}} \subseteq A_{\hat{x}}$ and $B_{\check{x}} \subseteq B_{\hat{x}}$. Now we have two alternatives again: either, we recursively explore *subclass of*-links to get more subject area-related Wikidata items for comparing Wikipedias (this leads to $\hat{B}_{\hat{x}}$ and $\hat{B}_{\check{x}}$) or not (generating $\check{B}_{\hat{x}} = B_{\hat{x}}$ and $\check{B}_{\check{x}} = B_{\check{x}}$). A test shows that recursive expansions on the side of B -hierarchies (generating $\hat{B}_{\hat{x}}, \hat{B}_{\check{x}}$) leads to overly large subject area representations that induce computationally hardly processable article networks. Take the example of the OECD category *Mathematics*: if we expand this category on the side of B -hierarchies, we get candidate items like *set*, which are dominated by *mathematical concept* as the target of a *study*-statement starting from *Mathematics*. But *set* has many instances in Wikidata, many of which are not mathematical concepts. Such examples, which occur frequently in the variants $\hat{B}_{\hat{x}}$ and $\hat{B}_{\check{x}}$, realize unwanted changes of subject area, so that B -sided expansions are mentioned here only as a theoretical alternative. To get processable networks, we alternatively represent the elements of item sets $C_x \in \{\hat{B}_{\hat{x}}, \hat{B}_{\check{x}}, \check{B}_{\hat{x}}, \check{B}_{\check{x}}\}$ as singletons, each containing a single Wikidata item that is finally used to extract Wikipedia article graphs:

$$\mathbb{B}_x = \{\{y\} \mid y \in C_x\} \tag{4}$$

Take the example of $\check{\mathbb{B}}_{\hat{x}}$, that is, the set of singletons, each containing a topic studied either directly or indirectly under OECD category $x - \cup_{x \in \mathbb{O}} \check{\mathbb{B}}_{\hat{x}}$ currently contains 172 labels of corresponding subject areas (see **Table 1**). Obviously, $\check{\mathbb{B}}_{\hat{x}} \subseteq \hat{\mathbb{B}}_{\hat{x}}, \hat{\mathbb{B}}_{\check{x}} \subseteq \hat{\mathbb{B}}_{\hat{x}}, \hat{\mathbb{B}}_{\check{x}} \subseteq \hat{\mathbb{B}}_{\check{x}}$. Having identified all Wikidata *item sets* for each OECD category x , we get an expression for the set of candidate subject areas:

$$\mathbb{B} = \bigcup_{x \in \mathbb{O}} \{\hat{B}_{\hat{x}}, \hat{B}_{\check{x}}, \check{B}_{\hat{x}}, \check{B}_{\check{x}}\} \cup \hat{\mathbb{B}}_{\hat{x}} \cup \hat{\mathbb{B}}_{\check{x}} \cup \check{\mathbb{B}}_{\hat{x}} \cup \check{\mathbb{B}}_{\check{x}} \tag{5}$$

An element $B_x \in \mathbb{B}$ is a set of Wikidata items that are (in-)directly accessible from OECD category x via *studies*-links. Item sets like B_x serve to identify Wikipedia articles on the same subject across languages. This is achieved as follows: each item in B_x allows for identifying a corresponding set of instances (e.g., the painting *Bal du moulin de la Galette* by Pierre-Auguste Renoir) by exploring *instance of*-links in Wikidata. These instances are linked from Wikipedia articles and correspond to subjects in **Figure 6** as instances of subject area B_x , which in turn is derived from OECD category x to ensure that we address educationally relevant topics. At this point, an important difference to section 3 becomes clear: the approaches mentioned there operate on small, closed lists of abstract topic categories (similar to the OECD categories used here), whereas we use OECD categories only to address the level of subject areas and their subject instances (cf. **Figure 6**). The

reference to ILL does not solve the problem of these approaches, since ILLs merely define groups of articles on the same subject, to put it in our terminology. In contrast to this, we use Wikidata as a whole to extract arbitrary alignable article networks from different Wikipedias, differentiating between subject areas, subject instances and article-wise topic distributions—in this sense, our approach is both thematically stratified and open. In any event, we limit sampling to $\check{\mathbb{B}}_{\hat{x}}$ for extracting article networks. The reason is that $\check{\mathbb{B}}_{\hat{x}}$ and $\hat{\mathbb{B}}_{\check{x}}$ explore to few *studies*-links, while variants that expand B -level items induce unwanted topic changes (see above). From the 172 candidate subject areas belonging to $\check{\mathbb{B}}_{\hat{x}}$ (see **Table 1**), we select the 19 largest ones supplemented by six areas. **Figure 8** shows the boxplots of the article networks’ orders (i.e., number of their nodes) for each of these subject areas, which we derived from 35 Wikipedia language editions (see **Table 2**). For these 35 languages we trained topic models to detect the topic distributions of their articles (see section 4.2, Level 3 of **Figure 6** and Step 2 of **Figure 7**).

Now, let $\mathbb{W} = \{\mathcal{W}_1 \dots, \mathcal{W}_i\}$ be the set of all Wikipedias $\mathcal{W}_i = (V_i, A_i) \in \mathbb{W}$ each represented as a directed graph with the set of vertices (i.e., articles) V_i and arc set A_i . Further, let $a \in \mathbb{B}$ be a subject area, then we induce for each Wikipedia \mathcal{W}_i the *Item Network* (IN) $I_a(\mathcal{W}_i)$ of all articles on subjects that are directly or indirectly studied under the OECD category corresponding to a : it is defined as the subgraph of \mathcal{W}_i ’s article graph that consists only of articles which by their Wikidata links are mapped onto instances of elements of a . To generate this graph, we explore *instance of*-links from elements of a to Wikidata items addressed by Wikipedia articles. Let $\iota: \cup_{i=1}^l V_i \rightarrow V_{\mathcal{D}}$ denote the function that links articles as instances to Wikidata items in the latter sense, then we get the following expression for I_i^a (see Step 3 of **Figure 7**):

$$I_a(\mathcal{W}_i) = I_i^a = (V_i^a, A_i^a) = (\{v \in V_i \mid \iota(v) \in a\}, \{(v, w) \in A_i \mid v, w \in V_i^a\}) \tag{6}$$

This allows us to finally define the so-called *alignment set* of any pair of INs $I_i^x = (V_i^x, A_i^x), I_j^y = (V_j^y, A_j^y)$ which contains all pairs of articles of these INs that are alignable because of treating the same subject:

$$\mathcal{W}_{ij}^{xy} = \{(v, w) \mid v \in V_i^x \wedge w \in V_j^y \wedge \iota(v) = \iota(w)\} = (\mathcal{W}_{ji}^{yx})^{-1} \tag{7}$$

4.2. A Hybrid Approach to Measuring the Similarity of INs

So far, we described the object of measuring thematic dissimilarities of Wikipedia’s language editions. Now we define the aspects under which we measure these dissimilarities. To this end, we consider (1) *syntactic* (text-structural), (2) *semantic* (topic-related), and (3) text statistical aspects of the similarity of INs. While these *intratextual* aspects are all vertex content-related, a fourth *Similarity View* (SV) focuses on the network structure of INs and thus on *intertextual* aspects of similarity. Our starting point is the observation that each Wikidata item can be addressed by several Wikipedia articles, and vice versa,

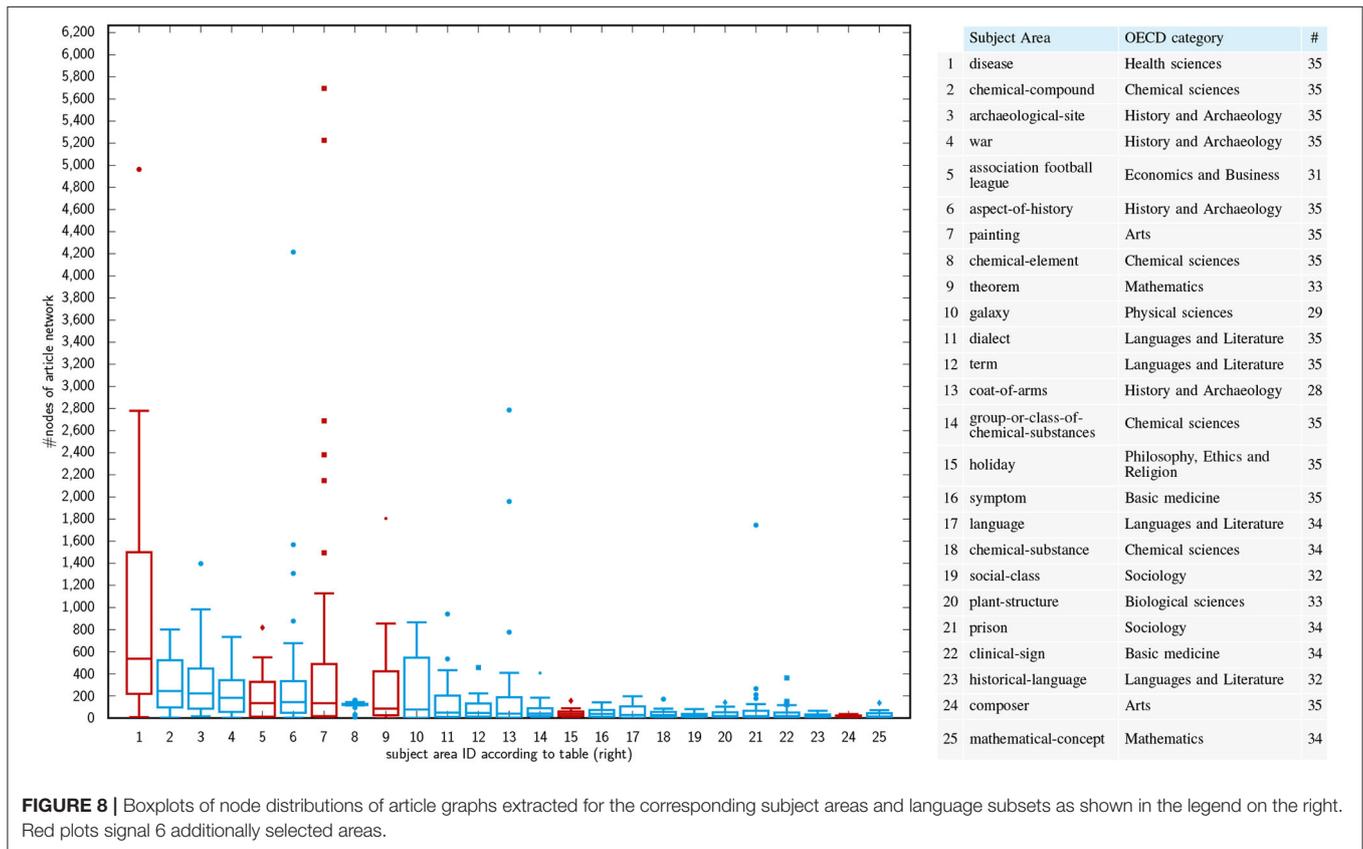


FIGURE 8 | Boxplots of node distributions of article graphs extracted for the corresponding subject areas and language subsets as shown in the legend on the right. Red plots signal 6 additionally selected areas.

the same article can describe several items, even if it is injectively mapped to Wikidata. The calculation of similarities of INs then requires the mapping of network similarities on the level of node content (intratextual) and network structure (intertextual). For this purpose, we vectorize the textual nodes of INs according to the similarity views SV1–4. That is, given a SV s , we assume that each node $v \in V_i^x$ of each IN $I_i^x = (V_i^x, A_i^x)$ is mapped onto a corresponding vector $\vec{v} \in \mathbb{R}^{k_s}$. This is expressed by the function

$$h_s: V_i^x \rightarrow \mathbb{R}^{k_s} \tag{8}$$

where k_s is the dimensionality associated with SV s . For INs $I_i^x = (V_i^x, A_i^x), I_j^y = (V_j^y, A_j^y)$ we calculate their similarity node- and network-wise using the alignment set \mathcal{W}_{ij}^{xy} (with a few exceptions, elements of vectors \vec{v} are non-negative): $\forall (v, w) \in V_i^x \times V_j^y$:

$$\sigma_{s_1}(I_i^x, I_j^y) = \begin{cases} \frac{\sum_{(v,w) \in \mathcal{W}_{ij}^{xy}} \sum_{n=1}^{k_s} h_s(v)[n]h_s(w)[n]}{\sqrt{\sum_{v \in V_i^x} \sum_{n=1}^{k_s} h_s(v)[n]^2} \sqrt{\sum_{w \in V_j^y} \sum_{n=1}^{k_s} h_s(w)[n]^2}} & V_i^x \neq \emptyset \neq V_j^y \in [-1, 1] \\ 0 & \text{else} \end{cases} \tag{9}$$

Computing Formula 9 is part of performing Step 6 of Figure 7. It requires defining vectorization functions h_s for SV $s = 1..3$ (cf. Step 4 of Figure 7) (SV4 is introduced below):

1. *SV1: Logical Document Structure (LDS)*: according to this SV, pairs of aligned articles are the more similar the more their LDS (Power et al., 2003) resemble each other, and the more such pairs of aligned articles of two INs resemble each other in this sense, the more similar these INs are. For measuring similarities of the LDS of aligned articles we vectorize articles along $k_1 = 11$ dimensions (cf. Callahan and Herring, 2011): *number of characters, number of sections, breadth of the table of content tree, depth of this tree, number of outgoing links to pages inside Wikipedia, number of outgoing links to pages outside Wikipedia, number of pictures, number of tables, number of links to the Integrated Authority File and related norm data, number of references, and number of categories*. In this way, we identify text pairs that address the same Wikidata item by similar document structures, e.g., with regard to the length of the presentation, the use of images, tables, or hyperlinks. Due to its orientation on surface structural features, this method can easily be calculated across languages.
2. *SV2: Thematic structure*: according to this SV, pairs of aligned articles are the more similar, the more similar the distributions of topics they address when describing their subjects (see Figure 6), and the more such pairs of aligned articles of two INs resemble each other in this sense, the more similar they are. SV2 is implemented with text2ddc (Uslu et al., 2019) (Step 2 in Figure 7), a neural network based on fastText (Joulin et al., 2017) which uses TextImager (Hemati et al., 2016) to preprocess texts (Step 1 in Figure 7). That is, topics are

TABLE 2 | Wikipedia language editions, which were analyzed thematically: "Topics" is the number of DDC-based topic classes trained for the corresponding language, "Train" is the number of training examples and "F-score" the harmonic mean of precision and recall of the corresponding test.

	Code	Language	Topics	Train	F-score
1	ar	Arabic	96	14,991	0.80
2	bs	Bosnian	87	5,599	0.83
3	ceb	Cebuano	68	2,069	0.87
4	ckb	Sorani	82	3,090	0.75
5	da	Danish	94	12,761	0.78
6	de	German	98	21,933	0.87
7	el	Greek	92	8,395	0.76
8	en	English	97	19,772	0.85
9	es	Spanish	95	16,951	0.85
10	fr	French	94	17,313	0.86
11	he	Hebrew	94	12,262	0.83
12	hi	Hindi	85	7,271	0.74
13	hu	Hungarian	91	10,854	0.85
14	id	Indonesian	93	11,265	0.81
15	it	Italian	94	15,894	0.85
16	ja	Japanese	93	16,390	0.84
17	ko	Korean	92	13,557	0.76
18	lv	Latvian	89	7,572	0.83
19	mk	Macedonian	88	5,750	0.76
20	ml	Malayalam	84	5,465	0.85
21	mr	Marathi	83	3,061	0.85
22	nl	Dutch	97	15,507	0.85
23	pl	Polish	96	16,356	0.84
24	pt	Portuguese	93	15,663	0.84
25	ro	Romanian	93	10,690	0.77
26	ru	Russian	97	17,302	0.85
27	sh	Serbo-cro.	94	9,536	0.82
28	si	Sinhala	81	2,521	0.83
29	simple	Simple English	93	10,882	0.83
30	sr	Serbian	91	10,607	0.82
31	sv	Swedish	95	16,458	0.80
32	te	Telugu	80	3,916	0.84
33	vi	Vietnamese	88	10,279	0.83
34	war	Waray	85	4,282	0.75
35	zh	Chinese	92	15,595	0.84

identified as elements of the 2nd level of the *Dewey Decimal Classification* (DDC), a topic model widely used in the field of libraries. To this end, each article is mapped onto a 98 dimensional topic vector, whose membership values encode the degree to which text2ddc estimates that the article deals with the topic corresponding to the respective dimension (cf. Mehler et al., 2019; Uslu et al., 2019). In this way, we identify text pairs that tend to describe the same subject of the same area in different languages under the perspective of similar topic distributions.

Since we do not have *Part of Speech* (POS) taggers for all target languages, we basically pursued a word-form-related

approach to train text2ddc: we generated training and test corpora by retrieving information from Wikidata, Wikipedia and the Integrated Authority File of the German National Library. Since Wikipedia is offered for a variety of languages, such corpora can be created for many languages. We optimized text2ddc with regard to selected linguistic features as a result of various pre-processing steps, such as lemmatization and disambiguation. In the last column of **Table 2** we show the *F*-values obtained for the corresponding tests. The highest *F*-value (87%) is achieved for German (where we also explored POS data), the lowest for Hindi (74%). Although this is a wider range of values, it is currently the only way to compare the content of texts in different languages in terms of the way their subjects are treated. And although text2ddc was trained for a larger number of languages⁷, we concentrated on those for which it achieves an *F*-value of at least 74%.

3. *SV3: Quantitative Text Structure (QTS)*: according to this SV, pairs of aligned articles are the more similar the more their QTSs resemble each other, and again, the more such pairs of aligned articles of two INs resemble each other in this sense, the more similar they are. To get comparable vector representations of the QTS of articles, we use a subset of 17 dimensions of quantitative linguistics as evaluated by Konca et al. (2020): *adjusted modulus, alpha, Gini coefficient, h-point, entropy, hapax legomena percentage, curve length, lambda, vocabulary richness, repeat rate, relative repeat rate, thematic concentration, secondary thematic concentration, type-token-ratio, unique trigrams, average sentence length, and number of difficult words*. Two text characteristics from Konca et al. (2020), which require POS tagging, are excluded from SV3, since POS-tagging tools were not available for all languages considered here. We additionally compute autocorrelations (lag 1–10) of consecutive sentence-related association probabilities with BERT (*Bidirectional Encoder Representations from Transformers*) (Devlin et al., 2019). BERT is a language model based on a bidirectional transformer (Vaswani et al., 2017), which uses encoder-decoder attention (Bahdanau et al., 2015) and self-attention (Cheng et al., 2016) mechanisms to calculate token representations conditioned on both left and right context. In total, we compute $k_3 = 27$ characteristics to vectorize the QTS of any article of the 35 languages considered here.

While SV1-3 are article content-related, graph similarity measures consider intertextual structures (SV4). We can calculate SV4 independently of SV1-3, or so that the contributions of SV1-3 are embedded at the node level of SV4. To elaborate the first variant, we consider four measures (thereby implementing Step 6 of **Figure 7**):

1. As a baseline, we compute the *Graph Edit distance-based Similarity* (GES) using the graph edit distance for labeled directed graphs (cf. Mehler et al., 2019). In this way we arrive

⁷See <https://textimager.hucompute.org/DDC/>

at a measure that calculates graph similarities as a function of overlaps of node and arc sets, respectively.

2. *Edge-Jaccard-Similarity (EJS)*: As a second baseline, we compute the arc set-based Jaccard similarity:

$$EJS(I_i^x, I_j^x) = \frac{|E_i^x \cap E_j^x|}{|E_i^x \cup E_j^x|} \in [0, 1]$$

$EJS(I_i^x, I_j^x) = 1$ if and only if I_i^x and I_j^x have identical non-empty arc sets (where the Jaccard similarity of empty sets is defined to be zero). EJS decreases with increasing symmetric difference in the arc sets.

3. *DeltaCon with Personalized PageRank*: Koutra et al. (2016) propose DeltaCon, a family of similarity measures for node-aligned graphs. It is based on the distance function

$$d(I_i^x, I_j^x) = \sqrt{\sum_{v \in V_i^x \cup V_j^x} \sum_{w \in V_i^x \cup V_j^x} \left(\sqrt{s_i(v, w)} - \sqrt{s_j(v, w)} \right)^2}$$

where the columns of $s_i, s_j \in \mathbb{R}^{|V_i^x \cup V_j^x| \times |V_i^x \cup V_j^x|}$ store *affinity* values between vertices $v, w \in V_i^x \cup V_j^x$ in the union graphs $I_i^x = (V_i^x \cup V_j^x, E_i^x)$ and $I_j^x = (V_i^x \cup V_j^x, E_j^x)$. As $d(I_i^x, I_j^x)$ is a metric, Koutra et al. (2016) propose to define the DeltaCon similarity as $sim(I_i^x, I_j^x) = \frac{1}{1+d(I_i^x, I_j^x)}$. DeltaCon can be parameterized by various node similarity measures, of which we consider a variant that is semantically meaningful in our context: Personalized PageRank (for vertices $v, w \in V(I_i^x)$) measures the probability of being in vertex w in the stable distribution of a random walk on I_i^x that has a probability of $1 - \alpha$ to be reset to v in each step. This models the probability of landing at an article w when starting at article v , randomly following links and going back to article v when one is “lost.” It thus models, considering all articles in turn as starting articles, a crowd of Wikipedia users following hyperlinks to navigate the information landscape. We proceed similarly to compute s_j and set $\alpha = 0.85$.

4. As both EJS and DeltaCon are highly sensitive to non-overlapping node sets, we restrict the input INs $I_i^x = (V_i^x, A_i^x)$ and $I_j^x = (V_j^x, A_j^x)$ for these two similarity measures to the subgraphs induced by the intersection of their vertex sets, that is, $I_i^x = I_i^x[V_i^x \cap V_j^x]$ and $I_j^x = I_j^x[V_i^x \cap V_j^x]$. This approach yields high similarities if the subgraphs on the aligned nodes are similar and disregards dissimilarity induced by unaligned nodes in both graphs. Hence, the latter two measures can be seen as “optimistic” and may give high similarities even if the overlap of two vertex sets is rather small. To provide an alternative to this view, we compute the arc set-based cosine graph similarity of Mehler et al. (2019). For this purpose we start from the following “axioms” concerning the similarity of INs:

A1 The higher the number of shared subject instances, the more similar the INs.

- A2 The lower the proportion of shared subject instances in the total number of instances of the underlying subject area, the less similar the INs.
- A3 The higher the number of shared paths, the more similar the INs.

A1 also concerns isolated vertices: two INs can be identical even if all their vertices are isolated. A2 damps this depending on the total number of shareable subjects: the smaller the orders of the INs in relation to this number, the less similar they are. A3 prefers pairs of INs that share many edges, but only in cases where equal paths start from aligned subjects. A3 essentially states that two networks are the more similar, the more similar they look like from the perspective of the more aligned nodes. A1 and A3 are measured with the apparatus of Mehler et al. (2019); to satisfy A2, we damp the resulting cosine similarity by the quotient of shared nodes and the total number of shareable items; we can do the same regarding the size of the networks, but refrain from this “pessimistic” variant in the present paper. We will refer to this variant by *Cosine Graph Similarity (CGS)*: it rates pairs of networks as similar that link large proportions of candidate Wikidata items in a similar way.

We calculate four graph similarity measures ranging from set-based to spherical measures, where EJS and DeltaCon weight similarities more optimistically and CGS more pessimistically.

4.2.1. Assessing Observed Similarities

Any similarity found between INs on the same subject area has to be evaluated according to how far it is higher than what is randomly expected and lower than what is ideally expected (see Step 5 of **Figure 7**). For this purpose, we consider the following bounds:

1. To get a *lower bound* we compute the similarities of random counterparts of INs: we consider Erdős-Rényi (Erdős and Rényi, 1959) graphs $R(I_i^x)$ chosen uniformly at random from all graphs of the same order and size as I_i^x . This randomization concerns SV4 without node-related similarities, assuming a bijection between the nodes of $R(I_i^x)$ and I_i^x based on their Wikidata items. Randomizations are performed 100 times per IN; similarity values are averaged accordingly.
2. As an *upper bound* we consider the maximum similarity of different language INs observed over all subject areas in our corpus of INs. Since we display these values as a function of the minimum $\min(|V_i^x|, |V_j^x|)$, we get an estimate of the maximum similarity observed for this minimum among all similarities observed in our experiment. Though the theoretical maximum is always 1 for identical graphs, this maximum is unlikely to be observable in practice even under the condition of comparatively similar INs. Therefore, by our re-estimation we achieve a more realistic upper bound that is actually observable.
3. To obtain the lower bound for SV1-3, the similarity of the non-aligned, randomly chosen articles from each pair of INs were calculated (using Formula 9). The articles in each pair were drawn from two independent random permutations,

whereupon the number of pairs (of articles) was kept the same as in the aligned case. To reduce the impact of possible outliers, the results of each calculation were averaged over 100 independent runs.

5. EXPERIMENT

Applying the methods of section 4 along the procedure of **Figure 7** to the Wikipedia editions of **Table 2** to generate language-specific INs for the subject areas of **Figure 8** produces the results of **Table 3**: based on the number of INs per subject area listed in **Figure 8** (whereby languages whose INs according to Formula 6 correspond to the empty graph for the given subject area are not listed), we arrive at 103,299 graph comparisons using our 7 similarity measures, three of which are vertex (SV1-3) and four hypertext structure-oriented⁸. Since we randomize each IN 100 times to perform the same procedure for each random setting, the final number of graph comparisons performed equals 10322900.

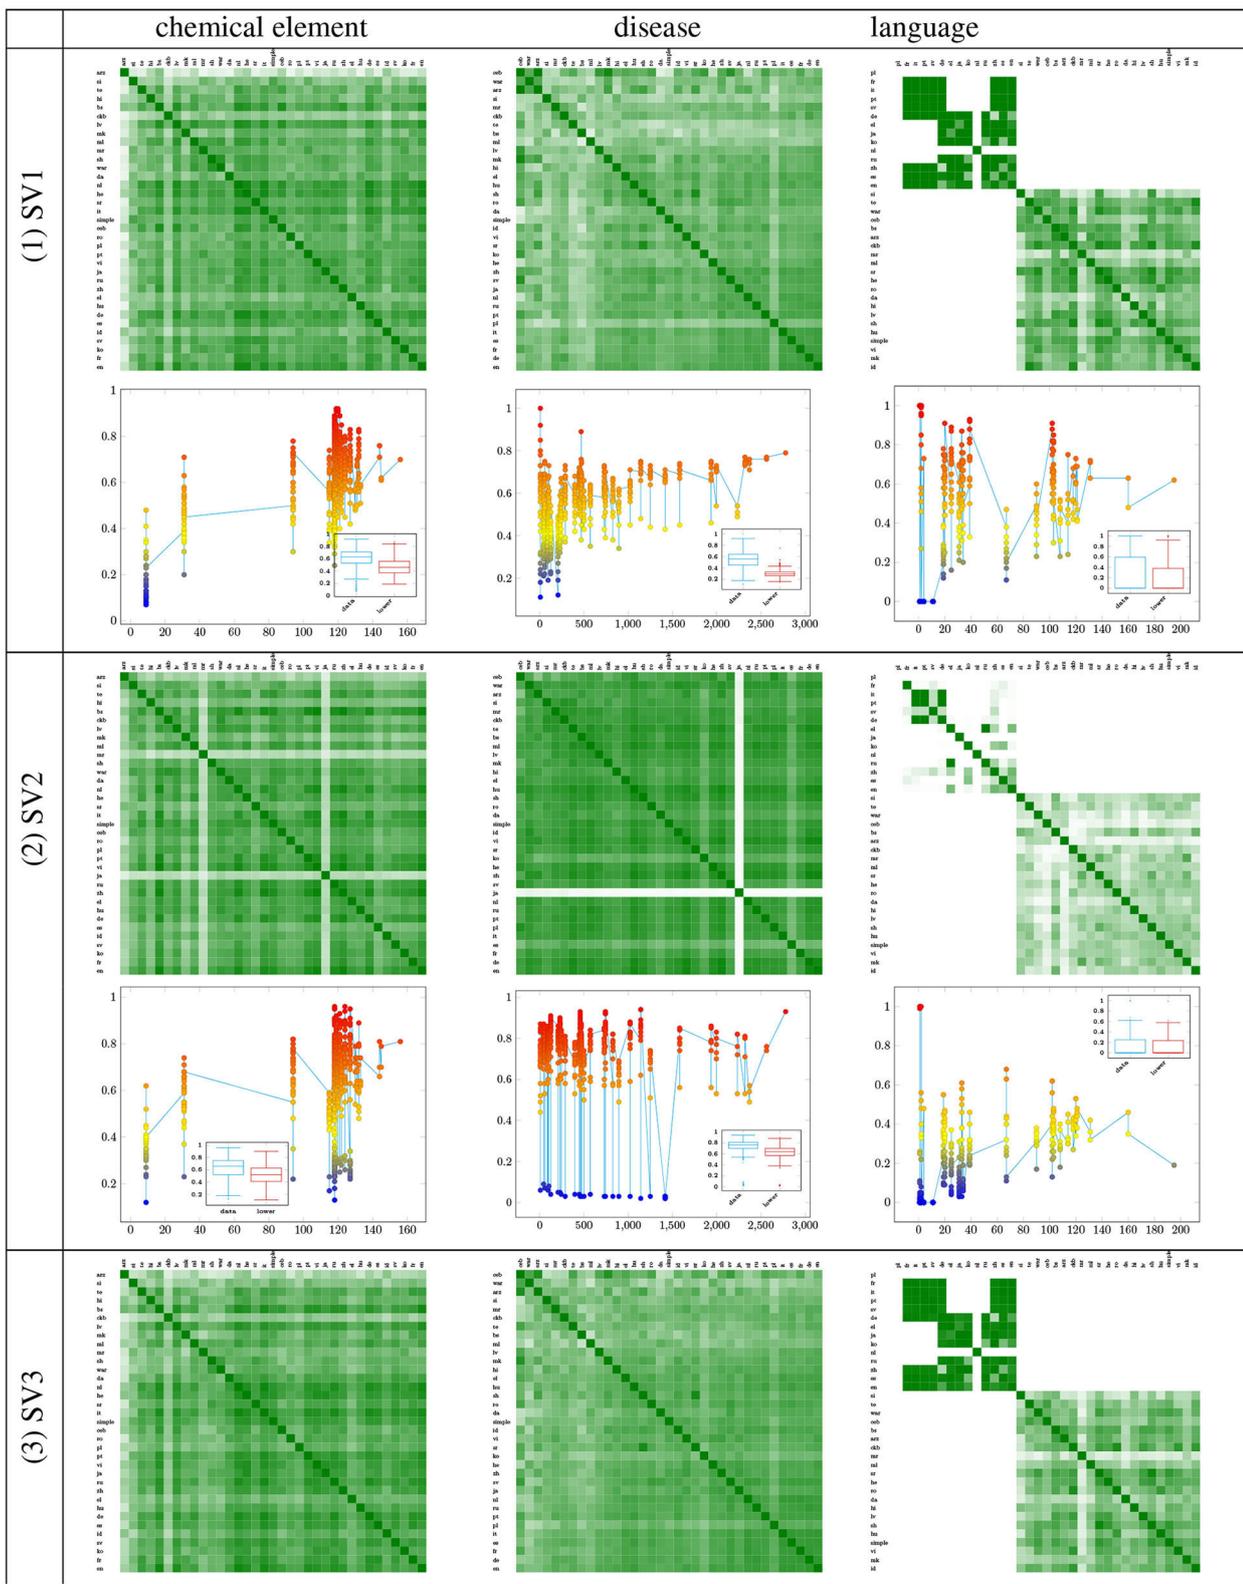
We start our experiment with three subject areas, whose analyses span the similarity spectrum observed by us: *chemical element* (OECD class *Chemical sciences*), *disease* (class *Health sciences*), and *language* (*Languages and Literature*). Let us first consider the set-based similarity measure GES. In the first column of Row 4 of **Table 3** we see the corresponding heatmap of the 35 languages' INs from **Table 2**: the greener, the more similar the INs of the respective language pair for this subject area. Apparently, we find that *chemical element* is treated very uniformly across languages in terms of hypertext structure. This subject area is the maximum of what we observed regarding this uniformity: it best approximates the ideal under the hypothesis that different language Wikipedias report uniformly on the same topic—note that the rows and columns of the heatmaps are ordered according to the orders $|V_i^x|$ of the INs I_i^x . This is confirmed by the curve displayed below the heatmap: it shows the similarity values listed in the heatmap as a function of $\min(|V_i^x|, |V_j^x|)$ (minimum of the orders of the input graphs I_i^x, I_j^x). In this way we see the influence of graph order: similarity values of comparisons with smaller graphs move to the left, those where both graphs are larger move to the right. In fact, we see for this subject area that most comparisons concern (equally) large graphs achieving high GES-values. This is contrasted by the last column of Row 4 which depicts the heatmap of subject area *language*: now, small graphs dominate the distribution with high similarity values, while pairs of significantly larger INs tend to have much lower values. This example demonstrates a specialization of a few language editions on a broad representation of this subject area, while the majority of editions tend to underrepresent it. All in all, we arrive at a zigzag curve of similarity values, which does not indicate a clear trend. In the mid of the range of these two examples in Row 4, we find the subject area *disease*: the heatmap now suggests that the language pairs are rather dissimilar from the perspective of

this subject area. This is particularly evident from the similarity curve, which shows an almost constant trend of small values with a small increase toward larger graphs. Let us now look at the same examples from the perspective of DeltaCon (Row 5); the situation remains essentially the same: *chemical element* induces a more homogeneous mass of simultaneously larger and similar graphs, *language* exhibits again a zigzag pattern, and *disease* reveals slightest similarities regardless of the INs' orders. Obviously, the latter subject area seems to contradict the ideal of homogeneous, equally informative Wikipedias very strongly. This is also confirmed by the randomizations according to section 4.2.1, which are presented as boxplots embedded in the similarity curves: blue is the boxplot of the similarity values shown in the curve itself and red is the boxplot of the similarity values of the randomizations. Interestingly, in the case of *language* randomized similarities are even higher than their empirical counterparts. This can be explained by the topology of ER graphs, which tend to have short diameters making shorter paths more probable. Anyhow, the situation is almost the same as in the case of Row 4: observed similarities seem to deviate only slightly from their randomized counterparts; similarities between these INs resemble those of corresponding random graphs. Now we look at CGS (Row 6), which evaluates graph similarities more “pessimistically” than DeltaCon. This is confirmed by the heatmaps: in the case of *chemical element*, the high similarity values change in favor of a “chessboard” view, while the binary regime disappears in the case of *language*, whose similarity progression now resembles that for *disease*. In the latter two cases, low similarities dominate, with the tendency of higher similarities only for pairs of larger graphs: based on CGS, pairs of small graphs are highly dissimilar (damping effect). Interestingly, the graph similarities as a function of the minimum of the orders of the input graphs no longer predominantly show a zigzag pattern as in the case of GES and DeltaCon: the impression of gradual transitions now prevails (2nd line of Row 6).

The latter assessments are confirmed by the set of 25 subject areas. **Table 4** shows the DeltaCon-based boxplots displayed as a curve: medians are represented by straight lines while the value ranges between the 25th and 75th percentile are colored accordingly; the blue curve represents observed similarities, the red one their random counterparts (“lower bound”) and the orange one the corresponding upper bounds (see section 4.2.1). **Table 4**, in which the subject areas are arranged according to decreasing median similarities of their INs, shows a clear trend: in almost all cases, observed similarities are below both the similarities of randomized INs and the upper bounds. That is, observed similarities are far away from ideally equally informative Wikipedias, which would cover the corresponding subject areas uniformly for all languages. Even more: as far as DeltaCon considers transitive dependencies of nodes along the same paths, it turns out that the INs' random counterparts even tend to have a greater hypertext-structural similarity—as explained above, this is partly due to their small diameters. Obviously, randomness makes networks seem more similar than if one follows existing walks in the real networks simultaneously. This even applies to the maximum values measured (upper bound). At the same time, we observe a broad spectrum of

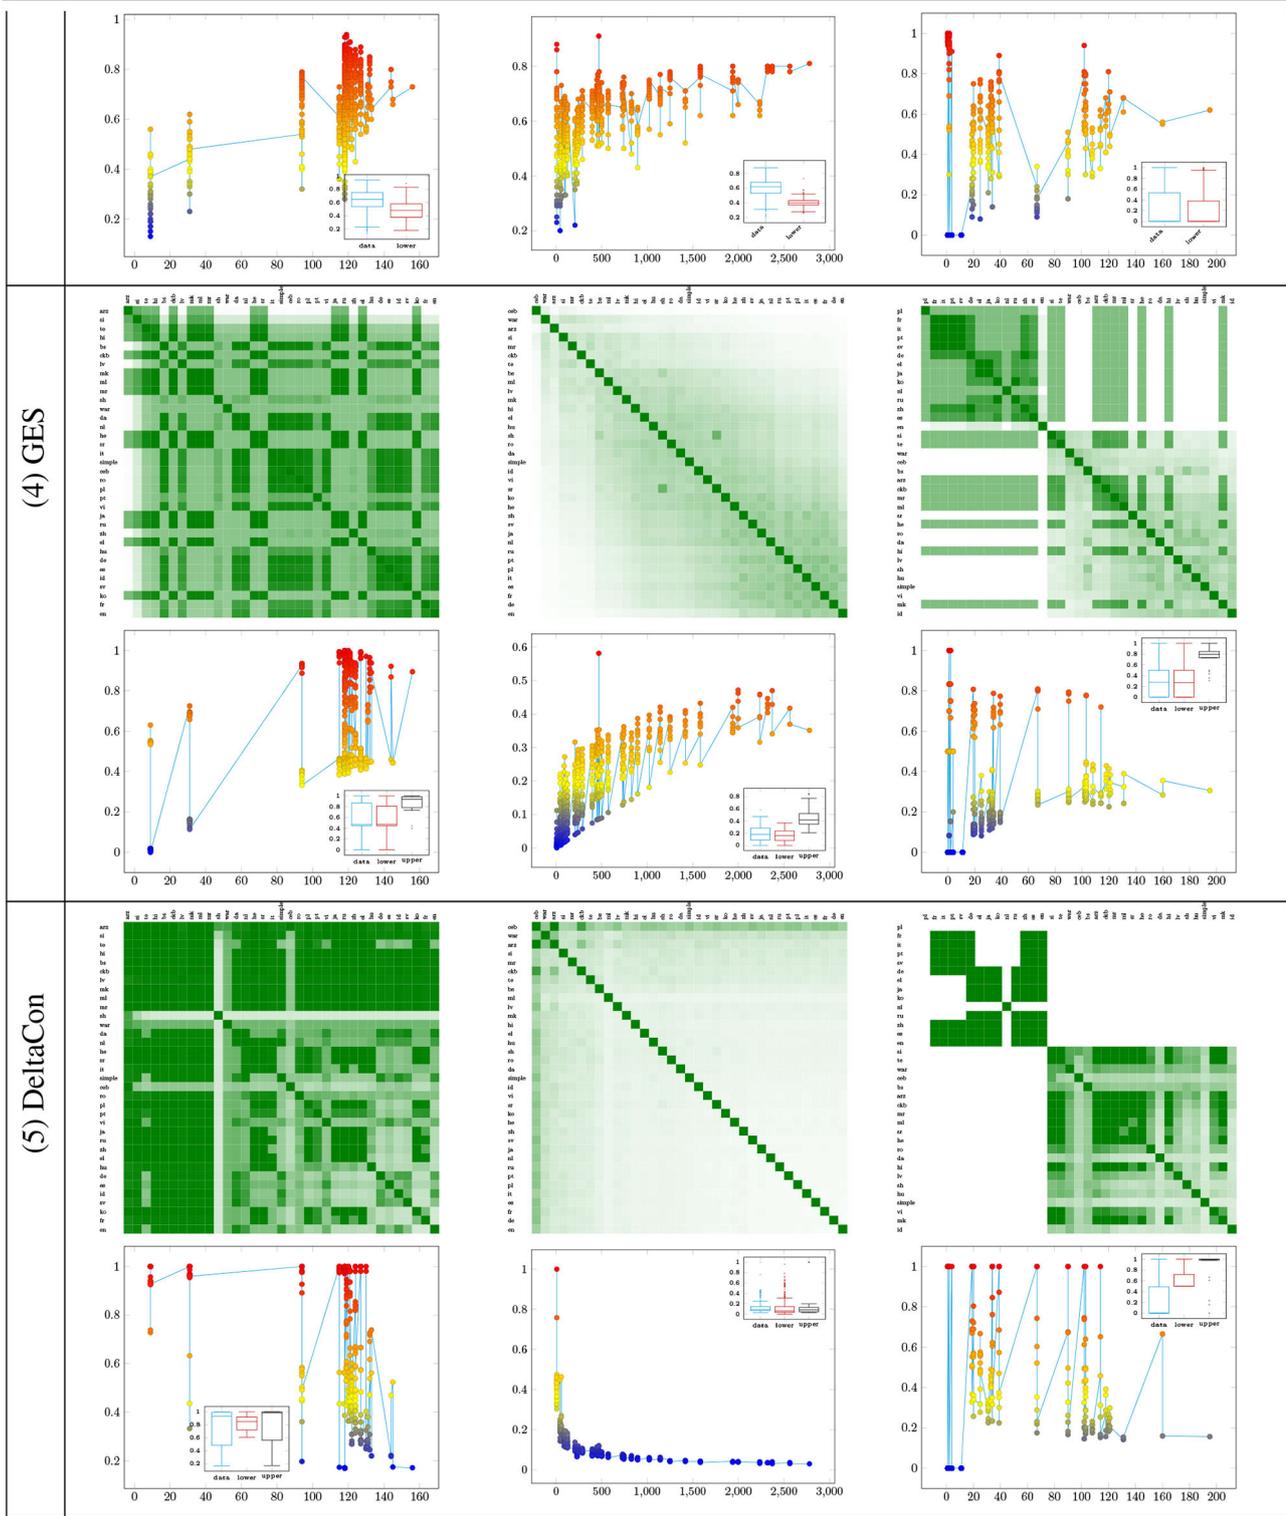
⁸If all INs were non-empty, this number would equal $\#_{\text{similarities}} \cdot \#_{\text{topics}} \cdot \#_{\text{comparisons/language}} = 7 \cdot 25 \cdot \binom{35}{2} = 110,250$.

TABLE 3 | Six measures (row) for computing the similarities of INs by example of three subject areas (column); rows and columns of the heatmaps correspond to languages (INs); curves below the heatmaps display similarity values as a function of the minimum order of the input graphs.



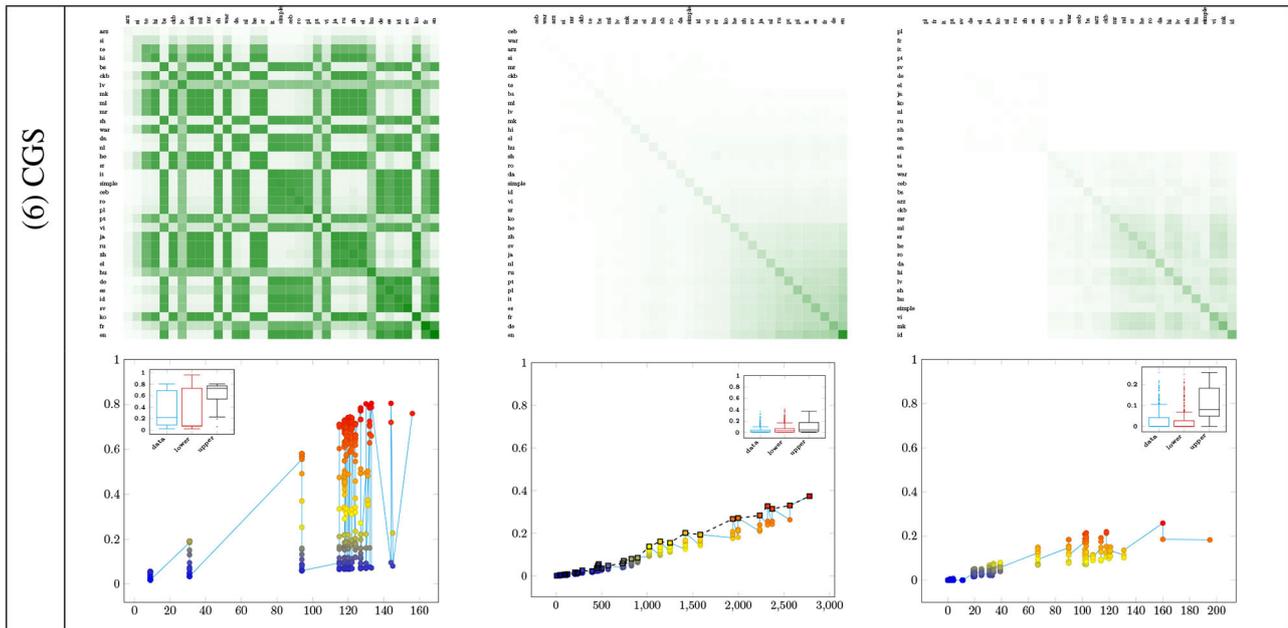
(Continued)

TABLE 3 | Continued



(Continued)

TABLE 3 | Continued



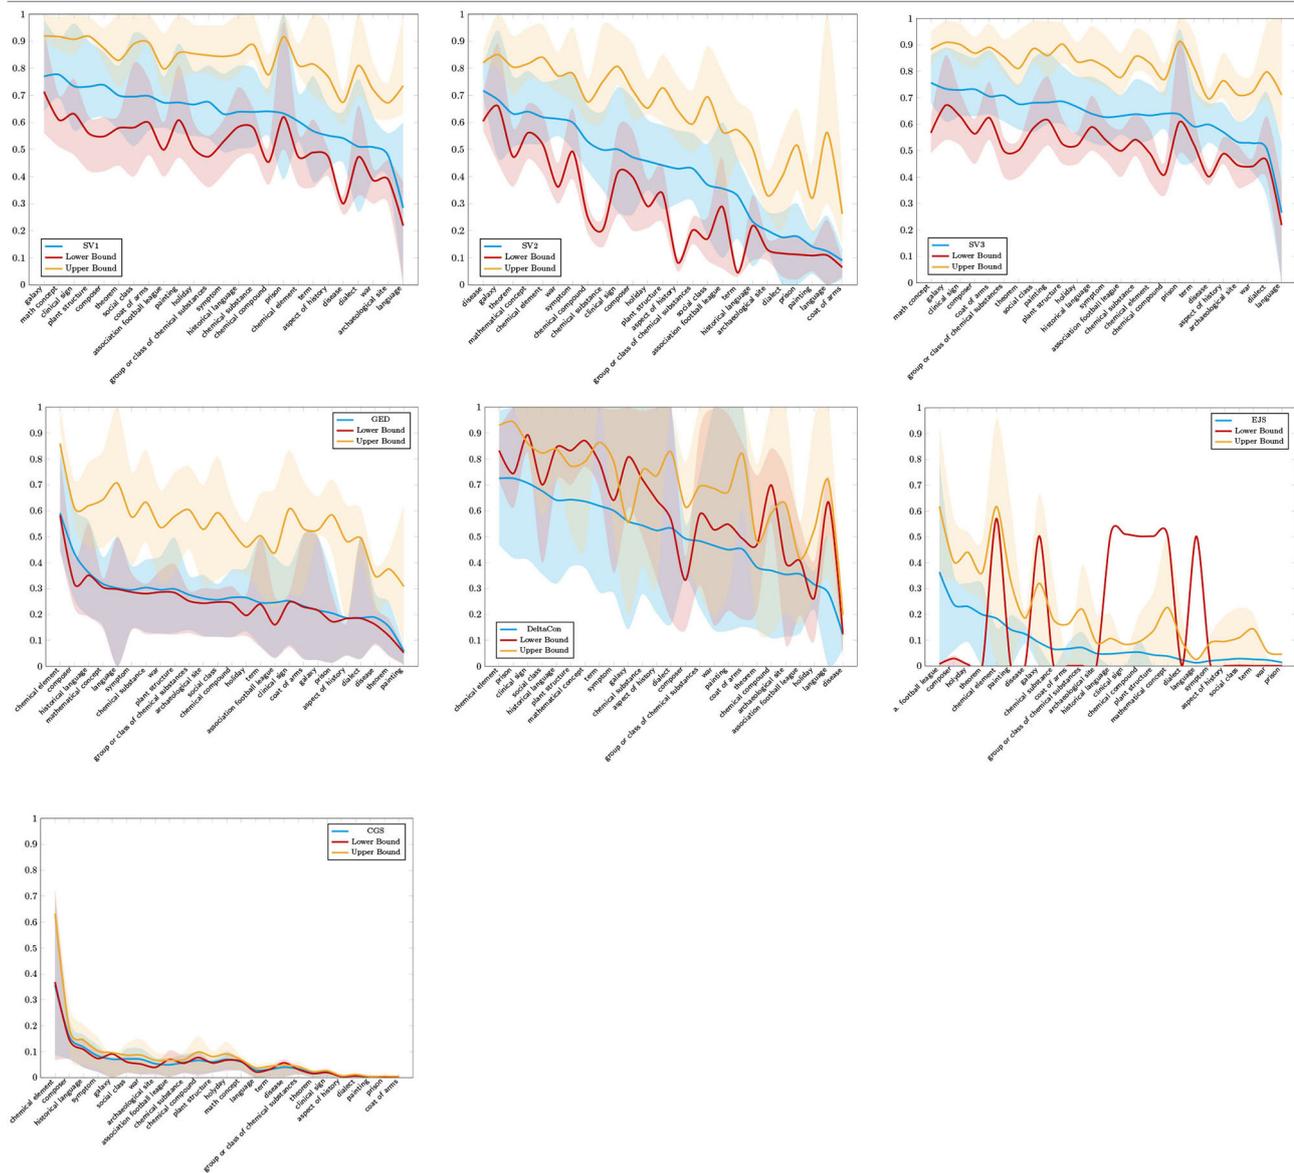
similarity values ranging from a minimum of about 20% (*disease*) to a maximum of about 75% (*chemical element*). Therefore it matters very much in which subject area one reads Wikipedia—our approach shows this in a fine-grained way for educationally relevant topics using a three-level topic model. The picture becomes even clearer when we look at the distribution of GES-based values in **Table 4**: observed similarities are now hardly distinguished from their random counterparts and far away from their upper bounds. A more chaotic picture, somewhat reminiscent of DeltaCon, results from the Jaccard-based variant EJS: observed similarities tend to be even smaller than their random counterparts, but not always. In **Table 4**, the least chaotic pattern is produced by CGS: we observe a monotonically decreasing function with small differences for minimum and maximum values, where the similarity values fall below those of the other measures.

From this lesson we learn that subject areas are rather unevenly distributed across Wikipedia's language editions, whether one measures their similarities using simpler (GES, EJS) or more complex measures (DeltaCon, CGS). But what picture do we get from looking at the content-related similarities of articles? In **Table 3** we start with SV1 (Row 1, LDS). Again, the picture is tripartite: *chemical element* marks the upper limit of observed similarities (the larger the INs, the higher their similarity), *language* the lower limit, and *disease* a middle case. *language* also shows that LDS-based similarities break down into two groups: for larger and for smaller graphs, while members of both groups show no or little similarities between each other. In any event, observed similarities tend to exceed their random counterparts (embedded blue (observed) and red (randomized) boxplots). This tripartition is basically confirmed by SV2 (thematic similarities measured by text2ddc

(Row 2 in **Table 3**). But now *disease* turns out as a subject area whose article graphs (INs) consist of thematically more homogeneous articles—more or less irrespective of the number of articles considered. In any event, thematic similarity is again concentrated more or less on large graphs in the case of *chemical element*. The third case (Row 3 in **Table 3**) concerns quantitative text structures. Here too, the general picture is confirmed: concentration on high values for large graphs in the case of *chemical element*, more evenly distributed, but high values in the case of *disease* and a bipolar picture in the case of *language*. As before, randomized counterparts are exceeded.

The rather exceptional case of SV2 (thematic similarities) is confirmed by the overall view for the 25 subject areas in **Table 3**: observed similarities along SV2 are below the values for LDS (SV1) and QTS (SV3), while the latter are more evenly distributed across the subject areas, possibly reflecting a law-like behavior as described by quantitative linguistics (Köhler et al., 2005). In any case, with few exceptions, the observed values are again within the interval spanned by their randomized variants and upper bounds. This points to similarity distributions far away from ideally homogeneously structured Wikipedia articles, which in addition would manifest almost the same topic distributions: actually, they do not. Apparently hypertextual dissimilarity is parallel to textual dissimilarity: *what is rather dissimilar in terms of intertextual structure, tends to be dissimilar in terms of intratextual structure as well*. However, we also observe the case of examples, such as *chemical element*, which has both high similarity values in terms of DeltaCon and SV2 (thematic similarity): this is, so to speak, the maximum of simultaneous inter- and intratextual similarity observed here. In any event, our study combines intra- and intertextual measurements where the former are based on three views, regarding LDS, QTS and

TABLE 4 | Distributions of similarity values of INs calculated by the similarity measures of section 4.2 and displayed per subject area. From left to right, top-down: SV1-3, GES, DeltaCon, EJS, CGS.

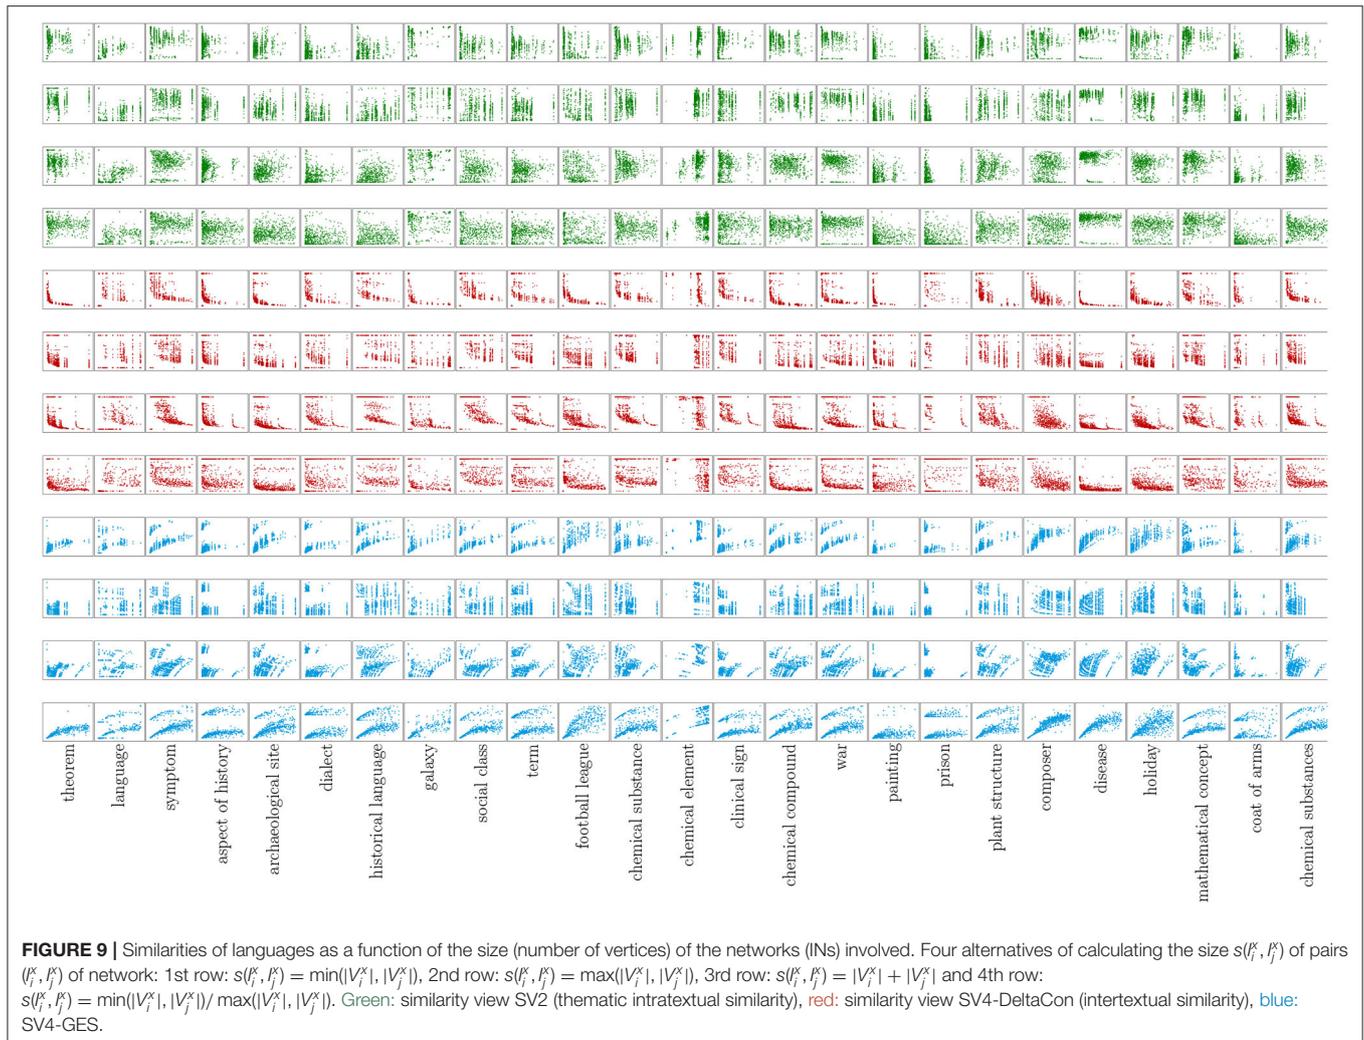


thematic text structures. In this way, we obtain a more precise, broader view of article content than has been possible with the methods of related research.

We now turn to correlation analysis and ask about the dependence of our similarity analysis on the size of the graphs involved, based on the hypothesis that the size (and thus indirectly the degree of development or activity) of a subject area explains our results. For this purpose, we calculate Spearman's rank correlation with respect to four data series, each of which is generated for two similarity views or three similarity measures (SV2, DeltaCon and GES).

The data series contrast measured similarities with the orders $|V_i^x|, |V_j^x|$ of the networks I_i^x, I_j^x involved.

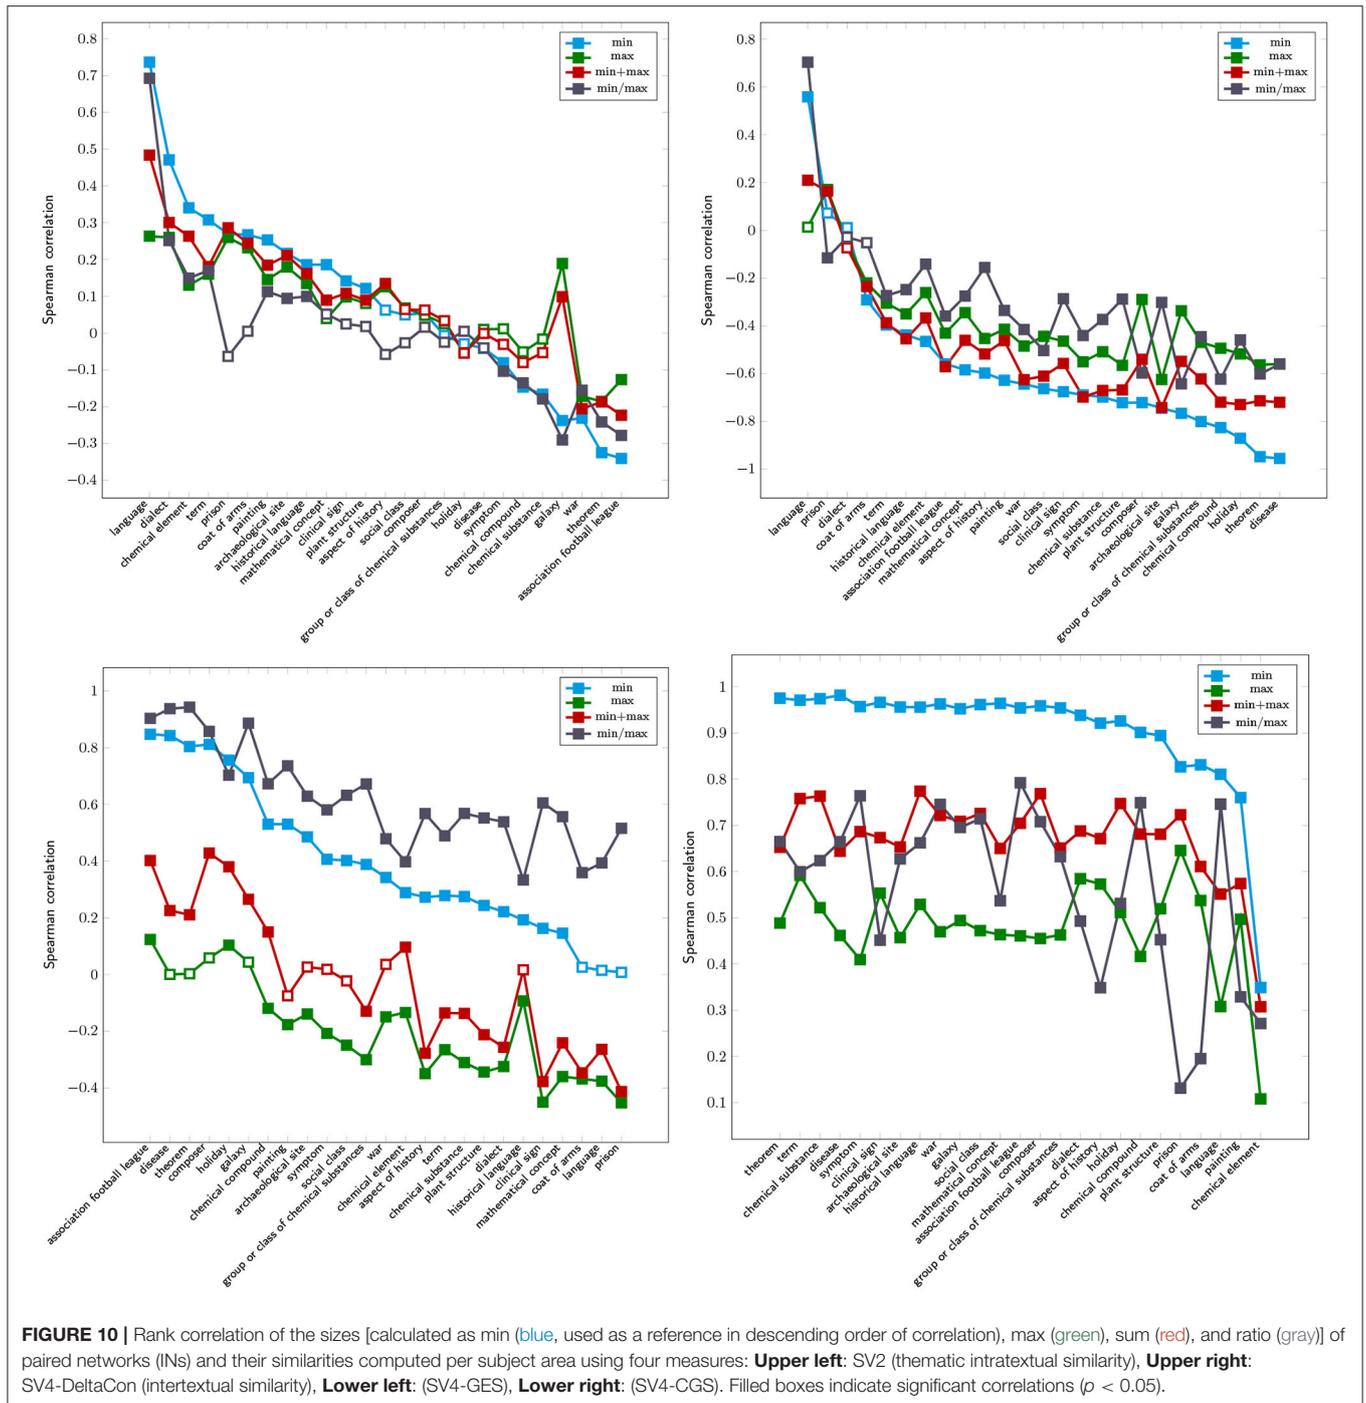
That is, we ask whether the rank of a pair of networks [the higher their similarity $\sigma(I_i^x, I_j^x)$, the higher the rank] correlates with its rank according to size $s(I_i^x, I_j^x)$ (the larger the networks, the higher the rank), distinguishing four alternatives: $s(I_i^x, I_j^x) = \min(|V_i^x|, |V_j^x|)$, $s(I_i^x, I_j^x) = \max(|V_i^x|, |V_j^x|)$, $s(I_i^x, I_j^x) = |V_i^x| + |V_j^x|$, and $s(I_i^x, I_j^x) = \min(|V_i^x|, |V_j^x|) / \max(|V_i^x|, |V_j^x|)$. The distributions are shown in **Figure 9**, the corresponding correlations in **Figure 10** (which additionally plots the values for CGS). For SV2, we observe a very strong effect regarding *language* ("visually" confirmed by **Table 3**) (note that the curves are ordered according to variant $\min(|V_i^x|, |V_j^x|)$, sorted in descending order); in most other cases correlations are



rather low (whether positive or negative). From this picture we conclude that observed thematic similarities of articles in different languages on the same subject cannot be attributed to the sizes of the networks involved. Remarkably, for SV2 we observe that the data series mostly coincide with variant $\min(|V_i^x|, |V_j^x|)$ regarding the correlations' order. This more or less also applies to DeltaCon in **Figure 10**. However, we now mostly observe higher negative correlations. Apparently, ranks in terms of structural similarity correlate negatively with size-related ranks. This means that if the INs are small, their similarities are likely to be large and vice versa. In the case of GES (**Figure 10**) we find this assessment more or less reversed: the correlations (blue) are almost all significantly higher than 0.2; that is, the larger the networks involved, the higher the similarity. From this perspective we can conclude that either both measures (GES and DeltaCon) contradict each other or (the more likely interpretation) they measure orthogonal aspects of graph similarity (the one set intersection-, the other walk-based). An exception is again CGS, which shows stable, high rank correlations not only for small graphs, for which it computes very

high values of dissimilarity, but also depends more on size than SV2 and DeltaCon—as motivated by the definition of CGS, size is a better predictor of it: the smaller the INs, the less their similarity and vice versa. In any case, the picture we get from this analysis is ambiguous, so that we hesitate to conclude that hypertextual similarity is reliably correlated with the orders of the networks involved: *intratextually, the similarity of INs does not depend on size and intertextually it does not show a clear trend.*

Next, we ask about the status of subject areas as a function of their analyses along SV1-4. Since each similarity measure induces a ranking of areas based on the average of the similarities observed for the language pairs (see **Table 4**), we can ask whether the rankings induced by different measures correlate or not. Lower rank correlations would then indicate unsystematic similarity relations in the sense that intra- and intertextual similarities do not point in the same direction. Lower rank correlations for either intra- or intertextual measures, in turn, would point to contradictory results. All in all, such findings would indicate that the INs under consideration exhibit incoherent similarities—contrary to the assumption of their



uniform similarity along intra- and intertextual dimensions. This is essentially what we find in **Table 5**: rank correlations are mostly low and not significant. One exception is the negative correlation of EJS and DeltaCon, which, apparently, measure different things. Another exception are the few examples of high positive correlations, such as those of GES and CGS. This correlation analysis shows that, with few exceptions, the similarity-based ordering of subject areas along one similarity

dimension (whether intra- or intertextual) does not allow us to infer their order along another dimension.

Next, we consider language networks whose edges correspond to the similarity values of the underlying language pairs (related network analyses have been conducted by Miquel-Ribé and Laniado, 2016; Samoilenko et al., 2016). We want to know in which subject areas which language clusters arise and whether different languages are center-forming in different subject areas

TABLE 5 | Rank correlations of the orderings of subject areas according to the average similarities obtained for the corresponding INs of language pairs where the similarities are computed by means of seven different similarity measures.

	SV1	SV2	SV3	GES	DeltaCON	EJS	CGS		SV1	SV2	SV3	GES	DeltaCON	EJS	CGS
SV1	1.00	0.27	0.95	-0.02	0.30	0.36	0.07	SV1	0.00	0.19	0.00	0.93	0.15	0.07	0.74
SV2	0.27	1.00	0.26	0.13	0.07	0.23	0.40	SV2	0.19	0.00	0.22	0.55	0.74	0.28	0.05
SV3	0.95	0.26	1.00	0.05	0.28	0.40	0.06	SV3	0.00	0.22	0.00	0.82	0.17	0.05	0.78
GES	-0.02	0.13	0.05	1.00	0.28	-0.06	0.72	GES	0.93	0.55	0.82	0.00	0.18	0.78	0.00
DeltaCON	0.30	0.07	0.28	0.28	1.00	-0.40	0.18	DeltaCON	0.15	0.74	0.17	0.18	0.00	0.05	0.38
EJS	0.36	0.23	0.40	-0.06	-0.40	1.00	0.17	EJS	0.07	0.28	0.05	0.78	0.05	0.00	0.41
CGS	0.07	0.40	0.06	0.72	0.18	0.17	1.00	CGS	0.74	0.05	0.78	0.00	0.38	0.41	0.00

(Left) Heatmap of correlation values, (Right) p-Values.

(see question Q3). Since the underlying similarity matrices (see Table 3) induce complete graphs, we filter out all edges whose similarity values are below the average similarity measured across all subject areas for the respective measure. Figure 11 illustrates the result by the example of three subject areas (*composer*, *painting*, and *language*) and two similarity views (SV2 and SV4-DeltaCon). The impression that we get already by this selection is rather confusing: the graphs look very different, from the perspective of the similarity measure and from the perspective of the subject area. A clear trend is not discernible (although English Wikipedia is usually prominently positioned). From this brief network analysis, we conclude that the different languages may have different degrees of salience depending on the subject area, while there is no single language that dominates in all these cases.

To conclude this analysis, we ask about the similarities of the positions of all 35 languages in the networks derived from the heatmaps of Table 3 and weight them by the ratio of these networks' strength centralities to the respective maximum. That is, we take the completely connected language similarity network for each of the 25 topics, in which each edge is weighted according to the similarity of the associated languages in the sense of the underlying similarity measure. Regardless of the actual structure of the INs, in the ideal state of equal INs for the different languages about the same topic, it is to be expected that the edges in the language networks are weighted by 1. Under this regime, the distance correlations of the strength centralities of languages should be maximal when being compared for each pair of topic-related language networks⁹: the languages then occupy identical network positions— independent of the topics. However, if all INs are maximally dissimilar, we likewise get maximum distance correlations of the languages according to their network positions. Thus, to differentiate between these two cases, we weight observed distance correlations by the quotient of observed network strength and maximum possible strength. For pairs of the fully connected language networks considered here,

this maximum is $35 * 34$, assuming a maximum edge weight of 1. To avoid unrealistically low weights, we multiply this maximum by the actually observed maximum value per similarity measure. The resulting heatmaps are shown in Figure 12: If INs on the same topic tend to be similar, both heatmaps are expected to be saturated, the weighted and the unweighted. Conversely, if these INs tend to be dissimilar, only the unweighted map is likely to be saturated. In Figure 12 we show the corresponding results for the optimistic variant of similarity measurement (DeltaCon) and its pessimistic counterpart (CGS). Obviously, with CGS we obtain higher distance correlations of the strength centralities of the languages in the 25 language networks than with DeltaCon, but very small weighted variants of these correlations, which indicate that the INs are extremely far from ideal similarities. The only exception is *chemical element*: this topic is described similarly in the different languages and more close to the ideal scope than any other topic. In the case of DeltaCon the weighted distance correlations are much higher, but still far from ideal. In Figure 13 we add the distance correlations generated by GES and SV2: for GES the unweighted correlations are similar to those of DeltaCon, for SV2 those of CGS. The situation is the opposite in the case of the weighted correlations. The general tendency is that the languages tend to have similar network positions, but at the price of lower similarities of the INs.

5.1. Discussion

Given the scope and growth dynamics of Wikipedia, it is not surprising that a reader who reads it in a particular language expects to be sufficiently broadly and deeply informed about the subject of his or her information search, provided it is one of the larger Wikipedias. We have shown that such an assumption about Wikipedia as a central building block of the IL does not apply—at least in the context of the subject areas examined here. Their different language versions differ so much in their treatment of the same subject area that it is necessary to know which area in which language someone is consulting if one wants to know how much the part of the IL he or she is traversing is biased. It may be the case that a reader's consultation of Wikipedia is accompanied by the assumption, that it is an open, dynamically

⁹For the notion of vertex strength (i.e., weighted degree values) see Barrat et al. (2004), for computing network centralities based on vertex indices see Feldman and Sanger (2007).

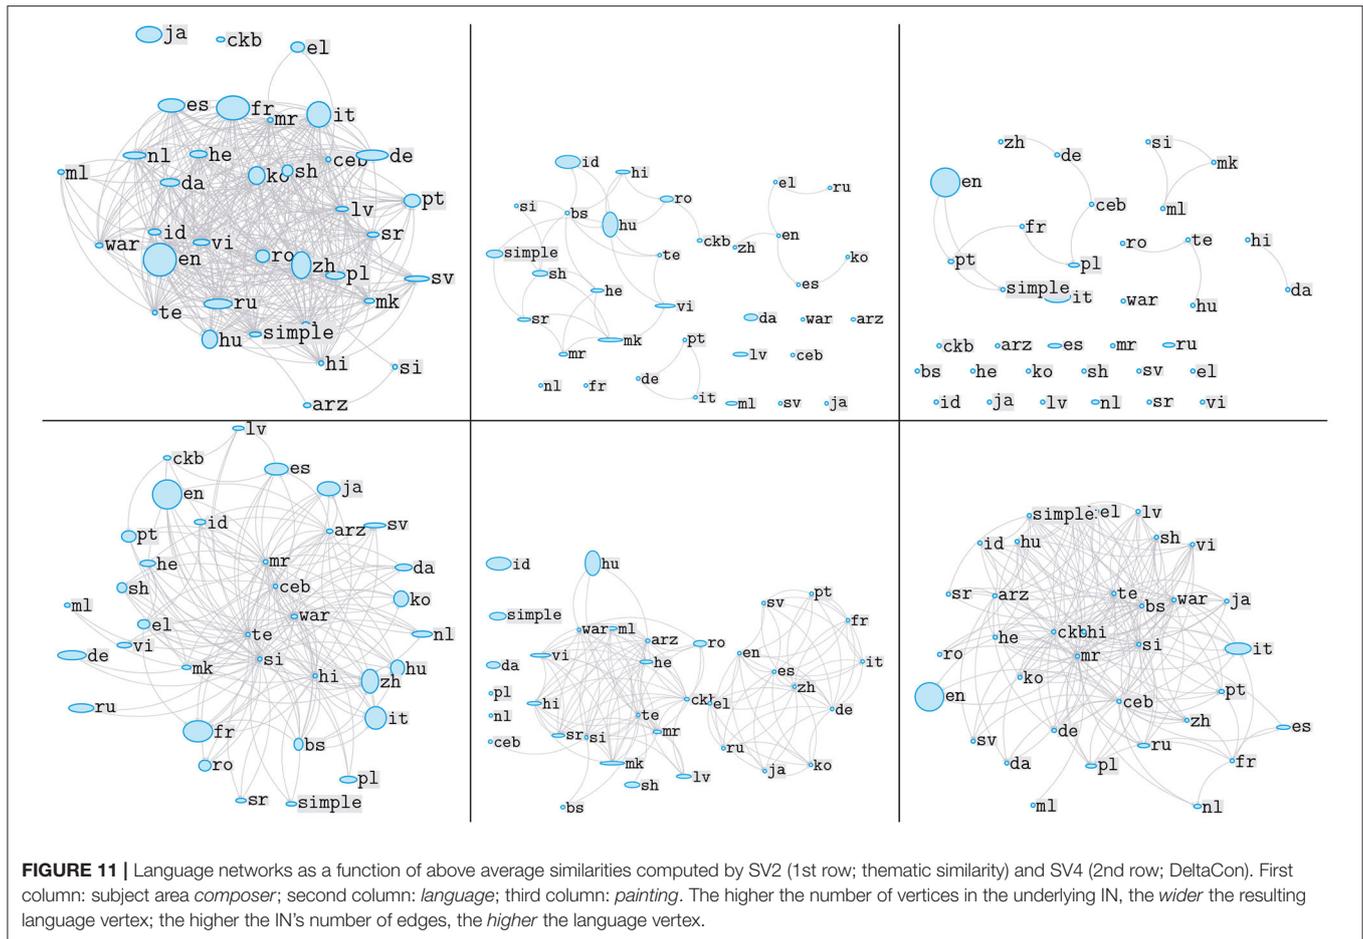


FIGURE 11 | Language networks as a function of above average similarities computed by SV2 (1st row; thematic similarity) and SV4 (2nd row; DeltaCon). First column: subject area *composer*; second column: *language*; third column: *painting*. The higher the number of vertices in the underlying IN, the wider the resulting language vertex; the higher the IN's number of edges, the higher the language vertex.

growing resource that largely covers relevant fields of knowledge, where one likely finds what one is looking for, and vice versa, that what is irrelevant is excluded. One might even assume that if a language version of Wikipedia is only sufficiently large, it will probably show this pattern of coverage—regardless of the underlying language. We have shown that at the current stage of development, such assumptions do not apply.

Our analysis has shown that subject areas covered by Wikipedia differ from language to language in a way that is hardly predictable by the size of the networks involved. Consequently, with regard to question Q1, we have to state that the similarities between the languages vary from subject area to subject area. It is therefore necessary to define the thematic reference (subject area) in order to say something about comparable languages. This may seem obvious, but it shows that Wikipedia's language editions are designed differently for different research fields and subject areas. With few exceptions (e.g., *chemical element*), we find that inequalities prevail within the same field. This makes it difficult to say which field is the more evenly distributed, language-independent one, which ensures that students can expect nearly equal information coverage regardless of language. Thus, in relation to question Q2, we conclude that the choice of the subject area has a major influence on the similarity assessment. And because of this influence, we do not discover a lingua franca

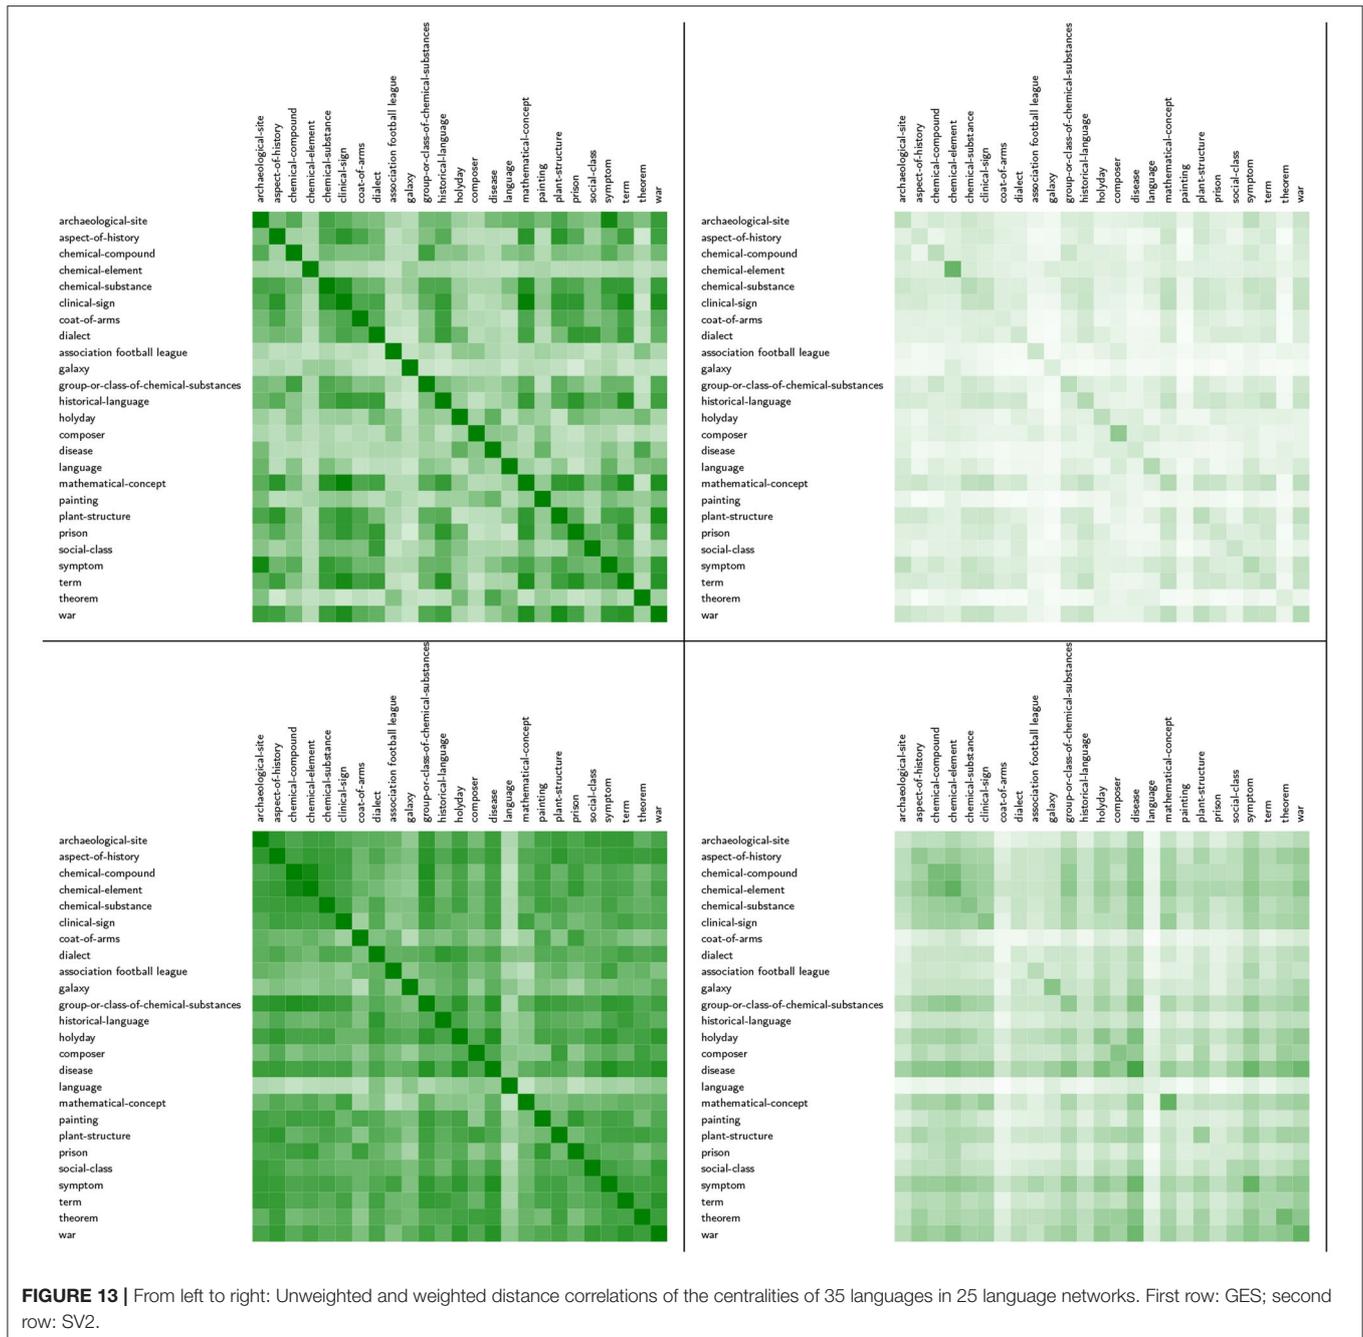
which, due to its size and coverage would serve as a kind of reference for the similarity relations of the different languages, so that although they may be dissimilar among themselves, they would be predominantly similar with respect to this central language. As things stand, the data do not support the attribution of this central role to English Wikipedia—at least from the point of view of the fields of knowledge and subject areas considered here (question Q3). In other words, knowing that Wikipedia is small in your first language and therefore probably does not cover the subject area of your task, it is hardly a way out of this dilemma to recommend reading English Wikipedia instead: it would be better to read at least both. However, we also see that depending on the subject area, other languages could play the role of primary sources (e.g., Hungarian in the case of subject area *language*—see **Figure 11**—or French in the case of subject area *composer*)—of course, this presupposes that one has the language skills for them.

Samoilenko et al. (2016) point out the influence of multiple dimensions on the commonalities, similarities and differences of Wikipedia's language editions, including language, culture, and geographical proximity. Apparently, we shed light on this view from the perspective of our four-part similarity analysis, which distinguishes between intratextual (text structural, quantitative, and thematic similarity) and intertextual (hypertextual) aspects,



where the latter are simultaneously examined by means of four measures. The important aspect is that these reference points may be influenced to varying degrees by linguistic, cultural or even geographical factors. In the case of SV3 (quantitative text profiles), for example, we observed rather evenly distributed similarities of the INs of the different subject areas, at a higher similarity level. The differences observed in this context may indicate the influence of the underlying languages, which their authors are less or barely able to “escape” through their own control. The extent to which SV3 models law-like behavior of

texts would then make INs of different or related languages look more similar—and differences could indicate a strong influence of the respective languages. SV1 (LDS) shows a similar pattern, although this may be due to the underlying webgenre, its editing rules and the way Wikipedia monitors compliance with them. However, such a finding could negate the role of this SV. In any case, one should not underestimate the source of information on which SV1 is based, as it concerns elements of documents that are likely to be the focus of reading: tables and especially figures and pictures. In this respect, a further development is



needed that integrates text analysis with image analysis and related approaches.

SV2 (thematic structure) differs from SV1 and SV3: It decreases rapidly and reaches a very low level of similarity (in case of subject area *disease* (see **Table 4**). Part of this dynamic is likely caused by the diverging *F*-values of our topic model (see **Table 2**). But let us assume its effectiveness. What could be the cause then, if we do not look for factors, such as size or age of a Wikipedia? Take the example of the subject area *disease*: could it be that it is cultural differences that determine which diseases

are described in which language edition by means of which intertextual structures (see the low similarity values measured with DeltaCon in **Table 3**)? Is there, so to speak, a cultural or any comparable disposition for the arrangement of intertextual contexts, while the corresponding intratextual similarities (SV2 in **Table 3**) tend to be much higher? In the case of subject area *holiday* (which according to SV2 and GES occupies a position in the middle, according to DeltaCon in the lower range and according to EJS in the upper range of the similarity spectrum), cultural references are rather likely. Whether true or

not, this gives rise to the question of which intertextual structures, which link-based factors for the production of multiple texts are culturally determined or have a cultural imprint. Is there, so to speak, a linguistic, cultural or knowledge area-related fingerprint that could be read from intertextual structures, from different parts of the IL, a fingerprint which could help to explain its dynamics beyond what is done by lexical-semantic analyses of link anchors? Even though our research does not answer this question, it does raise it and thus builds a bridge between the kind of data science we pursue and text-linguistic questions from the field of reading research.

To answer these questions, we need a multiple text model that includes the underlying IL as a limiting factor of what results from reading processes as multiple text; a model that considers linguistic, cultural, genre- and register-related (Halliday and Hasan, 1989) as well as social factors when asking about the function or meaning of a given or missing link. It is obvious that such a model-theoretical extension of reading research benefits considerably from browsing models, such as those developed by more recent hypertext research (Dimitrov et al., 2017; Lamprecht et al., 2017): it may not be surprising that readers are more likely to select links in the initial sections of Wikipedia articles according to text-structural criteria. However, this research has also highlighted the importance of semantic criteria for link selection. It is now necessary to further explore this system of motives and to extend it to the generation of multiple texts as a whole. In this context, our approach to DeltaCon is of interest, which evaluates random walks as a source of information to assess vertex affinities. The reason for this is that it can easily be linked to empirical research on the reading or navigation behavior of users in Wikipedia and comparable resources. In this way, it would integrate intertextual structural analysis with the pragmatics of real hypertext use. By additionally integrating the article-content-related similarity views SV1-3 as developed here, it would open up a very broad spectrum of information sources for the analysis and comparison of multiple texts, namely syntactic, semantic and pragmatic sources. This could ultimately pave the way to go beyond the detection of limiting factors, as the IL in the form of Wikipedia imposes on reading, to gain models of how this IL is actually represented by its readers and entire reader communities through distributed reading processes, that is, as a distributed cognitive map of the IL.

Irrespective of these findings, conclusions and prospects for future work, a number of boundary conditions must be considered with regard to our research, which at the same time affect its limitations:

1. *Coverage*: Although we have implemented an extensible procedure for extracting INs as intertextual manifestations of subject areas, which in principle can take into account any of the millions of Wikidata items to conduct cross-linguistic studies, we have only analyzed a subset of 25 such areas. This approach could be extended by asking about the convergence of similarities/dissimilarities between languages, as a result of studying a much wider range of subject areas. In this sense, we provided the starting point for a more detailed examination of the thematic biases of Wikipedia, compared to what has been

studied so far. This level of detail originates from our three-level topic model (**Figure 6**), which should be expanded into a model of thematic-rhematic intertextuality (cf. Mehler et al., 2019).

2. *Similarity analysis*: We implemented a hybrid approach to measure the similarities of INs, the structure-oriented part of which includes four approaches to graph similarity measurement. However, the spectrum of relevant measures is much wider (Emmert-Streib et al., 2016), so that their expressiveness and significance for intertextuality research should be examined more thoroughly. One may even think of a multiple source similarity measure that simultaneously maps various structural and other informational sources to assess the similarity of multiple texts and hypertexts.
3. *Network analysis*: The modeling of reading requires the modeling of cognitive processes, which in the case of distributed reading means the modeling of processes in social networks. As explained in our introduction, such an endeavor requires modeling the “fluent alignment” of two processes: on the one hand, the multi-authorial writing of Wikipedia and its embedding in the larger IL and, on the other hand, the diversity of reading situations, their task contexts and contextual resources by which they are conditioned. In this way, our approach could be further developed by integrating models of web-based writing research¹⁰ and social network analysis.
4. *Content mining*: Although we integrate a semantic model regarding the thematic structure of single texts, we do not yet consider their content beyond this level. This means, for example, that although our approach can find strong similarities between different Wikipedias on the same topic, the hypertexts in question can inform about this topic in very different ways, for example by making different statements or by embedding them in different argumentative contexts. By extending the semantic part of our text representation model (in particular by including knowledge graphs and representations based on semantic role labeling (Palmer et al., 2010) to detect, e.g., assertions or claims) we would be able to more reliably detect semantic (dis-)similarities between texts and overcome the corresponding limitation of our model.

A more comprehensive model of distributed reading and writing that meets these extensions to overcome the related limitations is certainly a challenge for future research, especially if it is to be based on thorough linguistic analysis. Currently, we see no way out of such a research direction for education science, that is, for studying learning in the age of information.

5.2. How Does Our Approach Relate to Linguistic Relativity?

Starting from a selected set of topics, we have shown that different language Wikipedias produce quite different networks for informing about the same topics. That is, we detected a

¹⁰For a recent network theoretical sentiment analysis of online writing see, for example, Stella et al. (2020). For a review of network theoretical approaches to knowledge networks in education science see Siew (2020).

bias: extent and organization of a topic's representation depend on the underlying Wikipedia—the former are biased by the latter. Since we related this bias to the languages in which the Wikipedias are written, we spoke of a *linguistic bias*. This raises the question what our approach contributes to research on what is known as *Linguistic Relativity* (LR) (Lakoff, 1987; Lucy, 1997) or *Cultural Relativity* (CR) (Gumperz and Levinson, 1991, 1996) (cf. also Sharifian, 2017). Are the differences we observe caused by differences in the underlying languages (LR) or even by cultural differences (CR) between the communities of writers producing these Wikipedias? Our approach does not allow for a direct answer to this question, as it is not based on the linguistic or social data required for such an undertaking. Nevertheless, it is worth explaining how it relates to this research.

To clarify this we utilize the distinction of *structure-*, *domain-*, and *behavior-centered* approaches according to Lucy (1997): the first start from observed semantic differences between languages to examine their influence on thought, the second from “domain[s] of experienced reality” (Lucy, 1997, p. 298) to ask how languages represent them, the third from language-specific practices of use. According to Lucy (1997, p. 298), the distinguishing feature of domain-specific approaches is that they characterize domains independently of the target languages. This is what we intend to do when deriving topic representations from Wikidata items and their relations (cf. Mehler et al., 2011): we explore this data to gain access to conceptual representations of parts of experienced reality, ask how they are described in different Wikipedias and whether the networked descriptions based thereon are commensurable or not. Our first assessment is that most approaches to Wikipedia's LR or CR (cf. Massa and Scrinzi, 2012; Laufer et al., 2015; Miquel-Ribé and Laniado, 2016; Miz et al., 2020) are domain-centered in such a way. Before assessing what they can say about LR, we go one step further in characterizing our approach, this time with the help of Lakoff (1987, p. 322) who discusses conceptual organization as a reference point for assessing the commensurability of conceptual systems, with which our approach to graph similarity is apparently compatible. In particular, Lakoff (1987, p. 334) concludes that organizational differences point to different conceptual systems and thus to LR. Let us assume that INs *manifest* such conceptual systems hypertextually. The differences we find in these INs could then reflect conceptual differences of the underlying languages and, if these differences are culturally determined, CR: conceptualizations of the subject areas investigated here would then be linguistically relative and ultimately culture-specific. If languages encode world views (Gumperz and Levinson, 1996, 2), Wikipedias can then be seen as manifestations of parts of such views, and since languages differ in encoding them, their Wikipedias are consequently non-trivially different. Moreover, due to the distributed authorship of Wikipedia, there is a direct link to the concept of distributed cognition (Hollan et al., 2000), to which approaches to CR are connected (Gumperz and Levinson, 1991; Sharifian, 2017). From this point of view, it seems plausible to assume effects of different communities, each of which produces representations of conceptual systems

or world views in the form of hypertexts, that are more or less incommensurable. From this it follows that by examining such differences we should get access to the differences of the underlying worldviews and their encodings. To sum up: we use Wikidata to identify cross-linguistic conceptual units (as a bridge to experienced reality), examine their language-specific lexicalizations (article names) and interconnections (hyperlinks) using corresponding excerpts from Wikipedia, interpret their differences as evidence of LR and speculate on CR as its cause (see above).

Apparently this consideration already brings us to the end of the analogy. The reason for this assessment concerns the research objects of the areas compared here. More precisely, the question is to which entities or systems the observed differences are finally ascribed. While research on LR aims to make statements about *language systems*—beyond the level of lexis—or *conceptual systems* (by asking whether the same parts of reality (e.g., the color spectrum) are conceptualized and coded in a language-specific way), we and our relatives in Wikipedia make statements about *textual instances* of such systems. The former deal with differences in language systems, while we study hypertextual differences without directly drawing conclusions about the underlying systems—beyond the lexical level. That is, we do not directly contribute to research on LR at the level of *language as system* but rather at the level of *language as text* (Hjelmslev, 1969). Moreover, beyond the question of whether or not a concept is manifested and networked in a language, we do not consider cases of conceptual splitting, fusion, etc. From this point of view, it is certainly no overstatement to claim that we found evidence of the type of linguistic bias described above. But it would be an overstatement to claim that we thereby measure LR on the level of language systems and the underlying cultures. Though we can speculate that cultural differences are responsible for the differences we measure, we cannot yet prove this with our apparatus. To consider language as a system, linguistic analyses are required, such as those provided by comparative linguistics (Bisang and Czerwinski, 2019). For example, we can ask about differences in the linguistic manifestation of topics (e.g., information structure, density and uncertainty, relevance, salience, etc.) and what effects this has on the organization and linking of articles. Based on this, we could ask for language-specific text patterns and attribute them to author communities and thus to cultural differences, but we would need independent data to support such conclusions. At least the linguistic part of this task now lies within the interdisciplinary reach of comparative and computational linguistics, so that we can put research on *reading/writing multiple texts* on a broader methodological basis. This is work for the future.

5.3. Information Processing and Online Reasoning

With regard to information processing and online reasoning in higher education, our approach has at least two implications:

1. As reviewed in section 3, Wikipedia is one of the most important resources in education. In accordance with the

Matthew principle, it is the primary target of first (and often final) information searches, either directly or via Google searches. This prominent position is contrasted by (i) the skewness of the thematic similarity relations we found, (ii) the unevenly distributed depths and widths of thematizations of the same topics across different Wikipedias, and (iii) the contextual dependence of the thematic similarities of these Wikipedias. This tripartite skewness and dependence, which is associated with a certain conception of linguistic relativity (see above), is now in turn in conflict with a way of using Wikipedia according to a reformulation of the *Closed World Assumption* (CWA) (Reiter, 1978); this reformulation assumes that Wikipedia is (almost) complete in the sense that it describes (almost) everything that is “relevant” to a particular subject area, while leaving out everything that is “irrelevant” in that sense¹¹: *what is relevant is in Wikipedia, or it is not relevant*. Such an assumption, which is not seriously supported by anyone, is obviously wrong, and we have given three more reasons for its rejection. The question is, however, whether even learners in the field of higher education search for and process information in *implicit* agreement with such an assumption or a similar view (possibly only for reasons of effort reduction). From the perspective of reading scenarios corresponding to this view, our research reveals a considerable potential of “false negatives,” that is, of knowledge units or components that have not been described but could have been described in Wikipedia. Our similarity analyses show that the same topics are described in different Wikipedias in different densities, and thus what is described in one Wikipedia is likely omitted in others (leading to “false negatives” in the latter sense); otherwise we would have observed much higher similarities of these Wikipedias. Thematic relevance (and consequently accessibility for the online reading process) is then either language-specific—which, given the topics selected here as examples, is unlikely to be the case—or one must acknowledge that resources, such as Wikipedia strictly contradict the latter variant of the CWA. Reading habits in the sense of this variant and actual information offer stand then in fundamental contradiction to each other. This assessment should be seen from the perspective that online accessibility attracts readers and that the status of Wikipedia as a widely accepted, unrivaled learning resource increases its likelihood of being read in higher education in whatever scenario (thus confirming a variant of the Matthew principle).

2. With the amount of information available online on numerous fields of knowledge and its accessibility through a medium with a unique selling point, such as Wikipedia, we can observe another phenomenon that is relevant to higher education but closed to the first glance. Starting from the concept of domain-specific learning and reading skills (List and Alexander, 2019), the question arises as to the online transparency and visibility of domain boundaries, which allow a learner to notice, in the context of his domain-specific task completion, at which

points he or she leaves the content area of his or her domain while switching to a domain that is possibly unfamiliar to him or her and for which he or she lacks the knowledge background and corresponding learning skills. This question is related to problems on two scales: on the micro-level, it concerns the cognitive load on the part of the learner due to an increased need for the integration of domain-external knowledge; on the macro-level, it corresponds to the blurring of domain boundaries, of established disciplinary differences, which are characterized by the development of the aforementioned domain-specific learning and reading skills. In speculative terms, a distributed reading process takes place as a result of numerous such events, in which each topic can be contextualized by any other topic of any domain (small-world effect). As a consequence of such a process, the imparting of domain differences, peculiarities and boundaries can be complicated even in higher education. At first glance, our research does not contribute to studying such processes, since we focused on specific topics for our similarity analyses, for which we have developed an extraction procedure, since Wikipedia does not make the relevant topic-related subnetworks visible. At second glance, however, our research also provides insights in this respect, as readers will hardly always follow the paths on which our similarity analyses were based. But if dissimilarity is already the predominant diagnosis for the latter paths, then this applies all the more to the free, unbound search for information (*what must already be considered dissimilar in our sense can only appear as more dissimilar under the latter regime*), so that the assessment made under the previous bullet point is reinforced. This implies the simplified consequence that with the increasing importance of learning resources like Wikipedia, the transfer of domain-specific knowledge can become more difficult and thus gains in importance.

We may add a third point concerning the divergence of time scales, since while the importance of Wikipedia is likely to grow steadily, this does not necessarily apply to the closing of the knowledge gaps diagnosed here: the idea of a convergence of all Wikipedias with regard to the increasingly similar presentation of the same subject areas is likely to be as daring as it is unrealistic. In any case, our approach shows that when it comes to investigating web-based resources with regard to their effects on (higher) education, the computational linguistic approach developed here appears indispensable if concrete models of the affected sections of the corresponding IL are required. That aims certainly beyond Wikipedia, but remains methodically in the context of what has been elaborated here.

6. CONCLUSION

We introduced a three-level topic model in combination with a graph-theoretical model for measuring the intra- and intertextual similarities of article networks from different language editions of Wikipedia. In this way we built a bridge between reading

¹¹This learning resource-related reformulation of the CWA replaces the concept of *truth* with the concept of *relevance with respect to specific information needs*.

research, educational science and computational linguistics. To this end, we described a new perspective for reading research that focuses more on the information landscape as a limiting factor of online reading. We have continued research showing that Wikipedia exhibits a topical coverage bias. However, we have done this using a much more elaborate topic and text structure model in conjunction with a quantitative model of hypertext structure, a hybrid model that is more realistic from the perspective of hypertext linguistics (Storrer, 2002). Finally, we have derived two implications from our findings; the first concerns a variant of the closed world assumption and the related prediction that learning processes based on shallow reading and quick search processes are likely to be affected by the thematic dissimilarities and contextual dependencies of online information resources as we have observed in Wikipedia, and that this effect may also have negative effects on the acquisition of domain-specific knowledge and corresponding learning skills. In future work we will continue elaborating our computational hypertext model. This will be done with special attention to hypertext usage to obtain graph models that integrate browsing behavior into graph similarity analysis. In this way, we will try to model students' online learning by combining information about *how they read* with information about *what they read* and how this is networked. This will possibly give an outlook on two things: the construction and integration of knowledge on the part of single students, learners, or readers and its distributed counterpart regarding larger communities of students, learners or readers and their distributed reading processes.

REFERENCES

- Álvarez, G., Oeberst, A., Cress, U., and Ferrari, L. (2020). Linguistic evidence of in-group bias in English and Spanish Wikipedia articles about international conflicts. *Discourse Context Media* 35:100391. doi: 10.1016/j.dcm.2020.100391
- Bahdanau, D., Cho, K., and Bengio, Y. (2015). "Neural machine translation by jointly learning to align and translate" in *Proceedings of the International Conference on Learning Representations* (San Diego, CA), 1–15.
- Bao, P., Hecht, B., Carton, S., Quaderi, M., Horn, M., and Gergle, D. (2012). "Omnipedia: bridging the Wikipedia language gap," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, TX), 1075–1084. doi: 10.1145/2207676.2208553
- Barrat, A., Barthélemy, M., Pastor-Satorras, R., and Vespignani, A. (2004). The architecture of complex weighted networks. *Proc. Nat. Acad. Sci. U.S.A.* 101, 3747–3752. doi: 10.1073/pnas.0400087101
- Barzilai, S., and Zohar, A. (2012). Epistemic thinking in action: evaluating and integrating online sources. *Cogn. Instruct.* 30, 39–85. doi: 10.1080/07370008.2011.636495
- Bisang, W., and Czerwinski, P. (2019). "Performance in knowledge assessment tests from the perspective of linguistic typology," in *Frontiers and Advances in Positive Learning in the Age of InformaTiOn (PLATO)* (Cham: Springer), 207–235. doi: 10.1007/978-3-030-26578-6_16
- Braasch, J. L. G., Bråten, I., and McCrudden, M. T. (Eds.). (2018a). *Handbook of Multiple Source Use*. New York, NY; London: Routledge.
- Braasch, J. L. G., McCrudden, M. T., and Bråten, I. (2018b). "Reflections and future directions," in *Handbook of Multiple Source Use, Chapter 29*, eds J. L. Braasch, I. Bråten, and M. T. McCrudden (New York, NY; London: Routledge), 527–537. doi: 10.4324/9781315627496-29
- Britt, M. A., Rouet, J.-F., and Braasch, J. L. (2012). "Documents as entities: extending the situation model theory of comprehension," in *Reading - From Words to Multiple Texts*, eds M. A. Britt, S. R. Goldmann, and J. F. Rouet (New York, NY: Routledge), 161–179. doi: 10.4324/9780203131268
- Britt, M. A., Rouet, J.-F., and Durik, A. M. (2018). *Literacy Beyond Text Comprehension: A Theory of Purposeful Reading*. New York, NY: Routledge.
- Callahan, E. S., and Herring, S. C. (2011). Cultural bias in Wikipedia content on famous persons. *J. Am. Soc. Inform. Sci. Technol.* 62, 1899–1915. doi: 10.1002/asi.21577
- Cheng, J., Dong, L., and Lapata, M. (2016). "Long short-term memory-networks for machine reading," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (Austin, TX: Association for Computational Linguistics), 551–561. doi: 10.18653/v1/D16-1053
- Cho, B.-Y., and Afflerbach, P. (2015). Reading on the internet: realizing and constructing potential texts. *J. Adolesc. Adult Liter.* 58, 504–517. doi: 10.1002/jaal.387
- Coiro, J., Sparks, J. R., and Kulikowich, J. M. (2018). "Assessing online collaborative inquiry and social deliberation skills as learners navigate multiple sources and perspectives," in *Handbook of Multiple Source Use*, eds J. L. Braasch, I. Bråten, and M. T. McCrudden (New York, NY; London: Routledge), 485–501. doi: 10.4324/9781315627496-27
- Conde, A., Arruarte, A., Larrañaga, M., and Elorriaga, J. A. (2020). How can Wikipedia be used to support the process of automatically building multilingual domain modules? A case study. *Inform. Process. Manage.* 57:102232. doi: 10.1016/j.ipm.2020.102232
- Denning, P., Horning, J., Parnas, D., and Weinstein, L. (2005). Wikipedia risks. *Commun. ACM* 48, 152–152. doi: 10.1145/1101779.1101804

DATA AVAILABILITY STATEMENT

Code and datasets used for this study can be found on github: <https://github.com/texttechnologylab/WikiSim>.

AUTHOR CONTRIBUTIONS

AM has written sections 1, 2, 3, and 5 (together with PW), generated **Figures 1–4, 6, 7, 11–13** and **Table 1**. AM, MK, PW, TU, and WH have written section 4. WH extracted INs from Wikidata and Wikipedia and generated **Figures 8–10** as well as **Tables 3–5**, together with AM, he generated **Figure 5**. TU generated **Table 2**, trained and tested text2ddc. AM (CGS), MK (SV1, SV3), PW (EJS, DeltaCon), TU (SV2), and WH (GES) implemented, computed, and described the similarity measures. AM, MK, PW, and WH proofread and revised the final manuscript. All authors contributed to the conception of the paper, its algorithmization, and the interpretation of the results.

FUNDING

WH was partly financed by the State of Hessen through the LOEWE research focus on *minority studies: language and identity*. PW was partially funded by the Federal Ministry of Education and Research of Germany as part of the competence center for machine learning ML2R (01|S18038C) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC 2070-390732324.

- DeStefano, D., and LeFevre, J.-A. (2007). Cognitive load in hypertext reading: a review. *Comput. Hum. Behav.* 23, 1616–1641. doi: 10.1016/j.chb.2005.08.012
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2019). “BERT: pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (Minneapolis, MN), 4171–4186.
- Dimitrov, D., Singer, P., Lemmerich, F., and Strohmaier, M. (2017). “What makes a link successful on wikipedia?” in *Proceedings of the 26th International Conference on World Wide Web, WWW 17* (Geneva: International World Wide Web Conferences Steering Committee), 917–926. doi: 10.1145/3038912.3052613
- Downs, R. M., and Stea, D. (1977). *Maps in Minds: Reflections on Cognitive Mapping*. New York, NY: Harper & Row.
- Emmert-Streib, F., Dehmer, J., and Shi, Y. (2016). Fifty years of graph matching, network alignment and network comparison. *Inform. Sci.* 346–347, 180–197. doi: 10.1016/j.ins.2016.01.074
- Erdős, P., and Rényi, A. (1959). On random graphs. *Publ. Math.* 6, 290–297.
- Feldman, R., and Sanger, S. C. (2007). *The Text Mining Handbook. Advanced Approaches in Analyzing Unstructured Data*. Cambridge: Cambridge University Press.
- Goldman, S. R., Braasch, J. L., Wiley, J., Graesser, A. C., and Brodowinska, K. (2012). Comprehending and learning from internet sources: processing patterns of better and poorer learners. *Read. Res. Q.* 47, 356–381. doi: 10.1002/RRQ.027
- Graham, M., Straumann, R. K., and Hogan, B. (2015). Digital divisions of labor and informational magnetism: mapping participation in Wikipedia. *Ann. Assoc. Am. Geograph.* 105, 1158–1178. doi: 10.1080/00045608.2015.1072791
- Gumperz, J. J., and Levinson, S. C. (1991). Rethinking linguistic relativity. *Curr. Anthropol.* 32, 613–623. doi: 10.1086/204009
- Gumperz, J. J., and Levinson, S. C. (1996). “Introduction: Linguistic relativity re-examined,” in *Rethinking Linguistic Relativity* (Cambridge: Cambridge University Press), 1–18.
- Halavais, A., and Lackaff, D. (2008). An analysis of topical coverage of Wikipedia. *J. Comput. Mediat. Commun.* 13, 429–440. doi: 10.1111/j.1083-6101.2008.00403.x
- Halliday, M. A. K., and Hasan, R. (1989). *Language, Context, and Text: Aspects of Language in a Sociosemiotic Perspective*. Oxford: Oxford University Press.
- Hargittai, E., and Dobransky, K. (2017). Old dogs, new clicks: digital inequality in skills and uses among older adults. *Can. J. Commun.* 42, 195–212. doi: 10.22230/cjc.2017v42n2a3176
- Hartman, D. K., Hagerman, M. S., and Leu, D. J. (2018). “Toward a new literacies perspective of synthesis: multiple source meaning construction,” in *Handbook of Multiple Source Use, Chapter 4*, eds J. L. Braasch, I. Bråten, and M. T. McCrudden (New York, NY; London: Routledge), 55–78. doi: 10.4324/9781315627496-4
- Head, A. (2013). Project information literacy: what can be learned about the information-seeking behavior of today’s college students? *SSRN Electron. J.* doi: 10.2139/ssrn.2281511. [Epub ahead of print].
- Hecht, B., and Gergle, D. (2009). “Measuring self-focus bias in community-maintained knowledge repositories,” in *Proceedings of the Fourth International Conference on Communities and Technologies* (University Park, PA), 11–20. doi: 10.1145/1556460.1556463
- Hecht, B., and Gergle, D. (2010a). “The Tower of Babel meets Web 2.0: user-generated content and its applications in a multilingual context,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10* (Atlanta, GA), 291–300. doi: 10.1145/1753326.1753370
- Hecht, B. J., and Gergle, D. (2010b). “On the “localness” of user-generated content,” in *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, CSCW '10* (New York, NY: ACM), 229–232. doi: 10.1145/1718918.1718962
- Hemati, W., Uslu, T., and Mehler, A. (2016). “TextImager: a distributed UIMA-based system for NLP,” in *Proceedings of COLING 2016: System Demonstrations* (Osaka), 59–63.
- Hjelmslev, L. (1969). *Prolegomena to a Theory of Language*. Madison, WI: University of Wisconsin Press.
- Hollan, J., Hutchins, E., and Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Trans. Comput. Hum. Interact.* 7, 174–196. doi: 10.1145/353485.353487
- Holloway, T., Bozicevic, M., and Börner, K. (2007). Analyzing and visualizing the semantic coverage of Wikipedia and its authors: research articles. *Complexity* 12, 30–40. doi: 10.1002/cplx.20164
- Hsieh, Y. P. (2012). Online social networking skills: the social affordances approach to digital inequality. *First Monday* 17. doi: 10.5210/fm.v17i4.3893
- Jiang, Y., Bai, W., Zhang, X., and Hu, J. (2017). Wikipedia-based information content and semantic similarity computation. *Inform. Process. Manage.* 53, 248–265. doi: 10.1016/j.ipm.2016.09.001
- Joulin, A., Grave, E., Bojanowski, P., and Mikolov, T. (2017). “Bag of tricks for efficient text classification,” in *Proceedings of the 15th Conference of the EACL: Volume 2, Short Papers* (Valencia: Association for Computational Linguistics), 427–431. doi: 10.18653/v1/E17-2068
- Karimi, F., Bohlin, L., Samoilenko, A., Rosvall, M., and Lancichinetti, A. (2015). Mapping bilateral information interests using the activity of Wikipedia editors. *Palgrave Commun.* 1, 1–7. doi: 10.1057/palcomms.2015.41
- Kendeou, P., Robinson, D. H., and McCrudden, M. T. (2019). “Modeling the dissemination of misinformation through discourse dynamics,” in *Misinformation and Disinformation in Education: An Introduction* (Charlotte: Information Age Publishing), 1–4.
- Kintsch, W. (1998). *Comprehension. A Paradigm for Cognition*. Cambridge: Cambridge University Press.
- Kittur, A., Chi, E. H., and Suh, B. (2009). “What’s in Wikipedia?: mapping topics and conflict using socially annotated category structure,” in *Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI '09* (New York, NY: ACM), 1509–1512. doi: 10.1145/1518701.1518930
- Köhler, R., Altmann, G., and Piotrowski, R. G. (Eds.). (2005). *Quantitative Linguistics. An International Handbook*. Berlin; New York, NY: Mouton de Gruyter.
- Konca, M., Mehler, A., Baumartz, D., and Hemati, W. (2020). *From Distinguishability to Informativity: A Quantitative Text Model for Detecting Random Texts*.
- Konieczny, P. (2016). Teaching with Wikipedia in a 21st-century classroom: perceptions of Wikipedia and its educational benefits. *J. Assoc. Inform. Sci. Technol.* 67, 1523–1534. doi: 10.1002/asi.23616
- Koutra, D., Shah, N., Vogelstein, J. T., Gallagher, B., and Faloutsos, C. (2016). DeltaCon: principled massive-graph similarity function with attribution. *ACM Trans. Knowl. Discov. Data* 10, 28:1–28:43. doi: 10.1145/2824443
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. Chicago, IL: University of Chicago Press.
- Lamprecht, D., Lerman, K., Helic, D., and Strohmaier, M. (2017). How the structure of wikipedia articles influences user navigation. *New Rev. Hypermed. Multimed.* 23, 29–50. doi: 10.1080/13614568.2016.1179798
- Laufer, P., Wagner, C., Flöck, F., and Strohmaier, M. (2015). “Mining cross-cultural relations from Wikipedia: a study of 31 European food cultures,” in *Proceedings of the ACM Web Science Conference* (Oxford), 1–10. doi: 10.1145/2786451.2786452
- Lemmerich, F., Sáez-Trumper, D., West, R., and Zia, L. (2019). “Why the world reads Wikipedia: beyond english speakers,” in *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining* (Melbourne, VIC), 618–626. doi: 10.1145/3289600.3291021
- List, A., and Alexander, P. A. (2019). Toward an integrated framework of multiple text use. *Educ. Psychol.* 54, 20–39. doi: 10.1080/00461520.2018.1505514
- Loh, K. K., and Kanai, R. (2016). How has the internet reshaped human cognition? *Neuroscientist* 22, 506–520. doi: 10.1177/1073858415595005
- Lorini, V., Rando, J., Saez-Trumper, D., and Castillo, C. (2020). Uneven coverage of natural disasters in Wikipedia: the case of flood. *arXiv* 2001.08810.
- Lucassen, T., and Schraagen, J. M. (2010). “Trust in Wikipedia: how users trust information from an unknown source,” in *Proceedings of the 4th Workshop on Information Credibility, WICOW '10* (Raleigh, NC), 19–26. doi: 10.1145/1772938.1772944
- Lucy, J. A. (1997). Linguistic relativity. *Annu. Rev. Anthropol.* 26, 291–312. doi: 10.1146/annurev.anthro.26.1.291
- Massa, P., and Scrinzi, F. (2012). “Manypedia: comparing language points of view of wikipedia communities,” in *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration, WikiSym 12* (New York, NY: Association for Computing Machinery), 1–9. doi: 10.1145/2462932.2462960

- McMahon, C., Johnson, I., and Hecht, B. (2017). "The substantial interdependence of Wikipedia and Google: a case study on the relationship between peer production communities and information technologies," in *Eleventh International AAAI Conference on Web and Social Media* (Montréal, QC), 142–151.
- Mehler, A., Gleim, R., Gaitsch, R., Uslu, T., and Hemati, W. (2019). From topic networks to distributed cognitive maps: Zipfian topic universes in the area of volunteered geographic information. *Complexity*. 4, 1–47. doi: 10.1155/2020/4607025
- Mehler, A., Hemati, W., Uslu, T., and Lücking, A. (2018). "A multidimensional model of syntactic dependency trees for authorship attribution," in *Quantitative Analysis of Dependency Structures*, eds J. Jiang and H. Liu (Berlin; New York, NY: De Gruyter), 315–348. doi: 10.1515/9783110573565-016
- Mehler, A., Pustynnikov, O., and Diewald, N. (2011). Geography of social ontologies: testing a variant of the Sapir-Whorf hypothesis in the context of Wikipedia. *Comput. Speech Lang.* 25, 716–740. doi: 10.1016/j.csl.2010.05.006
- Mehler, A., and Ramesh, V. (2019). "TextInContext: on the way to a framework for measuring the context-sensitive complexity of educationally relevant texts—a combined cognitive and computational linguistic approach," in *Frontiers and Advances in Positive Learning in the Age of InformaTiOn (PLATO)*, ed O. Zlatkin-Troitschanskaia (Cham: Springer International Publishing), 167–195. doi: 10.1007/978-3-030-26578-6_14
- Mehler, A., Sharoff, S., and Santini, M. (Eds.). (2010). *Genres on the Web: Computational Models and Empirical Studies*. Dordrecht: Springer.
- Mesgari, M., Okoli, C., Mehdi, M., Nielsen, Årup Nielsen F., and Lanamäki, A. (2015). The sum of all human knowledge: a systematic review of scholarly research on the content of Wikipedia. *J. Assoc. Inform. Sci. Technol.* 66, 219–245. doi: 10.1002/asi.23172
- Miquel-Ribé, M., and Laniado, D. (2016). "Cultural identities in wikipedias," in *Proceedings of the 7th 2016 International Conference on Social Media & Society* (London), 1–10. doi: 10.1145/2930971.2930996
- Miz, V., Hanna, J., Aspert, N., Ricaud, B., and Vanderghenst, P. (2020). "What is trending on Wikipedia? Capturing trends and language biases across Wikipedia editions," in *Companion Proceedings of the Web Conference 2020* (Taipei), 794–801. doi: 10.1145/3366424.3383567
- Nagel, M.-T., Schäfer, S., Zlatkin-Troitschanskaia, O., Schemer, C., Maurer, M., Molerov, D., et al. (accepted). How do university students' web search behavior, website characteristics, and the interaction of both influence students' critical online reasoning? *Front. Educ.*
- Oeberst, A., von der Beck, I., Back, M. D., Cress, U., and Nestler, S. (2018). Biases in the production and reception of collective knowledge: the case of hindsight bias in Wikipedia. *Psychol. Res.* 82, 1010–1026. doi: 10.1007/s00426-017-0865-7
- Oeberst, A., von der Beck, I., Matschke, C., Ihme, T. A., and Cress, U. (2019). Collectively biased representations of the past: Ingroup bias in Wikipedia articles about intergroup conflicts. *Br. J. Soc. Psychol.* 59, 791–818. doi: 10.1111/bjso.12356
- OECD (2007). *Revised Field of Science and Technology (FOS)*. Available online at: www.oecd.org/science/inno/38235147.pdf
- Okoli, C., Mehdi, M., Mesgari, M., Årup Nielsen, F., and Lanamäki, A. (2014). Wikipedia in the eyes of its beholders: a systematic review of scholarly research on wikipedia readers and readership. *J. Assoc. Inform. Sci. Technol.* 65, 2381–2403. doi: 10.1002/asi.23162
- Okoli, C., Mehdi, M., Mesgari, M., Årup Nielsen, F., and Lanamäki, A. (2012). The people's encyclopedia under the gaze of the sages: a systematic review of scholarly research on Wikipedia. *SSRN Electron. J.* 1–138. doi: 10.2139/ssrn.2021326
- Palmer, M., Gildea, D., and Xue, N. (2010). *Semantic Role Labeling*. Morgan & Claypool Publishers.
- Pentzold, C., Weltevrede, E., Mauri, M., Laniado, D., Kaltenbrunner, A., and Borra, E. (2017). Digging wikipedia: the online encyclopedia as a digital cultural heritage gateway and site. *J. Comput. Cult. Herit.* 10, 1–19. doi: 10.1145/3012285
- Perfetti, C. A., Rouet, J.-F., and Britt, M. A. (1999). "Toward a theory of documents representation," in *The Construction of Mental Representations During Reading*, eds H. van Oostendorp and S. R. Goldman (Mahwah, NJ: Erlbaum), 99–122.
- Pickering, M. J., and Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behav. Brain Sci.* 27, 169–226. doi: 10.1017/S0140525X04000056
- Power, R., Scott, D., and Bouayad-Agha, N. (2003). Document structure. *Comput. Linguist.* 29, 211–260. doi: 10.1162/089120103322145315
- Primor, L., and Katzir, T. (2018). Measuring multiple text integration: a review. *Front. Psychol.* 9:2294. doi: 10.3389/fpsyg.2018.02294
- Reiter, R. (1978). "On closed world data bases," in *Logic and Data Bases*, eds H. Gallaire and J. Minker (Boston, MA: Springer US), 55–76. doi: 10.1007/978-1-4684-3384-5_3
- Salmerón, L., Kammerer, Y., and Delgado, P. (2018). "Non-academic multiple source use on the internet," in *Handbook of Multiple Source Use*, eds J. L. Braasch, I. Bråten, and M. T. McCrudden (New York, NY; London: Routledge), 285–302. doi: 10.4324/9781315627496-17
- Samoilenko, A., Karimi, F., Edler, D., Kunegis, J., and Strohmaier, M. (2016). Linguistic neighbourhoods: explaining cultural borders on Wikipedia through multilingual co-editing activity. *EPJ Data Sci.* 5:9. doi: 10.1140/epjds/s13688-016-0070-8
- Samoilenko, A., Lemmerich, F., Weller, K., Zens, M., and Strohmaier, M. (2017). "Analysing timelines of national histories across Wikipedia editions: a comparative computational approach," in *Eleventh International AAAI Conference on Web and Social Media*, 210–219.
- Scaffidi, M. A., Khan, R., Wang, C., Keren, D., Tsui, C., Garg, A., et al. (2017). Comparison of the impact of wikipedia, UpToDate, and a digital textbook on short-term knowledge acquisition among medical students. *JMIR Med. Educ.* 3:e20. doi: 10.2196/mededu.8188
- Sharifian, F. (2017). Cultural linguistics and linguistic relativity. *Lang. Sci.* 59, 83–92. doi: 10.1016/j.langsci.2016.06.002
- Siew, C. S. Q. (2020). Applications of network science to education research: quantifying knowledge and the development of expertise through network analysis. *Educ. Sci.* 10:101. doi: 10.3390/educsci10040101
- Singer, P., Lemmerich, F., West, R., Zia, L., Wulczyn, E., Strohmaier, M., et al. (2017). "Why we read Wikipedia," in *Proceedings of the 26th International Conference on World Wide Web* (Perth, WA), 1591–1600. doi: 10.1145/3038912.3052716
- Smith, D. A. (2020). Situating Wikipedia as a health information resource in various contexts: a scoping review. *PLoS ONE* 15: e0228786. doi: 10.1371/journal.pone.0228786
- Stella, M., Restocchi, V., and Deyne, S. D. (2020). #lockdown: Network-enhanced emotional profiling in the time of COVID-19. *Big Data Cogn. Comput.* 4:14. doi: 10.3390/bdcc4020014
- Storrer, A. (2002). Coherence in text and hypertext. *Document Des.* 3, 156–168. doi: 10.1075/dd.3.2.06sto
- Sweller, J. (1994). Cognitive load theory, learning difficulty, and instructional design. *Learn. Instruct.* 4, 295–312. doi: 10.1016/0959-4752(94)90003-5
- Tobler, W. R. (1970). A computer movie simulating urban growth in the Detroit region. *Econ. Geogr.* 46, 234–240. doi: 10.2307/143141
- Uslu, T., Mehler, A., and Baumartz, D. (2019). "Computing classifier-based embeddings with the help of text2ddc," in *Proceedings of the 20th International Conference on Computational Linguistics and Intelligent Text Processing (CICLing 2019)* (La Rochelle).
- van Dijk, T. A. (1980). *Macrostructures. An Interdisciplinary Study of Global Structures in Discourse, Interaction, and Cognition*. Hillsdale, NJ: Erlbaum.
- van Dijk, T. A., and Kintsch, W. (1983). *Strategies of Discourse Comprehension*. New York, NY: Academic Press.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). "Attention is all you need," in *Advances in Neural Information Processing Systems* (Long Beach, CA), 5998–6008.
- Wagner, C., Graells-Garrido, E., Garcia, D., and Menczer, F. (2016). Women through the glass ceiling: gender asymmetries in Wikipedia. *EPJ Data Sci.* 5:5. doi: 10.1140/epjds/s13688-016-0066-4
- Warncke-Wang, M., Uduwage, A., Dong, Z., and Riedl, J. (2012). "In search of the ur-wikipedia: universality, similarity, and translation in the Wikipedia inter-language link network," in *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration* (Linz), 1–10. doi: 10.1145/2462932.2462959
- Wolf, M. (2018). *Reader, Come Home: The Reading Brain in a Digital World*. New York, NY: Harper.

- Zlatkin-Troitschanskaia, O., Bisang, W., Mehler, A., Banerjee, M., and Roeper, J. (2019a). "Positive learning in the internet age: developments and perspectives in the plato program," in *Frontiers and Advances in Positive Learning in the Age of InformaTiOn (PLATO)*, ed O. Zlatkin-Troitschanskaia (Cham: Springer International Publishing), 1–5. doi: 10.1007/978-3-030-26578-6_1
- Zlatkin-Troitschanskaia, O., Brückner, S., Molerov, D., and Bisang, W. (2019b). "What can we learn from theoretical considerations and empirical evidence on learning in higher education? Implications for an interdisciplinary research framework," in *Frontiers and Advances in Positive Learning in the Age of InformaTiOn (PLATO)* (Cham: Springer), 287–309. doi: 10.1007/978-3-030-26578-6_21
- Zlatkin-Troitschanskaia, O., Schmidt, S., Molerov, D., Shavelson, R. J., and Berliner, D. (2018). "Conceptual fundamentals for a theoretical and empirical framework of positive learning," in *Positive Learning in the Age of Information*

(PLATO)-A Blessing or a Curse? eds O. Zlatkin-Troitschanskaia, G. Wittum, and A. Dengel (Springer), 29–50. doi: 10.1007/978-3-658-19567-0_4

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Mehler, Hemati, Welke, Konca and Uslu. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.