



On How Cognitive Computing Will Plan Your Next Systematic Review

Maisie Badami¹(✉), Marcos Baez², Shayan Zamanirad¹, and Wei Kang³

¹ University of New South Wales (UNSW), Kensington, Australia
{m.badami, shayan.zamanirad}@unsw.edu.au

² LIRIS – University of Claude Bernard Lyon 1, Villeurbanne, France
marcos.baez@liris.cnrs.fr

³ Data61, CSIRO, Eveleigh, Australia
wei.kang@data61.csiro.au

Abstract. Systematic literature reviews (SLRs) are at the heart of evidence-based research, setting the foundation for future research and practice. However, producing good quality timely contributions is a challenging and highly cognitive endeavor, which has lately motivated the exploration of automation and support in the SLR process. In this paper we address an often overlooked phase in this process, that of planning literature reviews, and explore under the lenses of cognitive process augmentation how to overcome its most salient challenges. In doing so, we report on the insights from 24 SLR authors on planning practices, its challenges as well as feedback on support strategies inspired by recent advances in cognitive computing. We frame our findings under the cognitive augmentation framework, and report on a prototype implementation and evaluation focusing on further informing the technical feasibility.

Keywords: Systematic review · Cognitive process · Web services · Word embedding

1 Introduction

Systematic Literature Reviews (SLRs) are valuable research contributions that follow a well-known, comprehensive, and transparent research methodology. It is at the heart of evidence-based research, allowing researchers to systematically collect and integrate empirical evidence regarding research questions. Given their demonstrated value, SLRs are becoming an increasingly popular type of publication in different disciplines, from medicine to software engineering [11].

Despite the valuable contributions of systematic reviews to science, producing good quality timely reviews is a challenging endeavor. Studies have shown that SLRs might fail to provide a good and complete coverage of existing evidence, missing up to 40% of relevant papers [6], and even end up being outdated by the time of publication [6, 22] – this without considering those never published.

The reasons behind these challenges have been documented in several studies [7, 8, 20], which attribute them to the demanding nature of the involved tasks,

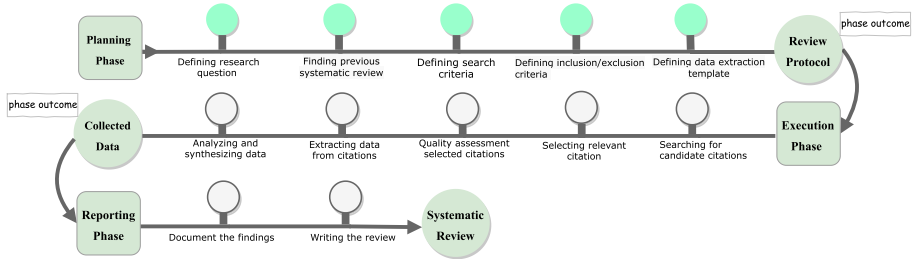


Fig. 1. The SLR process, defined by planning, execution and reporting activities

lack of expertise, limitations of support technology, and issues with primary studies. Recent advances in cognitive computing and collaborative technology offer an opportunity to address these challenges, and support researchers in planning, running and reporting SLRs (see Fig. 1 for an overview of the SLR process). We have seen new techniques and platforms enabling large-scale collaboration [13, 25, 26], and automation opportunities [9, 16, 19], offering promising results in different research activities relevant to the SLR process. Most of these efforts however are centered around the screening and identification of relevant scientific articles – and rightly so as it is one of the most time-consuming phases – but leaving other critical tasks largely unexplored.

In this paper we address a much less explored phase of the SLR process, that of planning the reviews. Guidelines and recommendations (e.g., [11]) define the main activities in this phase as i) identifying the need for undertaking the review, ii) defining the research questions (RQs), iii) defining the search and eligibility criteria, and iv) the data extraction template. These tasks are fundamental to guiding the SLR process and setting the foundation to having meaningful and original contributions, good coverage of the literature and a process free of bias. Yet, as we will see, they are often poorly performed, if at all.

In what follows we investigate how cognitive augmentation can support the planning phase of SLRs. We build on the insights and feedback from SLR authors to identify challenges and support strategies inspired by recent advancements in cognitive computing, framing the results under the framework for cognitive process augmentation. We also report on our early prototype and evaluation runs, showing the potential of augmentation in identifying the need for undertaking a review by leveraging word embeddings to find relevant SLRs from an input RQs.

2 Challenges in Planning SLRs

The challenges in running an SLR can be found throughout the entire process. These have been observed in the literature [7, 8, 20] as well as in our preliminary work, where we run an open-ended survey with more than 50 SLR authors tapping on their experience running SLRs. The results indicated that planning tasks are generally perceived as difficult to manage, requiring higher level of expertise and domain knowledge compared to more labor-intensive tasks.

Motivated by these insights, we run a second survey with 24 authors who published SLRs in top software engineering outlets in the last two years. This survey focused on their experience in planning SLRs, and inquired about i) whether planning tasks were properly addressed in their last SLR project, ii) the importance of addressing some salient challenges, and asked for iii) feedback on some support strategies to address the emerging challenges. In Fig. 2 we summarise the feedback on the first two points.

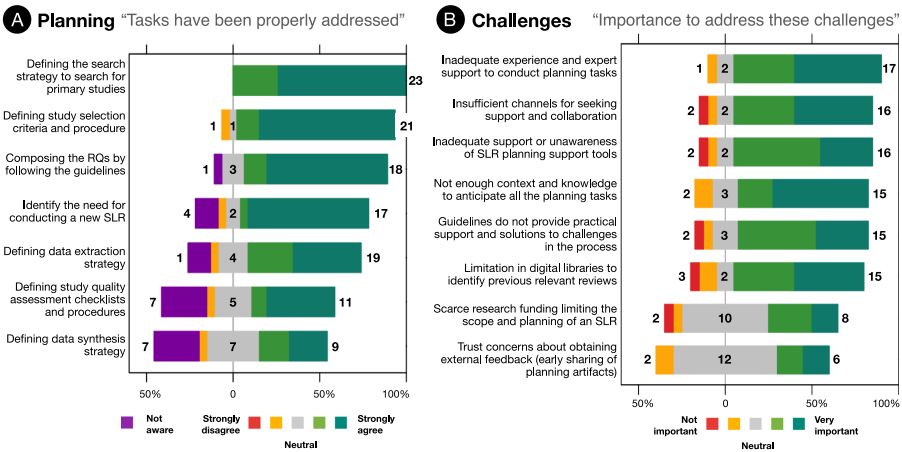


Fig. 2. Summary of feedback from SLR authors. A) Feedback on how planning tasks are addressed. B) Importance of addressing salient challenges in the planning.

While most authors reported positively to having addressed the tasks properly – not surprisingly given the quality of the outlets – there was still a significant number of researchers reporting neutral to negative, and in some cases even not being aware of certain tasks. More illuminating is to observe the challenges that we identified and the importance authors put in addressing them (Fig. 2), which we summarise and group below and address in the next section:

- C1** Inadequate experience and support by current tools and guidelines.
- C2** Insufficient context and knowledge to anticipate tasks in the planning phase.
- C3** Limitations of digital libraries to identify relevant SLRs.
- C4** Inadequate expert support and trust concerns in obtaining external feedback.

3 Cognitive Support in Planning SLRs

In this section, we present the conceptual design of a platform to support the planning of SLRs. In the following we describe the solutions and strategies addressing the main challenges, and the feedback obtained from SLR authors.

Strategies to Address Planning Challenges. The strategies were derived in brainstorming sessions among the co-authors, taking as input the insights from the first round of interviews with SLR authors on workarounds and strategies employed in the process, our own experience and prior work. The resulting strategies leverage techniques from machine learning and data-mining to address the main challenges. Below we describe the strategies for each of the challenges (C#).

- **Chatbot assistant** that allows authors to ask questions about the SLR process and best practices (C1). The chatbot provides a natural language interface to all information encoded in guidelines and recommendations.
- **Step-by-step guides** for each of the tasks (C1). The interface provides a conversational interface that assists authors in the preparation of each “artefact”, (e.g., RQs), by providing step-by-step prompts based on guidelines.
- **Incremental and iterative process** that can be adapted as more information becomes available to the researcher (C2). For example, RQs and inclusion/exclusion criteria can be refined as we identify similar literature reviews and learn more about the topic.
- **Incorporating context and knowledge** (C2). The system leverages information available in seed papers and similar literature reviews by extracting SLR-specific metadata relevant to the protocol (e.g., RQs, search strategy) and making them available to the authors as a reference point at each step.
- **Search focused on similar SLRs** (C3). Instead of defining complex queries and terms to identify similar reviews, the search focuses only on SLRs. The search results provide SLR-specific information, including RQs, inclusion/exclusion criteria, search strategy, etc.
- **From RQs to similar SLRs** and papers (C3). The system allows users to go from their (partial) RQs directly to similar SLRs by enriching information with extra data from seed papers.
- **Expert networking** (C4). We presented strategies including artefact-specific evaluation criteria, improved discovery of experts, and leveraging groups of trust. However, we limit the discussion on human-human collaboration.

We illustrate how the above strategies can be combined in a the concept tool in Fig. 3, depicting the iteration over the first two steps of the planning.

Feedback from SLR Authors. We requested feedback on the concept tool and strategies in our second wave of survey with 24 SLR authors. Authors were shown a mockup and descriptions,¹ and asked to i) rate the strategies on whether they addressed the specific challenges, ii) provide feedback on potential barriers for adopting them, as well as iii) other alternatives not foreseen.

The strategies to address the inadequate experience and support (C1) received positive feedback by almost all participants (15/16). For example, a

¹ All materials related to the study can be found at https://github.com/maisieb01/Cognitive_SLR.git.

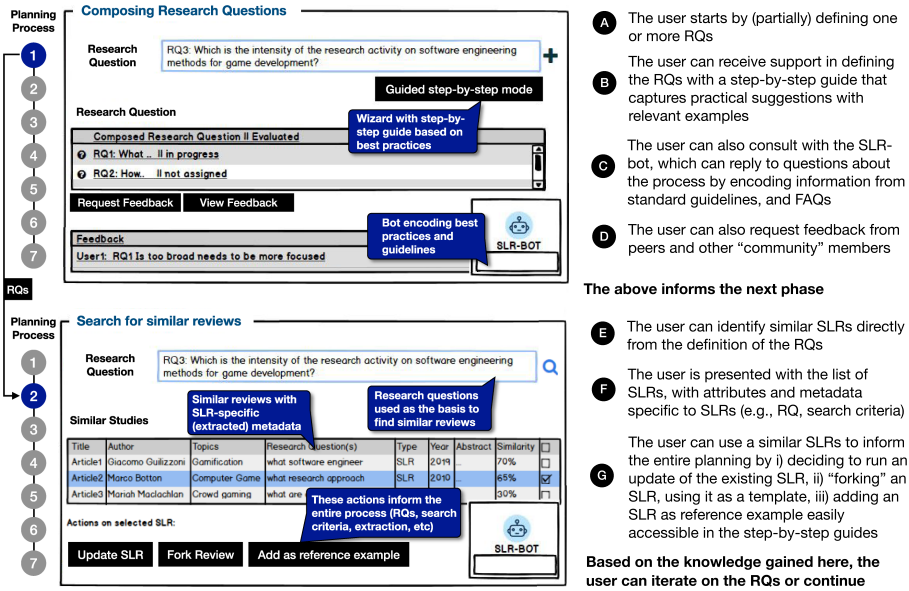


Fig. 3. Interaction illustrating the first two steps in the planning: defining the RQs and identifying the need for the SLR. The same strategies apply to the rest of the steps.

participant stated “I had to search for hours to understand how to get tasks done, so I think a Bot that helps with the common question would be really helpful”. Another one stated “Support from peers is the most important help that I’ve been missing when conducting my study. [...] I didn’t know much about the process and where to start, the step by step guideline would have helped me a lot”.

Participants reported potential barriers being, i) mismatch between the assistance provided by the tool and practices of the target domain of the SLR, and ii) the quality of the underlying algorithms, including the actual bot recommendations.

Regarding the lack of context and knowledge (C2), all participants provided positive feedback (17/17). The only concern was about extracting information from multiple SLRs following different approaches and the tool mixing them up.

The strategies to address the limitation of digital libraries (C3) also received positive feedback (15/17). The argument against it came from an author not sure about the effectiveness “I’m not sure if searching by [RQs] will help identify an area which needs more research or just other SLRs .. [it] just seems like a ‘type’ filter as e.g. in Scopus or adding ‘review’ in Title from Google Scholar”.

Among the suggestions we can mention: i) adapt the builtin guidelines to the target domain of the SLR, and consider other frameworks such as PICO, ii) extend the SLR extraction capabilities to recommend highly reputable venues, infer the guidelines to follow and create a to-do list authors could follow, iii) expand the search to suggest research that is very relevant, and iv) include explanations for as to why papers are suggested.

4 Conceptual Architecture

In delivering on the vision of cognitive support in the planning of SLRs, we rely on the framework for cognitive process augmentation by Barukh et al. [2].

As we will see, the strategies we devised inform this architecture at different layers.

Foundation: Existing technology provides support for coordinating *data*, *tasks* and *collaboration* that we can leverage to build our vision of a cognitively-augmented planning process. Starting from the process itself, we have seen the planning of SLRs to be an incremental and iterative process leading to a review protocol. The process management in this context could rely on lightweight artifact-centric systems (e.g., Gelee [1]) where the researcher drives the process while the system advises on the steps to take based on community-specific guidelines. Along the process, some tasks are already supported by current online services, such the search and access to scientific articles. Digital libraries and search engines provide access to article data and metadata, but under the limitations pointed out in the previous section, requiring researchers to engage in significant manual effort. Thus, although the data and knowledge required to elaborate the research protocol and inform the process is available, identifying, curating and adopting such knowledge is a challenging endeavour.

Enablement: The next layer leverages existing data sources and services to apply domain-specific data extraction and enrichment that will enable cognitive augmentation. Components such as *article recommendation*, enabling search for similar SLRs and papers, *article augmentation*, enriching SLRs with domain-specific metadata, *activity recommendation*, recommending steps based on process definitions and progress, and *knowledge graph*, aggregating knowledge about the process in queryable format, are among the enabling components. In this context, SLRs, primary research articles, and guidelines on how to run the SLR process, are the main sources of information. Lower-level algorithms such as named-entity recognition, word-embeddings and similarity serve as building blocks for these higher level components.

Delivery: The researcher finally experiences the cognitive augmentation in the planning through Conversational AI as well as intelligent GUIs. Conversational AI helps in delivering assistance in the process, providing a natural language interface to query the vast knowledge encoded in guidelines, and receive practical assistance in each step of the process. In the form of more general conversational interfaces, guided prompts would provide step-by-step guidance to assist researchers in knowledge-intensive tasks (e.g., defining RQs). We also recognise the need for serving more traditional delivery systems such as GUIs, to dote complex tasks with intelligent features (e.g., domain-specific search).

5 Prototype Implementation and Evaluation

In this section we describe our ongoing exploration into the technical feasibility of our approach.

We started with *article recommendation* as it emerged as a promising component based on the feedback from SLR authors, but raised concerns in terms of feasibility. We note however that providing support for identifying related SLRs has all the potential benefits but fewer of the concerns (e.g., recall and accuracy) with respect to assisting the identification and screening of relevant literature [12]. In the planning, bringing up the most relevant SLRs and papers can provide the additional context in a task that some people are not even aware of.

Prototype. The prototype is our initial exercise into understanding the technical requirements of the system, as well as a working tool to serve evaluation purposes. On the surface, the current version of the prototype provides a set of REST APIs (and an accompanying user interface) that given a set of RQs in input, it returns the most relevant SLRs, along with the relevance score computed based on the available models. Figure 4 presents the pipeline of the implementation.

The source of information is currently a curated database of SLRs, where domain-specific information has been manually extracted so as to evaluate the recommendation component in isolation. The end goal is to have a data layer that can interface with existing services to access structured and unstructured data, which can then be processed to automatically extract relevant domain-specific information. While we currently use a MySQL database to store and retrieve raw and curated data, the concept of *Data Lake* [3] emerges as a promising direction to store and query structured and unstructured SLR data.




Fig. 4. Architecture pipeline

In entering the augmentation layer, the data is pre-processed in two steps: (a) normalizing text corpus (e.g., removing special characters and stop words, all to lowercase); and (b) lemmatizing and converting each word to its base form. We leverage Stanford CoreNLP toolkit² to perform this process.

Following the data cleaning, the next step is to extract meaningful information from the data. The inspiration behind our proposed approach is leveraging a similar approach to the word embedding’s model [18] that represents words in a Vector Space Model (VSM). We have extended the idea of considering a “word”

² <https://nlp.stanford.edu/software/tagger.html>.

as a vector to represent the SLR-related corpus (e.g., RQs and SLR abstracts) in a vector space. To create the VSMSs, we employ an N-gram selector component to extract all the keywords (nouns and verbs) from sentences of the given context. We leverage Stanford Part-Of-Speech (POS) Tagger [15] to achieve this. Then, we create a list of n-grams out of these keywords and transform them into vector representations for each corpus. After encoding the given corpora into vectors, these vectors are used to calculate the similarity between desired corpora. Information is then augmented based on a pool of word embedding models. Our work leverages state-of-the-art algorithms widely used in NLP communities. Several such algorithms (e.g. GloVe [21], Word2Vec [17], Numberbatch [24], WikiNewsFast [4], and GoogleNews [5]) come with efficient implementations that are readily available as libraries to use.

The REST APIs then expose the functionality of the article recommendation component for programmatic access. A front-end application takes these services and wraps them up in a user interface.

Planned Evaluation. The goal of the initial evaluation is to inform specific design decisions regarding the algorithmic support for recommending papers. Among the main design decisions we consider: i) *What models will better serve the specific task?* The idea is to identify among the embedding models and architectures the most promising candidates to build on, and understand whether investing in domain-specific embedding models is required. Then, ii) *What information should we leverage when assessing the relevance calculations?* (e.g., title, title-abstract, RQs, full-text). The aim is to understand what (combinations of) information to focus when assessing the relevance of SLRs from an input RQs, and therefore to consider in the extraction process.

The prototype supports these two dimensions, models and selective information, so as to serve the evaluations. The dataset of SLRs is being manually constructed to incorporate for each SLR a set of related SLRs (as reported in the reviews) and not relevant SLRs as judged by human experts. Armed with the human-annotated dataset, we evaluate the quality of word embedding models by assessing how well the similarity scores of the word vectors correlate with human judgment [23]. The similarity is calculated as the distance between the vectors representing RQs and SLRs, using cosine similarity as measure, which has been found suitable for SLRs in prior work [14].

We rely on Spearman's rank (r_s) correlation between the word embedding models similarity score and researchers annotations to evaluate how well the similarity of the given pairs (e.g., RQs and abstracts) agrees with human judgments [10]. We performed a preliminary test run to tune the experimental setup, and the early results are encouraging. Results under limited settings (e.g., dataset of 160 SLRs, only abstract-RQ comparisons) already show good level of agreement ($r_s = 0.67$), for the best performing model, although at this point this is anecdotal since more comprehensive tests are required.

6 Conclusion and Future Work

We have seen that planning is a challenging endeavor, requiring resources, expertise or context that is often missing when undertaking a new topic, when performing it for the first time, or when resources are lacking – all typical scenarios in research. This paper shows that cognitive processes provide the ingredients to address these issues and support researchers in this often overlooked but impactful phase. As for ongoing and future work, we are in the process of refining the technical details of the first experiment, and planning an evaluation with end-users so as to assess the actual benefits of the approach when compared to standard tools. We also continue with our human-centered design approach to the development of the overall tool and algorithms, which will inform all components of the platform.

Acknowledgments. We acknowledge CSIRO Data61 for funding scholarship on this research.

References

1. Báez, M., et al.: Gelee: cooperative lifecycle management for (composite) artifacts. In: Baresi, L., Chi, C.-H., Suzuki, J. (eds.) *ICSOC/ServiceWave -2009*. LNCS, vol. 5900, pp. 645–646. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10383-4_50
2. Barukh, M.C., et al.: Cognitive augmentation in processes. In: Aiello, M., Bouguet-taya, A., Tamburri, D.A., van den Heuvel, W.-J. (eds.) *Next-Gen Digital Services. A Retrospective and Roadmap for Service Computing of the Future*. LNCS, vol. 12521, pp. 123–137. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-73203-5_10
3. Beheshti, A., Benatallah, B., Nouri, R., Chhieng, V.M., Xiong, H.T., Zhao, X.: Coredb: a data lake service. In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pp. 2451–2454 (2017)
4. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. *Trans. Assoc. Comput. Linguist.* **5**, 135–146 (2017)
5. Kenneth Ward Church: Word2vec. *Nat. Lang. Eng.* **23**(1), 155–162 (2017)
6. Créquit, P., Trinquart, L., Yavchitz, A., Ravaud, P.: Wasted research when systematic reviews fail to provide a complete and up-to-date evidence synthesis: the example of lung cancer. *BMC Med.* **14**(1), 8 (2016)
7. Garousi, V., Felderer, M.: Experience-based guidelines for effective and efficient data extraction in systematic reviews in software engineering. In: *Proceedings of the 21st International Conference on Evaluation and Assessment in Software Engineering*, pp. 170–179 (2017)
8. Hassler, E., Carver, J.C., Hale, D., Al-Zubidy, A.: Identification of SLR tool needs—results of a community workshop. *Inf. Software Technol.* **70**, 122–129 (2016)
9. Howard, B.E., et al.: Swift-review: a text-mining workbench for systematic review. *Syst. Rev.* **5**(1), 87 (2016)
10. Huang, E.H., Socher, R., Manning, C.D., Ng, A.Y.: Improving word representations via global context and multiple word prototypes. In: *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 873–882 (2012)

11. Kitchenham, B.: Procedures for performing systematic reviews. Keele, UK, Keele University **33**(2004), 1–26 (2004)
12. Kontonatsios, G., et al.: A semi-supervised approach using label propagation to support citation screening. *J. Biomed. Inform.* **72**, 67–76 (2017)
13. Krivosheev, E., Casati, F., Baez, M., Benatallah, B.: Combining crowd and machines for multi-predicate item screening. *Proc. ACM Hum.-Comput. Interact.* **2**(CSCW), 1–18 (2018)
14. Lopes, A.A., Pinho, R., Paulovich, F.V., Minghim, R.: Visual text mining using association rules. *Comput. Graph.* **31**(3), 316–326 (2007)
15. Manning, C.D., Surdeanu, M., Bauer, J., Finkel, J.R., Bethard, S., McClosky, D.: The Stanford CoreNlp natural language processing toolkit. In: *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pp. 55–60 (2014)
16. Marshall, I.J., Kuiper, J., Wallace, B.C.: Robotreviewer: evaluation of a system for automatically assessing bias in clinical trials. *J. Am. Med. Inform. Assoc.* **23**(1), 193–201 (2016)
17. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. *arXiv preprint [arXiv:1301.3781](https://arxiv.org/abs/1301.3781)* (2013)
18. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pp. 3111–3119 (2013)
19. Ouzzani, M., Hammady, H., Fedorowicz, Z., Elmagarmid, A.: Rayyan-a web and mobile app for systematic reviews. *Syst. Control Found. Appl.* **5**(1), 210 (2016)
20. Palomino, M., Dávila, A., Melendez, K.: Methodologies, methods, techniques and tools used on SLR elaboration: a mapping study. In: Mejia, J., Muñoz, M., Rocha, Á., Peña, A., Pérez-Cisneros, M. (eds.) *CIMPS 2018. AISC*, vol. 865, pp. 14–30. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-01171-0_2
21. Pennington, J., Socher, R., Manning, C.D.: Glove: global vectors for word representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543 (2014)
22. Sampson, M., Shojania, K.G., Garritty, C., Horsley, T., Ocampo, M., Moher, D.: Systematic reviews can be produced and published faster. *J. Clin. Epidemiol.* **61**(6), 531–536 (2008)
23. Shi, B., Lam, W., Jameel, S., Schockaert, S., Lai, K.P.: Jointly learning word embeddings and latent topics. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 375–384 (2017)
24. Speer, R., Chin, J., Havasi, C.: Conceptnet 5.5: an open multilingual graph of general knowledge. *arXiv preprint [arXiv:1612.03975](https://arxiv.org/abs/1612.03975)* (2016)
25. Sun, Y., et al.: Crowdsourcing information extraction for biomedical systematic reviews. *arXiv preprint [arXiv:1609.01017](https://arxiv.org/abs/1609.01017)* (2016)
26. Vaish, R., et al.: Crowd research: open and scalable university laboratories. In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, pp. 829–843 (2017)