

null • Business of Knowing, summer 2021

Sustainable Authorship Models for a Discourse- Based Scholarly Communication Infrastructure

Joel Chan

Published on: Aug 23, 2021

DOI: 10.21428/6ffd8432.8b4aad0c

License: [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

Introduction

Today's scholarly communication infrastructure is not designed to support scholarly synthesis. When gathering sources for a literature review, researchers need to answer questions about theories, lines of evidence, and claims, and how they inform, support, or oppose each other. This information cannot be found simply in the titles of research papers, in groupings of papers by area, or even in citation or authorship networks (the sole focus of most scholarly communication infrastructure).

This limitation is a serious impediment to knowledge building and synthesis. Consider that, [by some estimates](#), approximately 50% or more of the time cost of systematic reviews is devoted to workarounds for this infrastructural limitation: screening papers (title, abstract, and full-text) to determine if it actually contains a claim that is relevant and worth checking, then extracting the claims and metadata for analysis; worse, other scholars do not get to benefit from this intermediate product and must start all over again. With this in mind, it's not so surprising that [systematic reviews are rarely updated even when they need to be](#). Many [doctoral dissertations also lack coverage and synthesis of literature](#), and published papers are [not much better](#). It's plausible, too, that this limitation contributes substantially to [slowdowns in research progress](#) via the [growing burden of knowledge](#).

How might we build an alternative scholarly communication infrastructure that can overcome this core limitation?

II. Discourse graphs: the promise and the authorship bottleneck

For decades, [researchers across a range of disciplines](#) have been [developing a vision](#) of an [alternative infrastructure](#) centered on a more appropriate core information model: knowledge *claims*, linked to supporting *evidence* and their context through a network or *graph* model. For conciseness here, I call this model a "**discourse graph**", since the graph encodes discourse relations between statements, rather than ontological relationships between entities.

Much crucial conceptual and technical progress has been made at the level of formal standards, and [several proof-of-concept implementations](#) have [demonstrated the promise of this concept](#). However, [adoption, particularly in terms of authorship, remains a hard open problem](#). In general, coverage of the literature and breadth of sustained contributors remains far lower than we would like. As one data point, contributions to servers for the nanopublications standard for discourse graphs are

almost all within bioinformatics and contributed by tens of authors. Tobias Kuhn, a lead on this standard, [puts it well](#): we want an ocean of such "micropublications", but "[a]t the moment, this is no more than a puddle" (p. 492)

I believe the UX problems (broadly construed beyond just usability) that contribute to this bottleneck are both high leverage and relatively neglected. First, **contributing to shared discourse graphs is currently disconnected from the intrinsic practices of scholarship**, both in terms of toolsets (separate specialized tools and webapps/platforms), and practices (often more formal and unable to mix with the informal speculative notes that are the lifeblood of research work). This disconnect creates significant opportunity costs for authorship. It also leaves the work that scholars already do as a substantial untapped source of potential sustainable contributions. Consider that [by some estimates](#), full-time faculty self-report reading about 200 articles per year; there were an [estimated 700k full-time faculty in 2018](#). So we can estimate time spent reading ~100M articles per year as a lower bound on untapped resources, since students, part-time faculty, research scientists, and citizen scientists also spend significant time reading articles. This matches (and likely exceeds) the scale of the [total number of published research papers](#). Further, the **intended audience/beneficiaries of this authoring work are most often some unknown others**. This is problematic because, all things being equal, scholars are likely to choose activities that directly contribute to their own work and their direct responsibilities (collaborators, trainees, students, etc.), even if they value benefits to society.



III. Sustainable authorship of discourse graphs by integrating into scholarly practices

Based on this analysis, I believe a promising but underexplored solution path for this authorship bottleneck is to **build tools that integrate authoring contributions to discourse graphs into the intrinsic tasks of effective scholarship practices**.


Here I describe one example point of integration: reading and sensemaking for literature reviews.

A user story

Consider Curie, a researcher, who is studying the role of analogies in cross-boundary innovation. She writes notes about the papers she reads in a [digital outliner notebook](#),

in which she is also drafting a literature review for her research project.

Let's take a look at her notebook and how she might be able to integrate authoring and usage of a discourse graph.

• @sanerAnalogiesOutBlue1999 for  QUE - How might domain distance modulate the effects of analogies on creative output? ¹

- General notes
- Method
 - very similar to @dunbarHowScientistsThink1997's In vivo method
 - sample: three ~70-90 min colloquia sessions (one speaker + ~30-60 audience members from department; all 3 speakers were senior researchers, well-known, two from other universities) in social psych, cognitive psych, and cognitive neuropsych + five ~70-90 min lab group meetings from two different lab groups (one in developmental psych with 10 scientists: two faculty, three postdocs, three graduate students, and two research assistants; one in cognitive psych, with 19 scientists: three faculty, six postdocs, five graduate students, and five research assistants), all at major research university in North America
 - measures
- Results
 - Analogies were relatively frequent: ~3-15 in each lab group meeting, and ~4-15 in each colloquium
 - Far analogies were rare and never used in psychology lab group meetings for reasoning (vs. mere mentions); far analogies were rare in colloquia as well, but were frequently used for reasoning
 - Overall: within-domain analogies dominated (~80% of analogies)
 - Table 2: Percent of analogies categorized as within domain, used, similarity based, asserted, and presenter given (separately by setting).

Category	Coll.	Lab Group	p
Within Domain	77%	81%	n.s.
Used	43%	27%	.2
Similarity Based	70%	68%	n.s.
Asserted	93%	97%	n.s.

Curie's reading notes for a paper (Saner, Analogies Out Blue, 1999; shown here as a reference to a “@sanerAnalogiesOutBlue1999” resource, in grey) that informs a question about analogical distance (shown here as a reference to a “QUE note about domain distance”, in gold). Note the mixture of formal and informal notes. Some portions of her notes here are bullets with content “collapsed” underneath (shown here as purple circles), a common feature in outliner software.

Curie selects a portion of her notes where she has written up a core result about far analogy use, and marks it as formal evidence (denoted in the system in pink) for later use in synthesis.

Leaving aside the particularities of the software, the general *content* structure of her notes is similar to a Google Doc of reading notes: a mix informal and formal observations and structure, including general notes about related ideas, key details about methods, and the core results of the paper.

But there is one crucial difference: while writing notes for a paper, Curie has marked out a key piece of evidence (EVD) from the paper that might inform her synthesis for

her focal question about how domain distance modulates the effects of analogies on creative output. This marking creates a new document (or page) in the software with that evidence note as a title, and allows Curie to reference that specific piece of evidence elsewhere in her notebook (similar to Wikipedia), such as while drafting an outline.

As Curie begins to need more contextual details while comparing and making sense of multiple EVD notes, she can elaborate the EVD notes with more details over time, such as by migrating in screenshots of key tables and figures, or methodological details like participants and measures.

EVD - Far analogies were rare and never used in psychology lab group meetings for reasoning (vs. mere mentions); far analogies were rare in colloquia as well, but were frequently used for reasoning - @sanerAnalogiesOutBlue1999


• **Summary:**

- Overall: within-domain analogies dominated (~80% of analogies)

- Table 2: Percent of analogies categorized as within domain, used, similarity based, asserted, and presenter given (separately by setting).

Category	Coll.	Lab Group	p
Within Domain	77%	81%	n.s.
Used	43%	27%	.2
Similarity Based	70%	68%	n.s.
Asserted	93%	97%	n.s.
Presenter	90%	70%	.1

(p. 4)

- This is similar to  **EVD** - Molecular biologists with a reputation for innovation rarely used very far analogies in their lab meetings while generating novel scientific concepts; instead, they relied mainly on analogies to the same or other biological organisms - @dunbarHowScientistsThink1997
- **No differences** in percentage raw occurrence of **far analogies** between psychology lab group and colloquia **settings**, but **stark difference in "use" of far analogies**: none of the far analogies in lab group were actually **used**, but 71% of far analogies in colloquia were (vs. equal likelihood, ~30%, of usage for near analogies in both settings)

• **Grounding Context**

- setting: two psychology lab groups and one colloquium series at a major research university in North America (over 100 faculty, students, postdocs, and research assistants; colloquium series presenters from around the world)
- sample: three ~70-90 min colloquia sessions (one speaker + ~30-60 audience members from department; all 3 speakers were senior researchers, well-known, two from other universities) in social psych, cognitive psych, and cognitive neuropsych + five ~70-90 min lab group meetings from two different lab groups (one in developmental psych with 10 scientists: two faculty, three postdocs, three graduate students, and two research assistants; one in cognitive psych, with 19 scientists: three faculty, six postdocs, five graduate students, and five research assistants)
- measure: analogy yes/no - yes if comparing similarity/dissimilarity of two things in a structural or functional sense; no if x is attribute of y, or x is-a (instance of) y (e.g., "x is a connectionist system"). IRR is .90
- measure: analogy within-domain vs. not (**far analogies**): domain is defined as psychology, bc couldn't reason through systematic way of distinguishing between domains?
 - There were two dimensions that were most critical with respect to our research questions. The first was the "Source" dimension. Here we coded for the distance between source and target according to whether the target came from within the same domain as the source or from a different domain altogether. This dimension was taken from Dunbar (1996) and modified to suit the current domain of psychology. Dunbar distinguished between within organism, between organism, and non-biological analogies. While these distinctions are important and appropriate for molecular biology, they are not applied so easily to psychology. For

Curie elaborates the body of the evidence note — for the result from Saner 1999 about far analogy use — with contextual details, such as an excerpt of a key table of results, and grounding context such as methods and setting details.

Let's take a closer look at an outline Curie is drafting for her literature review.

- I think the first subclaim is the most obvious one: 🟢 **CLM - Analogical distance of inspirations for an idea are positively related to the idea's creativity.** This is kind of a weak sauce claim, to be honest. A bit too general for my taste.
 - **Case for the claim**
 - There is some weak **support** from an early experiment from @dahlInfluenceValueAnalogical2002 that everyone cites in the creativity/engineering design literature. Two core results relating to **analogical distance** come from the third experiment.
 - There's also some **evidence** from citation level data that looks at combinations of "knowledge domains" (with journals as a proxy) and finds correlations between having atypical or highly novel combinations of knowledge domains being related to being outliers in terms of citation impact
 - Some complicated, qualified **support** also comes from studies that find contingent benefits of far analogies
 - Also some **support** from **outsider innovation** effects
 - And some **support** in the form of negative effects of near analogies
 - **Counterarguments**
 - There is some **counter-evidence** from m/s/In vivo studies of science labs and design teams
 - For example, Dunbar traced the discovery process of four top molecular biology labs, and found that they rarely used far analogies to generate novel scientific concepts: 🟡 **EVD - Molecular biologists with a reputation for innovation rarely used very far analogies in their lab meetings while generating novel scientific concepts; instead, they relied mainly on analogies to the same or other biological organisms - @dunbarHowScientistsThink1997**
 - This is a strong result, even though it isn't experimental, due to its high external validity. Despite lack of experimental controls, it at least puts a cap on the claim about far analogies being important/necessary for creative breakthroughs, at least in molecular biology.
 - A very similar result was found by @sanerAnalogiesOutBlue1999, who found that far analogies were basically never used in reasoning in psychology lab meetings, but were used in colloquia while ideas were being communicated with new audiences: 🟡 **EVD - Far analogies were rare and never used in psychology lab group meetings for reasoning (vs. mere mentions); far analogies were rare in colloquia as well, but were frequently used for reasoning - @sanerAnalogiesOutBlue1999**
 - Related to this is 🟡 **EVD - In a design team, concepts tended to be more similar to their immediately preceding concepts after far analogy use compared to using near or no analogies - @chanImpactAnalogiesCreative2015.** This isn't *quite* a measure of novelty as we normally think of it (the reference for similarity is concepts generated in the last 10 or 5 lines). But it at least opposes the idea that you might be primarily using far analogies to get some distance from where you currently are.
 - Dunbar also claims that 🟢 **CLM - far analogies are systematically overrated in their importance for creative breakthroughs due to memory bias - @dunbarHowScientistsThink1997**
 - One **support** for this is a specific anecdote of tracing a significant discovery through a series of (near) analogies, and then asking the scientist to recall/recount the origins of that discovery: Dunbar found that 🟡 **EVD - A molecular biologist who had made a major scientific conceptual change did not recall any of the spontaneous analogies used to enact that change - @dunbarHowScientistsThink1997**
 - And some experimental **counter-evidence** showing equal effects of near analogies on novelty

Curie drafts an outline that references special claim and evidence notes extracted from papers (shown here as green CLM and pink EVD notes, respectively). The system recognizes patterns of argumentation in her writing and outlining and automatically creates discourse links between relevant CLM and EVD notes that are referenced in the outline.

It is similar to a normal scholarly outline, with a mixture of formal and informal notes, and links to resources and references. Again, there is a small but crucial difference: Curie can reference *specific results* (evidence notes) while making sense of the case for and against a focal claim.

This enables her to access contextual details for comparing/contrasting claims and evidence a hover or click away without breaking the flow of writing, in contrast to a paper-level citation. In this way, Curie benefits directly from having marked out these CLM and EVD notes.

- **Counterarguments**

- There is some [counter-evidence](#) from [m/s/In vivo](#) studies of science labs and design teams

- For example, Dunbar traced the discovery process of four top molecular biology labs, and found that they rarely used far analogies to generate novel scientific concepts: 🦋 [EVD - Molecular biologists with a reputation for innovation rarely used very far analogies in their lab meetings while generating novel scientific concepts; instead, they relied mainly on analogies to the same or other biological organisms - @dunbarHowScientistsThink1997](#)

- This is a strong result, even though it isn't experimental, due to its high external validity. Despite lack of experimental controls, it at least puts a cap on the claim about far analogies being important/necessary for creative breakthroughs, at least in molecular biology.

- A very similar result was found by Saner et al (1999), who found that far analogies were basically never used in reasoning in psychology lab meetings, but were used in colloquia while ideas were being communicated with new audiences: 🦋 [EVD - Far analogies were rare and never used in psychology lab group meetings for reasoning \(vs. mere mentions\); far analogies were rare in colloquia as well, but were frequently used for reasoning - @sanerAnalogiesOutBlue1999](#)

- Related to this is 🦋 [EVD - In a design team, concepts tended to be more similar to their immediately preceding concepts after far analogy use compared to using near or no analogies - @chanImpactAnalogiesCreative2015](#). This isn't *quite* a measure of novelty as we normally think of it (the reference for similarity is concepts generated in the last 10 or 5 lines). But it at least opposes the idea that you might be primarily using far analogies to get some distance from where you currently are.

Hovering over the titles of evidence notes from Saner, Dunbar, and Chan, calls up contextual details (stored inside each evidence note) for easy comparison without breaking the flow of writing.

Finally, consider what happens when Lovelace, a new student, joins the project. To onboard her, Curie runs a graph query to collect claim and evidence notes that inform the focal question, and exports and emails them to Lovelace. She can choose to share just the claim and evidence notes, or also the narrative context of their use in the body of a question note, or the discussions in the reading notes, as appropriate.

Alternatively, she could also share hyperlinks if she has an extension to her notebook that auto-publishes only discourse-graph subsets of her notes to a shared repository.

The screenshot shows a web interface for a synthesis tool. At the top, it says "Synthesis". Below that is a "Match" section with an "Evidence" dropdown. Two conditions are listed: "That Informs" and "That Opposes", both with toggle switches. The "That Opposes" condition is selected, showing a query: "[[CLM]] - Analogical distance of inspirations for an idea are positively related to the idea". Below this is a "Results" section with a "Query" input and a "Results" button. The results are displayed as a list of 4 items, each starting with "[[EVD]]" and followed by a description of evidence.

Synthesis

Match

Evidence ▾

☒ That Informs ▾ [[QUE]] - How might domain distance modulate the effects of analogies on creative outp 🗑

☒ That Opposes ▾ [[CLM]] - Analogical distance of inspirations for an idea are positively related to the idea 🗑

Add Condition +

Query

Results

Found 4 results

- [[EVD]] - In general, MechE students' ideas when given near analogies were at least as novel, and sometimes moreso, then ideas generated with far analogies, or seeing nothing - [[@tsengRoleTimingAnalogical2008]]
- [[EVD]] - In a design team, concepts tended to be more similar to their immediately preceding concepts after far analogy use compared to using near or no analogies - [[@chanImpactAnalogiesCreative2015]]
- [[EVD]] - Molecular biologists with a reputation for innovation rarely used very far analogies in their lab meetings while generating novel scientific concepts; instead, they relied mainly on analogies to the same or other biological organisms - [[@dunbarHowScientistsThink1997]]
- [[EVD]] - Far analogies were rare and never used in psychology lab group meetings for reasoning (vs. mere mentions); far analogies were rare in colloquia as well, but were frequently used for reasoning - [[@sanerAnalogiesOutBlue1999]]

Curie runs a graph query to find evidence that informs the focal question for which she wants to onboard Lovelace, a new student. She focuses on evidence the opposes their current hypothesized claim about the benefits of analogical distance, to better focus discussion and planning for the next set of studies. She exports these notes to a directory of markdown notes to share with Lovelace.

The graph query works because the notebook Curie is using has an underlying extension that recognizes the argument structure that she is using in the outline, through a mixture of indentation patterns and keywords. Here, for instance, Curie can query for opposing evidence for a claim because the system has formalized an "Opposed By" relation between the CLM and the EVD by recognizing a pattern of writing in her outline.

Over the next few weeks, Lovelace spends her time modifying, elaborating, and integrating these notes into her own notebook (instead of laboriously extracting claims and evidence from a long list of papers!), and writes up some notes on new evidence from recently published work that Curie hasn't yet read. She shares these updates with Curie, and the resulting updates to the synthesis outline sparks a novel hypothesis that the project team decides to test for their next set of experiments.

Some observations

This user story illustrates how the work of authoring a discourse graph can be integrated into familiar, intrinsically useful scholarly practices of reading, note-taking, and writing, to the direct benefit of scholars and their colleagues.

But it also demonstrates the technical feasibility of this vision! These screenshots are not mockups: they are snapshots of my own notes, which I have written for my own work (for a literature review), and actually shared with students and collaborators. The [digital notebook I am using](#) did not require me to do a lot of other extra work like setting up an environment or deploying a personal server; the only thing I had to do was install an extension — [an active research project](#) — to the notebook with a single click.

This notebook is also but one of a Cambrian explosion of [similarly extensible hypertext-enabled digital notebooks](#) that can technically accomplish this same basic shape of a workflow. These tools are quickly growing in their userbase, significantly extending beyond [older more niche/homegrown tools](#) that have similar basic capacities, and also spawning new sets of [technical and cultural practices](#) for easily structuring and sharing notes.

IV. Conclusion

I am excited to imagine a world where anyone who cares about understanding the frontiers of knowledge are equipped with tools that enable them to annotate and write notes that better benefit themselves *and* share discourse graph subsets of their notes to enrich scholarship practices with their immediate colleagues. I want to broaden the lens of scholars to include nonprofit research institutions compiling nonpartisan literature reviews to inform [policymaking](#), and highly motivated communities of patients and their families who are seeking to understand and contribute to research on diseases that personally affect them.

Can this bottom-up, decentralized, peer-to-peer infrastructure help advance original visions around a single universal shared discourse graph? I believe the answer is not directly, but this may actually be a feature rather than a bug. Distributed knowledge graphs are notoriously hard to achieve consensus on, especially as they scale, and there is emergent evidence that local [contextualization, ambiguity and contestation may be crucial for scholarly progress](#).

Therefore, I am excited about institutional structures that can steward local federations of discourse graphs (e.g., at the level of labs, centers, or institutions), enabled by technical mechanisms for dynamic interoperability, such as [Project Cambria](#). If institutions and local collaborations institute methods of consensus, error-checking, and editing for integrating (as an analog to, say, pull requests to open-source projects), there could also be a natural check and balance that is appropriately scaled

for bad actors peddling misinformation. As these local federations gather critical mass, we can direct existing technical and institutional structures — repositories, collections, and search databases — or emergent distributed infrastructures —such as [distributed knowledge graphs](#) — to curate and index subsets of them for sharing beyond lab groups, for conversations with policymakers and practitioners, facilitating larger centers and research consortia, and so on.

I believe a future with this shape would be marked by sustained, growing contributions to shareable discourse graphs. By substantially lowering the overhead to synthesis, such infrastructures could in turn power more sustainable, accelerated scientific progress across disciplines.



□ [read more essays in the “Business of Knowing” summer series](#) □