



Words as Bridges: Exploring Computational Support for Cross-Disciplinary Translation Work

Calvin Bao
University of Maryland
College Park, Maryland, USA
csbao@umd.edu

Marine Carpuat
Department of Computer Science
University of Maryland
College Park, Maryland, USA
marine@umd.edu

Yow-Ting Shiue
Department of Computer Science
University of Maryland College Park
College Park, Maryland, USA
Artificial Intelligence Foundation
San Francisco, California, USA
ytshiue@umd.edu

Joel Chan
College of Information Studies
University of Maryland
College Park, Maryland, USA
joelchan@umd.edu

Abstract

Scholars often explore literature outside of their home community of study. This exploration process is frequently hampered by field-specific jargon. Past computational work often focuses on supporting translation work by removing jargon through simplification and summarization; here, we explore a different approach that preserves jargon as useful bridges to new conceptual spaces. Specifically, we cast different scholarly domains as different language-using communities, and explore how to adapt techniques from unsupervised cross-lingual alignment of word embeddings to explore conceptual alignments between domain-specific word embedding spaces. We developed a prototype cross-domain search engine that uses aligned domain-specific embeddings to support conceptual exploration, and tested this prototype in two case studies. We discuss qualitative insights into the promises and pitfalls of this approach to translation work, and suggest design insights for future interfaces that provide computational support for cross-domain information seeking.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Information systems** → **Specialized information retrieval**; • **Computing methodologies** → *Machine translation*.

Keywords

cross-domain information seeking, information foraging, scholarly term translation

ACM Reference Format:

Calvin Bao, Yow-Ting Shiue, Marine Carpuat, and Joel Chan. 2025. Words as Bridges: Exploring Computational Support for Cross-Disciplinary Translation Work. In *30th International Conference on Intelligent User Interfaces (IUI '25, Cagliari, Italy*



This work is licensed under a Creative Commons Attribution 4.0 International License. *IUI '25, Cagliari, Italy*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1306-4/25/03
<https://doi.org/10.1145/3708359.3712110>

'25), March 24–27, 2025, Cagliari, Italy. ACM, New York, NY, USA, 26 pages.
<https://doi.org/10.1145/3708359.3712110>

1 Introduction

Scholars often explore literature outside of their home community of study. When crossing domains or disciplines, jargon can be a major challenge [39, 46, 48]. Researchers may struggle to formulate effective queries to search for papers related to a specific concept they have in mind, because the same or similar concept may be referred to with very different terms. For instance, while research on creativity in Psychology talks about “examples” and “structured imagination”, researchers in Engineering may talk about similar things in terms of “inspirations” and “precedents”. The same surface form may also map to very different underlying concepts: for instance, the term “stimulus” in Psychology often refers to stimuli used in experiments, or that are perceived by participants, which is a distinct concept from usage in the Management and Organization science literature: specific fiscal interventions to induce more economic activity. To navigate these terminology barriers and build understandings of new domains, researchers must do difficult and time-consuming “translation work” [48].

Some research has approached addressing these terminology barriers as a text-simplification problem, exploring techniques such as text summarization/simplification to *remove* jargon. For instance, [5] generates plain language summaries and key questions to help lay readers understand medical research papers. While simplification is an appropriate way to make research more accessible to lay users, we worry that it might hamper the creative potential of juxtaposing the nuances of specialized concepts from a different field with concepts you may already be familiar with. Other approaches involve finding analogical relationships between papers, focusing on overlaps in problems or solutions between authors [50] or research papers [7, 25]. This approach is well-suited for cross-domain *problem solving*, but may be less suited for more conceptual or theoretical explorations.

In this paper, we frame the term mapping problem in scholarly cross-domain information seeking as a *translation* problem, taking

our cue from research on translation work in interdisciplinary research [48]. Rather than removing jargon, the goal in this framing is to identify novel and fruitful conceptual alignments between specialized concepts that might lead to novel insights. With this framing, we investigate how to adapt Natural Language Processing (NLP) techniques that seek to align distinct conceptual spaces, whether across languages, or between distinct time periods or sub-communities using the same language. For instance, multilingual NLP techniques such as machine translation have been used to not only translate between two very different languages, but also quite similar ones such as dialects [59]. It may also be applied to “translation” within the same language but diachronically (e.g. ancient to modern Chinese, [9]) or synchronically (e.g. across partisan divides [29]). We hypothesize that we might similarly be able to treat different research communities as speaking different languages and apply multilingual NLP techniques to assist with translating concepts between scholarly domains.

As a first step, we explore how to adapt techniques from unsupervised cross-lingual alignment of word embeddings to explore conceptual alignments between domain-specific word embedding spaces. This family of techniques aligns a source monolingual embedding space to a target monolingual embedding space. This alignment then allows for computing word similarities across the two embedding spaces. Here, we build “monolingual” embedding spaces for individual research communities, and then construct computational alignments between the communities. We explore two well-known techniques for cross-lingual embedding alignment — MUSE [12], and VecMap [3].

To investigate the promise and challenges of this approach, we developed a prototype cross-domain search engine that uses aligned research-domain-specific embeddings to support conceptual exploration, and tested this prototype across two case studies with PhD-level interdisciplinary scholars exploring concepts across scholarly domains. In these case studies, cross-lingual alignments returned more novel and relevant alignments of concepts between domains compared to monolingual baselines, and comparable performance to a strong LLM baseline.

These results suggest that the analogy between multilingual NLP and cross-disciplinary translation work is fruitful and worth exploring further. We summarize computational and design implications from our design iterations and case studies for future interfaces that provide computational support for scholarly cross-domain information seeking.

In sum, in this paper, we contribute:

- **A novel computational approach for supporting cross-domain translation work**, based on an adaptation of unsupervised cross-lingual alignment of word embeddings from multilingual NLP
- **A prototype cross-domain search engine** that supports cross-domain translation work using cross-lingual alignment techniques
- **Empirical proof-of-concept from two case studies** with PhD-level interdisciplinary scholars that validate the promise of adapting multilingual NLP techniques for cross-domain translation work, and suggest future directions for developing this solution approach

We make our code and data publicly available.¹

2 Related Work

Our work advances research on intelligent systems for cross-domain and cross-disciplinary scholarly information seeking by conceptualizing cross-disciplinary information seeking as translation work. This conceptualization motivates us to adapt distributional semantics techniques for aligning terms across domains to support cross-domain information seeking.

2.1 Cross-Domain Information Seeking as Translation Work

Cross-domain information seeking has been extensively studied in the field of library and information science [47]. Early work documented how information *scatter* — the fragmentation of information across sources, organizations, and disciplinary clusters — poses substantial challenges to scholars and librarians who seek relevant information via browsing or directed search in databases [6]. The boundaries that need to be crossed to discover and integrate new ideas may be drawn at disciplinary boundaries, but are also often defined by “invisible colleges” [13] of scholars and/or works that cite each other [54]. These “invisible colleges” may partially cross disciplinary boundaries but cohere around common topics of interest and ways of approaching those topics [22], such as ethnic-racial studies. In this paper, we use the term “cross-domain information seeking” as an umbrella term to refer to the range of information seeking activities that cross disciplinary, sociological, paradigmatic, and other significant boundaries between research communities.

Palmer’s [45, 46] seminal studies of the information work of interdisciplinary scientists documented how scientists “import” relevant information from outside their specialties using a complex mix of processes and practices, such as browsing and consulting a variety of journals, cultivating networks of personal contacts across disciplines, and even embedding graduate students as intermediaries to other labs to identify and transfer tacit knowledge. Of particular interest is the framing of interdisciplinary information work as “translation work”: drawing on actor-network theory [34], Palmer describes how interdisciplinary scholars “define, interpret, and redefine new information, retaining essential elements of the original context while revising and reapplying it for their own purposes” [48, p. 107]. In this conceptualization of translation work, jargon can be a barrier, and less specialized language can facilitate broader cross-discipline engagement with ideas [39, 41], but jargon may also act as *bridges* to novel perspectives worth importing and integrating into new insights. We see a potentially fruitful analogy here to the ways in which variations in semantic overlap between languages can signal opportunities for conceptual expansion: for instance, Dörk and colleagues [16] drew on the French noun *flâneur* — which refers to an urban man who wanders through the boulevards and cafes of Paris without a driving goal other than experiencing city life — to develop a novel conceptual vision of information seeking that focuses on open-ended, enjoyable construction of explorable information spaces.

¹<https://github.com/oasisresearchlab/crossdisciplinary-embeddings>

We are motivated by this conceptualization of translation work to investigate a jargon-preserving computational approach to facilitating cross-domain information seeking, potentially revealing new and undiscovered connections between scholarly domains.

2.2 Intelligent Interactive Systems for Supporting Scholarly Information Seeking

The field of HCI and intelligent user interfaces is home to substantial, long-running bodies of work on intelligent systems for supporting scholarly information seeking, from citation recommendation and information retrieval approaches [7, 25, 26, 35, 44, 50], to novel search and sensemaking interfaces [10, 11, 15, 24, 27, 51]. Work in this area often deals with scholarly information seeking in a domain-centric or domain-agnostic manner: users are either explicitly or implicitly modeled as searching in a single search space of “the scientific literature”.

Our work is situated in a smaller body of work that aims to directly support information seeking that crosses domain or disciplinary boundaries. Some of this work uses citation-based approaches to identify partitions in the literature and enable users to bridge those partitions: for example, literature-based discovery systems [56] aim to discover conceptual chains that bridge disjoint citation networks. In a classic motivating success story for this approach, Swanson [58] identified a citation network A where “fish oil” often co-occurs with “blood viscosity”, and another citation network B where “blood viscosity” and “Raynaud’s syndrome” often co-occur; A and B lacked citation connections between them, and Swanson proposed that there may be a novel opportunity to explore how fish oil might treat Raynaud’s syndrome. Another example is Portenoy’s [50] system that also uses citation network gaps to recommend novel *authors* whose work a scholar might find inspirational.

Another line of work develops content-based approaches to bridging domain and disciplinary gaps. Some, for example, focus on simplifying or removing jargon, often to bridge gaps between expert and lay audiences [4]. Other systems draw on conceptualizations of analogical problem solving to identify *analogical* relationships between research papers that focus on shared problems in spite of or in opposition to domain or topical similarity, enabling discovery of cross-domain papers that may be relevant to a research problem [7, 25, 52]. Finally, recent work has begun to explore the use of Large Language Models (LLMs) to support cross-disciplinary information seeking; these approaches prompt LLMs to generate research questions or queries from the perspective of other disciplines or domains, which can then be used as inspiration or queries for information seeking [37, 61].

Our work is most closely related to content-based approaches to supporting cross-domain information seeking. We extend this body of work by focusing on identifying novel and interesting specialized *concepts* (vs. authors or papers) across domain boundaries.

2.3 Distributional Semantics Methods for Aligning Terms across Domains

We explore how our goal of identifying interesting conceptual alignments between scholarly domains can draw on distributional

semantics methods from NLP that learn computational representations for words based on the context of how they co-occur with other words. This process is motivated by the “distributional hypothesis,” [53] which hypothesizes that *distributional similarity* is correlated with *meaning similarity*, so one could be approximated by the other. The learned distributional representations of words as vectors enables the use of measures like cosine similarity between word vectors to discover related terms.

For our purposes, building a single representation of word meaning is likely suboptimal, since it would fail to account for variations in language use across specific domains or communities. We thus draw on methods for inducing *alignments* between two distinct vector spaces [3, 12, 31, 40, 49, 60]. These methods have been applied to discover cross-language translations of words or produce mappings for bilingual lexicon induction [2, 23, 33, 43], which power a range of applications, such as cross-lingual information retrieval systems [36], where a user can input a query in one language, but receive relevant results from documents in other languages, and studying semantic change of terms over time within the same language [19, 30, 32, 55], or between different sub-communities of the same language, such as across partisan divides [29, 42].

Because we rarely expect there to be abundant parallel data between pairs of scholarly domains, and do expect domain-specific corpora to be relatively small (if the domains in question are quite granular), we specifically draw on *unsupervised* alignment methods that are robust to small corpora concerns in low-resource settings, such as MUSE [12] and VecMap [3].

3 System Design

Here, we discuss the design and implementation of a prototype cross-domain retrieval system that leverages embedding alignment techniques. The goal of these systems is to enable a user to input a term from a “home” domain, and explore sets of terms from a different “target” domain that may map to the same or related concepts as the home term. We overview the functionality of the retrieval system, and then describe how we train and align domain-specific embeddings to facilitate target-community retrieval.

3.1 Cross-Domain Search Engine based on Aligned Domain-Specific Embeddings

Figure 1 shows a prototype search engine that could provide this functionality. It enables users to select a “home” and “target” research domain, to explore how a specific user-defined term from that home domain might align with specialized terms from the target domain.

The search for related terms from the target community draws on the vocabulary from the target community’s domain-specific corpus, and ranks the terms from that vocabulary based on their cosine similarities between their embeddings and the query term in the home-target aligned space. To facilitate sensemaking of results, we retrieve links anchored with context for the term.

Query

Select your Home Community: psychology / creativity research

Select a Target Community: management / org science

Enter a term in your Home Community: stimulus

Query

Main results (debug)

People in [management / org science] may use a different term when they talk about something conceptually close to [stimulus]. We don't even know if such a term exists! But these may be close...

trajectories Toggle context

- [Trajtenberg et al. 1997, abstract] We find also evidence of the existence of 'technologi trajectories'.
- [ALKAERSIG et al. 2017, abstract] Consistently with recent advances in the literature, we argue that multiple technological trajectories can coexist in a field; therefore, firms may contribute to technological development by recombining in novel ways the capabilities that are widespread in the field, or by building novel and rare capabilities.
- [ALKAERSIG et al. 2017, Introduction] Therefore, the efforts that individual firms undertake to deviate from a given technological path, may generate a variety of persisting technological trajectories.
- [Gilsing et al. 2008, Control variables] However, they usually experience problems in diversifying into new technological areas inhibiting experimentation and favoring specialization along existing technological trajectories (Levinthal and March, 1993; March, 1991; Ahuja and Lampert, 2001).
- [Gilsing et al. 2008, Insert figure 2 here] This finding is also in line with the organizational learning literature: large established organizations have difficulties in diversifying into new technological areas, inhibiting experimentation and favoring specialization along existing technological trajectories (Levinthal and March, 1993; March, 1991; Ahuja and Lampert, 2001).

spillovers Toggle context

varies Toggle context

cultures Toggle context

domestic Toggle context

histories Toggle context

Figure 1: Term search user interface for exploring potential home → target bridges

3.2 Training and Aligning Domain-Specific Word Embeddings

To train domain-specific embeddings, we start by constructing separate corpora for different research domains. Motivated by conceptualizations of research domains as not just disciplines, but potentially “invisible colleges” of researchers cohering around common approaches to shared research interests [13], we define domain-specific corpora in terms of citation networks, rather than overall subject headings. Specifically, we seed a research domain by collecting a set of “seed” papers for that domain collaboratively with the researchers, and then define a corpus for that research domain by sampling additional papers through citation expansion.

Once we have defined domain-specific corpora, we train domain-specific word embeddings for each domain corpus. For each domain $d_s \in D$, where D refers to all the domains encompassing the set of papers collected for a specific researcher, we train a word embedding model, resulting in an embedding space E_{d_s} , which maps words in that domain to vectors in a continuous space. Mathematically, the embedding for domain d_i is a mapping $E_{d_s} : V_{d_s} \rightarrow \mathbb{R}^k$ from V_{d_i} , the vocabulary of domain d_i to a k -dimensional vector space, where each $E_{d_i}(w)$ is a vector in \mathbb{R}^k that embeds the usage patterns of the word $w \in V_{d_i}$.

Our goal is to learn an alignment function a such that $E_{d_t} = a(E_{d_s})$, effectively transforming the embeddings from one domain to align with those of another. In this work, we consider alignment functions a that are linear transformations of the original embedding space and can be intuitively interpreted as rotations of the embedding space E_{d_s} to better align with the space in E_{d_t} . Figure 3

provides an illustrative example of how similar terms from different domains may become closer in the aligned embedding space.

In our experiments, we trained fastText² embeddings with the skipgram model, a neural network model that learns word embeddings by learning to predict words in a context window around a given input target word.

We then apply computational techniques to learn our alignment function a to align pairs of domain-specific embeddings into a shared embedding space. Within this space, we can then explore how terms from one domain-specific corpus — defined by its embedding in the aligned space — relates other terms in another domain-specific corpus via metrics such as cosine similarity to the embeddings for those terms in the aligned space.

We draw our techniques from work on building cross-lingual embeddings (CLE). Specifically, we adapt VecMap [3] and MUSE [33], both of which learn orthogonal transformation matrices that maintain the internal structure of each monolingual space while aligning them within a shared cross-lingual space. Figure 3 shows that after applying cross-lingual alignment between domain-specific embeddings for “Psychology” and “Management / Organization Science”, the embedding for the term “examples” from Psychology could be moved to be closer to the term “substitutes” from Management / Org Science in the aligned space. This alignment could inspire new conceptualizations of the process of creating from examples, by considering dynamics of negotiation, competition, and strategy (which are central to the concept of strategic substitutes). We visualize this entire process in Figure 2.

²<https://fasttext.cc>

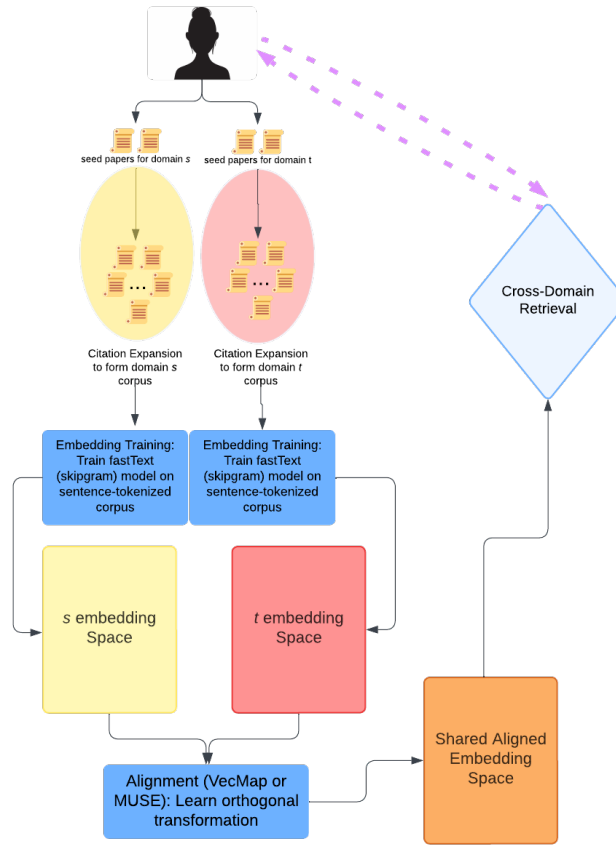


Figure 2: Overview of the pipeline from researcher-provided seed papers, through corpus development, mono-domain embedding training, and alignment to support retrieval across domains.

4 Case Study 1

To explore the potential of our alignment approach, we first investigate the novelty and relevance of terms from a target domain retrieved in response to a query term from a “home” domain using a MUSE-aligned shared space between the domains. We compare these results to terms retrieved by other (monolingual) baseline approaches. We chose an example setting of a researcher studying creativity from the perspective of Psychology, seeking cross-disciplinary insights from Management and Organization Science.

4.1 Scenario and Domain Corpora Preparation

We build a corpora of papers with the S2ORC dataset [38]. In this study, a domain is defined by a set of papers relevant to “creativity” research in Psychology. To build the corpora, the fourth author provided seed papers in Psychology and in Management. A list of all seed papers can be found in the accompanying supplemental materials. We then enlarged the pool of papers for each community through citation expansion: we retrieved papers from a single level of both incoming and outgoing citations from each seed paper. Some statistics of the corpora are provided in Table 1.

We constructed a scenario where the psychology researcher is seeking interesting parallels to the concepts of “stimulus” and “examples” from their “home” community of Psychology in the “target” community of Management and Organization Science. The researcher would review a set of terms (in the context of the sentences they were used in) and assess their degree of *relevance* (alignment with their home concept) and *novelty* (degree to which it expanded their thinking about the topic). We chose this initial setting because the fourth author of this work has a background in psychology and is studying creativity, and would thus be able to judge the novelty and relevance of retrieved terms.

4.2 Comparison Points and Evaluation

To evaluate how our MUSE adaptation might perform compared to existing approaches to search, we constructed an evaluation set consisting of aligned terms and their context from our MUSE approach, as well as 3 baselines that implement some variant of “monolingual” embeddings:

- (1) *fasttext-target*: using monolingual fastText embeddings trained on the whole corpus (across communities), but retrieving only from the target corpus. This baseline is a direct

Community	# seed papers	# papers after expansion	# sentences	# tokens
psychology / creativity research	7	659	77.9k	47.1k
management / org science	8	280	30.4k	27.3k

Table 1: Corpus statistics for the community corpora used in case study 1




Figure 3: Visual example using a cross-lingual alignment method to align two research communities (Psychology and Management / Organization Science research on creativity and innovation). The intuition is to find a “rotation” that can align two domain-specific embedding spaces such that terms that are semantically similar across the two embedding spaces are close together in the aligned space. Here we show an example from our data, where the MUSE alignment mapped the term “examples” (from Psychology) to the conceptually related term “substitutes” (from Management) in the aligned space. Visual of rotating spaces adapted from [12].

comparison to the MUSE pipeline, with the only difference being the embeddings

- (2) fasttext-combined: using monolingual fastText embeddings trained on the whole corpus, and retrieving from both the home and target corpus, and
- (3) SBERT: using an off-the-shelf contextual word embedding, retrieving from both the home and target corpus

We expect fasttext-target and fasttext-combined to perform poorly relative to MUSE, but for different reasons: we expect fasttext-target to return results that are contaminated by terms that are similar on the surface, but actually refer to different concepts in the respective scholarly domain; in contrast, we expect fasttext-combined to retrieve mostly direct matches from the home community. SBERT is a strong baseline, to explore whether existing contextual embedding approaches might be sufficient for teasing out cross-disciplinary nuances in the use of terms.

Each retrieval pipeline returned 50 (term, context sentence) pairs, where the term occurs in the context sentence. Details of each pipeline follow.

4.2.1 MUSE, target corpus. In this pipeline, we leverage the MUSE aligned embedding space. Given the user’s query term, the 10 most similar terms in the target space are obtained. Then, for each of the 10 expanded terms, 5 context sentences are retrieved from the target community corpus. Each paper is allowed to contribute up to 2 context sentences; once the maximum is reached, the pipeline would

move on to the next paper that has sentences with occurrences of the expanded term.

4.2.2 BASELINE 1: fastText, target corpus. This condition differs from MUSE in that the similar terms are obtained from fastText embeddings. The context sentence retrieval part is almost the same, except that we may go beyond the top 10 expanded term and use the 11th, 12th, ... similar terms, until we get 50 context sentences. The context sentence search is over the target community corpus.

4.2.3 BASELINE 2: fastText, combined corpus. This is another fastText baseline, which is the same as above except that the context sentence search is over the combined corpus containing papers of all communities.

4.2.4 BASELINE 3: Off-the-shelf, contextual (SBERT), combined corpus. This is an off-the-shelf search baseline which does not have the term expansion step. In contrast, both the query term and all the context sentences are encoded with contextualized embeddings. We retrieve a list of sentence contexts from the target community, then sort it by SBERT cosine similarity with the query term. To fit the (response term, context sentence) format, we assign the word in the context sentence that has the highest universal fastText embedding similarity to the query term in the “retrieved sentence” as the “response term”. We use the SBERT toolkit³ and adopt the all-mpnet-base-v2⁴ model. We consider this as a suitable contextualized embedding model for comparison for this task because SBERT’s training data contains the S2ORC dataset [38].

4.3 Relevance and Novelty Judgments

To enable judgments of the results while being naive to their source, we mixed the results from each pipeline into a spreadsheet. The author with expertise in creativity then rated each of the 200 term occurrences (i.e., a term and its context sentence), blind to pipeline source, along two dimensions: 1) *relevance* (0 to 2) and *novelty* (0 to 2).

4.4 Results

Figure 4 and Figure 5 show scatterplots of the relevant and novelty ratings for the query results from each pipeline, with the corresponding terms plotted on the scatterplot for clarity. Each term’s relevance or novelty rating is averaged over ratings provided to its specific occurrences in context sentences. Note that the number of terms for SBERT is often less than the other pipelines, since retrieval is focused on *sentences* (and then extracting the most relevant terms), rather than terms directly. These plots show that, on average, the results from MUSE were rated as more novel, while maintaining a high degree of relevance. In contrast, the fastText and SBERT baselines primarily returned relevant results with low novelty. This

³<https://www.sbert.net>

⁴<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

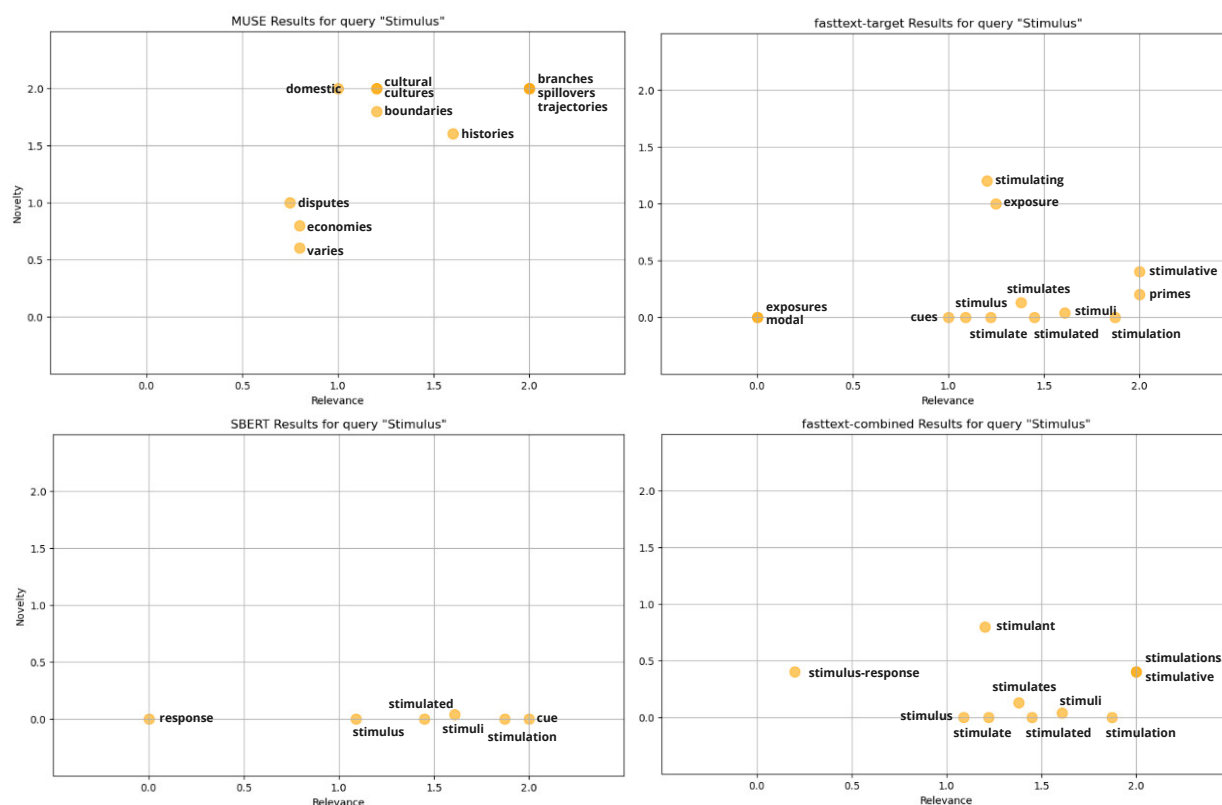


Figure 4: Results and corresponding relevance and novelty ratings from each pipeline for the query: Psychology[stimulus] → Management. More terms in the MUSE pipeline were rated as having both higher relevance and novelty, compared to other pipelines, which were highly relevant but not as novel.

advantage of MUSE over the monolingual baselines was strongly pronounced for the “stimulus” query. For the “examples” query, MUSE results are higher variance, with high novelty/relevance terms like (evolutionary) “pockets”, and (strategic) “substitutes”, but also noisy matches to irrelevant terms like “notebook” and “buyers”. But the quantitative pattern of high relevance but low novelty holds for the monolingual baselines.

Qualitatively, the monolingual baselines generally returned results that are very similar, and in many cases, linguistic variants of the original query (e.g., stimuli, stimulus, cue). In some cases, these surface form variations were actually less relevant to the original query term sense: for instance, a substantial number of the results from the monolingual baselines returned context sentences that used the term “examples” in a generic sense, not the specific sense of examples used to stimulate creativity. In contrast, the most relevant results from MUSE were frequently specialized terms that are quite different conceptually from the original query term. For example, for the query “stimulus”, MUSE returned specialized terms like “spillovers” (referring to transfer or R&D knowledge across industries), “cultures” (outlining a different dimension of knowledge transfer), and “trajectories” (emphasizing the conceptual history of knowledge that can be transferred across domains). For the query

“examples”, key results included specialized terms like (coevolutionary) “pockets” (connecting the idea of example-based inspiration to genetic algorithms and evolutionary dynamics), and (strategic) “substitutes” (evoking processes of negotiating strategic interests around substitutes for assets and resources in management science).

5 Case Study 2

To further explore the potential of our alignment approach, we conducted a second case study with an interdisciplinary researcher in Applied Linguistics, Human-Computer Interaction, and Education, interested in target disciplines of Ethnic-Racial Identity, Professional Identity, and Immigrant Studies. We report the community corpora in Table 2.

We assess the validity of our tool with a think-aloud study with our prototype cross-domain search engine. In addition to using the prototype to evaluate term alignments, we also use it as a design probe to exploring how our research participant uses the tool to assess related concepts in a target domain. As in Case Study 1, we constructed a scenario where this researcher would see parallels to predefined concepts; here, we use the concepts of “Education[third

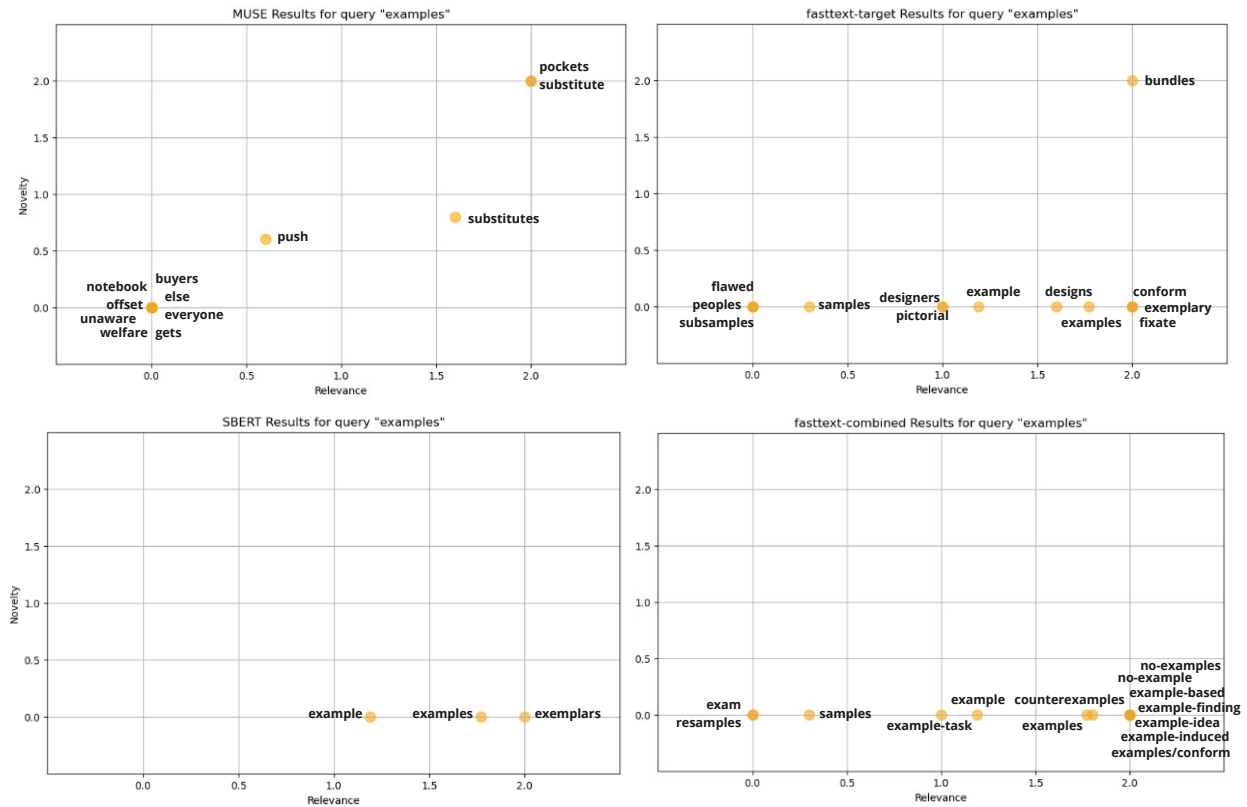


Figure 5: Results and corresponding relevance and novelty ratings from each retrieval pipeline for the query: Psychology[examples] → Management. Across retrieval systems, many of the terms were relevant but not very novel.

space]" and "Applied Linguistics[translanguaging]", to explore conceptual alignments with research communities of Ethnic-Racial Identity, Professional Identity, and Education.

5.1 Scenario and Domain Corpora Preparation

We developed the information-seeking scenario and domain corpora for this case study in collaboration with Jane Foster⁵, a PhD student in Human-Computer Interaction with strongly interdisciplinary research interests. The broad question that drove her research was *"How to communicate across different languages for different social and communicative purposes, while being mindful of linguistic dominance dynamics"*? Subthemes and areas of interest for this broad question included how language use relates to dynamics of identity expression, developing intersubjectivity, and transferring knowledge across different language communities.

In a series of initial conversations with Jane Foster, we identified three research communities that she considered her "home communities": 1) HCI, language technologies, and machine translation (hci), 2) applied linguistics (ling), and 3) language education (edu). Each of these home communities contained core specialized

concepts that were foundational to her research: for instance, "AI-mediated communication" was a central concept of interest in HCI, while concepts like "translanguaging" and "linguistic repertoire" were core ideas in applied linguistics, and "funds of knowledge" and "third space" were central ideas from (language) education. She described being interested in exploring connections of these concepts and her research questions to other research domains, such as ethnic and racial studies (eth), immigrant studies (imm), and professional identity development (pro) in interdisciplinary settings. These research communities thus served as the "target domains" for this case study.

We seeded Jane Foster's home domain corpora by asking her to identify a set of seed papers for the key concepts she identified in each home domain (along with any other background readings she found important). We repeated this process for the target domains of ethnic and racial studies and immigrant studies in a similar manner: seed papers for ethnic and racial studies were drawn from an annotated bibliography for an advanced PhD student whose dissertation was on identity development amongst immigrant teens; and seed papers for immigrant studies were drawn from an annotated bibliography for a research project on immigrant acculturation, led by an Assistant Professor whose research focuses on information ethics and justice, and social inclusion, for immigrants from the

⁵Pseudonym to protect participant anonymity

African, Afro-Caribbean, and Afro-Latinx diaspora in the US. The PhD student and Assistant Professor are departmental colleagues of the fourth author. The fourth author, whose core expertise includes interdisciplinarity, then created a reading list on professional identity development in interdisciplinary settings. They used the LitMaps⁶ tool to build out a set of readings in the successively expanded citation networks for a seed set of papers on professional identity in interdisciplinary teams. The seed papers for all domains studied here are provided in the supplementary materials.

From these seed papers, we then developed domain corpora for each target community by gathering forward and backward citations from each community’s seed papers using the Semantic Scholar Academic Graph API⁷. Table 2 shows the final corpus statistics for these home and target domain corpora.

We computed normalized modularity score [18] and present them in Table 3 for several community pairings explored in Case Studies 1 and 2.

5.2 Design Iterations

We expand on the methodology used in Case Study 1 (section 4) to further explore the application of cross-lingual embeddings for facilitating cross-disciplinary translation. Using MUSE, we started by aligning these embeddings to create the shared vector space where concepts from one domain can be translated into related concepts in a target domain.

Qualitative Sanity Check. To understand how well the MUSE system’s term mappings aligned concepts between scholarly domains, we conducted a qualitative sanity check. This involved generating a set of term translations using MUSE for several key concepts extracted from both domains. After manually reviewing the generated term mappings we discovered that the mappings did not exhibit the level of coherence we were aiming for, and many terms seemed to lack meaningful counterparts in the target domain. For instance, terms central to one domain’s field of study were often mapped to generic terms and concepts in the target domain, rather than particularly interesting or novel concepts. For example, we observed that *Education[funds of knowledge] → Ethnic-Racial Identity* would map to terms like “social psychology”, “socio-cognitive”, and “developmental processes”. We also observed broken mappings, where target-mapped terms would not be meaningful, e.g. *Human-Computer Interaction[AI-mediated communication] → Education[j.]*, where “j.” is a response term.

Pivot to VecMap. We decided to leverage the same monolingual fastText embeddings but experiment with VecMap [3]. VecMap has been shown in cross-lingual contexts to be more robust than MUSE [14], which can suffer from instability and stochasticity. Qualitatively, we saw more potentially relevant and interesting conceptual mappings with VecMap, and fewer obviously irrelevant mappings. For example, *ling[repertoire] → pro[lifeworld]* and *ling[repertoire] → pro[sensemaking]*. Both “sensemaking” and “lifeworld” can be highly related conceptually to “repertoire” because these concepts deal with the way people understand and interpret the world informally. We detail complementary quantitative metrics to assess

the strength of the alignment through two reference-free metrics: Normalized Modularity [18] and Normalized Cosine Similarity, as suggested by [20].

Normalized Modularity. Normalized modularity is a reference-free evaluation metric for assessing cross-lingual word embeddings. It works by measuring how well word vectors from different languages mix in a shared embedding space; the motivation is to detect if intra-lingual similarity (if words are more similar within-language) is greater than cross-lingual similarity (if words are more similar across languages), which is an indication that the embeddings fail to align the semantic meanings of words across languages.

The cross-lingual word embeddings are represented as a graph where words are nodes and edges represent similarity between word vectors using cosine similarity. Modularity then measures the difference between the proportion of edges that connect nodes within the same language, and the random expected proportion of these edges. A high modularity value indicates the presence of distinct language-specific clusters, suggesting poor cross-lingual alignment. A low modularity value indicates weaker separation between languages, suggesting better cross-lingual alignment.

$$Q = (e_s - a_s^2) + (e_t - a_t^2) \quad (1)$$

We use Equation 5.2 from [18] to compute normalized modularity scores between pairings of domains, where e_i refers to the fraction of edges that connect words in the embedding space of domain i , and a_i refers to the expected number of edges within the embedding space of domain i . This score is then normalized:

$$Q_{norm} = \frac{Q}{1 - (a_s^2 + a_t^2)}$$

Table 3 presents the results for VecMap and MUSE embeddings across a range of source and target pairings. First, we see that HCI $\rightarrow ?$ tends to have a stronger alignment with MUSE. For Applied Linguistics $\rightarrow ?$, VecMap generally shows better alignment, with the exception of \rightarrow HCI. Finally, for Education $\rightarrow ?$, the alignment is roughly equivalent between VecMap and MUSE. All source-target pairings in the queries we test in Case Study 2 indicate lower normalized modularity with VecMap.

Salient Cosine Similarity. We also computed cosine similarity of domain-salient terms and their translations in a target domain as an intrinsic evaluation metric, while normalizing similarities by the average cosine similarity of the embedding space following suggestions from [20]. This allows us to assess how tightly terms are mapped to similar terms in the shared embedding space. Our goal was to ensure that the learned representations were tightly aligned to enable similarity computations between embeddings in the domain pairing.

To support this evaluation, we extracted the top-100 salient, topical terms from each domain’s corpus using Latent Dirichlet Allocation (LDA) to use as tests. These terms reflect key concepts within each domain’s scholarly discourse. The system then translated these terms into their nearest equivalents in the target domain space using either MUSE or VecMap embeddings.

We compute the cosine similarity for each source and target term pair. We normalize this score by the average similarity score of all source-target translation pairs for all vocabulary terms in

⁶<https://www.litmaps.com/>

⁷<https://www.semanticscholar.org/product/api>

Community	# seed papers	# papers after expansion	# sentences	# tokens
hci	28	1.7k	9.8k	15.7k
ling	28	8.4k	43.7k	53.4k
edu	46	4.5k	22.7k	32.1k
eth	21	6.4k	34.1k	36.5k
pro	55	4.4k	24.4k	30.7k
imm	43	8.2k	45.85k	47.8k

Table 2: Corpus statistics for the community corpora used in case study 2




Figure 6: Pairwise domain evaluation of MUSE and VecMap using normalized cosine similarity on the salient topical terms from the source domain. Salient similarity scores are notably higher with the VecMap alignments, than with the MUSE alignments.

the source domain. This quantifies the overall performance of the system’s framework in creating semantically meaningful mappings.

We found consistently higher cosine similarity scores for VecMap compared to MUSE (see Figure 6). This complements the other qualitative and quantitative metrics to suggest that the MUSE alignment approach may have struggled to generate reliable mappings compared to VecMap, and provide complementary validation of our choice to run Case Study 2 with alignments from VecMap.

Qualitative Assessments and Query Selection for Think-Aloud Study. To get a sense of how effectively the system handles the queries we test across different target domains, we extracted and qualitatively assessed 10 response terms for each of the following queries, for all target domains:

- Education: funds of knowledge
- Education: third space
- Applied Linguistics: translanguaging
- Applied Linguistics: repertoire
- HCI: AI-mediated communication

Two authors of this paper independently rated each response term using discrete labels: **-1: Not relevant at all**, **0: Some relevance but generic or unsure**, **1: Very relevant and sensible**.

The Pearson correlation of ratings was 0.53, indicating moderate agreement. Per source domain, target domain, and query, we computed the average of those ratings. We also marked each result as a “potential hit” if either rater judged it as potentially very relevant and sensible (i.e., a rating of 1). We then summed the number of “potential hits” in each home-query-target domain triple to identify the top four most relevant queries to test in Case Study 2 with the interdisciplinary researcher. We did this to prioritize “recall” of potentially useful alignments to validate and discuss with Jane Foster in the case study. As Table 4 shows, these were:

- Education (edu): third space → Ethnic-racial Identity (eth) (5 potential hits)
- Education: third space → Professional Identity (pro) (4 potential hits)
- Applied Linguistics (ling): translanguaging → Education (5 potential hits)
- Applied Linguistics: translanguaging → Ethnic-racial Identity (7 potential hits)

We selected domain-query pairs with a higher proportion of potential hits because terms that were irrelevant or too broad were quite straightforward to recognize, but we wanted expert validation of the terms that we saw as potentially meaningful analogs in the

Source Domain	Target Domain	VecMap Modularity ↓	MUSE Modularity ↓
hci	ling	0.289	0.231
	edu	0.265	0.223
	imm	0.333	0.228
	eth	0.297	0.228
	pro	0.229	0.188
ling	hci	0.276	0.259
	edu	0.189	0.269
	imm	0.237	0.279
	eth	0.279	0.282
	pro	0.225	0.281
edu	hci	0.268	0.228
	ling	0.189	0.182
	imm	0.250	0.254
	eth	0.204	0.240
	pro	0.127	0.274
psy	mgmt	0.469	0.261

Table 3: Normalized Modularity Comparison between VecMap and MUSE. Lower modularity scores indicate better cross-lingual alignment, and is correlated with higher downstream performance [18].

target domains, and could seed meaningful discussion on cross-domain exploration.

5.3 Study Design

System Comparison. In our second case study, we introduced a baseline comparison using GPT-4o-mini alongside the VecMap system. Our goal in comparing to GPT is not to assess if our system could outperform an LLM, but rather to use it as a point of comparison for an alternative approach that relies on much larger scale data and compute, and for which we already have some evidence of multilingual capabilities [1, 28].

We conducted a blind comparison between the two systems using a repeated measures experimental design with our interview participant, aiming to assess their identification of sensible, and sensible+new terms. For each query, our participant viewed the results of each system in a counterbalanced order to mitigate ordering effects. This experiment design allows us to directly compare the performance of the two systems in generating meaningful and relevant terms across different interdisciplinary domains.

The results for each query were shown in our prototype cross-domain search engine interface Figure 1: both GPT and VecMap results could be displayed in an identical fashion, showing a ranked list of terms and a set of context sentences for each term, along with URLs to the research papers that contain those sentences, for more context. For each set of results for a query, we asked the participant to identify and discuss any terms that were sensible and sensible+new. After this stage, we allowed for free play of the VecMap system to spur needfinding discussion on the cross-domain information seeking task. The entire case study was conducted in-person, and lasted approximately 60 minutes. During this time, the participant’s screen and discussion when navigating the system

was recorded using Zoom. The recordings were later transcribed using MacWhisper⁸ for analysis.

We discuss in more detail how the retrieval pipeline for the two systems are implemented below, and then jump into the results.

VecMap Alignment System. With this system, we leverage the VecMap-aligned embedding space. Given the user’s query term, the K=10 most similar terms in the target space is obtained. Then, for each of the 10 response terms, up to 5 context sentences are retrieved from the target domain corpus of abstracts. Each paper is constrained to at most 2 context sentences.

GPT-4 Baseline. We initially prompted GPT4o-mini to generate related concepts from the target community, but it sometimes produced out-of-vocabulary terms or combined words in unintended ways. Since our search engine relies on adherence to the target community’s terminology, we shifted from open-ended generation to constrained selection. Prompt:

You are a researcher in the field of {SOURCE_COMMUNITY}, researching the concept of {HOME_CONCEPT}. Find 10 similar concepts to this in the field of {TARGET_COMMUNITY}, separated by string comma. You are ONLY allowed to select from the following terms: {candidate_terms}. Do not combine different words in this vocabulary, you are ONLY allowed to select from them.

We note however that this does require larger context windows, to fit the entire vocabulary into the prompt.

5.4 Results: Think-Aloud

We discuss the results from the case study as described in subsection 5.3. In tables 5 to 8, we present as a reference the list of terms

⁸<https://goodsnooze.gumroad.com/l/macwhisper>

Home	Query	Target	Avg. Rating	# Potential Hits	Potential Hits
edu	funds of knowledge	eth	-0.45	1	group__identity
		hci	-0.65	0	
		imm	-0.20	2	conceptual__framework, tridimensional
		ling	-0.39	1	multimodality
		pro	-0.11	1	habitus
	third space	eth	0.2	5	social__identity__theories, social__identity__theory, social__identity__approach, constructionist, identity__negotiation
		hci	-0.67	1	intersectional
		imm	-0.30	2	TDE, information__worlds
		ling	-0.25	3	multi/plurilingual continua, intertextuality
		pro	-0.05	4	CoPs, brokerage, boundary-crossing, brokering
hci	AI-mediated communication	edu	-0.75	1	nepantla
		eth	-1.00	0	
		imm	-0.50	1	e-diaspora
		ling	-0.45	2	transnational, transnationalism
		pro	-0.75	0	
ling	repertoire	edu	-0.35	3	repertoires, repertoire, tacit
		eth	-0.15	4	heritage__culture, worldview, subculture, heritage
		hci	-0.95	0	
		imm	-0.45	1	host__culture
		pro	-0.22	3	repertoire, lifeworld, sensemaking
	translanguaging	edu	0.03	5	translanguaging, teacher-talk, transliteracy, translingual
		eth	0.45	7	transcultural, cultural__identity, EICM, multicultural__education, bicultural__identity, ethno-cultural, acculturation__process
		hci	-0.17	3	mixed-initiative, cooperative, co-creative
		imm	-0.45	0	
		pro	0.00	3	situated__learning, CID, situated__learning__theory

Table 4: Qualitative assessment results terms returned for all home-domain[query]-target-domain pairs tested for Case Study 2, assessed by two authors of this paper. To prioritize recall for validation in the think-aloud, potential hits are defined as results where at least one rater gave a 1 (very relevant and sensible).

that the participant indicated were sensible and sensible + new, broken down by GPT and VecMap-based systems. We discuss further insights from each query in sections 5.4.1 to 5.4.4.

5.4.1 Query Insights: Education[third space]→ Ethnic-racial Identity. (Table 5)

Jane Foster began by mapping the concept of “third space” from Education to Ethnic-racial Identity. She emphasized that VecMap’s results were more focused and contextually relevant. Terms like “social identity”, “identity development”, and “cultural socialization” were sensible and aligned well with her understanding of “third

spaces” as informal settings of daily life interactions. In contrast, the terms returned by the GPT system, while relevant, focused more on “racial” and “ethnic” aspects of identity formation. While intrigued by the term **hybrid**, she felt that its association with identity was too distinct from the concept of third spaces:

*I think this one surprised me a little. I never considered **hybrid** to modify identities. . . . if it’s used in a general sense, let’s say social interaction, it may potentially map to “third space.” Because it can be a **hybrid** mode of interaction, like it’s not entirely in-person or it’s not entirely online... Or it’s a **hybrid** between family space and other, non-family space. Here, what I’m seeing is*

Education[third space] → Ethnic-racial Identity		
	VecMap	GPT
Sensible + New	<ul style="list-style-type: none"> The article concludes that national character is a sociocultural and psychological indicator that influences identity construction in political autobiographies. 	<ul style="list-style-type: none"> Parental cultural socialization and general parenting quality are important predictors of ethnic identity (EI) development in adolescents.
Sensible	<ul style="list-style-type: none"> It offers a nuanced understanding of the author’s career development journey to an authentic work identity. This analytic autoethnography, situated in multicultural, democratic South Africa, describes how historic moments in the country’s political evolution influenced the author personally . . . I am a Black South African; I am a gay professional and so, who am I at work? This research reexamines common conceptualizations of ethnic identity as a relatively stable, dispositional-like construct—in contrast to constructionist theories that emphasize the contextual, situational, and transitory nature of ethnic identity—and integrates the two frameworks. The current study explored ways schools facilitate supportive or marginalizing experiences for first generation Arab heritage youth in the United States and investigated how these experiences impact acculturative experiences and identity negotiation for these students. The social identity framework suggests that exposure to high-status ingroup or low-status outgroup portrayals enhances self-esteem through positive ingroup distinctiveness. 	<ul style="list-style-type: none"> The social identity framework suggests that exposure to high-status ingroup or low-status outgroup portrayals enhances self-esteem through positive ingroup distinctiveness. When non-native English speakers enter the classroom, immediate differences in language, cultural values, and peer interactions impact identity development and negatively influence self-perceptions.

Table 5: Terms selected by Jane Foster as sensible, and sensible+new for Education[third space] → Ethnic-racial Identity. Terms supplied by the retrieval system (either GPT or VecMap) are provided in-context, using just one representative line for demonstration purposes.

Education[third space] → Professional Identity		
	VecMap	GPT
Sensible + New	<ul style="list-style-type: none"> Rooted in a sociological model illustrating interacting levels of power at macro-, meso-, and microlevels, we argue that authentic research-practice partnerships and clinician-researcher collaborations can mitigate this power differential. 	...
Sensible	<ul style="list-style-type: none"> Drawing on social learning theory, and in particular on communities of practice (CoPs), the results of the study unpack how the digital nature of online peer discussion groups enables three interrelated learning mechanisms (sharing, symphonizing and shaping). This study explores how doctors perform this bridging work, conceptualized as cultural brokerage, on the job. Using a boundary-crossing lens, the article analyses the workplace arrangement, in which students’ learning across professional boundaries and their negotiations around a boundary object were prerequisites to coordinate their interprofessional knowledge and manage emerging challenges while being in charge of care on the ward. This article argues that time plays a crucial role of brokering disparate knowledge spaces in Indigenous discourses of postgenomics, with both Indigenous cosmological frames and DOHaD/epigenetics centring a circular temporal model. 	...

Table 6: Terms selected by Jane Foster as sensible, and sensible+new for Education[third space] → Professional Identity. Terms supplied by the retrieval system (either GPT or VecMap) are provided in-context, using just one representative line for demonstration purposes.

a **hybrid** type of identity. So in the identity sense, I wouldn’t use this term... But for modality of interaction, yeah.

5.4.2 Query Insights: Education[third space] → Professional Identity. (Table 6)

In the second query, Jane Foster again explored the concept of “third space” in education, but this time in Professional Identity. She

noted that to her, she would not map “third spaces” to professional identity in general, because “third spaces” are life-relevant, situated in informal, everyday environments, while professional identity is shaped in more structured, work-related contexts.

Nevertheless, VecMap yielded terms such as **research-practice**, **brokering**, **brokerage**, and **boundary-crossing**, which resonated with her idea of third spaces as transitional zones:

Applied Linguistics[translanguaging] → Education

	VecMap	GPT
Sensible + New	<ul style="list-style-type: none"> Seven bilingual teacher candidates implemented a strategic transliteracy approach to writing instruction, observing student assets and teaching cross-linguistically. Students can develop a translingual awareness that perceives languages as plural and fluid, and this awareness can help students succeed outside of monolingual norms. 	<ul style="list-style-type: none"> Seven bilingual teacher candidates implemented a strategic Transliteracy approach to writing instruction, observing student assets and teaching cross-linguistically. The present paper aims at finding out what modifications in translation of visitors' rules may be necessary if the target readership is to be considered, and at explicating the translation process through applying a transdisciplinary perspective of ELF studies, linguistic landscape (LL) studies, cross-field studies on conceptualization, translanguaging and translation studies. For this purpose, the text addresses, first, the theoretical developments of the decolonial turn; secondly, it discusses the notion of critical interculturality; and finally, it presents a reflection on the links between critical interculturality, decolonial pedagogy, and a vision of teacher education as a dynamics of personal development.
Sensible	<ul style="list-style-type: none"> The mothers and their daughters drew on their everyday ways of being (e.g., playfulness, translanguaging) in their mathematical interactions. We argue that our cases show how it is possible to bridge the gap between heteroglossic conceptions of language and language in applied linguistics, and what is conceived as legitimate language practices in the classroom. 	<ul style="list-style-type: none"> The mothers and their daughters drew on their everyday ways of being (e.g., playfulness, translanguaging) in their mathematical interactions.

Table 7: Terms selected by Jane Foster as sensible, and sensible+new for Applied Linguistics[translanguaging] → Education. Terms supplied by the retrieval system (either GPT or VecMap) are provided in-context, using just one representative line for demonstration purposes.

Applied Linguistics[translanguaging] → Ethnic-racial Identity

	VecMap	GPT
Sensible + New	<ul style="list-style-type: none"> This study assessed the ability of a set of practice-based measures to identify indicators of positive transcultural parenting. Becoming multicultural through early immersive culture mixing (EICM)—i.e., growing up with a mix of cultures that coexist and interact to form an emergent hybrid culture within one's home—is a rapidly rising phenomenon in many parts of the world. 	...
Sensible	<ul style="list-style-type: none"> The acculturation aspect examined unique manifestations of conflict related to the acculturation process of immigrant families. 	...

Table 8: Terms selected by Jane Foster as sensible, and sensible+new for Applied Linguistics[translanguaging] → Ethnic-racial Identity. Terms supplied by the retrieval system (either GPT or VecMap) are provided in-context, using just one representative line for demonstration purposes.

*Ah, I see ... This is sort of a third space, it's between research and practice, somewhere in between ... I have not seen **research-practice** before ... If it's **brokering** or **boundary-crossing**, then it's between two different workplaces...*

GPT, however, returned terms closely tied to professional settings that were less relevant because they did not capture the nuances of the life-relevant aspects central to “third spaces”, contrasting them with the terms she found from the VecMap system:

I think because these terms focus more on the professional... I'd consider that most of them in the “workplace,” rather than the “third space.” Unless it touches upon “research-practices” where it maps to another space or in-between... Or if it's “brokering” or “boundary-crossing,”

then it's between two different... [work]spaces. I think if it's professional identity-related, I wouldn't consider these sensible.

I think the only thing is identity formation, so I want to see more in how it is used, and whether it really has an analogical relationship with third space.

5.4.3 Query Insights: Applied Linguistics[translanguaging] → Education. (Table 7)

For the third query, Jane Foster explored how “translanguaging” mapped to terms in Education. Both systems contributed useful terms, but with different perspectives: VecMap focused more on learning-related concepts, while GPT focused on multilingualism.

While neither system was fully comprehensive, both offered inspiring concepts for the participant. Jane Foster highlighted that

translanguaging relates to integrating separate linguistic systems within a person's repertoire, which maps closely to terms like "transliteracy", "transdisciplinary", and "interculturality". She noted that related terms like "code-switching", although relevant, refers to a distinct linguistic concept from translanguaging.

Same-term responses prompted Jane Foster to consider it in different contexts:

It returned translanguaging... but it might be a different context!... Oh it's the same one. I think it's sensible because it's the same one so I'll still put it down... It didn't change my thoughts on translanguaging.

Jane Foster was particularly interested in **transliteracy**, **transdisciplinary**, and **interculturality**, which she felt captured the "translanguaging" concept, contrasting those terms with *distinct* terms such as **code-switching**:

*I see a very strong analogical connection, because everything is integrated ... people are just drawing from this one repertoire. When I am considering **transliteracy**, **transdisciplinary**, **interculturality**, it's really: how do we pool everything together and draw connections between them? I didn't include **bilingual** / **pluralism**. These are relevant terms but are different concepts from "translanguaging"... People would think it equates to **code-switching**, but I don't see it that way. These terms in general emphasize: you speak two languages, and how do you switch back and forth? Instead of seeing everything as one unit.*

Jane Foster also assessed the nuances of the "translanguaging" term, noticing that some of the response terms were too broad to capture the original concept:

*I think they're losing the core idea of "translanguaging". "Translanguaging" COULD be a strategy used to **scaffolding** learning, but "scaffold" would map to "translanguaging" as a "pedagogy" instead of "translanguaging" itself.*

*I could see "translanguaging" as a practice of **teacher-talk**, but it doesn't map there. It's hierarchical... it's **teacher-talk** then underneath it's "translanguaging."*

Jane Foster found both systems useful but in different ways. The VecMap system returned terms that were more focused on learning, while the GPT-based system selected terms that were more aligned with language and multi-linguism. While neither system was fully comprehensive, both offered inspiring concepts for the participant.

5.4.4 Query Insights: Applied Linguistics[translanguaging] → Ethnic-racial Identity. (Table 8)

In the last query, Jane Foster explored how translanguaging maps to Ethnic-racial Identity. While VecMap did not surface many interesting terms, **EICM** (early immersive culture mixing) in particular captured her interest. She perceived the term as suggesting an environment where cultural elements are deeply intertwined, making it difficult to separate distinct cultural identities. Within the sentence context of **EICM**, "hybrid" resurfaced, connecting back to her first query, which prompted her to reflect further on how hybrid cultures manifest:

I feel like I need to do more reading on "hybrid". I'm intrigued by this term in a neutral way. I can't see how

it would be related, but because I have seen this... Hybrid of multiple identities; that is a clear delineation. But this is an emergent hybrid. It's hard to see one person's cultural root and then separate it; for example, that part belongs to US culture, that belongs to Chinese culture; so they are using a different.. way [of looking at it].

As for GPT, thematically, these terms were considered too broad:

This focuses a lot on identity in general. Because it's general, it has lost the nuances of "translanguaging"...

When prompted to compare the two systems, Jane Foster consistently favored the VecMap-based system, noting that the system was able to deliver more research-focused terms.

5.4.5 Free Play. For the free play portion of the system, Jane Foster explored **Linguistics[translanguaging]** → **HCI** using just the VecMap system. Her goal was to understand how the HCI community discusses the concept of translanguaging and how it might be applied in design contexts.

Some of the response terms did not capture the fully-formed use cases the participant had in mind. For example, **speculative** appeared instead of **speculative design**. Moreover, the system did not offer direct, concrete applications for design. While these terms do not directly map to the core concept of "translanguaging", they still could be useful to describe how one might approach studying it in HCI, serving as a guide to consider different ways to approach "translanguaging" methodologically.

*I feel like it gives me a methodological approach, rather than a design application. I don't necessarily think [**speculative**] maps to the concept [of "translanguaging"], but it's useful to consider my approach... **co-creative** could also be a participatory design method. Reading this list, it seems that HCI would consider "translanguaging" as an open question. That there is not a specific or concrete claim yet, but it maps to a series of features of this term that guide the methodological choices to study this term in HCI.*

5.4.6 Design Feedback. Discussion with the participant after reviewing queries yielded insights about how a system that enables cross-domain alignment of concepts might support specific translation work scenarios.

One use case the participant mentioned for exploring unfamiliar domains was to investigate how theories they are familiar with are specifically applied in different target contexts.

If I have good familiarity on a term, I want to see what are the other claims surrounding this topic just so I would understand this term more directly. My perspective is this: this is how people in Linguistics lay claims on this term, and this is the social science theory to understand HCI problems.

*I want to know how "translanguaging" could be implemented. I know well about the theory, but outside of the classroom, how do we do that? With these terms like **transliteracy**, I want to see how these sort of actions of pooling resources together would work in other contexts,*

and to see if there are any strategies or behaviors or interventions I could sort of mimic.

In this scenario, the scholar would query the term “translanguaging” from Applied Linguistics → Education to get a sense of how the concept of translanguaging is studied in education contexts. This use case can be handled in current search systems in the case that the specific term is used in the exact same way in the target domain.

A second use case focuses on identifying concepts that are *similar* to a specific concept of interest. As an example, the participant described the need to find theoretical backings cited in their home community.

So now I would want to see a connection of how I could link the social science theory back to HCI, and I want to get a sense of... “have similar things been done in HCI?” Or explore how the HCI community thinks of this concept, for example, translanguaging [from Linguistics].

This is akin to the idea of searching for similar theories that inform work done in applied domains, and writing related works to get a sense of what similar things were done in the past. Importantly, this use case of finding alignments between similar concepts may not be directly supported by techniques that gloss over jargon or domain differences in language use. Instead, methods like cross-domain alignment, as we explored here, might be more useful.

6 Discussion

In this paper, we explored how we might adapt techniques for cross-domain alignment of word embeddings to support translation work in cross-domain scholarly information seeking. We then tested this concept across two case studies with interdisciplinary scholars’ real cross-domain information seeking scenarios. In the first case study, training and aligning domain-specific word embeddings with MUSE enabled retrieval of more relevant **and** novel response alignments between terms in home and target research domains, compared to either the contextual SBERT approach and the combined-corpora modeling approaches (more relevant, less novel). In the second case study, training and aligning domain-specific word embeddings with VecMap enabled retrieval of equally or more relevant and novel alignments between terms in home and target research domains, compared to responses from GPT. Because LLMs have been trained over diverse, large-scale sources of data, we analogize this setting to the combined-corpora modeling approaches we tested in Case Study 1. However, the results of Case Study 2 point to off-the-shelf LLMs being unable to produce terms suitable for this cross-domain exploration task for scholars. We note the effectiveness of our prototype approach, despite the small size of the underlying corpora.

Together, these results suggest that modeling communities separately and then aligning them into a shared embedding space can preserve specialized knowledge and produce opportunities for interesting concept explorations. Overall, we view the quantitative and qualitative example results as a promising proof of concept that multilingual NLP approaches might be fruitfully applied to the scenario of cross-domain information seeking, by framing scholarly domains as distinct language-using communities.

6.1 Limitations and Future Work

Our case studies provide high ecological validity and qualitative depth, due to their grounding in real cross-domain information seeking scenarios and direct involvement of interdisciplinary scholars; but more precise quantitative estimates of efficacy across additional settings are needed to establish generalizability and robustness.

Another important next step will be to systematically investigate the conditions under which different alignment-based approaches are likely to be fruitful. Our investigation here suggested that VecMap was a more effective alignment technique for the Case Study 2 corpora compared to MUSE, but we have little direct insight into precisely why this is the case. Further, having seen the stronger results for VecMap in Case Study 2, we attempted to replicate and improve on the results of Case Study 1 with VecMap replacing MUSE as the alignment technique: unfortunately, the resulting alignment failed to produce similarly relevant and novel terms as MUSE. Investigating the reasons for these performance differences is an important next step. This could begin by investigating specific assumptions made by alignment techniques about the size, distribution, and overlap between domain corpora, but also investigate how these interact with specific alignment technique hyperparameters, such as the choice of vocabulary size and dimensionality of embeddings to keep.

These investigations will be greatly enhanced by the availability of reliable intrinsic, automated evaluation metrics to correlate with the extrinsic cross-domain translation task we are interested in. In our investigations, we adapted evaluation metrics that show moderate correlation with cross-lingual tasks, but there is no prior evidence on their effectiveness in this specialized task. Thus, our choice to pivot to VecMap for Case Study 2 was primarily driven by our qualitative manual inspection of query results: the cosine similarity of salient terms and normalized modularity metrics [18] were interpreted overall as a complementary validation of our choice, rather than a fine-grained evaluation of the quality of specific domain alignments. We also observed some correlations between these metrics and the qualitative assessment ratings, in theoretically predicted directions: normalized modularity was negatively correlated with average qualitative rating scores, $r = -.37$, and salient cosine similarity was positively correlated, $r = .57$. These observations suggest these metrics may be useful for automatically evaluating alignment quality for this task. However, we are hesitant to interpret these as strong quantitative results, given the small N (15 total domain pairs, 1-2 queries per pairing), and the lack of gold standard ratings for the query results (since the domain expert only directly validated four sets of query results). Thus, there is still a need to directly validate alignment metrics for this specific application of aligning different scholarly domains.

Finally, while both our case studies benefit from the involvement of PhD-level scholars with the relevant interdisciplinary expertise, the selection process introduces certain limitations. We chose these two participants based on their professional relevance to different case study domains, limiting our ability to generalize the findings from the qualitative evaluation of the cross-disciplinary term alignments to the diversity of scholars in these fields.

6.2 Additional Multilingual NLP Approaches to Explore for Cross-Domain Scholarly Translation Work

In this paper, we tested the concept of alignment-based approaches with the simplest instantiation: linear alignment of monolingual embeddings, and computing similarities between word embeddings in the shared space. Despite the promising examples of useful alignments, the quantitative trends (e.g., 50-60% potential hits at 10 results on the upper bound, but a number of cases with little to no strong potential results) indicate the difficulty of the task. Thus, we are excited to explore additional extensions and multilingual NLP approaches that could support cross-domain scholarly translation work.

One possible direction is to move away from a strictly unsupervised setting, and adapt additional techniques from multilingual alignment to improve alignment quality, both at the alignment stage and the retrieval stage. First, we might be able to enhance the quality of our cross-domain scholarly alignment by injecting supervision into learning better representations. For example, labeled pairs of specialized terms from different domains can help bootstrap better alignments. Additionally, semi-supervised techniques such as self-training could help expand coverage of cross-domain mappings, allowing us to benefit from the flexibility of the unsupervised approach and the effectiveness of more supervised techniques. From a retrieval perspective, we could integrate techniques from informational retrieval, such as learning-to-rank, or ensembling our distributional semantics-based approach with traditional retrieval methods (e.g. BM25 and TF-IDF). Another direction might be to use provided or learned [17] hierarchies from ontologies or controlled vocabularies, to focus retrieval on more specific (vs. broader) domain-specific terms.

Finally, it may be fruitful to explore multilingual approaches powered by LLMs. In Case Study 2, we tested a prototype translation-like prompting approach with GPT-4o to select potentially relevant and novel terms within the constraints of our setup. While initial results seemed promising, the think-aloud results revealed that terms were still too imprecise and broad. Natural extensions could be to refine the prompt design to explore how to select novel and sensible terms; some prompting-based methods could range from: few-shot prompting, chain-of-thought, to prompting pipelines on top of embedding-based retrieval.

6.3 Implications for Intelligent Interactive Systems for Supporting Cross-Domain Scholarly Information Seeking

From a systems perspective, we are keen to explore new interface designs for reading and searching that draw on multilingual techniques to support cross-domain information seeking. One example might be a “translate” button in a reading interface: after providing information about their “home” domain, scholars could select text portions that are unfamiliar to them, right click to show the menu, and then click “translate” to see term mapping to their home domain. This could be implemented in a user interface similar to the ScholarPHI system for augmenting scientific papers with position-sensitive *definitions* of mathematical terms and symbols

[21]. Another approach might be to augment existing search systems with a list of cross-aligned terms that act as filters for exploring portions of a large unfamiliar corpus of search results that might most closely align with their home domain concepts. In either approach, we could consider exposing relevant vocabulary terms in both source and target domains to guide natural user queries.

In keeping with the richness of translation work as potentially navigating and integrating multiple related novel concepts from a new domain, rather than translating individual terms into their direct equivalents, we also wonder about alternative designs that operate over *collections* of alignments, rather than individual lists of search results. Such designs may align better with our Case Study 2 participant’s preference for making sense of collections of terms, rather than individual terms. In particular, the participant wanted to get a sense of all the response terms at once, before being able to assess how relevant or novel a term was. Judging in-silo versus all-at-once is certainly a design consideration. The former might lead to more criticality and premature rejection, which could lead to frustration if many returned terms are not relevant even if a minority are fruitful alignments to specialized concepts (akin to the precision versus recall tradeoff). The latter might give a user more of a mental sketch of grouping and clustering thematic conceptual mappings, better rewarding the experience given the presence of fruitful alignments, even if they are situated next to irrelevant concepts. The framing of the tool as a search interface might also lead users to desire more directly relevant terms or 1-to-1 mappings, given broader themes in HCI around how affordances of a search interface can substantially shape how people assess the relevance of cross-domain inspirations [25, 57], or use/adapt them for novel insights [8].

Overall, we believe framing interdisciplinary foraging in terms of *translation work*, in alignment with multilingual NLP, has strong potential to yield useful new designs for intelligent interactive systems that support cross-domain information-seeking.

References

- [1] Jordi Armengol-Estapé, Ona de Gibert Bonet, and Maite Melero. 2022. On the Multilingual Capabilities of Very Large-Scale English Language Models. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declercq, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Jan Odijk, and Stelios Piperidis (Eds.). European Language Resources Association, Marseille, France, 3056–3068. <https://aclanthology.org/2022.lrec-1.327>
- [2] Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2016. Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Jian Su, Kevin Duh, and Xavier Carreras (Eds.). Association for Computational Linguistics, Austin, Texas, 2289–2294. <https://doi.org/10.18653/v1/D16-1250>
- [3] Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Iryna Gurevych and Yusuke Miyao (Eds.). Association for Computational Linguistics, Melbourne, Australia, 789–798. <https://doi.org/10.18653/v1/P18-1073>
- [4] Tal August, Lucy Lu Wang, Jonathan Bragg, Marti A. Hearst, Andrew Head, and Kyle Lo. 2022. Paper Plain: Making Medical Research Papers Approachable to Healthcare Consumers with NLP. *ToCHI Volume 30* (Dec. 2022). <https://dl.acm.org/doi/10.1145/3589955>
- [5] Tal August, Lucy Lu Wang, Jonathan Bragg, Marti A. Hearst, Andrew Head, and Kyle Lo. 2023. Paper Plain: Making Medical Research Papers Approachable to Healthcare Consumers with Natural Language Processing. *ACM Transactions on Computer-Human Interaction* (April 2023). <https://doi.org/10.1145/3589955> Just Accepted.

- [6] Marcia J Bates. 1996. Learning about the information seeking of interdisciplinary scholars and students. *Library trends* 45, 2 (1996), 155–164. tex.publisher: Johns Hopkins University Press.
- [7] Joel Chan, Joseph Chee Chang, Tom Hope, Dafna Shahaf, and Aniket Kittur. 2018. SOLVENT: A Mixed Initiative System for Finding Analogies between Research Papers. *Proceedings of ACM Human-Computer Interaction: CSCW 2, CSCW* (2018), 31. <https://doi.org/10.1145/3274300>
- [8] Joel Chan, Zijian Ding, Eesh Kamrah, and Mark Fuge. 2024. Formulating or Fixating: Effects of Examples on Problem Solving Vary as a Function of Example Presentation Interface Design. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3613904.3642653>
- [9] Ernie Chang, Yow-Ting Shiu, Hui-Syuan Yeh, and Vera Demberg. 2021. Time-Aware Ancient Chinese Text Translation and Inference. In *Proceedings of the 2nd International Workshop on Computational Approaches to Historical Language Change 2021*, Nina Tahmasebi, Adam Jatowt, Yang Xu, Simon Hengchen, Syrielle Montariol, and Haim Dubossarsky (Eds.). Association for Computational Linguistics, Online, 1–6. <https://doi.org/10.18653/v1/2021.lchange-1.1>
- [10] Joseph Chee Chang, Amy X. Zhang, Jonathan Bragg, Andrew Head, Kyle Lo, Doug Downey, and Daniel S. Weld. 2023. CiteSee: Augmenting Citations in Scientific Papers with Persistent and Personalized Historical Context. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3544548.3580847>
- [11] Duen Horng Chau, Aniket Kittur, Jason I. Hong, and Christos Faloutsos. 2011. Apolo: Making Sense of Large Network Data by Combining Rich User Interaction and Machine Learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 167–176. <https://doi.org/10.1145/1978942.1978967> 00000.
- [12] Alexis Conneau, Guillaume Lample, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word Translation Without Parallel Data. *ArXiv abs/1710.04087* (2017).
- [13] Diana Crane. 1972. *Invisible Colleges; Diffusion of Knowledge in Scientific Communities*. Chicago: University of Chicago Press.
- [14] Yerai Doval, Jose Camacho-Collados, Luis Espinosa Anke, and Steven Schockaert. 2020. On the Robustness of Unsupervised and Semi-supervised Cross-lingual Word Embedding Learning. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis (Eds.). European Language Resources Association, Marseille, France, 4013–4023. <https://aclanthology.org/2020.lrec-1.495>
- [15] Cody Dunne, Ben Shneiderman, Robert Gove, Judith Klavans, and Bonnie Dorr. 2012. Rapid understanding of scientific paper collections: Integrating statistics, text analytics, and visualization. *Journal of the American Society for Information Science and Technology* 63, 12 (Dec. 2012), 2351–2369. <https://doi.org/10.1002/asi.22652>
- [16] Marian Dörk, Sheelagh Carpendale, and Carey Williamson. 2011. The information flaneur: a fresh look at information seeking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Vancouver BC Canada, 1215–1224. <https://doi.org/10.1145/1978942.1979124>
- [17] Lior Forer and Tom Hope. 2024. Inferring Scientific Cross-Document Coreference and Hierarchy with Definition-Augmented Relational Reasoning. <https://doi.org/10.48550/arXiv.2409.15113> arXiv:2409.15113 [cs].
- [18] Yoshinari Fujinuma, Jordan Boyd-Graber, and Michael J. Paul. 2019. A Resource-Free Evaluation Metric for Cross-Lingual Word Embeddings Based on Graph Modularity. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Anna Korhonen, David Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, Florence, Italy, 4952–4962. <https://doi.org/10.18653/v1/P19-1489>
- [19] William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2018. Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change. <http://arxiv.org/abs/1605.09096> arXiv:1605.09096 [cs].
- [20] Katharina Hämmerl, Jindřich Libovický, and Alexander Fraser. 2024. Understanding Cross-Lingual Alignment—A Survey. In *Findings of the Association for Computational Linguistics ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikrumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand and virtual meeting, 10922–10943. <https://doi.org/10.18653/v1/2024.findings-acl.649>
- [21] Andrew Head, Kyle Lo, Dongyeop Kang, Raymond Fok, Sam Skjonsberg, Daniel S. Weld, and Marti A. Hearst. 2021. Augmenting Scientific Papers with Just-in-Time, Position-Sensitive Definitions of Terms and Symbols. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. <https://doi.org/10.1145/3411764.3445648> 00006.
- [22] Birger Hjørland. 2002. Epistemology and the socio-cognitive perspective in information science. *Journal of the American Society for Information Science and Technology* 53, 4 (2002), 257–270. <https://doi.org/10.1002/asi.10042> 00509 _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/asi.10042>
- [23] Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. 2018. Loss in Translation: Learning Bilingual Word Mapping with a Retrieval Criterion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, Brussels, Belgium, 2979–2984. <https://doi.org/10.18653/v1/D18-1330>
- [24] Hyeonsu Kang, Joseph Chee Chang, Yongsung Kim, and Aniket Kittur. 2022. Threddy: An Interactive System for Personalized Thread-based Exploration and Organization of Scientific Literature. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3526113.3545660>
- [25] Hyeonsu B. Kang, Xin Qian, Tom Hope, Dafna Shahaf, Joel Chan, and Aniket Kittur. 2022. Augmenting Scientific Creativity with an Analogical Search Engine. *ACM Transactions on Computer-Human Interaction* (March 2022). <https://doi.org/10.1145/3530013> Just Accepted.
- [26] Hyeonsu B Kang, Nouran Soliman, Matt Latzke, Joseph Chee Chang, and Jonathan Bragg. 2023. ComLittee: Literature Discovery with Personal Elected Author Committees. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–20. <https://doi.org/10.1145/3544548.3581371>
- [27] Hyeonsu B Kang, Tongshuang Wu, Joseph Chee Chang, and Aniket Kittur. 2023. Synergi: A Mixed-Initiative System for Scholarly Synthesis and Sensemaking. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–19. <https://doi.org/10.1145/3586183.3606759>
- [28] Tannon Kew, Florian Schottmann, and Rico Sennrich. 2023. Turning English-centric LLMs Into Polyglots: How Much Multilinguality Is Needed? *ArXiv abs/2312.12683* (2023). <https://api.semanticscholar.org/CorpusID:266374637>
- [29] Ashiqur R. KhudaBukhs, Rupak Sarkar, Mark S. Kamlet, and Tom Michael Mitchell. 2020. We Don't Speak the Same Language: Interpreting Polarization through Machine Translation. In *AAAI Conference on Artificial Intelligence*.
- [30] Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015. Statistically Significant Detection of Linguistic Change. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 625–635. <https://doi.org/10.1145/2736277.2741627>
- [31] Saurabh Kulshreshtha, Jose Luis Redondo Garcia, and Ching-Yun Chang. 2020. Cross-lingual Alignment Methods for Multilingual BERT: A Comparative Study. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 933–942. <https://doi.org/10.18653/v1/2020.findings-emnlp.83>
- [32] A Kutuzov, L Øvrelid, T Szymanski, and E Velldal. 2018. Diachronic word embeddings and semantic shifts: a survey. (2018). <https://www.sciencedirect.com/science/article/pii/S1319157823001234#b0240>
- [33] Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc'Aurelio Ranzato. 2018. Unsupervised Machine Translation Using Monolingual Corpora Only. arXiv:1711.00043 [cs.CL] <https://arxiv.org/abs/1711.00043>
- [34] Bruno Latour. 2007. *Reassembling the Social: An Introduction to Actor-Network-Theory*. OUP Oxford. Google-Books-ID: BgJREAAQBAJ.
- [35] Yoonjoo Lee, Hyeonsu B Kang, Matt Latzke, Juho Kim, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2024. PaperWeaver: Enriching Topical Paper Alerts by Contextualizing Recommended Papers with User-collected Papers. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–19. <https://doi.org/10.1145/3613904.3642196>
- [36] Robert Litschko, Goran Glavaš, Simone Paolo Ponzetto, and Ivan Vulić. 2018. Unsupervised Cross-Lingual Information Retrieval using Monolingual Data Only. arXiv:1805.00879 [cs.CL] <https://arxiv.org/abs/1805.00879>
- [37] Yiren Liu, Pranav Sharma, Mehul Jitendra Oswal, Haijun Xia, and Yun Huang. 2024. PersonaFlow: Boosting Research Ideation with LLM-Simulated Expert Personas. <http://arxiv.org/abs/2409.12538> arXiv:2409.12538 [cs].
- [38] Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld. 2020. S2ORC: The Semantic Scholar Open Research Corpus. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 4969–4983. <https://doi.org/10.18653/v1/2020.acl-main.447>
- [39] Li Lucy, Jesse Dodge, David Bamman, and Katherine A. Keith. 2023. Words as Gatekeepers: Measuring Discipline-specific Terms and Meanings in Scholarly Publications. <https://doi.org/10.48550/arXiv.2212.09676> arXiv:2212.09676 [cs].
- [40] Kelly Marchisio, Neha Verma, Kevin Duh, and Philipp Koehn. 2022. IsoVec: Controlling the Relative Isomorphism of Word Embedding Spaces. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 6019–6033. <https://doi.org/10.18653/v1/2022.emnlp-main.404>

- [41] Peter McMahan and James Evans. 2018. Ambiguity and Engagement. *Amer. J. Sociology* 124, 3 (Nov. 2018), 860–912. <https://doi.org/10.1086/701298.00010>.
- [42] Jeremiah Milbauer, Adarsh Mathew, and James Evans. 2021. Aligning Multidimensional Worldviews and Discovering Ideological Differences. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 4832–4845. <https://doi.org/10.18653/v1/2021.emnlp-main.396>
- [43] Smaranda Muresan and Judith L. Klavans. 2013. Inducing terminologies from text: A case study for the consumer health domain. *Journal of the American Society for Information Science and Technology* 64, 4 (mar 2013), 727–744. <https://doi.org/10.1002/asi.22787>
- [44] Srishti Palani, Aakanksha Naik, Doug Downey, Amy X. Zhang, Jonathan Bragg, and Joseph Chee Chang. 2023. Relatedly: Scaffolding Literature Reviews with Existing Related Work Sections. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–20. <https://doi.org/10.1145/3544548.3580841>
- [45] Carole L. Palmer. 1999. Structures and strategies of interdisciplinary science. *Journal of the American Society for Information Science* 50, 3 (1999), 242–253. [https://doi.org/10.1002/\(SICI\)1097-4571\(1999\)50:3<242::AID-ASI7>3.0.CO;2-7](https://doi.org/10.1002/(SICI)1097-4571(1999)50:3<242::AID-ASI7>3.0.CO;2-7)
- [46] Carole L. Palmer. 2001. *Work at the Boundaries of Science: Information and the Interdisciplinary Research Process*. Springer Science & Business Media. Google-Books-ID: P20cBQAAQBAJ.
- [47] Carole L. Palmer. 2017. Information Research on Interdisciplinarity. In *The Oxford Handbook of Interdisciplinarity* (2 ed.), Robert Frodeman (Ed.). Oxford University Press, 429–442. <https://doi.org/10.1093/oxfordhb/9780198733522.013.35>
- [48] Carole L. Palmer and Neumann. 2002. The Information Work of Interdisciplinary Humanities Scholars: Exploration and Translation. *The Library Quarterly: Information, Community, Policy* 72, 1 (2002), 85–117. www.jstor.org/stable/4309582 Publisher: University of Chicago Press.
- [49] Barun Patra, Joel Ruben Antony Moniz, Sarthak Garg, Matthew R. Gormley, and Graham Neubig. 2019. Bilingual Lexicon Induction with Semi-supervision in Non-Isometric Embedding Spaces. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Anna Korhonen, David Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, Florence, Italy, 184–193. <https://doi.org/10.18653/v1/P19-1018>
- [50] Jason Portenoy, Marissa Radensky, Jevin D West, Eric Horvitz, Daniel S Weld, and Tom Hope. 2022. Bursting Scientific Filter Bubbles: Boosting Innovation via Novel Author Discovery. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3491102.3501905>
- [51] Napol Rachatasumrit, Jonathan Bragg, Amy X. Zhang, and Daniel S Weld. 2022. CiteRead: Integrating Localized Citation Contexts into Scientific Paper Reading. In *Proceedings of the 27th International Conference on Intelligent User Interfaces (IUI '22)*. Association for Computing Machinery, New York, NY, USA, 707–719. <https://doi.org/10.1145/3490099.3511162>
- [52] Marissa Radensky, Simra Shahid, Raymond Fok, Pao Siangliulue, Tom Hope, and Daniel S. Weld. 2024. Scideator: Human-LLM Scientific Idea Generation Grounded in Research-Paper Facet Recombination. <https://doi.org/10.48550/arXiv.2409.14634> arXiv:2409.14634 [cs].
- [53] Magnus Sahlgren. 2008. The Distributional Hypothesis. *The Italian Journal of Linguistics* 20 (2008), 33–54. <https://api.semanticscholar.org/CorpusID:23750999>
- [54] Pamela E. Sandstrom. 2001. Scholarly communication as a socioecological system. *Scientometrics* 50, 3 (Jan. 2001), 573–605. <https://doi.org/10.1023/A:1012786609066>
- [55] Dominik Schlechtweg, Anna Hatty, Marco del Tredici, and Sabine Schulte. 2019. A Wind of Change: Detecting and Evaluating Lexical Semantic Change across Times and Domains. *ACL* (2019). <https://aclanthology.org/P19-1072.pdf>
- [56] Yakub Sebastian, Eu-Gen Siew, and Sylvester O. Orimaye. 2017. Emerging approaches in literature-based discovery: techniques and performance review. *The Knowledge Engineering Review* 32 (Jan. 2017). <https://doi.org/10.1017/S0269888917000042.00000>
- [57] Arvind Srinivasan and Joel Chan. 2024. Improving Selection of Analogical Inspirations through Chunking and Recombination. In *Proceedings of the 16th Conference on Creativity & Cognition (C&C '24)*. Association for Computing Machinery, New York, NY, USA, 374–397. <https://doi.org/10.1145/3635636.3656207>
- [58] Don R. Swanson. 1986. Undiscovered Public Knowledge. *The Library Quarterly: Information, Community, Policy* 56, 2 (1986), 103–118. <https://doi.org/10.1086/601720.00000> Publisher: University of Chicago Press.
- [59] Yu Wan, Baosong Yang, Derek F. Wong, Lidia S. Chao, Haihua Du, and Ben C.H. Ao. 2019. Unsupervised Neural Dialect Translation with Commonality and Diversity Modeling. *ArXiv abs/1912.05134* (2019).
- [60] Ruochen Xu, Yiming Yang, Naoki Otani, and Yuexin Wu. 2018. Unsupervised Cross-lingual Transfer of Word Embedding Spaces. *ArXiv abs/1809.03633* (2018). <https://api.semanticscholar.org/CorpusID:52186890>
- [61] Chengbo Zheng, Yuanhao Zhang, Zeyu Huang, Chuhan Shi, Minrui Xu, and Xiaojuan Ma. 2024. DisciLink: Unfolding Interdisciplinary Information Seeking Process via Human-AI Co-Exploration. <https://doi.org/10.48550/arXiv.2408.00447>

arXiv:2408.00447 [cs].

A Qualitative Assessments of Mappings

A.1 Assessments for Case Study 1

query_term	cross_term	cossim	relevance	novelty
examples	gets	0.6043121219	0	0
examples	gets	0.6043121219	0	0
examples	gets	0.6043121219	0	0
examples	gets	0.6043121219	0	0
examples	gets	0.6043121219	0	0
examples	substitutes	0.5895254016	2	2
examples	substitutes	0.5895254016	2	2
examples	substitutes	0.5895254016	1	0
examples	substitutes	0.5895254016	1	0
examples	substitutes	0.5895254016	2	0
examples	notebook	0.5892814994	0	0
examples	notebook	0.5892814994	0	0
examples	pockets	0.5892053246	2	2
examples	pockets	0.5892053246	2	2
examples	welfare	0.5747747421	0	0
examples	welfare	0.5747747421	0	0
examples	welfare	0.5747747421	0	0
examples	welfare	0.5747747421	0	0
examples	welfare	0.5747747421	0	0
examples	substitute	0.5725717545	2	2
examples	substitute	0.5725717545	2	2
examples	substitute	0.5725717545	2	2
examples	substitute	0.5725717545	2	2
examples	substitute	0.5725717545	2	2
examples	unaware	0.5722236037	0	0
examples	unaware	0.5722236037	0	0
examples	unaware	0.5722236037	0	0
examples	unaware	0.5722236037	0	0
examples	unaware	0.5722236037	0	0
examples	push	0.5718544126	1	1
examples	push	0.5718544126	1	1
examples	push	0.5718544126	0	0
examples	push	0.5718544126	1	1
examples	push	0.5718544126	0	0
examples	offset	0.5687605739	0	0
examples	offset	0.5687605739	0	0
examples	offset	0.5687605739	0	0
examples	offset	0.5687605739	0	0
examples	else	0.5636897683	0	0
examples	else	0.5636897683	0	0
examples	else	0.5636897683	0	0
examples	else	0.5636897683	0	0
examples	else	0.5636897683	0	0
examples	buyers	0.5636130571	0	0
examples	buyers	0.5636130571	0	0
examples	buyers	0.5636130571	0	0
examples	buyers	0.5636130571	0	0
examples	buyers	0.5636130571	0	0
examples	everyone	0.5618861914	0	0
stimulus	trajectories	0.6212699413	2	2
stimulus	trajectories	0.6212699413	2	2
stimulus	trajectories	0.6212699413	2	2
stimulus	trajectories	0.6212699413	2	2

stimulus	trajectories	0.6212699413	2	2
stimulus	spillovers	0.6189801693	2	2
stimulus	spillovers	0.6189801693	2	2
stimulus	spillovers	0.6189801693	2	2
stimulus	spillovers	0.6189801693	2	2
stimulus	spillovers	0.6189801693	2	2
stimulus	varies	0.6030950546	2	2
stimulus	varies	0.6030950546	1	1
stimulus	varies	0.6030950546	1	0
stimulus	varies	0.6030950546	0	0
stimulus	varies	0.6030950546	0	0
stimulus	cultures	0.5986304879	2	2
stimulus	cultures	0.5986304879	1	2
stimulus	cultures	0.5986304879	1	2
stimulus	cultures	0.5986304879	1	2
stimulus	cultures	0.5986304879	1	2
stimulus	domestic	0.5788759589	1	2
stimulus	domestic	0.5788759589	1	2
stimulus	domestic	0.5788759589	1	2
stimulus	domestic	0.5788759589	1	2
stimulus	domestic	0.5788759589	1	2
stimulus	histories	0.5694535375	2	2
stimulus	histories	0.5694535375	2	2
stimulus	histories	0.5694535375	2	2
stimulus	histories	0.5694535375	1	1
stimulus	histories	0.5694535375	1	1
stimulus	cultural	0.563362658	2	2
stimulus	cultural	0.563362658	1	2
stimulus	cultural	0.563362658	1	2
stimulus	cultural	0.563362658	1	2
stimulus	cultural	0.563362658	1	2
stimulus	disputes	0.561938107	1	1
stimulus	disputes	0.561938107	1	1
stimulus	disputes	0.561938107	1	2
stimulus	disputes	0.561938107	0	0
stimulus	economies	0.5614743233	2	2
stimulus	economies	0.5614743233	1	1
stimulus	economies	0.5614743233	1	1
stimulus	economies	0.5614743233	0	0
stimulus	economies	0.5614743233	0	0
stimulus	boundaries	0.5581468344	2	1
stimulus	boundaries	0.5581468344	1	2
stimulus	boundaries	0.5581468344	1	2
stimulus	boundaries	0.5581468344	1	2
stimulus	boundaries	0.5581468344	1	2
stimulus	branches	0.5577203035	2	2

Table 9: Response terms for query terms of interest from Psychology → Management / Organizational Science, paired with cosine similarity and qualitative assessments from one annotator.

A.2 Assessments for Case Study 2

source	target	query_term	cross_term	similarity	annotator1	annotator2	avg
ling	hci	translanguaging	deliberative	0.8743163943	-1	0	-0.5
ling	hci	translanguaging	mixed-initiative	0.8533095717	1	1	1
ling	hci	translanguaging	initiative	0.8460384011	-1	0	-0.5
ling	hci	translanguaging	cooperative	0.8315920234	1	0	0.5
ling	hci	translanguaging	speculative	0.8281762004	-1	-1	-1
ling	hci	translanguaging	proactive	0.8109608293	-1	-1	-1
ling	hci	translanguaging	co-creative	0.8079094887	1	1	1
ling	hci	translanguaging	Creative	0.8072760105	-1	0	-0.5
ling	hci	translanguaging	adaptive	0.8055086136	-1	0	-0.5
ling	edu	translanguaging	translanguaging	0.790014565	0	1	0.5
ling	edu	translanguaging	teacher-talk	0.784806788	1	1	1
ling	edu	translanguaging	transliteracy	0.782607913	1	1	1
ling	edu	translanguaging	classroom-based	0.7754539251	-1	-1	-1
ling	edu	translanguaging	language-learning	0.7680151463	-1	0	-0.5
ling	edu	translanguaging	scaffold	0.7634933591	-1	0	-0.5
ling	edu	translanguaging	literacy	0.7614750266	-1	0.5	-0.25
ling	edu	translanguaging	translingual	0.7595765591	0	1	0.5
ling	edu	translanguaging	languaging	0.7590286136	0	1	0.5
ling	edu	translanguaging	teaching/learning	0.7577226758	-1	-1	-1
ling	imm	translanguaging	adapts	0.8113001585	-1	0	-0.5
ling	imm	translanguaging	machine	0.7610801458	-1	-1	-1
ling	imm	translanguaging	EBIs	0.7520438433	1	-1	0
ling	imm	translanguaging	mentoring	0.7489615083	-1	-1	-1
ling	imm	translanguaging	adaptations	0.747436583	-1	0	-0.5
ling	imm	translanguaging	impetus	0.7393995523	-1	-1	-1
ling	imm	translanguaging	adaption	0.7382391691	-1	0	-0.5
ling	imm	translanguaging	crosscultural	0.7377848029	0	0	0
ling	imm	translanguaging	intercultural	0.7375655174	0	0	0
ling	imm	translanguaging	multicultural	0.7371763587	0	0	0
ling	eth	translanguaging	transcultural	0.8437370062	1	1	1
ling	eth	translanguaging	cultural	0.8020462394	-1	0	-0.5
ling	eth	translanguaging	cultural_identity	0.7983759642	1	1	1
ling	eth	translanguaging	EICM	0.7838115692	1	1	1
ling	eth	translanguaging	multicultural_education	0.7721879482	0	1	0.5
ling	eth	translanguaging	bicultural_identity	0.771873951	1	1	1
ling	eth	translanguaging	ethno-cultural	0.7650336623	1	0	0.5
ling	eth	translanguaging	multi-cultural	0.7632349133	0	0	0
ling	eth	translanguaging	acculturation_process	0.7588860989	1	1	1
ling	eth	translanguaging	best_practices	0.757483542	-1	-1	-1
ling	pro	translanguaging	CLIL	0.8253515363	0	0	0
ling	pro	translanguaging	learning	0.8214411736	-1	0	-0.5
ling	pro	translanguaging	situated_learning	0.8092653155	1	0	0.5
ling	pro	translanguaging	CID	0.8076415658	1	1	1
ling	pro	translanguaging	pedagogic	0.7999457121	-1	0	-0.5
ling	pro	translanguaging	M-learning	0.7973175049	-1	0	-0.5
ling	pro	translanguaging	pedagogical	0.7970983386	-1	0	-0.5
ling	pro	translanguaging	e-learning	0.7967173457	0	0	0
ling	pro	translanguaging	student_learning	0.7891011238	0	0	0
ling	pro	translanguaging	situated_learning_theory	0.7856515646	1	0	0.5
ling	hci	repertoire	vision	0.6829593182	-1	-1	-1
ling	hci	repertoire	participation	0.6826735139	-1	0	-0.5
ling	hci	repertoire	Real-time	0.674511373	-1	-1	-1
ling	hci	repertoire	envision	0.6718433499	-1	-1	-1

ling	hci	repertoire	freedom	0.6680535674	-1	-1	-1
ling	hci	repertoire	path	0.6620223522	-1	-1	-1
ling	hci	repertoire	puppet	0.6607789993	-1	-1	-1
ling	hci	repertoire	self-control	0.6590454578	-1	-1	-1
ling	hci	repertoire	partner	0.6556383371	-1	-1	-1
ling	hci	repertoire	influential	0.6525364518	-1	-1	-1
ling	edu	repertoire	grounds	0.6947515607	-1	0	-0.5
ling	edu	repertoire	flows	0.6836689115	-1	0	-0.5
ling	edu	repertoire	unique	0.6813056469	-1	-1	-1
ling	edu	repertoire	repertoires	0.6811887026	0	1	0.5
ling	edu	repertoire	repertoire	0.6778647304	0	1	0.5
ling	edu	repertoire	resource	0.6778272986	-1	0	-0.5
ling	edu	repertoire	accumulated	0.6653056741	-1	0	-0.5
ling	edu	repertoire	reorient	0.6639546752	-1	-1	-1
ling	edu	repertoire	truths	0.6492972374	-1	-1	-1
ling	edu	repertoire	tacit	0.6480041742	0	1	0.5
ling	imm	repertoire	new__culture	0.7574108243	0	0	0
ling	imm	repertoire	hence	0.7198644876	-1	-1	-1
ling	imm	repertoire	adapts	0.7161327004	-1	0	-0.5
ling	imm	repertoire	successfully	0.7103006244	-1	-1	-1
ling	imm	repertoire	dominant__culture	0.7074056268	0	-1	-0.5
ling	imm	repertoire	s	0.7032558918	-1	-1	-1
ling	imm	repertoire	different__cultures	0.6992135644	0	0	0
ling	imm	repertoire	adaption	0.6937709451	-1	0	-0.5
ling	imm	repertoire	host__culture	0.6921414733	0	1	0.5
ling	imm	repertoire	navigation	0.6919963956	-1	0	-0.5
ling	eth	repertoire	self	0.7625374794	-1	0	-0.5
ling	eth	repertoire	heritage__culture	0.751629293	1	1	1
ling	eth	repertoire	worldview	0.7122991085	0	1	0.5
ling	eth	repertoire	sense	0.711098969	-1	0	-0.5
ling	eth	repertoire	subculture	0.7054489255	0	1	0.5
ling	eth	repertoire	tangible	0.704145968	-1	-1	-1
ling	eth	repertoire	heritage	0.6999636889	0	1	0.5
ling	eth	repertoire	socio-cultural	0.6926341653	-1	0	-0.5
ling	eth	repertoire	markers	0.69101125	-1	0	-0.5
ling	eth	repertoire	belongs	0.6903909445	-1	-1	-1
ling	pro	repertoire	ideology	0.7007747889	-1	0	-0.5
ling	pro	repertoire	repertoire	0.6773253679	0	1	0.5
ling	pro	repertoire	lifeworld	0.6770594716	1	1	1
ling	pro	repertoire	different	0.6749215126	-1	-1	-1
ling	pro	repertoire	frames	0.6726770997	-1	0	-0.5
ling	pro	repertoire	sensemaking	0.6652895808	1	0	0.5
ling	pro	repertoire	competent	0.6618661284	-1	0	-0.5
ling	pro	repertoire	manifest	0.6616282463	-1	-1	-1
ling	pro	repertoire	core	0.6598476768	-1	0	-0.5
edu	hci	funds__of__knowledge	platform	0.781198144	-1	0	-0.5
edu	hci	funds__of__knowledge	platforms	0.7583290339	-1	0	-0.5
edu	hci	funds__of__knowledge	forms	0.7142107487	-1	-1	-1
edu	hci	funds__of__knowledge	formats	0.7071314454	-1	-1	-1
edu	hci	funds__of__knowledge	links	0.7066740394	-1	0	-0.5
edu	hci	funds__of__knowledge	norm	0.7046322227	-1	0	-0.5
edu	hci	funds__of__knowledge	grief	0.6988732815	-1	-1	-1
edu	hci	funds__of__knowledge	format	0.697663188	-1	-1	-1
edu	hci	funds__of__knowledge	Algorithms	0.6881314516	0	0	0

edu	hci	funds_of_knowledge	theoretically	0.6832361221	-1	0	-0.5
edu	ling	funds_of_knowledge	theorize	0.7694742084	-1	0	-0.5
edu	ling	funds_of_knowledge	knowledge-rich	0.7588395476	-1	0	-0.5
edu	ling	funds_of_knowledge	conceptualized	0.7509846687	-1	0	-0.5
edu	ling	funds_of_knowledge	lens	0.7487373352	-1	0	-0.5
edu	ling	funds_of_knowledge	conceptualize	0.747408092	-1	0	-0.5
edu	ling	funds_of_knowledge	knowledges	0.725109756	-1	1	0
edu	ling	funds_of_knowledge	deploys	0.7211380601	-1	-1	-1
edu	ling	funds_of_knowledge	knowledge	0.7191357613	-1	0	-0.5
edu	ling	funds_of_knowledge	multimodality	0.711112082	0	1	0.5
edu	imm	funds_of_knowledge	conceptualize	0.7420441508	-1	0	-0.5
edu	imm	funds_of_knowledge	contextualize	0.730682373	-1	0	-0.5
edu	imm	funds_of_knowledge	contextualizes	0.7160271406	-1	0	-0.5
edu	imm	funds_of_knowledge	multidimensional	0.7149877548	0	0	0
edu	imm	funds_of_knowledge	acculturation_theory	0.7143118382	1	-1	0
edu	imm	funds_of_knowledge	conceptual_framework	0.714176178	1	0	0.5
edu	imm	funds_of_knowledge	tridimensional	0.7110076547	1	0	0.5
edu	imm	funds_of_knowledge	conceptualizes	0.7109881043	-1	0	-0.5
edu	imm	funds_of_knowledge	conceptual_model	0.7070842385	-1	0	-0.5
edu	imm	funds_of_knowledge	multidimensionality	0.7054414153	-1	0	-0.5
edu	eth	funds_of_knowledge	derives	0.7214701772	-1	-1	-1
edu	eth	funds_of_knowledge	conceptualize	0.7162780762	-1	0	-0.5
edu	eth	funds_of_knowledge	theoretical_approaches	0.7159905434	-1	0	-0.5
edu	eth	funds_of_knowledge	synthesize	0.7152888179	-1	0	-0.5
edu	eth	funds_of_knowledge	group_identity	0.7140477896	1	1	1
edu	eth	funds_of_knowledge	derived	0.7053064108	-1	-1	-1
edu	eth	funds_of_knowledge	existing	0.7047510743	-1	-1	-1
edu	eth	funds_of_knowledge	bridge	0.7008541226	-1	0	-0.5
edu	eth	funds_of_knowledge	theorize	0.6947138906	-1	0	-0.5
edu	eth	funds_of_knowledge	conceptualized	0.6915852427	0	0	0
edu	pro	funds_of_knowledge	conceptualised	0.7170065045	0	0	0
edu	pro	funds_of_knowledge	concepts	0.7028673291	0	0	0
edu	pro	funds_of_knowledge	conceptualize	0.7015826106	0	0	0
edu	pro	funds_of_knowledge	integrate	0.6972215772	0	0	0
edu	pro	funds_of_knowledge	concept	0.6937665939	-1	0	-0.5
edu	pro	funds_of_knowledge	habitus	0.6908137798	1	1	1
edu	pro	funds_of_knowledge	conceptualise	0.6904374957	-1	0	-0.5
edu	pro	funds_of_knowledge	understandings	0.6889657378	-1	0	-0.5
edu	pro	funds_of_knowledge	theory	0.6889147758	-1	0	-0.5
edu	hci	third_space	insights	0.7838318944	-1	-1	-1
edu	hci	third_space	intersectional	0.7603951097	0	1	0.5
edu	hci	third_space	lights	0.7572006583	-1	-1	-1
edu	hci	third_space	CBE	0.7424249053	0	-1	-0.5
edu	hci	third_space	future	0.7412995696	-1	-1	-1
edu	hci	third_space	puppet	0.7375090122	-1	-1	-1
edu	hci	third_space	instructors	0.734104991	-1	-1	-1
edu	hci	third_space	interdisciplinary	0.7302479744	0	0	0
edu	hci	third_space	discuss	0.728939414	-1	-1	-1
edu	ling	third_space	multi/plurilingual	0.7795246243	1	1	1
edu	ling	third_space	continua	0.766535759	1	0	0.5
edu	ling	third_space	readership	0.7573106289	-1	-1	-1
edu	ling	third_space	reconceptualise	0.7523953319	-1	0	-0.5
edu	ling	third_space	intertextuality	0.7505890131	1	1	1
edu	ling	third_space	conceptualized	0.7432819605	-1	-1	-1

edu	ling	third__space	concept	0.7429624796	-1	-1	-1
edu	ling	third__space	reconceptualization	0.7417594194	1	-1	0
edu	ling	third__space	conceptualizes	0.7408658266	-1	-1	-1
edu	ling	third__space	nexus	0.7386583686	-1	0	-0.5
edu	imm	third__space	TDE	0.7607435584	1	1	1
edu	imm	third__space	construction	0.7474842072	-1	-1	-1
edu	imm	third__space	theorised	0.7378210425	-1	-1	-1
edu	imm	third__space	concept	0.7370481491	-1	-1	-1
edu	imm	third__space	conceptual_framework	0.728639245	1	-1	0
edu	imm	third__space	proposal	0.7261104584	-1	-1	-1
edu	imm	third__space	information_worlds	0.7233423591	1	1	1
edu	imm	third__space	enrich	0.7202795744	-1	-1	-1
edu	imm	third__space	acculturation_theory	0.7174776196	0	0	0
edu	imm	third__space	reconstruction	0.716917634	1	-1	0
edu	eth	third__space	identity.	0.7945914865	-1	0	-0.5
edu	eth	third__space	identity__construction	0.7858144045	0	0	0
edu	eth	third__space	social_identity_theories	0.773149848	1	0	0.5
edu	eth	third__space	social_identity_theory	0.7703753114	1	0	0.5
edu	eth	third__space	social_identity_approach	0.7689462304	1	0	0.5
edu	eth	third__space	identity_theory	0.7616072297	0	0	0
edu	eth	third__space	constructionist	0.7581090927	1	0	0.5
edu	eth	third__space	hybrid	0.7478024364	-1	1	0
edu	eth	third__space	identity_negotiation	0.7414662242	1	0	0.5
edu	eth	third__space	social_identity	0.7385947704	0	0	0
edu	pro	third__space	CoPs	0.7541846633	1	1	1
edu	pro	third__space	entrepreneurship	0.7500813603	-1	0	-0.5
edu	pro	third__space	research-practice	0.7392860055	-1	1	0
edu	pro	third__space	conceptualized	0.7345022559	-1	-1	-1
edu	pro	third__space	distinctive	0.7331488132	-1	-1	-1
edu	pro	third__space	conceptualised	0.7331167459	-1	-1	-1
edu	pro	third__space	brokerage	0.7313123345	1	1	1
edu	pro	third__space	conceptualise	0.7266747952	-1	-1	-1
edu	pro	third__space	boundary-crossing	0.7237938046	1	1	1
edu	pro	third__space	brokering	0.7229532003	1	1	1
hci	ling	AI-mediated__communication	transnational	0.8087528944	1	1	1
hci	ling	AI-mediated__communication	transnationalism	0.8066139817	1	1	1
hci	ling	AI-mediated__communication	Everyday	0.7917954922	-1	-1	-1
hci	ling	AI-mediated__communication	Transnational	0.7847146392	-1	1	0
hci	ling	AI-mediated__communication	Looking	0.7837681174	-1	-1	-1
hci	ling	AI-mediated__communication	racialization	0.78137362	0	-1	-0.5
hci	ling	AI-mediated__communication	emergency	0.7741872072	-1	-1	-1
hci	ling	AI-mediated__communication	Masculinities	0.7690383196	-1	-1	-1
hci	ling	AI-mediated__communication	Identities	0.7651382089	-1	-1	-1
hci	edu	AI-mediated__communication	makerspaces	0.77935642	0	-1	-0.5
hci	edu	AI-mediated__communication	Partnerships	0.7601863146	-1	-1	-1
hci	edu	AI-mediated__communication	Inclusive	0.7550260425	-1	-1	-1
hci	edu	AI-mediated__communication	Immigrant	0.7508330941	-1	-1	-1
hci	edu	AI-mediated__communication	interrogates	0.7342031002	-1	-1	-1
hci	edu	AI-mediated__communication	Tongan	0.7277263999	-1	-1	-1
hci	edu	AI-mediated__communication	anti-immigrant	0.7198304534	-1	-1	-1
hci	edu	AI-mediated__communication	nepantla	0.7160634995	1	0	0.5
hci	edu	AI-mediated__communication	compassionate	0.7143685222	-1	-1	-1
hci	edu	AI-mediated__communication	Collaboration	0.71210289	-1	0	-0.5
hci	imm	AI-mediated__communication	placemaking	0.7857161164	1	-1	0

hci	imm	AI-mediated__communication	place-making	0.7733535171	1		1
hci	imm	AI-mediated__communication	HRM	0.7627471685	-1	-1	-1
hci	imm	AI-mediated__communication	Divakaruni	0.7560439706	-1	-1	-1
hci	imm	AI-mediated__communication	e-diaspora	0.7536911964	1	0	0.5
hci	imm	AI-mediated__communication	globalisation	0.7535432577	-1	0	-0.5
hci	imm	AI-mediated__communication	facilitating	0.7498921752	-1	0	-0.5
hci	imm	AI-mediated__communication	TDE	0.7489527464	0	0	0
hci	imm	AI-mediated__communication	talent	0.7473618984	-1	-1	-1
hci	imm	AI-mediated__communication	advent	0.7453843355	-1	-1	-1
hci	eth	AI-mediated__communication	wider	0.8147138357	-1	-1	-1
hci	eth	AI-mediated__communication	Chagrin	0.8083021641	-1	-1	-1
hci	eth	AI-mediated__communication	Falls	0.808240056	-1	-1	-1
hci	eth	AI-mediated__communication	landscape	0.7854062319	-1	-1	-1
hci	eth	AI-mediated__communication	landmark	0.7750949264	-1	-1	-1
hci	eth	AI-mediated__communication	agenda	0.7665138245	-1	-1	-1
hci	eth	AI-mediated__communication	BAM	0.7627230883	-1	-1	-1
hci	eth	AI-mediated__communication	scientists	0.7586004138	-1	-1	-1
hci	eth	AI-mediated__communication	institutional	0.7544520497	-1	-1	-1
hci	eth	AI-mediated__communication	higher__education	0.7536155581	-1	-1	-1
hci	pro	AI-mediated__communication	APRT	0.7847443819	-1	-1	-1
hci	pro	AI-mediated__communication	VAC	0.763710022	-1	-1	-1
hci	pro	AI-mediated__communication	SGM/DSD	0.7588301301	-1	-1	-1
hci	pro	AI-mediated__communication	Brazil	0.7540581226	-1	-1	-1
hci	pro	AI-mediated__communication	BGI	0.7479191422	-1	-1	-1
hci	pro	AI-mediated__communication	RCMI	0.7463548183	-1	-1	-1
hci	pro	AI-mediated__communication	SBE	0.7405706048	1	-1	0
hci	pro	AI-mediated__communication	GEM-SLP	0.7375261188	-1	-1	-1
hci	pro	AI-mediated__communication	wellbeing	0.7351579666	0	0	0
hci	pro	AI-mediated__communication	Delivery	0.7338174582	-1	0	-0.5

Table 10: Response terms for query terms of interest, paired with cosine similarity and qualitative assessments from two annotators.