

Beyond Text: Characterizing Domain Expert Needs in Document Research

Sireesh Gururaja Nupoor Gandhi Jeremiah Milbauer Emma Strubell
Language Technologies Institute, Carnegie Mellon University, Pittsburgh, PA, USA
{sgururaj, nmgandhi, jmilbaue, estrubel}@cs.cmu.edu

Abstract

Working with documents is a key part of almost any knowledge work, from contextualizing research in a literature review to reviewing legal precedent. Recently, as their capabilities have expanded, primarily text-based NLP systems have often been billed as able to assist or even automate this kind of work. But to what extent are these systems able to model these tasks as experts conceptualize and perform them now? In this study, we interview sixteen domain experts across two domains to understand their processes of document research, and compare it to the current state of NLP systems. We find that our participants processes are idiosyncratic, iterative, and rely extensively on the social context of a document in addition its content; existing approaches in NLP and adjacent fields that explicitly center the document as an object, rather than as merely a container for text, tend to better reflect our participants' priorities, though they are often less accessible outside their research communities. We call on the NLP community to more carefully consider the role of the document in building useful tools that are accessible, personalizable, iterative, and socially aware.

1 Introduction

From contextualizing scientific research in literature reviews, to understanding the functioning of complex organizations, experts conduct a wide variety of tasks that depend on document research. Document research, i.e. reading, understanding, and otherwise working with collections of documents is a process that underlies almost all knowledge work. As such, there is a rich body of literature that aims to understand how experts in various fields read documents, characterizing goals, processes, and how their experiences and knowledge inform what and how they read (Bazerman, 1985; Hillesund, 2010; Mysore et al., 2023, *inter alia*).

More recently, primarily text-based NLP tools

such as LLMs have been proposed as "solutions" to document research. General purpose commercial models are billed as being able to "understand" both single documents and even whole corpora, context length limits willing, and there are a growing number of purpose built tools targeted at particular professions, from legal document tasks (Wiggers, 2024; Merken, 2024; Ravaglia, 2024) to aiding in the process of scientific discovery (AI4Science and Quantum, 2023), with some claims going as far as to argue that some parts of the scientific process could soon be wholly automated (Lu et al., 2024).

But to what degree are LLMs able to model document research and understanding as currently carried out by experts? In this study, we interview 16 domain experts working in materials science, law, and policy, to understand their processes of document research: Their goals, the uses they have for these documents, and how they evaluate the documents' content for relevance and quality.

In our analysis, we derive tasks common to the experts we interview, and assess the degree to which modern tools from NLP and adjacent fields address those tasks. We find that our experts' processes are highly personal and varied, involve iteratively constructing mental model, and are consistently informed not only by the content of the documents, but by the material and social context of their production. Tools that design around this context, such as citation-aware tools for scientific support (He et al., 2019; Heimerl et al., 2016), better reflect experts' processes, though they tend to be limited to domains that explicitly center publication structure.

By contrast, modern, general-purpose systems from NLP tend to reflect a common, information-centric view: Documents are merely containers for information, and that information within documents can therefore be segmented, decontextualized and displayed without regard for the source

document. This is reflected in the affordances of the systems themselves, which often operate either on individual sentences, or on segments whose lengths are determined by an underlying model's context length or a chunking algorithm (Asai et al., 2023), which may or may not align with semantic boundaries (Qu et al., 2024).

The experts we interviewed view documents in ways that cohere much more strongly with theory found in the Science and Technology Studies (STS) literature: that documents are not merely conduits for information, but are traces of social processes whose details are crucial to the documents' interpretation and use. For our experts, these details can inform evaluating a document's provenance, assessing a document against background knowledge of a field, or considering the global context in which a document exists. In other words, a document "serves not simply to communicate, but also to coordinate social practices" (Brown and Duguid, 1996). We therefore call for the NLP community to develop **accessible** systems and methodologies that are more **personalizable**, **iterative** and **socially aware** (§6) in order to more accurately reflect the views and priorities of their users, and the rich social context in which documents are produced and consumed.

2 Related Work

Understanding Reading. We locate this work in the tradition of work that seeks to understand expert readers and their processes for working with the documents they read, which often take the form of interview studies. Mysore et al. (2023), for example, conducts semistructured interviews and think-aloud sessions with data scientists for how they conduct literature review. We also see many similarities with the findings in Bazerman (1985), who finds that physicists rely on "purpose-laden schemas", which include models of both content and authorship and other metadata, similar to what we find with materials scientists.

Document Theory. We draw on the STS literature for theories of document-centric views of knowledge work. Lund (2009) provides a broad overview of document studies beginning in the early 20th century, focusing on the materiality and social production of documents through the digital age. He points out the shift in focus away from documents towards disembodied information in library sciences in the 1960s. Frohmann (2004)

corroborates this shift (albeit with slightly different dates) and places it contemporaneously with "discourses of...artificial intelligence and informatics," while arguing for an understanding of the contingent, social role of the scientific publication. These two sources illustrate the a possible origin for the elision of the document in the framing of contemporary NLP. Brown and Duguid (1996) similarly argue that a document, rather than being a "conduit" for information, serve as a mode of social coordination and control in their production and distribution.

Document-aware Reading Support. Document awareness as a principle of system design is an active area of research, especially in human-computer interaction and information retrieval. The Semantic Reader project (Lo et al., 2024) incorporates a number of these features into a reader that shows an enriched view of a PDF document, and works like He et al. (2019) and Heimerl et al. (2016) allows exploration through citations and other metadata. Work in NLP that accounts for metadata like citations, such as Viswanathan et al. (2021), is less common. We note, however, that metadata is seldom considered in reading support work outside of scientific documents, and personalized reading support, i.e. work that considers the context of the reader, is also rare.

Challenges to NLP. There is a growing ambivalence in NLP towards the practice of benchmark-based evaluation (Gururaja et al., 2023). This paper joins a growing number which call for benchmarks to be more closely aligned to end-user needs. Newman-Griffis et al. (2021) call for what they term "translational NLP;" which proposes an application focus as the driver of scientific progress. Katz et al. (2023) propose a new benchmark that presumes the strength of LLMs at traditional NER tasks as the basis for proposing a much more difficult benchmark that better reflects information seeking needs.

3 Methods

We recruited 16 participants, beginning with a convenience sampling method (Galloway, 2005), in which the authors began by interviewing existing non-computer science collaborators across the projects they worked on, and then by snowball sampling (Parker et al., 2019), in which participants were asked to recommend other interview candi-

dates. Our participants, six women and ten men, were drawn from collaborations in the materials science, law, and policy communities, were all based in the U.S. and had a wide age distribution, with five participants between ages 25 and 34, six between 35 and 44, three between 45 and 54, and two 55-64. Ten of our participants were associated with materials science, though many identified themselves as belonging to other disciplines, such as chemical or mechanical engineering. All of our participants in this group were either professors or postdoctoral researchers whose primary focus was the synthesis, characterization, or modeling of materials. As such, the document research that they described to us was primarily the process of literature review: keeping up to date with the subfields they already worked in, or learning about new subfields that became interesting to their work. The remaining six participants were academics and professionals whose jobs involve reading, researching, interpreting, or otherwise engaging with law, public policy, or governmental records. This population's document research operated on many more kinds of documents, but the informational goals were largely similar, e.g. staying abreast of relevant policy, legal precedent, or public reactions to policy.

These populations are neither very similar nor dissimilar; we use them as a way of understanding what themes in how experts work with documents might generalize across groups with nominally different tasks, and what themes might be specialized to a single domain. In essence, we aim to establish a loose lower bound of the variety of tasks that professionals across different domains carry out.

We conducted semi-structured interviews (Weiss, 1995) with our participants that lasted between 27 and 73 minutes, with the median interview lasting 53 minutes. Interviews were conducted with a dedicated notetaker, and we recorded the interview audio with participant consent. We followed an interview guide (included in Appendix A), that developed four broad themes: The participant's current work and positionality, their current process for document research including how it fit into their work, how they evaluate documents for relevance and/or quality, and finally what existing tools they use to perform document research. We developed this set of questions in collaboration with materials scientists on one author's project, and was evaluated with a test interview before wider interviewing, with only minor changes to the wording of some

questions. The guide did not change between the two subpopulations we interviewed. We conducted the interviews between February and July 2024.

Following the procedures of grounded theory (Strauss and Corbin, 1990), at the conclusion of each interview, authors produced an analytical memo detailing the themes from that interview (Glaser et al., 2004). After a sufficient number of interviews for recurring, coherent themes to emerge, authors began a process of independently open coding the data, developing a thematic taxonomy that generalized across interviews. After this, the authors met regularly to discuss and refine the open codes into a preliminary closed coding frame (Miles and Huberman, 1994). Authors then annotated each interview with the themes from this closed coding frame. In analysis meetings, the authors further refined this closed coding frame by adding, merging, and removing codes, then iteratively re-coding the data. The analysis of this paper emerged from closed-coded versions of the data, and was validated against the original transcripts for appropriate context.

4 Summary of Interviews

Document research implied a wide variety of activities to our participants. Despite the variety of tasks, however, several common threads emerged that bridged the disciplinary gap. We conceptualize the tasks carried out by our participants across domains to broadly fall into three categories: Local context tasks, global context tasks, and corpus construction. In the following sections, we first characterize the logistics of the documents that our participants work with, describe some of the examples of each of these three types of task, as well as some broader themes across participants.

4.1 Document Characteristics

The documents our participants worked with were primarily, though not always, in PDF format (with exceptions including maps and raw data files), and originated across a wide temporal range. Several participants described having to work with scanned documents which may or may not have had OCR applied to them. In the case of materials scientists, many often consulted technical reports from the 1950s to 70s that were originally typewritten; some of our policy experts looked at digitized government documents that were hand-scanned. Even in cases where documents had been produced digi-

tally, understanding rich visual content, like page layouts, tables, and charts was a key area of focus for nearly all of our participants.

4.2 Task types

Local context tasks. Local context tasks, which only involve the content of a single document, resemble common information extraction tasks. In the materials science context, this often manifested as extracting information about how experiments were conducted, the materials that resulted, and their associated properties, in keeping with the principles of data-driven design of materials (Himanen et al., 2019; Olivetti et al., 2020). For instance, participant 21 described how their students “*work on collecting information from the articles, and...build the models that can predict the material properties*” (21). Information extraction tasks were also present in the policy domain, with one participant describing “*trying to extract policy data from these plans, including...different entities, different policy parameters*” (1). Local context tasks are only carried out once a researcher has already developed a mental model for the content of the documents they are searching, and are conducted only in corpora and subfields that the researcher knows well.

Global context tasks. By contrast, global context tasks, which rely on signals from other documents, or the focus document’s connections to them, are often much more exploratory. Common to all of our participants was the task of coming to understand a new corpus or subdomain with which they were previously unfamiliar; our experts were frequently reading and working with documents outside their core area of expertise. In materials science, this was described as “*building my own intuition of a classic material science thing...if I vary this, it goes up or it goes down.*” (2); in the law and policy domain, this could be understanding a corpus of government communications obtained through a freedom of information request, or the ramifications of a new policy through the public response to it. In these cases, information extraction-like approaches are insufficient: as one participant put it, “*it wouldn’t be...sufficient to say...we’re just searching for a needle in a haystack...we’re interested in understanding the haystack.*” (39). Researchers described constructing a mental model that was subdomain- or corpus-specific, progressively refining that mental model through the encounter with more documents. Over time, this al-

lowed them to develop intuitions and expectations of the corpus, which also provided signal when they were subverted. This process parallels the background knowledge integral to evaluating the novelty and epistemic status for materials scientists’ keeping up-to-date with their existing interests. Global context tasks were seen as a precursor to local context tasks: only after building a reliable mental model for an area did our participants feel comfortable looking at individual documents one at a time.

Corpus construction. While materials scientists consistently described a standard workflow that relied on academic search engines, these types of resources were only available to our participants in law and policy in the case of firm-internal documents or legal precedent. More commonly, participants explained that it was not trivial to collect corpora, beginning with “*putting in phone calls to various libraries...to sort of find out what material is available*” (39). Our participants frequently described having assemble their corpora from documents “*chopped up into different chapters*” (8), or that contained a “*reference to some other document that contains the relevant information*” (1). Constructing corpora in this way often involves a great deal of expertise, both in knowing what to include, by conducting “*manual verification of the documents we retrieve to understand if they actually are the documents we’re seeking.*” (1), and verifying “*the degree to which it is complete or extensive, that’s an important consideration.*” (1).

4.3 Broader themes

Awareness of contemporary technology. We asked about experts’ current processes, including the tools they had used, or whether they had incorporated AI tools into their workflow. While many of our participants were technically sophisticated, with some training their own BERT-based models to do topic classification, or writing an emacs-based tool for paper discovery, usage of LLM tools like ChatGPT was largely constrained to non-document tasks like writing or code assistance. By contrast, when attempting to use them for document research tasks, they described a number of pitfalls. Some expressed doubts about the lack of specific background knowledge, or concern about ceding control of the research process; others described trying to get models to work for their process and facing challenges. One participant, for instance, said “*can*

we just dump a bunch of PDFs into a GPT and get a summary? And it turns out it wasn't that easy. It wasn't like plug and play. But it also was showing some potential." (17). Though some of our participants had evaluated contemporary tools, they still considered them models that had to be customized, as in earlier machine learning, rather than useful as drop-in tools.

Access to technology. Of the tasks that we heard described, many are the subjects of active research in the NLP and NLP-adjacent communities. However, very few of our participants had access to these technologies, primarily because only a few of them could or chose to in furtherance of their tasks. This was perhaps nowhere more evident than in the case of digitization and OCR, where one participant described scoping a project based on *"whether the material is digitized or whether we're going to need to digitize it."* (39), later discussing how a digitization system that considered each page a separate document prevented them from using keyword searches effectively: *"if we had reliable high quality text, and we had our document organization [taken] into consideration...then we could have used a keyword based search as like a candidate classifier."* (39)

Personalization. Regardless of field, a consistent theme that we observed in our interviews was how idiosyncratic each researcher's process of document research was. There seemed to be no agreement, even within fields, on what makes a document relevant to a given search or what cues a researcher might use to assess the epistemic status of individual claims. We view this as a major challenge for NLP systems.

5 Task Analysis

In this section, we identify key tasks that our participants shared across fields, and compare the current state-of-the-art NLP tools with the needs that our experts outlined.

5.1 Information Extraction

Traditional IE, where we tag entities of concrete, well-defined types, would be useful for many of our participants. However, our participants' needs were for concepts that would be specific to their research objectives (i.e. not available in off-the-shelf models) and difficult to define succinctly, like tagging spans that provide evidence of 20-25 core political

values in an argument, like *"equality and justice, liberties, security, safety"* (42). Participants also described the acceptable granularity of extracted information varying per-project: *"we just used like more bag-of-words-based, keywords-based approaches...but the whole research project in that case assumed a level of bluntness that wouldn't be appropriate for other projects."* (39) This is different to standard benchmark datasets for IE where the types are more concrete and well-defined (Ding et al., 2021; Tjong Kim Sang and De Meulder, 2003). Further, supervised neural systems require a large amount of expensive annotations for each new tag set (Li et al., 2020), though recent work on few-shot IE with LLMs has aimed to reduce the potential annotation burden (Hofer et al., 2018; Ashok and Lipton, 2023; Huang et al., 2020).

Our participants also emphasized that not all IE tasks of interest to them involve only local information, conflicting with traditional IE focused on within-document and short document settings. They describe settings where IE would be applied to long scientific papers and policy briefs: *"I don't think there's a way for me to get the information that I need without having the full document, but once I have that document, I will only use the specific portions that I need."* (11). These documents may exceed the context window of LLM-based IE systems and result in poorer performance (Dagdelen et al., 2024).

Evaluating IE tools extrinsically can reveal the significance of existing vulnerabilities. The value of better aligning evaluation principles with user needs was highlighted by one of our participants: *"it would still dramatically reduce the amount of time a researcher would need to spend... But the conclusion was that these tools... were inadequate on their own"* (39) For example, Adams et al. (2024) shows that LLMs do not perform well enough on long document clinical notes for reliable clinical use in question-answering. In relation extraction, there has been work to soften evaluation metrics to accommodate for the use of generative models, reflecting a shift towards evaluation methods that are aligned with downstream utility where it is often enough to recover spans with overlapping span boundaries (Jiang et al., 2024). These types of evaluations would better highlight the most impactful open problems in IE.

5.2 Multimodality and OCR

Understanding visual and layout features of documents was a priority for nearly all of our participants, the vast majority of whom worked with documents prominently featuring tables, charts, or which conveyed information through layout. One materials scientist characterized their needs as “*looking for a statement backed up by data and the data can be just numbers. It can be graphs. It can be pictures of microstructures or just all*” (0), and another described extracting “*a ton of data and figures...because they’re usually X, Y plots.*” (5). These tasks are inherently multimodal: while some LLM approaches recommend linearizing tables into text, charts, images, and layout information must be handled visually. Compounding this is the tendency for these documents to be scanned or photocopied instances of paper documents, but which still necessitate the processing detailed above, resulting in it being “*tough to really search through them, and so they might not...appear when you’re doing a lit review on Google Scholar*” (2).

While visual understanding of documents is not a solved problem, dedicated multimodal layout understanding models like the LayoutLM series (Huang et al., 2022, *inter alia*), which serve both as visual segmentation models as well as representations for downstream document reasoning tasks remain an active area of research, general-purpose models like Qwen-2-VL (Wang et al., 2024) and LLaVA-Next (Liu et al., 2024) now include visual document understanding and OCR benchmarks like DocVQA (Mathew et al., 2021) and TextVQA (Singh et al., 2019), and proprietary models like Claude now have modes designed explicitly around PDF processing¹.

5.3 Iterative Search and Exploration

One common theme repeated across multiple participants was that the search process is inherently iterative. Rather than rely on a single set of results identified for a particular information need, researchers will iteratively expand their search across multiple stages, using results to inform each successive step, progressively building a mental model of the search space.

Researchers iterate for different reasons. Some seek to identify the provenance of the information they find: “*I go and read a current paper and I*

find where they cited that they got information from. And I read where that person got the information from and I read where that person got the information from. I usually try and find the original sources to everything.” (5), with one even stating “*sometimes I read a paper and the main use of it is the references in the paper*” (4). Others iterate precisely to establish the global context surrounding a particular paper: “*you have to like build up this context around the paper. What came before it? What is it citing? And how does the content of that paper relate to the ideas that came before it?*” (4).

Iterative approaches do exist within the information retrieval literature. Initial retrieved documents can be used as a source of information on related lexical terms (Attar and Fraenkel, 1977), which additional can help address issues like terminology drift (section 5.4), and the process of learning as you search, where “*reading and coding...helps you generate, develop contextual knowledge*” (39). More recent work has focused on iterative approaches which search through the structure of documents (Min et al., 2019; Zaheer et al., 2022; Hsu et al., 2024), albeit with a focus on the task of multi-hop reasoning (Yang et al., 2018) rather than information exploration.

Iterative construction of mental models could also be supported by the task of ontology induction. While there has been recent work to induce concepts from individual documents (Matos et al., 2024), our participants highlighted the need for concepts detected at the corpus level that are tailored to a specific research question. Despite initial work in the scientific domain (Katz et al., 2024), abstracting concepts across documents has been shown to be challenging for state-of-the-art LLMs (Guo et al., 2024), even without the per-question adaptation.

5.4 Terminology

One of the consistent difficulties that our participants faced across fields, similar to those in Mysore et al. (2023), was terminology. Our participants described three types of terminology shift. The first, temporal, is when the meaning of a word or a term drifts over time. This might be because of shifts in community usage, like one materials scientist pointed out, “*in like 2008-ish, the Chinese community decided that the word shock [testing] also means high rate Kolsky bar testing*” (5). It could also be because the referents of the words themselves have changed, as in the case of committee

¹<https://docs.anthropic.com/en/docs/build-with-claude/pdf-support>

organization in a local government: “if you’re looking for committee reports from before 2018 and you’re looking for hospitals, the committee on hospitals didn’t exist until 2018.” (40). The second, domain-specific change, is when different fields use different words or terms for very similar things: “I call it a surrogate model. If you ask a statistician what they’re going to call it, they’re going to call it an emulator. If you ask someone in the reliability community...they’re going to call it a response surface.” (17). These two types of terminology shift are partially addressed in the existing literature. Periti and Montanelli (2024) survey approaches to track new usage and senses for existing vocabulary; Lucy et al. (2022) quantifies domain-specific terminology usage and synonymy across fields; Head et al. (2021) provide context-sensitive definitions of technical terms and mathematical symbols.

However, while existing systems work on either identifying terms as near-synonyms or providing definitions, our participants emphasized that understanding the differences in meaning and why one term might be used instead of another was also important. For example, one materials scientist outlined how while “oxidation” and “aqueous corrosion” meant similar processes, the keywords used to search for one vs the other, and the numbers that would characterize those properties would be different: “if it’s aqueous corrosion...they might care about the atmosphere or basically the liquid concentration a lot more...But if you were then going to oxidation in high temperature, they are mostly looking at mass gain data.” (2) The simultaneous focus on the similarities and differences in mostly-synonymous words reflects a process in which our participants tended to jointly model the semantic content of documents alongside the documents’ social context, like its authorship or intended audience. This was especially true with the final type of terminology shift, political, where people describe similar things, but may use different language to convey different valence, or signal which aspects of an issue are being prioritized, as one participant gestured to in the case of climate policy: “rural communities will talk about micro grids, not as a part of climate action, but as a way to get off the dependency on investor owned utilities like PG&E.” (8).

In these cases, our participants not only had to engage in an iterative process to find the synonyms, they also had to reason about the positionality of the authors and why they might use a different

term. Understanding identity, its presence in corpora, and its interaction with LLMs is still a new area of research: Kantharuban et al. (2024) and Li et al. (2024) demonstrate the sensitivity of models like ChatGPT to implicit markers of identity in responses and refusals, respectively, and Lucy et al. (2024) investigated author positionality and community belonging, and Milbauer et al. (2021) demonstrated cross-community lexical differences linked with ideology. Reading support that addresses the needs of our participants would need to unify several running themes of work in an accessible interface: providing definitions, enriching those definitions with understanding of where assumptions and practice across communities might differ and the implications of different terminology across domains.

5.5 Corpus Construction

For corpus construction, which we describe in section 4.2, our participants most often described starting with collections of library resources and Google searches, only eventually moving to write web scrapers if the kinds of documents were similar enough. Automated tools, such as custom scrapers that use LLMs to explore documents (Huang et al., 2024; Ma et al., 2023) would be extremely useful for expanding corpora beyond what is feasible to manually collect. However, our participants stressed the need to reconstruct documents from chapters, understand document versioning, and how to iteratively build “some checks of what’s missed by [a search], false positives, false negatives” (42) in constructing corpora, implying that document scrapers for constructing corpora would need to accommodate an iterative, exploratory style to truly function for this purpose.

5.6 Global and Social Context

When discussing how they handled the uncertainty of dealing with potentially contradictory information from multiple documents, our participants repeatedly made clear that content alone was rarely sufficient either to search for documents, or to assess the results of such a search, with one participant stating that document content accounted for “at best...40 to 45%” (40) of what made a document trustworthy. In order to determine the degree to which statements and document content could be trusted, our participants relied on a number of complementary contextual signals. In this section, we discuss the signals our participants used when

working with corpora. We note, however, that most of the issues outlined in this section remain underexplored in language technologies; as such, we primarily describe desiderata for future efforts.

Perhaps the most commonly used signal for our participants was consistency checks against contextual knowledge of their field and “standard practices” within it. Even in the domain of scientific publications, where published, peer-reviewed papers are commonly taken as “vetted and approved by the scientific establishment” (Cronin, 1984), materials science participants reiterated the need to critically evaluate papers and their results based on “how well rooted [they are] on the first principles of the field,” (28) and whether a given document passed their notion of a “self-consistency test” (0), asking “Does this graph support comments and assertions made in the text?” (0), even though they were the intended audience of the paper. If numerical results seemed unlikely (by intuition), researchers saw this as an indicator that a paper required additional scrutiny: “it looks like they’re kind of cherry picking their data to make it look good. And you want...full disclosure. Are they highlighting cases where maybe the thing doesn’t work?” (17) These standards were also seen as changing over time: one materials scientist pointed out that “These days, titanium alloys are used for very specific applications, which is then going to be biasing the data that’s collected on them” (28), and a policy researcher noted that “planning history is full of these best practices, which turn out to be failures 10, 20 years later.” (8) This poses a challenge to current NLP systems, which do not model the “unexpectedness” of a particular statement or document within the context of similar documents, instead treating most content in documents as propositional and uniform in importance, making them vulnerable to spin (Yun et al., 2025). Liu et al. (2023) and Milbauer et al. (2023) have recently developed approaches for linking information across documents, however the field still works within a framework of resolving conflicts in propositional knowledge (see e.g. Xu et al. (2024), who survey approaches to mitigating knowledge conflicts in models and conclude that current approaches remain insufficient), rather than dealing with the contingent, temporally and socially bound truth values found in documents.

Participants also often used metadata like authorship as part of assessing document reliability across all fields. However, we found that the signals that

participants chose to rely on were highly individual, with two participants rarely agreeing on how to evaluate a document. We heard contradictory opinions on the value of metadata like publication venue, citations, the principal investigator of the lab group that published the paper, and even the quality of the diagrams and charts within the paper from materials scientists. Even in cases where a participant used a certain signal, they emphasized its contextual nature. In the case of citations, for example, the age of the subfield was an important qualifier for one participant, who said “There’s an enormous variation in citation rate depending on the field and how specialized it is. So anything that’s old and highly specialized like steel research, numbers are tiny. There’s just so few people active.” (0) Both of these signals — background knowledge and metadata — can be seen as authors using a mental model of how their field functions in order to evaluate a document. Further, the document is the unit of analysis that allows both of signals to be used - consistency checks operate at the level of individual document components interacting, and the document is the unit which can be published, cited, and so on. Document-aware models and tools tend to be restricted to science, where metadata pertaining to documents is readily available through resources like S2ORC (Lo et al., 2020), but the necessity for document-based understanding is not: signals from the social context of document production like authorship and provenance were also important for our law and policy participants.

Having a social model of the production of a document is even more important for our participants in cases where they are not the original or intended audience for a document, such as when examining planning documents, or communications obtained through government transparency programs. In these cases, participants often report having to piece together the information they seek by comparing multiple documents, including differences between versions of the same document, to get a better picture. For example, one participant details inferring political positioning for potential appointments in local government by “looking for a pattern of how someone may not be saying like their views, but in actuality are sort of expressing it through something that is legally binding.” (40) Another participant detailed a case in which the differences between versions of a document indicated important information about the honesty of the information: “they write a first draft, and then

they literally edit out some of the material that's too honest for public consumption...I can't treat both of those pieces of information as equally important...the first one I know is more forthcoming, and the second one is more whitewashed." (39), where "a concept that's articulated infrequently relative to other concepts can be much more important for understanding the, like, underlying story." (39). As discussed in 5.4, modeling the types of social processes underlying document production is still new to NLP. Further, while document similarity is a robust and well-studied topic, finding meaningful differences in documents that might indicate information being hidden is much less common, with some work being published in legal NLP venues (Li et al., 2022).

6 Conclusion

In this paper, we interview 16 domain experts from the domains of materials science and law/policy to understand how they conduct document research. We find that their self-described processes rely on understanding and actively modeling the social processes by which text is produced, through the construct of a document and its associated metadata. As exemplified by the success and usage of layout understanding models and the HCI work on scientific reading support, document-aware tools match the processes of our experts. More concretely, however, there are four qualities of systems that we call on the NLP community to build into new document research tools for expert users outside the field:

Accessible Though we have characterized the state of existing research as it addresses many of the concerns of our experts, we note that the most accessible NLP tools today by far are primarily text-based and rarely consider documents a first-class citizen. By contrast, NLP work that does center documents is far less accessible. Many of the issues faced by our participants, whether working with older, undigitized documents, understanding terminology drift, or leveraging metadata to assess the provenance of a document are already the focus of existing research, but are not as accessible to them as web-based LLM systems. How can we both promote better modeling of how users engage with large quantities of text in accessible systems like LLMs, and make tools that do this modeling more accessible?

Personalizable We note in several places that our experts had idiosyncratic processes for finding, evaluating, and reading documents. From looking for different types of information both within and across fields, to having vastly different heuristics for assessing the provenance and reliability of a document, our participants, all successful experts in their fields, conduct research in highly personal, specific ways. Amid concerns of LLMs potentially stifling creativity and serving as a homogenizing force (Anderson et al., 2024; Kumar et al., 2024), NLP tools should encourage and support diversity of thought by allowing for experts to customize systems to their existing personal processes.

Iterative Almost all of our participants discussed the process of constructing a mental model as a series of iterative updates as they read new documents and reconciled the content with their expectations. This process of constructing a mental model was essential, and heavily used global context: experts relied extensively on signals like how any given document related to background knowledge, other documents in the corpora, standard practices in the field, as well as social signals like authorship and author positionality. As they read, their understandings of both the signals and how they pertained to the documents co-evolved. By contrast, NLP tools are often presented as either static or occasionally updated artifacts. To better support how experts work, NLP systems should malleably support their users' evolving mental models, including through easy schema updates, flexible data relabeling, and user-friendly retraining and evaluation loops.

Socially aware The social character of document production played an important role for our participants, whether in understanding terminology and why it might be used, modeling authors and participants in spheres of discourse they participated in, and evaluating information for reliability and provenance that included information beyond the propositional content of the text they worked with. Our participants frequently looked beyond what was written in text to how and why it was written, treating documents more as traces of social processes than containers or conduits for pure, disembodied information. NLP tools should be designed around the idea that social context is crucial for understanding text: authorship, audience, communicative intent, and format are all factors that readers already consider; as readers already know, models should be aware that no text is *just* text.

Limitations

Our interview study was evaluated and approved by an Institutional Review Board as STUDY2023_00000431.

While we constructed our sample deliberately to span two fields, we acknowledge that this was a convenience sample, and that it is neither specific to one field, nor representative of all the ways that people might work with documents. Our study also focuses exclusively on potential users of NLP tools who are already experts in their domains. Their concerns are not likely to be representative of non-domain expert users, especially non-experts reading technical language.

Acknowledgments

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-22-2-0121. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- Lisa Adams, Felix Busch, Tianyu Han, Jean-Baptiste Excoffier, Matthieu Ortala, Alexander Löser, Hugo JWL Aerts, Jakob Nikolas Kather, Daniel Truhn, and Keno Bressen. 2024. Longhealth: A question answering benchmark with long clinical documents. *arXiv preprint arXiv:2401.14490*.
- Microsoft Research AI4Science and Microsoft Azure Quantum. 2023. [The Impact of Large Language Models on Scientific Discovery: a Preliminary Study using GPT-4](#). *arXiv preprint*. ArXiv:2311.07361 [cs].
- Barrett R Anderson, Jash Hemant Shah, and Max Kreminski. 2024. [Homogenization Effects of Large Language Models on Human Creative Ideation](#). In *Proceedings of the 16th Conference on Creativity & Cognition*, C&C '24, pages 413–425, New York, NY, USA. Association for Computing Machinery.
- Akari Asai, Sewon Min, Zexuan Zhong, and Danqi Chen. 2023. [Retrieval-based language models and applications](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 6: Tutorial Abstracts)*, pages 41–46, Toronto, Canada. Association for Computational Linguistics.
- Dhananjay Ashok and Zachary C Lipton. 2023. Promptner: Prompting for named entity recognition. *arXiv preprint arXiv:2305.15444*.
- Rony Attar and Aviezri S Fraenkel. 1977. Local feedback in full-text retrieval systems. *Journal of the ACM (JACM)*, 24(3):397–417.
- Charles Bazerman. 1985. [Physicists Reading Physics: Schema-Laden Purposes and Purpose-Laden Schema](#). *Written Communication*, 2(1):3–23. Publisher: SAGE Publications Inc.
- John Seely Brown and Paul Duguid. 1996. [The Social Life of Documents; introduction by Esther Dyson](#). *First Monday*.
- Blaise Cronin. 1984. *The Citation Process: The Role and Significance of Citations in Scientific Communication*. T. Graham. Google-Books-ID: k1HN2uqm5LAC.
- John Dagdelen, Alexander Dunn, Sanghoon Lee, Nicholas Walker, Andrew S Rosen, Gerbrand Ceder, Kristin A Persson, and Anubhav Jain. 2024. Structured information extraction from scientific text with large language models. *Nature Communications*, 15(1):1418.
- Ning Ding, Guangwei Xu, Yulin Chen, Xiaobin Wang, Xu Han, Pengjun Xie, Haitao Zheng, and Zhiyuan Liu. 2021. [Few-NERD: A few-shot named entity recognition dataset](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3198–3213, Online. Association for Computational Linguistics.
- Bernd Frohmann. 2004. [Scientists and Other Information Users](#). In *Deflating Information*, From Science Studies to Documentation, pages 53–91. University of Toronto Press.
- Alison Galloway. 2005. [Non-probability sampling](#). In Kimberly Kempf-Leonard, editor, *Encyclopedia of Social Measurement*, pages 859–864. Elsevier, New York.
- Barney G Glaser, Judith Holton, et al. 2004. Remodeling grounded theory. In *Forum qualitative sozialforschung/forum: qualitative social research*, volume 5.
- Shaoru Guo, Yubo Chen, Kang Liu, Ru Li, and Jun Zhao. 2024. [NutFrame: Frame-based conceptual structure induction with LLMs](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 12330–12335, Torino, Italia. ELRA and ICCL.
- Sireesh Gururaja, Amanda Bertsch, Clara Na, David Widder, and Emma Strubell. 2023. [To build our future, we must know our past: Contextualizing paradigm shifts in natural language processing](#). In

- Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13310–13325, Singapore. Association for Computational Linguistics.
- Jianguan He, Qing Ping, Wen Lou, and Chaomei Chen. 2019. **PaperPoles: Facilitating adaptive visual exploration of scientific publications by citation links**. *Journal of the Association for Information Science and Technology*, 70(8):843–857.
- 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, CHI '21*, New York, NY, USA. Association for Computing Machinery.
- Florian Heimerl, Qi Han, Steffen Koch, and Thomas Ertl. 2016. **Citerivers: Visual analytics of citation patterns**. *IEEE Transactions on Visualization and Computer Graphics*, 22:190–199.
- Terje Hillesund. 2010. **Digital reading spaces: How expert readers handle books, the Web and electronic paper**. *First Monday*.
- Lauri Himanen, Amber Geurts, Adam Stuart Foster, and Patrick Rinke. 2019. Data-driven materials science: status, challenges, and perspectives. *Advanced Science*, 6(21):1900808.
- Maximilian Hofer, Andrey Kormilitzin, Paul Goldberg, and Alejo Nevado-Holgado. 2018. Few-shot learning for named entity recognition in medical text. *arXiv preprint arXiv:1811.05468*.
- Sheryl Hsu, Omar Khattab, Chelsea Finn, and Archit Sharma. 2024. **Grounding by trying: LLMs with reinforcement learning-enhanced retrieval**. *Preprint*, arXiv:2410.23214.
- Jiaxin Huang, Chunyuan Li, Krishan Subudhi, Damien Jose, Shobana Balakrishnan, Weizhu Chen, Baolin Peng, Jianfeng Gao, and Jiawei Han. 2020. Few-shot named entity recognition: A comprehensive study. *arXiv preprint arXiv:2012.14978*.
- Wenhao Huang, Zhouhong Gu, Chenghao Peng, Zhixu Li, Jiaqing Liang, Yanghua Xiao, Liqian Wen, and Zulong Chen. 2024. **AutoScraper: A Progressive Understanding Web Agent for Web Scraper Generation**. *arXiv preprint*. ArXiv:2404.12753 [cs] version: 1.
- Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. 2022. **Layoutlmv3: Pre-training for document ai with unified text and image masking**. In *Proceedings of the 30th ACM International Conference on Multimedia, MM '22*, page 4083–4091, New York, NY, USA. Association for Computing Machinery.
- Pengcheng Jiang, Jiacheng Lin, Zifeng Wang, Jimeng Sun, and Jiawei Han. 2024. **GenRES: Rethinking evaluation for generative relation extraction in the era of large language models**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2820–2837, Mexico City, Mexico. Association for Computational Linguistics.
- Anjali Kantharuban, Jeremiah Milbauer, Emma Strubell, and Graham Neubig. 2024. Stereotype or personalization? user identity biases chatbot recommendations. *arXiv preprint arXiv:2410.05613*.
- Uri Katz, Mosh Levy, and Yoav Goldberg. 2024. Knowledge navigator: Llm-guided browsing framework for exploratory search in scientific literature. *arXiv preprint arXiv:2408.15836*.
- Uri Katz, Matan Vetzler, Amir Cohen, and Yoav Goldberg. 2023. **NERretrieve: Dataset for next generation named entity recognition and retrieval**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3340–3354, Singapore. Association for Computational Linguistics.
- Harsh Kumar, Jonathan Vincentius, Ewan Jordan, and Ashton Anderson. 2024. **Human Creativity in the Age of LLMs: Randomized Experiments on Divergent and Convergent Thinking**. *arXiv preprint*. ArXiv:2410.03703 [cs].
- Jing Li, Aixin Sun, Jianglei Han, and Chenliang Li. 2020. A survey on deep learning for named entity recognition. *IEEE transactions on knowledge and data engineering*, 34(1):50–70.
- Victoria Li, Yida Chen, and Naomi Saphra. 2024. Chatgpt doesn't trust chargers fans: Guardrail sensitivity in context. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6327–6345.
- Xiang Li, Jiaxun Gao, Diana Inkpen, and Wolfgang Alschner. 2022. **Detecting Relevant Differences Between Similar Legal Texts**. In *Proceedings of the Natural Legal Language Processing Workshop 2022*, pages 256–264, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. **Llava-next: Improved reasoning, ocr, and world knowledge**.
- Nelson F Liu, Kenton Lee, and Kristina Toutanova. 2023. Anchor prediction: Automatic refinement of internet links. *arXiv preprint arXiv:2305.14337*.
- Kyle Lo, Joseph Chee Chang, Andrew Head, Jonathan Bragg, Amy X. Zhang, Cassidy Trier, Chloe Anastasiades, Tal August, Russell Authur, Danielle Bragg, Erin Bransom, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Yen-Sung Chen, Evie Yu-Yen Cheng, Yvonne Chou, Doug Downey, Rob Evans, Raymond Fok, Fangzhou Hu, Regan Huff,

- Dongyeop Kang, Tae Soo Kim, Rodney Kinney, Aniket Kittur, Hyeonsu B. Kang, Egor Klevak, Bailey Kuehl, Michael J. Langan, Matt Latzke, Jaron Lochner, Kelsey MacMillan, Eric Marsh, Tyler Murray, Aakanksha Naik, Ngoc-Uyen Nguyen, Srishti Palani, Soya Park, Caroline Paulic, Napol Rachatasumrit, Smita Rao, Paul Sayre, Zejiang Shen, Pao Siangliulue, Luca Soldaini, Huy Tran, Madeleine van Zuylen, Lucy Lu Wang, Christopher Wilhelm, Caroline Wu, Jiangjiang Yang, Angele Zamarron, Marti A. Hearst, and Daniel S. Weld. 2024. [The semantic reader project](#). *Commun. ACM*, 67(10):50–61.
- 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*, pages 4969–4983, Online. Association for Computational Linguistics.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. 2024. [The ai scientist: Towards fully automated open-ended scientific discovery](#). *Preprint*, arXiv:2408.06292.
- Li Lucy, Jesse Dodge, David Bamman, and Katherine A. Keith. 2022. [Words as Gatekeepers: Measuring Discipline-specific Terms and Meanings in Scholarly Publications](#). arXiv. Version Number: 2.
- Li Lucy, Suchin Gururangan, Luca Soldaini, Emma Strubell, David Bamman, Lauren Klein, and Jesse Dodge. 2024. [AboutMe: Using self-descriptions in webpages to document the effects of English pretraining data filters](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7393–7420, Bangkok, Thailand. Association for Computational Linguistics.
- Niels Windfeld Lund. 2009. [Document theory](#). *Annual Review of Information Science and Technology*, 43(1):1–55.
- Kaixin Ma, Hongming Zhang, Hongwei Wang, Xiaoman Pan, Wenhao Yu, and Dong Yu. 2023. [Laser: Llm agent with state-space exploration for web navigation](#). *arXiv preprint arXiv:2309.08172*.
- Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. 2021. [Docvqa: A dataset for vqa on document images](#). In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209.
- Emanuel Matos, Mário Rodrigues, and António Teixeira. 2024. [Towards the automatic creation of NER systems for new domains](#). In *Proceedings of the 16th International Conference on Computational Processing of Portuguese - Vol. 1*, pages 218–227, Santiago de Compostela, Galicia/Spain. Association for Computational Linguistics.
- Sara Merken. 2024. [Ai company luminance raises \\$40 mln as contracts tech investment booms](#). *reuters.com*. Accessed: 2024-12-15.
- Jeremiah Milbauer, Ziqi Ding, Zhijin Wu, and Tongshuang Wu. 2023. [NewsSense: Reference-free verification via cross-document comparison](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 422–430, Singapore. Association for Computational Linguistics.
- 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*, pages 4832–4845, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Matthew B Miles and A Michael Huberman. 1994. *Qualitative data analysis: An expanded sourcebook*. sage.
- Sewon Min, Danqi Chen, Luke Zettlemoyer, and Han-naneh Hajishirzi. 2019. [Knowledge guided text retrieval and reading for open domain question answering](#). *arXiv preprint arXiv:1911.03868*.
- Sheshera Mysore, Mahmood Jasim, Haoru Song, Sarah Akbar, Andre Kenneth Chase Randall, and Narges Mahyar. 2023. [How Data Scientists Review the Scholarly Literature](#). In *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval, CHIIR '23*, pages 137–152, New York, NY, USA. Association for Computing Machinery.
- Denis Newman-Griffis, Jill Fain Lehman, Carolyn Rosé, and Harry Hochheiser. 2021. [Translational NLP: A new paradigm and general principles for natural language processing research](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4125–4138, Online. Association for Computational Linguistics.
- Elsa A Olivetti, Jacqueline M Cole, Edward Kim, Olga Kononova, Gerbrand Ceder, Thomas Yong-Jin Han, and Anna M Hiszpanski. 2020. [Data-driven materials research enabled by natural language processing and information extraction](#). *Applied Physics Reviews*, 7(4).
- Charlie Parker, Sam Scott, and Alistair Geddes. 2019. [Snowball sampling](#). *SAGE research methods foundations*.
- Francesco Periti and Stefano Montanelli. 2024. [Lexical Semantic Change through Large Language Models: a Survey](#). *ACM Comput. Surv.*, 56(11):282:1–282:38.
- Renyi Qu, Ruixuan Tu, and Forrest Bao. 2024. [Is semantic chunking worth the computational cost?](#) *Preprint*, arXiv:2410.13070.
- Ray Ravaglia. 2024. [Donotpay: Ai agents improve student financial aid applications](#). *forbes.com*. Accessed: 2024-12-15.

Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards vqa models that can read. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Anselm Strauss and Juliet Corbin. 1990. *Basics of qualitative research*. Sage publications.

Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.

Vijay Viswanathan, Graham Neubig, and Pengfei Liu. 2021. [CitationIE: Leveraging the citation graph for scientific information extraction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 719–731, Online. Association for Computational Linguistics.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.

Robert S Weiss. 1995. *Learning from strangers: The art and method of qualitative interview studies*. Simon and Schuster.

Kyle Wiggers. 2024. [Openai-backed legal tech startup harvey raises \\$100m](#). *Techcrunch.com*. Accessed: 2024-12-15.

Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024. [Knowledge conflicts for LLMs: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8541–8565, Miami, Florida, USA. Association for Computational Linguistics.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*.

Hye Sun Yun, Karen Y. C. Zhang, Ramez Kouzy, Iain J. Marshall, Junyi Jessy Li, and Byron C. Wallace. 2025. [Caught in the Web of Words: Do LLMs Fall for Spin in Medical Literature?](#) *arXiv preprint. ArXiv:2502.07963 [cs]*.

Manzil Zaheer, Kenneth Marino, Will Grathwohl, John Schultz, Wendy Shang, Sheila Babayan, Arun Ahuja, Ishita Dasgupta, Christine Kaeser-Chen, and Rob

Fergus. 2022. Learning to navigate wikipedia by taking random walks. *Advances in Neural Information Processing Systems*, 35:1529–1541.

A Interview Guide

Begin by defining document research: we’re interested in processes you have for finding documents, or things within documents to help you in research tasks in a professional context.

Demographics/Positionality

1. Can you briefly describe your job, covering the kinds of research questions that you encounter in your line of work?
2. In your work, can you describe the cases where you have to look for documents, or things within documents? An example would be great.

How do you do document research now?

3. Can you describe the goals that you have when you do document research? What kinds of documents and information do you search for? If you have different kinds of searches with different goals, please describe them.
4. When searching for documents, are you searching for documents as a whole, or specific pieces of information/facts within those documents? How much of the document’s content as a whole do you end up using?
5. What purpose do those documents or facts have once you find them? (reference? Quotation material? Prior approaches to what you’re trying to solve?)
6. Is there an existing ontology to the kind of searching that you do? Are the things that you search for in documents part of a well-defined set of things, or is your approach to these documents creative?
7. To what degree is finding documents or facts an iterative process? Is there a mental model that you have of the space of possible documents that you update as you find new documents?
8. To what extent is the structure of documents relevant?

How do you evaluate the documents that you find?

9. When you search for documents or facts, what does it mean for a document or fact to be high quality to your purpose?
10. When evaluating a document or fact for relevance or quality, how much of that depends on the content of the document itself?
11. How much of that depends on knowledge of the field that you have that's not explicitly in the document - other important documents, standard practices, etc? Are there resources for that kind of domain knowledge?
12. How much depends on metadata, like citations, author affiliations, venue, etc?

Existing tools

13. What tools do you currently use to search for documents?
14. When executing a search, how quickly do you usually find the sort of thing you're looking for? Are there specialized keywords that get you to what you're looking for?
15. If you use specialized tools for your domain, what do they do differently from generic-domain tools, like Google search?
16. Do you ever write code to enable better searching? What are the tasks that code helps you with that existing tools are insufficient for?
17. If there was something that your tool could do differently or better, what would it be?
18. Have you used AI-based tools to aid in your work? How well have they suited your workflow and process?

Demographics

19. I am going to read some age brackets. Can you indicate when I read a bracket that your age falls into?
 - 18-24
 - 25-34
 - 35-44
 - 45-54
 - 55-64
 - 65+
20. Is there anything else in your background that you consider relevant?