acrl 2023

WELCOME TO THE MACHINE:
Ir/Responsible Use of Machine Learning in Research Recommendation Tools

Ali Krzton*

INTRODUCTION

Research practices continually adapt to developments in technology, and contending with the vast body of scholarly literature is no exception. As digital workflows have reduced the barriers to creating and publishing scholarship, researchers have scrambled to find ways to avoid "drinking from a firehose" and reduce information overload. Recommendations from colleagues about what sources are most important or worthwhile to read have always played a pivotal role in guiding researchers' attention, and they will continue to do so. But most of them have also become familiar with commercial services that help subscribers discover various forms of content tailored to their preferences and interests.

Machine learning, the technology that powers recommendation engines for Spotify, Netflix, Amazon and the like, machine learning, can likewise be applied to generate research recommendations, and demand for such services in the academic sector has increased.[1] Understanding how these tools function in a general sense has become critical for library and information professionals, both to competently advise researchers on literature management and to use them effectively in their own scholarly work.

In this paper I will give a brief overview of how machine learning algorithms work by pattern recognition. I then discuss how the use of machine learning influences the function of tools that recommend scholarly sources to researchers. Two general mechanisms are explored: one in which algorithms are designed to identify similarity and connection from the literature itself, and one in which they attempt to characterize the preferences and intentions of the user. I will compare their suitability for solving a particular research problem, ineffectual searches of unfamiliar literature in the course of interdisciplinary research, and outline their strengths and drawbacks. Finally, I discuss how librarians and other information professionals can leverage these insights when assisting researchers to reduce their search effort with machine learning responsibly.

RESEARCH RECOMMENDATIONS AS A MACHINE LEARNING APPLICATION

Research recommendation algorithms have been under active development for over twenty years,[2] but while previous approaches straightforwardly utilized document similarity or citation networks, this generation of research recommendation systems increasingly relies on machine learning to connect users with sources. The incorporation of this technology, a type of artificial intelligence, is a fundamental rather than an incremental change. The literature related to machine learning is highly technical, but an introductory data science text provides a compact definition in plain language: "Machine learning is the field of teaching machines and comput-

* *Ali Krzton, Research Data Management Librarian, Auburn University Libraries, alk0043@auburn.edu*

ers to learn from existing data to make predictions on new data without being explicitly programmed."[3] The key concept to note is *without being explicitly programmed*. Machine learning algorithms can reach the desired or specified outcome on the basis of any correlations detected in the input data, updating the weight given to those correlations based on feedback concerning how well the assigned task was performed. This is the "learning" element of machine learning.

With respect to navigating the literature by relating sources to one another, this means machine learning models will sometimes identify novel clusters of research in unforeseen ways. These clusters will not have a close parallel in shared topics, common authorship, or citation linkages. Taking full advantage of this phenomenon requires a shift in focus from how the clusters came about to whether people find such groupings useful or meaningful, but the potential for generating new ideas and lines of inquiry is clear. This can be particularly desirable when it comes to interdisciplinary research.

Interdisciplinary research challenges scholars to navigate unfamiliar literature. They may need to learn new sets of keywords or subject terms, or translate different terminology for similar concepts across fields, particularly in more exploratory stages of their project.[4] This makes searching for relevant sources more difficult and time-intensive. Tools that retrieve related works quickly and accurately can reduce the burden, and machine learning is theoretically suited to reduce the need for explicit description, or much description at all, of what users want to find. But practically speaking, how helpful are research recommendation services that employ machine learning in acquainting researchers with new subject areas faster? It depends on how the tool actually works; specifically, the source of the data it uses to determine relevance.

TOOLS BASED ON TEXTUAL CHARACTERISTICS

One class of tools that utilizes machine learning to identify relevant research leverages the text and/or the metadata of scholarly sources as data inputs. Then, models that evaluate the similarity of sources are developed from this data and guided to group or connect those with higher probabilities of being related. The class is referred to here as text-based tools, as opposed to content-based, since "content-based filtering" is a term of art related to how algorithms themselves work[5] and is not related to how tools are categorized in this paper. Those interested in trying out text-based research recommenders for themselves have a variety of choices at present (Table 1).

TABLE 1

A selection of research recommendation services that incorporate machine learning

Connected Papers	https://www.connectedpapers.com/
Research Rabbit	https://www.researchrabbit.ai
Semantic Scholar	https://www.semanticscholar.org/
Iris.ai	https://iris.ai/
Inciteful	https://inciteful.xyz
R Discovery	https://discovery.researcher.life/

A representative example is the Research Rabbit service,[6] which has been available to the general public since 2021.[7] Research Rabbit guides users through networks of associated sources after an initial paper is specified as a starting point. The browser-based interface allows users to map the literature in a particular area by bringing lists together with graphs of interconnected papers that adjust in real time depending on what is selected (Figure 1). Lists also update with suggestions after users identify more papers of interest, offering the option of continuing to build the collection. Of particular interest to interdisciplinary researchers, users can start multiple collections and keep them separate, permitting exploration within one topic or field at a time from a single account. While registration is required, anyone can set up an account for free, and the development team has pledged not to charge for the core discovery service.[8]

FIGURE 1

A sample project in the Research Rabbit interface. The contents of the Multilevel Society collection, one of two curated by this user, appear in the first list. To the right, information about the selected paper is summarized, followed by options for finding related work. The second list presents papers identified as similar, and a clickable graphical representation helps the user navigate this list as a cluster. Screen capture by the author.

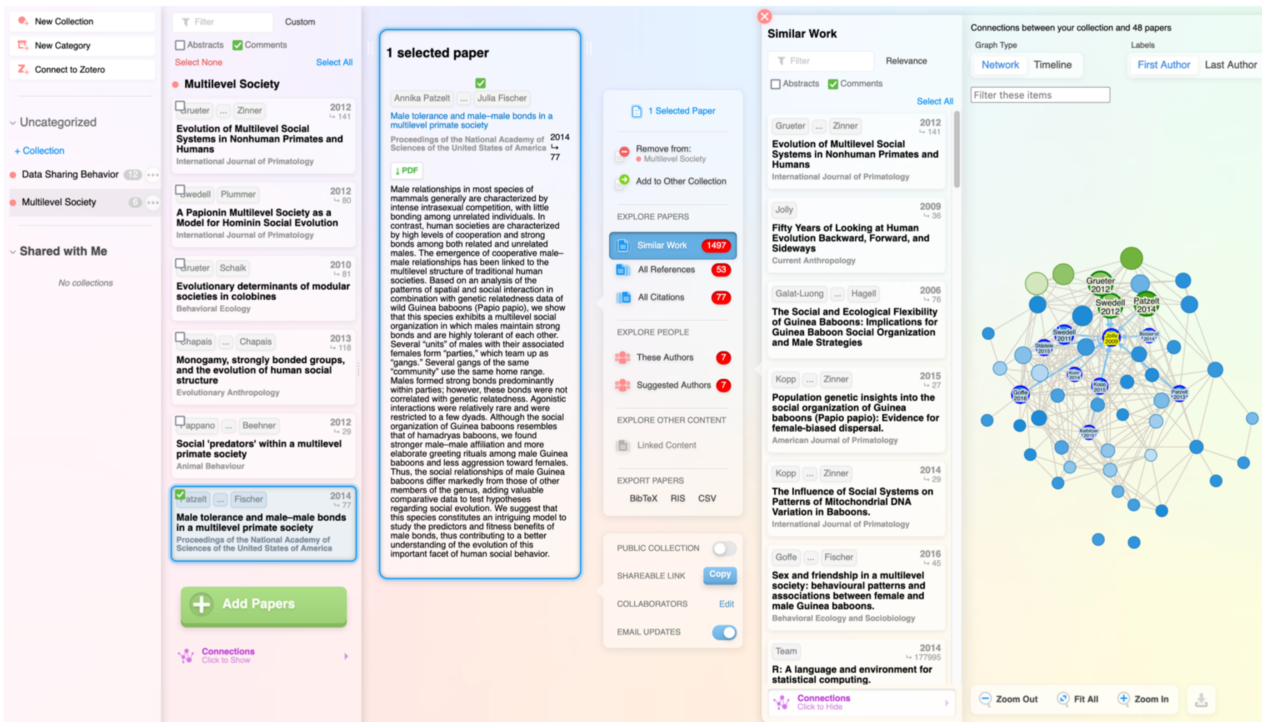

Text-based research recommendation tools are appealing for a variety of reasons, whether researchers want to move beyond traditional search to explore less familiar subject areas or keep up with recent developments. However, despite their promise, there are several limitations. One, general to machine learning applications, is that regardless of the amount of control a user has in selecting or ranking outputs, the training process is not reproducible, and in most cases, it is not even explainable.[9] For example, the developers of Research Rabbit have not provided any information on how their algorithms work.[10] There is also a tradeoff between greater freedom to tinker and ease of use; the learning curve of these tools could be a barrier to adoption for those unwilling to sink time into mastering yet another interface.

The scope of the dataset on which the service relies is another essential consideration. Some tools draw upon a massive corpus such as that maintained by Semantic Scholar, but others may be restricted to a smaller trove of open access papers or depend upon the extent of the user's institutional subscriptions.[11] Furthermore, many of these tools are offered on a "freemium" model or charge outright, putting those with limited funds at a disadvantage. Conversely, when a service has no apparent revenue source, it is reasonable to be concerned about sustainability. Sustainability of the underlying datasets can also be an issue. For example, Microsoft Academic Graph, the database on which Research Rabbit was built,[12] was discontinued at the end of 2021.[13] While it was incorporated into a successor database, OpenAlex, the fate of Microsoft Academic Graph serves as a reminder of how vulnerable the web's scholarly infrastructure can be.

TOOLS BASED ON USER PROFILES

A second class of tools aims the machine learning algorithms at users themselves to generate recommendations. These tools have an advantage over those discussed in the previous section in that they require less direct input, and therefore less effort, after the initial set-up. Once these tools have sufficient data to generate a model of what the user is presumably looking for, algorithms take over the heavy lifting and suggest additional resources to the user automatically. How users respond to those suggestions (i.e. marking relevant or not, saving, reading, or ignoring them) yields additional training data to further refine the model and improve its precision. Best-of-class services of this type give the impression of "knowing what you want before you do". Two such tools, ScienceDirect Recommendations and Google Scholar Recommended Articles, are notable due to the size and influence of the companies that offer them (Elsevier and Alphabet). For this reason, they are explored in greater detail here.

ScienceDirect Recommendations is an opt-in service available to anyone with an Elsevier user ID. The service's home page explains that it "uses machine learning and your online activity to suggest research tailored to your needs".[14] Once the user signs in on the ScienceDirect website, no further explicit action beyond typical engagement is necessary, unlike Research Rabbit and similar tools that require seed references or collections. Instead, data that trains the model which provides recommendations is generated from user activity, which is why the home page also suggests logging in and remaining so while on the ScienceDirect website to improve performance. While the nature of the activity that contributes to the user profile is not specified, examples almost certainly include the user's search terms, what sources they click on (and what sources they ignore), and what they view or download. However, more fine-grained data could also be incorporated into the model,[15] such as aspects of search behavior that indicate success or failure in finding relevant material, or even when they visit the website and for how long. The tool's designers are incentivized to gather any data that could improve the quality of the recommendations (and specifically, how well the recommendations maintain engagement with the ScienceDirect website and Elsevier publications).

From users' perspectives, it is not necessary to be aware what they did to prompt the recommendations they are given so long as those recommendations are relevant. However, such awareness may develop over time from repeated exposure to the content of those recommendations. For instance, a researcher may retrieve several papers to help a student learn about a topic only to see several items related to that subject appear in their recommendation feed. This may be perceived as diluting the usefulness of the service by reducing its personalization. This awareness could lead the user to make different choices while using the platform; some implications of this behavioral feedback loop are discussed in the next section.

One limitation ScienceDirect Recommendations shares with the literature-based tools discussed previously is that the sources it offers are constrained to a particular pool. Presumably, the library on which it draws consists primarily or exclusively of content published by Elsevier. While 15 million articles and book chapters[16] is an extensive body of scholarship, most researchers would find it too narrow to rely on exclusively regardless of how well-tailored ScienceDirect's recommendations seemed to be. In contrast, Google Scholar indexes a tremendously wide variety of sources.[17] It can, therefore, retrieve results from a much larger universe of possibilities than any specific publisher's recommendation engine.

Google Scholar is perhaps the best-known discovery tool among academics. While Google Scholar's interface puts simple keyword searching of the literature front and center, it also recommends particular sources. When a researcher with a Google Scholar profile is signed in, the Google Scholar homepage displays a list of recommended articles below the search box. The seed data for these recommendations is the user's own publication history as recorded in their Google Scholar profile. According to Google, relevance of articles is determined by "the topics of your articles, the places where you publish, the authors you work with and cite, the authors that work in the same area as you and the citation graph."[18] Unlike ScienceDirect Recommendations and most other recommendation engines, Google Scholar's recommendations are not influenced by user searches or any other user behavior within the platform other than keeping one's Google Scholar publication history up-to-date. However, researchers' publication and citation behavior is used to train the models regardless of whether they interact with Google Scholar thanks to publicly available metadata.

Although a key advantage of using Google Scholar to search the literature is its inclusiveness with respect to information sources, its recommendation service is far more restrictive than its competitors in the sense that it gives users no choice in where to begin; they may not even define their research interests apart from what they have published. Moreover, this particular tool leaves out new scholars with few or no publications. The way Google Scholar's recommendations are configured appears to be a reflection of Google's design philosophy more generally, with its emphasis on simplicity and determining results chiefly by algorithms with minimal user input.[19] The service also leverages social network effects due to the weight placed on author connections through co-authorship and citation. These relationships could ultimately determine which articles are recommended more than the subject matter of the research itself, but since Google's algorithms are proprietary, it is difficult for anyone to independently assess how the relevance of articles is ranked. Finally, an important aspect of this service model is it reminds users with profiles to sign in while they use Google Scholar so that the recommended articles appear on the home page. This makes it easier for Google to harvest data about the user's search history and behavior, and though it has no impact on the research recommendations connected to the Google Scholar profile, it may affect the user's results while using the general search engine as the algorithms adapt to new inputs.

Concerns about the privacy and integrity of personal data are reasonable when evaluating any services based around user profiles. Part of the appeal of profile-based research recommendations is that they are nearly always offered at no charge. However, these tools require extensive data collection in order to function as expected, including data passively generated by people's behavior within the platform (or outside of it) in addition to whatever information users explicitly enter for the purpose of building their profiles. Thus, while researchers do not pay for access in money, they may incur costs as a result of sharing data of strategic or personal importance with the providers.

IMPACTS OF PROFILING ON THE PRACTICE OF RESEARCH

With respect to the initial question of how machine learning tools can reduce the friction of searching within new subject areas for researchers working on interdisciplinary projects, it is apparent that research recommendations developed from user data will not accomplish this goal and can even be counterproductive. This is because machine learning models are meant to identify similar cases, but users who intend to explore new topics or methodologies are purposefully breaking from past patterns. Thus, their behavioral record is not an appropriate source of training data. Some developers appear to understand this problem and are engineering workarounds. For example, R Discovery is a relatively new research recommendation service that profiles users to create a personalized feed. At the time of this writing, the website advertises that multiple feeds are coming soon and provides the following rationale for implementing this feature: "We understand the needs of a modern interdisciplinary researcher like you and hence allow you to create multiple feeds so that your reads are organized the best."[20]

But the pernicious impacts of user profiling on the practice of research are not limited to interdisciplinary study. Due to the tendency of machine learning models to grow more precise after successive rounds of input and feedback, algorithmic extrapolations of research interests will become more focused with time. While there are contexts in which this could be desirable, it can also lock researchers into a particular track, discouraging them from branching off into other paths of inquiry or serendipitously encountering ideas from other fields of study. To the extent that social networks inform the algorithms that generate research recommendations, the set of sources that scholars working in a particular specialty are exposed to might converge on a subset of the most popular references, contributing to siloing and homogenous thinking. This is especially likely in the case of new publications, where algorithms may go beyond chasing trends to actively create them in response to the choices of the first or most connected readers of emerging research. In plain terms, an article of lower quality could be amplified over a higher-quality article on the same topic coming out at roughly the same time because a few users accessed it or interacted with it early. Because of the fast pace of research, the merits of the better article may not be recognized because the specialists in the field never see it in their feeds and move on to new recommendations before it can gain momentum.

Finally, it is possible for algorithmic profiles and users' self-awareness to become mutually reinforcing. While users' activities and responses feed data into the system, the system also transmits information back that gives users an impression of what their interests, or even their identities as scholars, are perceived to be. This sense of being "watched" can alter subsequent behavior, especially if researchers find that a particular service they come to depend on recommends highly relevant sources.[21] It is conceivable that, for instance, someone who knows views and clicks contribute to their profile could decide not to even look at an article less closely aligned with their current research project for fear of throwing their feed off course. Over time, small decisions such as these could add up, reducing overall willingness to engage in open-ended exploration. Recognizing the profound effects of such a loss provides an opportunity to consider how ubiquitous recommendation engines affect curiosity more broadly, within academic research as well as outside it.

RECOMMENDATIONS

Cost-benefit analysis of machine learning tools in the research discovery space has implications for professional practice. Librarians and others who provide information services to academic researchers need not know the inner workings of these tools in detail, but it is important to distinguish those that use aspects of sources vs. characteristics of the user (i.e. profiling) as data to train the models. These approaches can coexist, and it is frequently unclear which data streams are being incorporated into the recommendation engine.

In general, the more control the user has over defining the parameters of the starting points and what constitutes relevancy the better, and the less all user behaviors within the system determine what will be retrieved the better. Tools that encourage researchers to remain logged in at all times and attempt to build a single, comprehensive picture of their interests which follows them across projects are to be avoided. When consulting with researchers on technologies aimed at finding relevant sources, librarians and information professionals should keep an open mind and maintain perspective on the immense value to them of saving time and effort. Any concerns about a particular service should be raised in the simplest possible terms, with a clear explanation of the shortcomings (for example, excessive data harvesting or lack of transparency). Ideally, librarians and information professionals can suggest viable alternatives during these discussions to provide users with a way forward. For this and many other reasons, paying attention to developments in the application of machine learning tools to the practice of research is worthwhile.

NOTES

1. David Matthews, "Drowning in the Literature? These Smart Software Tools Can Help," *Nature* 597, no. 7874 (September 1, 2021): 141–42, https://doi.org/10.1038/d41586-021-02346-4.
2. Joeran Beel et al., "Research-Paper Recommender Systems: A Literature Survey," *International Journal on Digital Libraries* 17, no. 4 (November 1, 2016): 305–38, https://doi.org/10.1007/s00799-015-0156-0.
3. Prateek Gupta, *Practical Data Science with Jupyter*, 2nd ed. (New Delhi, India: BPB Publications, 2021): 145.
4. Rick Szostak, "The State of the Field: Interdisciplinary Research," *Issues in Interdisciplinary Studies* 31 (2013): 44–65, https://doi.org/10.7939/R3QB9V49Q.
5. Beel et al., "Research-Paper Recommender Systems."
6. "ResearchRabbit," ResearchRabbit, accessed February 28, 2023, https://www.researchrabbit.ai.
7. Aaron Tay, "ResearchRabbit Is out of Beta- My Review of This New Literature Mapping Tool," *Academic Librarians and Open Access* (blog), August 16, 2021, https://medium.com/a-academic-librarians-thoughts-on-open-access/researchrabbit-is-out-of-beta-my-review-of-this-new-literature-mapping-tool-3c593d061c63.
8. ResearchRabbit Team, "ResearchRabbit," Open Collective, February 26, 2023, https://opencollective.com/researchrabbit.
9. Lucie Beranová et al., "Why Was This Cited? Explainable Machine Learning Applied to COVID-19 Research Literature," *Scientometrics* 127, no. 5 (May 2022): 2313–49, https://doi.org/10.1007/s11192-022-04314-9.
10. Matthews, "Drowning in the Literature?"
11. Extance, "How AI Technology Can Tame the Scientific Literature."
12. Tay, "ResearchRabbit Is out of Beta- My Review of This New Literature Mapping Tool."
13. Dalmeet Singh Chawla, "Microsoft Academic Graph Is Being Discontinued. What's next?," *Nature Index*, June 15, 2021, https://www.nature.com/nature-index/news-blog/microsoft-academic-graph-discontinued-whats-next.
14. ScienceDirect, "Research Recommendations," ScienceDirect, accessed February 9, 2023, https://www.sciencedirect.com/research-recommendations.
15. Xiaomei Bai et al., "Scientific Paper Recommendation: A Survey," *IEEE Access* 7 (2019): 9324–39, https://doi.org/10.1109/ACCESS.2018.2890388.
16. ScienceDirect, "Research Recommendations."
17. Matthews, "Drowning in the Literature?"
18. Namit Shetty et al., "Scholar Recommendations Reloaded! Fresher, More Relevant, Easier," *Google Scholar Blog* (blog), February 12, 2021, https://scholar.googleblog.com/2021/02/scholar-recommendations-reloaded.html.

19. Matt Strom, "What Google Search Shows Us About the Future of Product Design," *UX Booth* (blog), September 10, 2019, https://www.uxbooth.com/articles/what-google-search-shows-us-about-future-of-product-design/.
20. Cactus Communications, "Academic Research Reading: Personalised Research Discovery," R Discovery, 2023, https://discovery.researcher.life/.
21. Minna Ruckenstein and Julia Granroth, "Algorithms, Advertising and the Intimacy of Surveillance," *Journal of Cultural Economy* 13, no. 1 (2020): 12–24, https://doi.org/10.1080/17530350.2019.1574866.