

Building a Better Mousetrap: Tools and Processes for Selling A Company

Chelsea Kerr
chelsea@zuva.ai
Zuva Inc
Toronto, Canada

Sasha Vtyurina
sasha@zuva.ai
Zuva Inc
Toronto, Canada

Adam Roegiest
adam@zuva.ai
Zuva Inc
Toronto, Canada

ABSTRACT

It is a fact of life for many start-ups that they must sell part of their company (i.e., fund raising) in order to have enough capital to grow the company to one day successfully exit the market. The unfortunate side effect of this necessity is that it places a large burden on start-ups to respond to information requests from potential buyers which then forces employees to step away from their day jobs to formulate responses. While it has been the norm to respond to such requests using manual review of contracts and other information sources, the increasingly competitive funding market has resulted in growing time pressure for all participants of start-up purchasing endeavours. Furthermore, current technological offerings often fall short of providing optimal support to the start-up and the buyer which continues to reinforce a process that is often cumbersome and chaotic.

In this work, we present an analysis of 19 interviews of regular participants on both sides of the sales process finding that the main pain points revolve around document management, request tracking, internal and external collaboration. Based upon this analysis, we describe an early-stage prototype to investigate a possible solution for efficient handling of buyers' information requests. We recruited 12 participants to test this prototype and found that the main issue was misalignment between language in the tool and participants' mental models. From these two sets of analyses, we present potential implications and considerations for building tools for infrequent but high-risk and high-reward information tasks.

ACM Reference Format:

Chelsea Kerr, Sasha Vtyurina, and Adam Roegiest. 2023. Building a Better Mousetrap: Tools and Processes for Selling A Company. In *ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '23)*, March 19–23, 2023, Austin, TX, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3576840.3578299>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHIIR '23, March 19–23, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0035-4/23/03...\$15.00

<https://doi.org/10.1145/3576840.3578299>

1 INTRODUCTION

"I don't understand it. Everyone spends their time trading shitty data rooms, and emails, and diligence trackers in Excel and nothing... Just this horrible, horrible, horrible situation and there's got to be a better mousetrap¹ out there but I haven't found one..."

A Study Participant

Selling a company in whole or part (e.g., fund raising) is often the route for many start-ups to achieve some form of long-term viability and success.² Despite recent economic downturns, the global market for merging or acquiring companies saw a volume of US\$642 billion between July and September 2022 indicating that this is still a common and lucrative activity [40]. For such deals to take place, however, the seller must go through a process, commonly called sell-side due diligence, which involves employees from the selling company (or the selling company's legal counsel) to interrupt their day-jobs to respond to information requests that will help the buyer determine what they believe is an appropriate valuation for the risks present in the company. Should this process be successful, the buyer will then engage in a no less arduous process, commonly called buy-side due diligence, where they verify the responses provided as well as look for more nuanced risks and conflicts relevant to the buyer. Sell-side due diligence is particularly impactful to the selling company since regular infusions of capital can help extend their runway to show viability and value which ideally leads to a subsequent complete acquisition. Combined with the fact that sell-side due diligence can be very disruptive to the regular flow of business and has the potential to negatively affect the bottom-line if not managed well (Section 3.2). Buy-side diligence has seen some prior investigation [24, 28, 42] but is also managed quite differently from sell-side diligence (i.e., almost entirely by a law firm). Due to this difference and particular importance on "getting this right" for the selling company, we provide a high-level comparison between the two processes in Section 2.1 and discuss sell-side diligence only for the remainder of this work.

When conducting sell-side due diligence, the selling company is tasked with responding to a *due diligence request list* (Figure 1) that contains multiple sections with each section typically corresponding to a subset of the organization (e.g., general policies and organizational chart, finance, legal, human resources). The exact information requested in any given request list will differ based upon the buying entity and the domain (e.g., healthcare, manufacturing, e-commerce) and so will the exact form of deliverable(s) (e.g.,

¹A better mousetrap" is an English idiom referring to the idea that any improvement to current mousetrap technologies would net the inventor much fame and fortune.

²Without loss of generality, we use "buy" and "sell" in this paper regardless of whether this is in the form of investment or outright acquisition due to some proportion of controlling interest in a company changing hands.

more mature companies would typically produce more thorough responses than a company with one or two customers). To investigate the high-level commonalities and the various perceptions of the sell-side due diligence process, we conducted 19 semi-structured interviews with individuals from both buy and sell sides focusing on their background and process involvement in sell-side due diligence, the collection and management of responses to the request list, and the outputs produced by the process that are then externally shared (Section 3.2).

Based upon our analysis of the interviews, we found that a serious lack of dedicated tools has seemingly impacted the ability for companies to conduct sell-side diligence in a less disruptive way (i.e., lots of manual tracking in spreadsheets). In an attempt to fill this gap, we describe our early stage machine learning powered tool that aims to assist sellers in more efficiently and effectively managing the diligence process by allowing them to create filters out of requests (Section 4). This tool allows users to leverage the results of machine learning to pull out relevant contract language, perform keyword filtering, and perform comparisons and boolean operations to perform this filtering. We evaluate the usability of this filter creation with 12 participants to attempt to determine where pain points have been remediated, where they remain, and if any are made worse.

We conclude this paper with a discussion of some implications of our user studies to sell-side diligence and other tasks that share similar properties (Section 5). This discussion focuses on tasks that are performed infrequently enough that participants do not create a fully fleshed out mental model and tasks that require high-level collaborative information seeking but with individual goals and motivations being at odds with other collaborators. In addition to these implications, we discuss how complex retrieval tasks, like sell-side diligence, limit how well researchers are able to model the tasks and perform user research on prototypes in a realistic task setting.

2 BACKGROUND

2.1 Buy-side and Sell-side Diligence

Selling a (part of a) company is very similar to selling a house (or a time share in one) in that the main goal of the seller is to do what is necessary to get the most money as quickly as possible [2]. On the other hand, buyers want to get the best value for their dollar while mitigating as much risk as possible. When selling a house, the buyer is will initially ask a small set of relatively informal questions (e.g., “when was the roof last done?”) to get a sense of what work, if any, might need to be done. In selling a company, buyers (e.g., investors, other companies) need to initially determine whether the company they want to buy is worth the asking price, whether the potential risks are worth the potential upside, and whether they believe a better price can be negotiated. In this *sell-side due diligence* process the seller provides responses to a list of requests supplied by the buyer that is based upon information in contracts, business software, policies, and, sometimes, the minds of employees. This list is called a *due diligence request list* (Figure 1) and corresponding responses provided by the seller is called a *disclosure schedule* which can be exceedingly complex or quite terse.

ITEM	
2. MATERIAL AGREEMENTS	
2.1	All banking documentation, loan agreements, promissory notes, indentures, guarantees, lines of credit, letters of credit, security agreements, pledges, contracts, commitments and other evidence of indebtedness. Verify all long term indebtedness of the Corporation, including original and current amounts outstanding, the due dates and interest rates.
2.2	Information about guarantees in respect of any indebtedness or any obligation of any other person with respect to the Corporation.
2.3	Any business combination, acquisition, merger, divestiture or share (or other security) purchase agreement involving the Corporation, as well as any correspondence, whether of a preliminary nature or otherwise, in connection therewith.
2.4	Copies of all letters of intent of the Corporation as well as any correspondence in connection therewith, whether of a preliminary nature or otherwise.
2.5	Government sponsored loans, government assistance programs, grants, subsidies, etc.
2.6	All significant marketing, management, service, consulting or any other similar type of contract.
2.7	All agreements, contracts or commitments limiting the freedom of the Corporation to engage in any line of business, to compete with any other person or to solicit clients or employees or imposing obligations of confidentiality.
2.8	Significant equipment and service leases.
2.9	Significant insurance policies and commitments (property, vehicles, liability, directors and officers, key-man, etc.).
2.10	All agreements, contracts or commitments which might reasonably be expected to have an adverse impact on the business or operations of the Corporation.
2.11	All contracts or agreements with or pertaining to the Corporation and to which directors, officers or owners of more than 5% of the shares of the Corporation are parties, including but not limited to promissory notes, receivables and/or payables.
2.12	All contracts, agreements or commitments entered into in the ordinary course of business, which cannot be terminated or cancelled without penalty within 30 days.

(a) Requests for all material agreements to a company.

ITEM	
5.9	A copy of the Corporation's most recent business plan.
6. REAL PROPERTY AND PERSONAL PROPERTY	
6.1	With respect to real property of the Corporation:
	(a) list all real property and interests in real property (including all leases) owned, leased or used by the Corporation and copies of all deeds, leases, subleases, sales contracts, appraisals, tenancy agreements, easements, licenses, royalty agreements, surveys, title reports and policies of title insurance or lawyer's abstract reports and any other documents pertaining to title of the real property
	(b) copies of any expert reports, appraisals or valuations prepared with respect to real property of the Corporation
	(c) copies of all outstanding mortgages and encumbrances on real property of the Corporation
	(d) copies of any title opinions previously rendered with respect to real property of the Corporation
	(e) copies of any outstanding offers to lease, purchase or sell any real property
	(f) details on any outstanding lease assignment or subleases
	(g) copies of any options affecting real property
	(h) copies of all contracts, arrangements and easements relating to services and utilities
	(i) copies of all property tax data including valuations with respect to all real property
6.2	Copies of all applicable zoning and building code information.
6.3	Copies of any notices or permits received from local zoning or building authorities and a list of any properties (owned or leased) that do not comply with local zoning or building codes.
6.4	A list of all machinery, equipment, furniture, fixtures, vehicles, and similar items of personal property owned or leased by the Corporation together with a schedule showing the terms of ownership/lease, the total cost, depreciation reserve and rates (if applicable), lease payments (if applicable) and present book value (if applicable)
6.5	A list of any other machinery, equipment, furniture or vehicles used in the business of the Corporation but not listed in 6.4 above.

(b) Requests for agreements relating to property.

Figure 1: Excerpts from a real due diligence request list.

In turn, the buyer performs more in-depth diligence to ensure that responses provided by the seller are true and factual as well as to ask for further elaboration on any newly revealed insights. This process is typically called *buy-side due diligence* (or some times mergers & acquisitions due diligence) and is typically handled by a law firm representing the buyer which performs an exhaustive review of the seller's documents. In buying a home, this is equivalent to hiring a home inspector to do a thorough inspection of the house and property history before making an offer. Interestingly, the market for buy-side diligence software grew over the last decade (e.g.,

Luminance, Seal Software, Kira Systems, Diligen, Eigen, Della) [1] due to high-cost associated with having highly paid law firms conduct this exhaustive review combined with entire departments in law firms focused on performing this function. Conversely, sell-side due diligence has not seen a corresponding maturity, we believe, due to the fact that it is a relatively infrequent task performed by different individuals in a company’s lifetime.

2.2 Collaboration

During a sell-side diligence process, multiple individuals interact with each other to accomplish a singular goal: “answer buyer’s requests for information as quickly and accurately as possible”. In some sense, this is a goal comprised of many “mini” tasks to find information or determine an answer to the request similar to the *individual search + no shared display* case in the work of Crescenzi et al. [15]. While sell-side diligence requires coordination of many individuals, much existing collaborative information seeking and retrieval research tends to focus on a small set of individuals working together to satisfy a singular information need but insights from this research will be applicable to sell-side diligence. For example, the seminal work Morris and Horvitz [37] identified three key aspects of collaborative web search: *awareness*, *division of labour*, and *persistence*. Following this work, Shah et al. [45] called for an addition of metadata to search results, and Shah [44] found that the number of coordination messages was decreased with the improved transparency into the partner’s activity. This observation extends to scenarios where a collaborator is an intelligent agent as observed in recent work by Avula et al. [4, 5] which indicates process tracking and visibility are important components of collaborative retrieval. As any single person is unable to complete the entire task due to the lack of knowledge, it is ever so important to facilitate effective interaction between people who are involved in the process [25].

As we discuss in Section 3.2 and 3.3, some of the main challenges faced by the sell-side diligence participants revolve around the need for efficient task distribution and visibility into the process. Ironically, one of the biggest similarities between the sell-side diligence process and collaborative web search might be the users “satisficing³ with current tools” [22] which manifests in participants using emails, instance messaging, and phone calls to coordinate and manage the work [11, 36] in both task domains.

2.3 Time Pressure

When investigating the effects of time pressure on information retrieval and seeking behaviours, there has generally been a split between time-critical (i.e., life or death) pressure [26, 35, 41] or more explicit time-constrained pressure (e.g., find as much as you can in N minutes) [16–18, 31, 32, 46]. We believe, sell-side diligence falls closer to time-constrained tasks due to implied constraints on individual time from others managing the process as well as a final deadline for diligence to be completed. The general trend observed in past work is that time pressure often results in a sacrifice of relevant material returned and some task-specific behavioural adaptations. In a professional setting, time-pressure in clinical decision

support has been observed to reduce accuracy and increase individual stress during the task [46] which aligns with observations discussed in Section 5.2. A sell-side diligence tool is similar to the search assistant in Crescenzi et al. [16] which was unused by some participants as they felt that they could find information quicker and more reliably themselves when under time pressure. This indicates that any tool for sell-side diligence should be perceived as low effort and highly reliable.

2.4 Task Complexity

As computing hardware and information access have expanded so too have the tasks that users of information systems attempt to complete. Ranging from factoid and non-factoid question answering [7] to timeline summarization of social media [30] to interacting with conversational agents [12] to recommending what products to buy [9], these tasks move beyond ad-hoc search to satisfy ever more complex information needs. Moreover, there has been an increasing investigation into “professional” search tasks such as electronic discovery [14, 38], systematic review [27, 29], and clinical decision support [21, 39]. In these cases, the task being investigated is often a (sizable) part of a much larger professional task (e.g., finding all relevant COVID-19 vaccine efficacy study and then leaving the meta-analysis related tasks to subsequent work). We find that the “search as learning” paradigm [13, 20] and other related exploratory search tasks are closest to sell-side diligence. In the sense that the seller is conducting this search and synthesis on behalf of the buyer in order for the buyer to learn about the company and make a determination to buy.

Prior work explored interfaces for complex search, sensemaking and document organisation activities [8, 23, 33] which could be compared to the process of due diligence.

3 UNDERSTANDING SELL-SIDE DILIGENCE IN PRACTICE

3.1 Methodology

To facilitate an understanding of the sell-side diligence process, we conducted semi-structured interviews with professionals that have been involved with either the sell-side or buy-side. Interviews covered three lines of questioning: (1) an initial discussion of the interviewee’s background and their experience of the sell-side diligence process; (2) how information is collected, curated, used in responses to requests, and collaborated on within an organization; (3) how outputs are produced and shared with the relevant parties. Interviews were kept to 30 minutes with some occasionally running to allow participants to complete any thoughts.

Participant selection was done by soliciting relevant members of our organization’s executive team’s professional network. This introduced an inherent bias to our sampling procedure but was largely unavoidable due to the nature of the research to be conducted. To mitigate this, we interviewed a thorough mix of professionals as possible to ensure an accurate representation of sell-side diligence in practice. In total, we interviewed 6 Mergers & Acquisitions lawyers, 7 corporate lawyers, 4 venture capital vice-presidents, 1 banker, 1 Chief Financial Officer, and 1 company co-founder. Interviewing stopped once no new insights were being uncovered

³Satisficing is a decision-making strategy that aims for a satisfactory or adequate result, rather than the optimal solution.

and no additional participants were available from different job positions to increase interview diversity.

Interviews were recorded with participants' consent, were automatically transcribed, and corrected for accuracy during review. Thematic analysis using MAXQDA was conducted to derive well-structured insights about key pain points and process hurdles as well as a valuable functionality that a sell-side diligence tool would require to be successful. Due to resource constraints, thematic analysis was conducted by a single researcher who may not have been present for the interview itself.

3.2 Analysis

We note that some interviews described the inciting incident for sell-side diligence to be either having an unsolicited bid arrive or to actively seek out a buyer. For obvious reasons, we omit discussion of this pre-diligence step and focus on aspects of the sell-side diligence process itself.

3.2.1 Document Management. While it would not be an unreasonable assumption that organizations are diligent in keeping good records about where important documents, contracts, and other company critical information exists. Nearly all of our interviewees that had assisted in sell-side diligence confirmed that such an assumption often does not hold in practice. Several participants expressed the idea that it was simply a matter of course to rely on the longest tenured employees as sources of truth for where this information resides (P4: *"I reached out to him and brought him in because he had been with the company since...not the very beginning but about a year or two into the company's history. So he knew where a lot of things were."*). This is perhaps a bit surprising until put into the context that companies taking part in sell-side diligence are not necessarily focused on best practices for data organization but on activities that support the bottom line (P21: *"...don't have very strong processes and controls, because they, you know, have only been in existence a few years. Things are moving quickly they don't have a big team..."*), P3: *"Early days companies don't even know they are supposed to keep track of it. Doesn't even occur to them ... They are all busy running around trying to build a company."*). One participant, in particular, believes it is critical that documents be digital during sell-side diligence. Otherwise, they opine *"[i]f all your contracts are on paper, you're screwed. If your contracts are digital, you're only slightly less screwed."* (P2). Accordingly, there is an initial hurdle in just *finding* all required materials and requires potentially non-trivial coordination to do so.

3.2.2 Managing Requests. Once a seller has performed the nearly Sisyphean task of finding all the necessary material, they begin the actual crux of the process in responding to the request list. Not only does responding involve the inputs of many departments (P2: *"...there are so many questions that aren't legal. Tons of finance related, tons of data related..."*), P4: *"...the remaining lists all focus on product and business and HR and retention and financials"*) but it involves what can be *"multiple rounds of information requests"* (P4) which creates a time and effort sink across the organization (P20: *"it could be a company [that we'd] maybe spend 20 hours a month with, suddenly becomes 40 or 50 hours a month"*), P4: *"we would literally be*

meeting every day in the evening and just sit down and go through [the request list] question by question and line by line.")

Request management becomes even more complex if there are several prospective buyers vying for the selling company since they have to manage multiple request lists that are phrased differently but are asking for, in many cases, the same information (P3: *"So, you have a standard list that goes across corporate, IP, employee benefits, tax."*, P1: *"...independent of the company, there's a relatively common list of questions that people would ask in the first instance. And then coming out of a call there will be clarification questions."*, P2: *"There's always a bunch of questions you want to ask that are the same, like 'Do you have any outstanding lawsuits?', 'Any environmental orders against your company?'"*). The difficulty lies in ensuring all similar questions have a consistent answer across bidders as it minimizes the work done and can expedite the process overall (P4: *"If I could go back in time and recreate the process, ...if I could get a list of possible questions that I know everyone's going to ask and then create a standard response for them. That would have eliminated 80% of the work. Because then you're just saying 'hey, I know everyone's going to ask this question, here's the answer' and it's the same answer for everyone."*, P21: *"what most people do is they just create one big data room, they accumulate all the questions from everyone else. And then you're kind of like almost like you're playing cards - you're dealing out the information as needed from that main data room."*) as well as reducing the risk of providing different answers. As one can imagine from this discussion, without rigorous processes and tools managing requests even from one bidder, let alone several, can be an arduous task.

3.2.3 Collaboration. Collaboration might be better characterized in the context of sell-side diligence as coordination of resources since each involved department needs access to their own part of the request list (P2: *"The techie folks don't need to see the contracts, they just need to have access to see their particular piece of it."*). In contrast to many collaborative information tasks, individuals will often perform their task then return to their day-jobs and leave the coordination of the overall endeavour to those assigned to that task. This is often a C-level executive and their reports⁴ (P13: *"The inefficiency of this process and the resource-intensive nature of this process is really problematic. Because you don't have bodies lying around. So you're taking people ... maybe your technology people instead of building product or fixing bugs, or your sales people who, instead of selling deals, have to worry about diligence stuff... The company's not being run. So you're having a real negative business impact."*, P17: *"But it's just trying to get your arms around everything. Because typically you are doing your normal job at the same time, so it's really keeping everything on track is difficult."*).

As mentioned above, collaboration can come in discussing what to provide to the buy-side and is where the sell-side law firm typically gets involved (P1: *"your law firm is helping to do some of the legal due diligence calls or guidance about how to do it and then strategic considerations and discussions, just advice about how to approach something from a sell perspective with your buyers."*). Accordingly, this framing of responses is critical to ensure that buyers are provided enough data to make risk assessments and, ultimately,

⁴We note that the very consistent term across interviewees for this role was "quarter-backing" but refrain from using it here to keep this sport-agnostic.

whether to continue the transaction but also that the right story is told about the data that is present (e.g., less than ideal terms were given for a contract that forms the backbone of the seller's annual revenue).

3.2.4 Tools. A consistent response from all interviewees that had conducted sell-side diligence was that there was a relatively small set of common tools: a virtual data room⁵ (e.g., Intralinks, HighQ, and sometimes Sharepoint, Google Drive, or Dropbox) and some form of spreadsheet software (e.g., Microsoft Excel, Google Sheets). This is not entirely surprising since spreadsheets are a straightforward tool to track and coordinate response while a virtual data room is the accepted way to share documents among legal professionals.

Out of all our interviewees, only 3 said that they had used contract review software tailored for buy-side diligence to aid in sell-side diligence. This was due to prior exposure to such tools and its effectiveness at expediting the review for key data points in contracts and other agreements. Such tools are tailored to professional law firms that do in-depth review during mergers and acquisitions and so come with a price tag and functionality that is representative of those tasks (i.e., it is tailored to legal professionals). It was no surprise that so few sell-side interviewees attempted to use such software since the ramp up and price tag would be prohibitive.

3.2.5 Non-Document Material. Much of the work on the sell-side consists of documents but participants frequently referenced the need to produce responses not based off documents. Such requests might correspond to the state of the business (P21: “go pull out all the stock based compensation amounts for last five years” or to historical elements (P3: “Explain to me about the history of the company, how did the IP get developed?”) to fiscal projections (P2: “What’s your projected revenue for next year? What was the earlier versions of that?”). Moreover, it may be that there are just generic questions that require a written answer but aren’t necessarily formalized in any document (P2: “You have any outstanding lawsuits? Any environmental orders against your company?”, P4: “...often there were just answers to questions, there weren’t really documents or anything that needed to be provided. It was like a “what is” question.”). There may also be a need to formalize existing business practices or views that may not have needed to be formalized (P21: “Things that didn’t need to be documented in a word document, so they were more questions of management, you know what’s your view on fraud, and how do you tackle fraud, you don’t have a policy around fraud.”). Consequently, when requests like these occur in sell-side diligence, it boils down to some number of humans figuring out an answer and producing a written response with any necessary evidence attached.

3.2.6 Buy-Side Considerations. A buyer’s focus in sell-side diligence is to quickly build an understanding of the contextual aspects of the selling company’s operations and any risks that may create should they buy (part of) the company. Much of the buyer’s interest is focused on finding answers about the material contracts and agreements of the company before then engaging in buy-side diligence (P18: “We would say, just look at the high level contracts, the big tickets, top few customers. And then we would ask for warranties about customer contracts.”, P3: “In the financing they don’t want to

know everything, they just want to know the top 10 things. <...> But when you do the sale of a company they’re basically saying ‘we want to know everything.’”).

There has been a recent trend, at least prior to the current economic climate, resulting from a very competitive funding market (i.e. venture capital and private equity investment) for buyers to curtail some of the more tedious and in-depth asks of sell-side diligence to prevent being scooped by a different buyer (P18: “And in the past few years it’s just gotten even more competitive on the European landscape. For VC it’s crazy. So the pressure is to sign deals. <...> you don’t have a lot of time. And you got a lot of ‘is this really important?’, ‘do we really care?’ We’ll iron it out later.”). By reducing the scope of sell-side diligence, buyers take additional risks and have less certainty about their purchases but are able to compete with their rivals. This means that buyers may become advocates for tools that allow them to return to more thorough request lists while maintaining competitiveness and keeping burdens on sellers low.

3.3 Elements of a Mousetrap

In this section, we synthesize the main elements that any tool seeking to be a “better mousetrap” for sell-side due diligence would need to tackle based upon our interviews. As many participants have suggested that companies seeking funding, especially early stage ones, will be more focused on running the business rather than having strict procedures and protocols in place for document management, we avoid making as many assumptions as possible about how an arbitrary sell-side company would operate in the following discussion.

3.3.1 Document/Evidence Repository. Managing where documents or any created evidence exists is useful for ensuring that the correct data is provided to buyers when the decision to turn over materials is made. This can encompass what some participants characterized as a “staging” area that allows internal employees to sort through different versions of documents before producing the final version for buyers. Additionally, ensuring that this repository can manage multiple on-going bids by interested providers would mean that sellers have more insight into what is being shared, with whom, and whether it is shared consistently (i.e., to prevent errors in reporting). While it is a daunting task, this repository would ideally connect with as many different data sources as possible to ensure easy importing of relevant materials without requiring manual intervention to move files around. Some connections, like those to virtual data rooms, are more likely to be consistently used but others that correspond to more mature organizations (e.g., Carta for stock option management) are beneficial to have but likely to be unused for a non-trivial portion of young companies. The more connections that can be made, the easier it is for sellers to provide correct and timely information to buyers but results in a tool that is more complex both for developers but also end users to manage.

3.3.2 Request and Response Management. Our participants very clearly expressed how time consuming and high touch managing sell-side diligence can be for a seller. Beneficially, a tool builder can leverage the success of contract review technology [1, 34] to aid in the review of contracts by automatically extracting relevant data

⁵A virtual data room is a shared digital location for all documents related to a particular deal and is the digital equivalent to a physical data room full of boxes.

points and using them to filter documents that meet the desired request criteria (e.g., “Provide all employment agreements where base salary exceeds \$150,000.”). Such technology is not perfect but can help focus any document review to ensure that manual review is kept to a minimum. Reducing manual review enables that manual effort to be better spent on handling requests that require human intervention (e.g., responding to the presence of any active lawsuits) and possible document creation. We expect that some written responses will also require evidence to support the answer (e.g., a copy of the capitalization table when reporting on unallocated stock options). The corresponding evidence, if any, would then be stored in the document repository. Regardless of how a request is responded to, we envision that being able to save requests and responses can go along way to reducing additional effort in the presence of multiple bids or in successive sales when doing multiple rounds of investment (i.e., some answers may be recyclable). To further reduce effort and potential errors, the tool can provide the ability to identify similar requests, possibly automatically.

3.3.3 Internal Collaboration. We view the goal of internal collaboration as ensuring that internal stakeholders are aware of what requests have been responded to, what new requests may have come in, who (if anyone) are these requests assigned to (e.g., to verify or manually answer), and being able to leave any necessary notes on requests or documents that may require further investigation or framing when presented externally (e.g., realizing that a contract in the repository is not a final version). Included in our vision of internal collaboration is the ability to manage access to arbitrary sections of the request list and data repository to facilitate focused input from different departments in the organization. This allows sell-side diligence managers to not worry about collating separate spreadsheets or providing overly broad access to an entire organization and the risks that could occur (e.g., accidental deletion of other responses). We might also expect that other common functionality in productivity software will be present in a sell-side diligence tool (e.g., notifying individuals when they are assigned a task) but refrain from digging deeply into the exact details since they are less relevant to this work. Overall, the described internal collaboration functionality aims to enable individuals to contribute as efficiently and effectively as possible and then return to their day-job to ensure the company remains desirable to buyers.

3.3.4 External Collaboration. Ensuring good external collaboration is probably one of the more crucial aspects of a sell-side diligence tool as it can colour the on-going relationship between buyer and seller. The seller is also likely to desire the most fine-grained control in this situation in order to be certain they provide the right information to the buyer. The existing alternative from our participants’ perspective is to perform the extra work of maintaining a “staging” virtual data room and sending spreadsheets of responses via email and then upload documents to final data room. With fine-grained control, sellers can reduce the amount of duplicative work in producing “safe to send” copies of work that they’ve already done and potentially given access to buyers sooner. In that vein, we can also envision the ability for buyers to submit new requests or clarifying questions in the tool itself to better manage who responds and when. The final component to external collaboration is producing the final set of deliverables to the buyer in order for

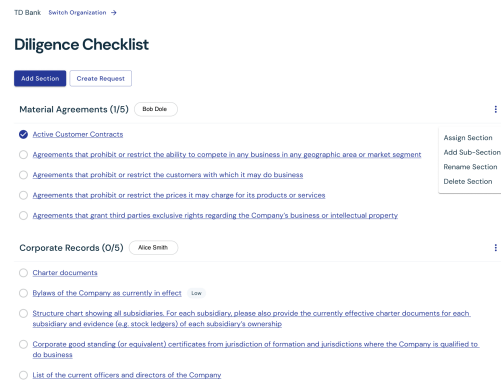


Figure 2: An example set of requests broken into sections for a fictional organization. Note that one request is marked as completed and others remain to be completed with sections being assigned to different individuals for verification.

them to perform any additional buy-side diligence. Ideally, the tool would connect with whatever virtual data room the buyer would like as well as facilitate the transmission of a disclosure schedule in the desired format. But practically speaking, we imagine that being able to export a single, well organized snapshot of the repository and the disclosure schedule will be seen as sufficient in most cases and likely a non-trivial improvement over the status quo.

4 A “BETTER MOUSETRAP”?

As discussed in Section 3.3, there are several broad areas that we believe a sell-side diligence needs to address to successfully improve the process from both sell-side and buy-side perspectives. Out of those areas, we think that the most practical area for human effort reduction is in identifying documents that are responsive to a particular request in the diligence request list. Especially as being able to formulate and create the necessary filters is much more of a user oriented problem rather than the more development heavy task of adding access controls and connectors to different data sources (e.g., Sharepoint, Salesforce). While there are certainly elements to be investigated in the other areas, prior research [16] suggests that ensuring an assistant to the search and filtering process is easy to use and not perceived as effort-heavy may improve the likelihood that it will be used. With that in mind, it was clear to us that helping users to easily find relevant documents to requests is the lowest-hanging fruit for improving the mousetrap.

In the remainder of this section, we describe a preliminary prototype revolving around the creation of search filters for requests in a due diligence request list and related management tasks. We tested this tool using new and repeat participants from the earlier study and present an analysis of their experience and what this implies for further iterations on our attempt to improve the mousetrap.

4.1 Request Management

Figure 2 depicts the initial prototype for request management. Users are able to create different sections which can correspond to the

Create a Request

Name:

Document Type: Shareholder purchase agreements | Select a Section: Material Agreements

Select document types you would like this filter to not automatically. The section this request belongs to.

Filters (Optional)

Add additional conditions to filter the documents added to the response to match the request.

Operator	Field	Value
AND	granted date	2022-07-30

If these conditions are met, supply the following outputs.

When these conditions are met, the document name and a link to the document will be added to the response for this request. You can choose to supply additional outputs such as specific field extractions or a text note.

Output Type	Field
Flag for additional review	granted date

Cancel Save

Figure 3: An example of a request filter creation screen that is meant to assist users in automatically filtering out documents that do not meet desired criteria for a particular request. In this example, this corresponds to any share or stock purchase agreements that have a stock granted date and that date is greater than July 30, 2022.

different sections in a due diligence request list but this organization is up to their discretion. Each section is assignable to one or more individuals in an organization to facilitate any checking of automatically generated responses or creating responses and associating them with the request (not pictured). Sections are assigned to individuals rather than to each request since our interview analysis indicated that sections are typically related to a particular organizational unit (e.g., legal, finance, IT) rather than grouped in some other manner. This screen also allows users to see at a glance the progress of the process and whether additional resources need to be made available to finish in a desired time frame (e.g., if one user is assigned to too many sections). We note that there isn't anything particularly nuanced about this screen and its functionality but helps to reinforce the experience that we expect users to take part in.

Our particular area of investigation is shown in Figure 3 which is an example request filter being created. The goal here is to allow users to easily combine pre-built machine learning models (referred to as "Field" in the prototype), document type restrictions, text search, and various logical and comparison operations to filter documents to only those that satisfy the desired request [19, 43]. Ensuring that filter creation be relatively painless and intuitive for users is where we believe the largest gains can be made and can pose the biggest barrier to adoption of a sell-side diligence tool (i.e., if this isn't easy then why use the tool). We note that this design was motivated with the understanding that not all documents would necessarily be present when such filters are created and that the results should be updated with new documents meeting the criteria.

4.2 Methodology

In order to validate our design choices, we invited 12 participants with various experience on buy and sell-sides (11 of whom we interviewed previously), and scheduled individual 30 minute remote video sessions. This still results in a biased sample of participants

Due Diligence Request List

In connection with the proposed investment by one or more affiliates of Awesome VC Fund into Spectacular Startup (the "Company"), please provide the following documentation and information for our review.

Where applicable, please provide the requested information for the Company, its predecessors (if any) and each subsidiary. If there are no documents that are responsive to a requested item, please indicate that there are no applicable documents for such item. We also request that until the proposed transaction is consummated, you submit for our review any new documents or other information of the type requested if they subsequently become available.

1. Material Agreements

- List of top 10 suppliers and customers and copies of their relevant contracts.
- Any other agreements which exceed \$100,000, individually or aggregated among affiliates
- Agreements that grant third parties exclusive rights regarding the Company's business or intellectual property
- All currently effective real estate property leases, including any amendments
- Agreements that prohibit or restrict (i) the ability to compete in any business in any geographic area or market segment, (ii) the customers with which it may do business or (iii) the prices it may charge for its products or services.
- Copies of all insurance policies covering the Company or its officers, directors or employees
- All agreements entered into by the Company out of the ordinary course of business or which are between the Company and any related party of the Company (including its officers and/or stockholders)

Figure 4: An modified real-world example of a request list that a seller could receive from a buyer. Participants of the second study were asked to build a filter for request 1.3.

but invited participants were generally those that we believed would use this part of the tool, at least some of the time.

During the call, participants were asked to imagine that their company was going through sell-side diligence and use the prototype in order to respond to the supplied request list. Figure 4 shows the example diligence request list that was presented to users and they were tasked with creating a filter that would seek to satisfy request 1.3. In the proposed scenario, all necessary documents have already been located and uploaded into the system, therefore the focus of the task was on using the prototype to select the correct documents to satisfy the request list. We note that the prototype users interacted with was a high-fidelity "paper" prototype that was created using HTML and JavaScript to emulate what would happen in a real tool. This was done to more readily get feedback on this part of the tool rather than waiting for a more complete end-to-end implementation.

During the sessions, we asked the participants to share their screen and to think aloud as they were attempting to complete the task. With the participants' permission, we recorded the sessions calls for later analysis. Notes were taken during the sessions and we continued recruiting until no new insights were being noticed. After all interviews were completed, one of the authors conducted a more thorough thematic analysis of the recorded videos using MAXQDA.

4.3 Analysis

The most prominent, and initially confusing, trend in participant recordings was around the use of the word "request" in the prototype. The majority of participants interpreted the "Create a Request" step as if they were requesting someone else to do the work of finding the documents. This interpretation is not entirely false as one may imagine requesting that the tool is tasked with sifting through the documents but it did not correspond to the idea of creating some kind of filter. A common alternative view suggested by participants was to frame this step more explicitly as search, whereby a user is establishing criteria to search the corpus of documents to then save them in a specific location in a data room.

There was also repeated confusion in how requests were named by participants. In the intended workflow, requests are treated as named entities mapping to a set of documents. However, our participants found it unintuitive and some interpreted the “name” input field (in Figure 3, this is “Share Ownership”) as an input for plain English search query. This is understandable considering that the mental model of most participants was parallel to ad-hoc search. They expected to complete their search, see the results, save them with an indicative name. Instead, the tool offered the inverse workflow: create the named entity, set search criteria, see the results. This misalignment is not unexpected due to the task scenario outlining that all documents are present in the tool rather than having to accommodate new documents over time.

One unexpected behaviour occurred around naming conventions for sections and requests. We expected that the names of sections and requests in the tool would mirror the wording in the example request list. Participants, however, consistently chose to use their own naming scheme throughout the study. Part of this may have manifested due to the artificial nature of the task but indicates to us that the tool may need to encourage closer alignment to the request list to ensure that users are certain that a particular request is reflected in the tool.

As shown in Figure 3, the prototype allowed participants to leverage several type of criteria, including text search and machine learning models (*Fields*). Participants were also able to limit the filters to certain pre-defined document types to more directly focus only on specific documents related to the request. During the study, we noticed that participants would often forego the filter creation step and skip to results only based on the document type. When asked, many explained this behaviour by wanting to set rough boundaries, examine the results and go back to add more filters if necessary. These “broad queries” can be explained by the fact that our participants were “novices” in using this tool [3]. We expect that after the initial introduction, users will spend more time setting up the search criteria to receive high quality results [6]. Further, as shown in prior research [10], since people are more likely to solicit the use of assistive tools for more complex tasks, it is plausible that requests requiring only a few documents can be fulfilled with simple filters.

One of the criteria types allowed participants to leverage a trained machine learning model to pull out specified concepts (e.g. all examples of non-compete clauses in a contract). Despite this option providing a potentially superior and efficient way to find the necessary documents, most of the participants issued text queries with the understanding that searching for an exact string match would require them to set up multiple conditions (e.g., different phrasing, handling capitalization) and may result in missing documents. Such reluctance to use machine learning models can be explained by the relative unfamiliarity of machine learning capabilities to participants (as mentioned in Section 3.2, tooling is limited). While interacting with a full-fledged tool over a high-fidelity prototype may change user behaviours, we plan to further improve discoverability and ease of use for this type of advanced functionality to further reduce barriers to use.

Regardless of filter type, participants expressed the desire to manually spot-check the output. This speaks to a “trust but verify”

approach towards assistive tools and signifies the fact that sell-side participants are not used to relying on sophisticated tools but rather are accustomed to conducting the work largely by hand. But we acknowledge that due to the stakes involved in sell-side diligence that some of this hesitance may remain even in a high-trust scenario.

5 BEYOND THE MOUSETRAP

5.1 Mental Models in Infrequent Tasks

In Section 4.3, one of the most surprising elements of the user study was that many participants were confused as to what a request was in the context of a due diligence *request* list. This was mildly frustrating for us since the participants themselves could not articulate a more appropriate name for “request” when prompted and none of our internal domain experts could either. We do note that “question” was perhaps the most consistent suggestion as replacement but it isn’t clear to us that this would be any better outright but may be consistent with the iterative process that can develop when buyers ask for further clarification.

After some lengthy discussions, we realized that for most participants, even those who have performed sell-side diligence several times, the process is not performed frequently enough to build a consistent and complete mental model. This is compounded by the fact that participants may not always be involved in all parts of the process whenever they participate in sell-side diligence. Accordingly, participants defaulting to ad-hoc search behaviours might result from consistent exposure in their own day-to-day job rather than suitability to the task (i.e., it is a fallback behaviour). Given that participants are focused on getting the work done as quickly as possible and likely with a focus on getting back to their day-job, it is no surprise that they may not fully develop a mental model for how they talk and think about the task. We find some indirect confirmation of this in the fact that, in our experience, lawyers engaging in buy-side due diligence often have a much stronger mental model about the task due to fact that it is their day-job.

We believe one of the most critical insights that may have broader implications is that mental model development for complex tasks is inversely correlated the frequency of the task. Moreover, when this task interferes with a more primary task even less attention is paid to it. It simply becomes something that stands in the way of getting other work done. The problem, of course, is that when the task also has a high-risk or high-reward component that lacking a mental model may mean errors or other problems can arise. To avoid this, we might look to help shape an innocuous mental model that encourages users to be consistent in their performance of tasks but not necessarily be cognizant of this (e.g., by using terminology that is generic enough to not pull them from the task).

5.2 Collaboration and Coordination

In Section 3.2, we saw that much of sell-side diligence is done individually and as quickly as possible to return to other day-to-day tasks. This means that while individual employees are supporting the overarching task of making sell-side diligence successful, they do so separate from the efforts of others and not having high-levels of collaboration. But there is a risk that interactions can be perceived “negative” when individuals are repeatedly get pulled back into the

process to provide clarification or provide additional responses. On the other hand, the interactions between buy-side and sell-side are more complex since both sides want the deal to complete but with different mitigating factors in mind. Sellers will want to do as little work as possible to secure the deal (as they have a business to run) and buyers want to obtain as much information as possible to judge risk and correct prices from sellers. This tension can obviously be manifest as adversarial interactions by either side since each have slightly conflicting underlying motivations in their work.

The problem in either case is that this can sour relationships either internally or externally between collaborators. As several participants expressed this being an undesirable (but not necessarily rare) outcome, avoiding such scenarios becomes a high priority since it can affect how successfully the company operates afterwards or may itself put the deal in jeopardy (P104: *“And what I found happens so many times is the diligence process becomes hostile, combative, and antagonistic for so many reasons...You go right from this really miserable diligence process into an integration process, and everyone still has the battle scars. And it really jeopardizes the success of a deal. Because the deal is not about ‘Okay, you bought the asset at the right valuation, you signed an M&A agreement, a stock purchase agreement, and now the deal is done.’ That’s the easy part. The hard part is integrating and working together. And if you spend six months beating each other up, you’re not necessarily going to have the best relationship going forward.”*). This can mean that there are interesting trade-offs in tools built to solve problems addressing a common overarching goal but competing with different individual goals from those completing the process. Balancing these different motivations can be further complicated by a less developed mental model and not understanding one’s role in the process. And so, tools for these processes may then need to be actively refined based upon different personas to ensure functionality does not exacerbate these tensions.

5.3 Limitations

In Section 4, we described a study focused only on a relatively small part of the sell-side diligence process in our high-fidelity prototype rather than a full fledged tool to perform sell-side diligence. The unfortunate reality is that creating a realistic simulation of such a task is not itself simple due to the inherent complexity of sell-side diligence resulting in needing enough representative documents and other information, finding enough participants willing to work together, and being able to justify the time commitment. A realistic sell-side diligence task is likely to require several hours worth of work by multiple individuals even if they were provided with a better mousetrap. Limiting the scope of a simulation would mean participants perform something more akin to a bare-bones usability test of the tool rather than actually performing the task. Especially since most participants may end up believing that the document set is so small that it would be easier just to do the work manually. Indeed, the best we could hope for (especially for comparative purposes) is to find a charitable company that had just completed a sale to “replay” their diligence process using the tool.

6 CONCLUSION

In this paper, we have presented the idea of sell-side due diligence and its importance to how companies, especially start-ups, sell themselves in whole or part. Through a diverse user study across the different perspectives in the sell-side diligence process, we found that document and information collection, managing buyer information requests, and internal and external collaboration are some of the biggest pain points in this process. Especially as sell-side diligence requires every day employees to dedicate non-trivial time away from their day-job to perform task related to selling the company, which can itself damage the company’s bottom line and productivity overall. Using request management as a starting point, we present a simple tool to create request filters that would allow users to automatically identify relevant documents from a collection when they meet desired criteria (e.g., employee contracts where base pay is over \$150,000). While our users were able to generally accomplish the task, the terminology used in the tool was consistently a source of consternation due to its incompatibility with their mental model. Interestingly, when prompted for suggestions to improve the terminology, participants would draw blanks. This fact combined with results from the first user study (e.g., work is compartmentalized, performed infrequently) indicated to us that participants often lacked a complete mental model of the task. Ultimately indicating to us that when building tools for similar high-risk/high-reward infrequent tasks that one should attempt to encourage consistency but to use terminology that will be seen as generic as possible and strike a “not wrong” balance (i.e., the language used is not wrong enough to pull them out of the task).

REFERENCES

- [1] 2017. Legal AI Co.s Seal, Kira + Leverton Show Buoyant Growth. <https://www.artificiallawyer.com/2017/09/15/legal-ai-co-s-seal-kira-leverton-show-buoyant-growth/>.
- [2] 2021. A Basic Primer on Selling a Business: Summary - An Overview of Business Sale Processes. <https://kirasystems.com/contract-central/a-basic-primer-on-selling-a-business-summary-an-overview-of-business-sale-processes/>.
- [3] Anne Aula and Klaus Nordhausen. 2006. Modeling successful performance in web searching. *Journal of the american society for information science and technology* 57, 12 (2006), 1678–1693.
- [4] Sandeep Avula, Gordon Chadwick, Jaime Arguello, and Robert Capra. 2018. SearchBots: User Engagement with ChatBots during Collaborative Search (*CHIIR '18*). Association for Computing Machinery, New York, NY, USA, 52–61. <https://doi.org/10.1145/3176349.3176380>
- [5] Sandeep Avula, Bogeum Choi, and Jaime Arguello. 2022. The Effects of System Initiative during Conversational Collaborative Search. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW1, Article 66 (apr 2022), 30 pages. <https://doi.org/10.1145/3512913>
- [6] Leif Azzopardi, Diane Kelly, and Kathy Brennan. 2013. How query cost affects search behavior. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. 23–32.
- [7] Valeriia Bolotova, Vladislav Blinov, Falk Scholer, W Bruce Croft, and Mark Sanderson. 2022. A Non-Factoid Question-Answering Taxonomy. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1196–1207.
- [8] George Buchanan, Ann Blandford, Harold Thimbleby, and Matt Jones. 2004. Supporting information structuring in a digital library. In *International Conference on Theory and Practice of Digital Libraries*. Springer, 464–475.
- [9] Wanling Cai, Yucheng Jin, and Li Chen. 2022. Impacts of Personal Characteristics on User Trust in Conversational Recommender Systems. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22*). Association for Computing Machinery, New York, NY, USA, Article 489, 14 pages. <https://doi.org/10.1145/3491102.3517471>
- [10] Robert Capra, Jaime Arguello, Anita Crescenzi, and Emily Vardell. 2015. Differences in the use of search assistance for tasks of varying complexity. In *Proceedings of the 38th international acm sigir conference on research and development in information retrieval*. 23–32.

- [11] Robert Capra, Gary Marchionini, Javier Velasco-Martin, and Katrina Muller. 2010. Tools-at-hand and learning in multi-session, collaborative search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 951–960.
- [12] Jason Ingyu Choi, Saar Kuzi, Nikhita Vedula, Jie Zhao, Giuseppe Castellucci, Marcus Collins, Shervin Malmasi, Oleg Rokhlenko, and Eugene Agichtein. 2022. Wizard of Tasks: A Novel Conversational Dataset for Solving Real-World Tasks in Conversational Settings. In *Proceedings of the 29th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, Gyeongju, Republic of Korea, 3514–3529. <https://aclanthology.org/2022.coling-1.310>
- [13] Kevyn Collins-Thompson, Preben Hansen, and Claudia Hauff. 2017. Search as learning (dagstuhl seminar 17092). In *Dagstuhl reports*, Vol. 7. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [14] Gordon V. Cormack and Maura R. Grossman. 2014. Evaluation of Machine-Learning Protocols for Technology-Assisted Review in Electronic Discovery. In *Proc. SIGIR 2014*.
- [15] Anita Crescenzi and Robert Capra. 2013. Collaboration in search processes and results sharing. In *Collaborative Information Seeking Workshop*.
- [16] Anita Crescenzi, Robert Capra, and Jaime Arguello. 2017. Time Limits, Information Search and the Use of Search Assistance. In *Proc. CHIIR 2017*.
- [17] Anita Crescenzi, Rob Capra, Bogeum Choi, and Yuan Li. 2021. Adaptation in Information Search and Decision-Making under Time Constraints. *Proc. CHIIR 2021*.
- [18] Anita Crescenzi, Diane Kelly, and Leif Azzopardi. 2016. Impacts of Time Constraints and System Delays on User Experience. *Proc. CHIIR 2016*.
- [19] Jonathan Donnelly and Adam Roegiest. 2020. The Utility of Context When Extracting Entities from Legal Documents. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2397–2404.
- [20] Souvik Ghosh, Manasa Rath, and Chirag Shah. 2018. Searching as Learning: Exploring Search Behavior and Learning Outcomes in Learning-Related Tasks. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval* (New Brunswick, NJ, USA) (CHIIR '18). Association for Computing Machinery, New York, NY, USA, 22–31. <https://doi.org/10.1145/3176349.3176386>
- [21] Travis R. Goodwin and Sanda M. Harabagiu. 2016. Medical Question Answering for Clinical Decision Support. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management* (Indianapolis, Indiana, USA) (CIKM '16). Association for Computing Machinery, New York, NY, USA, 297–306. <https://doi.org/10.1145/2983323.2983819>
- [22] Marti A Hearst. 2014. What's missing from collaborative search? *Computer* 47, 3 (2014), 58–61.
- [23] David G Hendry and David J Harper. 1997. An informal information-seeking environment. *Journal of the American Society for Information Science* 48, 11 (1997), 1036–1048.
- [24] Dan Hendrycks, Collin Burns, Anya Chen, and Spencer Ball. 2021. CUAD: An Expert-Annotated NLP Dataset for Legal Contract Review. <https://doi.org/10.48550/ARXIV.2103.06268>
- [25] James Hollan, Edwin Hutchins, and David Kirsh. 2000. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)* 7, 2 (2000), 174–196.
- [26] Spiros Koulouzis, Paul Martin, Huan Zhou, Yang Hu, Junchao Wang, Thierry Carval, Baptiste Grenier, Jani Heikkinen, Cees de Laat, and Zhiming Zhao. 2020. Time-critical data management in clouds: Challenges and a Dynamic Real-Time Infrastructure Planner (DRIP) solution. *Concurrency and Computation: Practice and Experience* 32, 16 (2020).
- [27] Matthew Lease, Gordon V Cormack, An T Nguyen, Thomas A Trikalinos, and Byron C Wallace. 2016. Systematic review is e-discovery in doctor's clothing. In *Proceedings of the 2nd SIGIR workshop on Medical Information Retrieval (MedIR)*.
- [28] Spyretta Leivaditi, Julien Rossi, and Evangelos Kanoulas. 2020. A Benchmark for Lease Contract Review. <https://doi.org/10.48550/ARXIV.2010.10386>
- [29] Hang Li, Harris Scells, and Guido Zuccon. 2020. Systematic review automation tools for end-to-end query formulation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2141–2144.
- [30] Jimmy Lin, Adam Roegiest, Luchen Tan, Richard McCreadie, Ellen M Voorhees, and Fernando Diaz. 2016. Overview of the TREC 2016 Real-Time Summarization Track. In *TREC*.
- [31] Chang Liu and Yiming Wei. 2016. The impacts of time constraint on users' search strategy during search process. *Proc. ASIST 2016*.
- [32] Chang Liu, Fan Yang, Yu Zhao, Qin Jiang, and Lu Zhang. 2015. What Does Time Constraint Mean to Information Searchers? *Proc. IliX 2015*.
- [33] Catherine C Marshall, Frank M Shipman III, and James H Coombs. 1994. VIKI: Spatial hypertext supporting emergent structure. In *Proceedings of the 1994 ACM European conference on Hypermedia technology*. 13–23.
- [34] Sara Merken. 2021. Global M&A just suffered its worst Q3 in a decade. <https://www.reuters.com/legal/transactional/legal-tech-dealmaking-continues-litera-scoops-up-kira-systems-2021-08-10/>.
- [35] Nina Mishra, Ryen W. White, Samuel leong, and Eric Horvitz. 2014. Time-Critical Search. In *Proc. SIGIR 2014*.
- [36] Meredith Ringel Morris. 2013. Collaborative search revisited. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 1181–1192.
- [37] Meredith Ringel Morris and Eric Horvitz. 2007. SearchTogether: An Interface for Collaborative Web Search. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology* (Newport, Rhode Island, USA) (UIST '07). Association for Computing Machinery, New York, NY, USA, 3–12. <https://doi.org/10.1145/1294211.1294215>
- [38] Douglas W Oard, Jason R Baron, Bruce Hedin, David D Lewis, and Stephen Tomlinson. 2010. Evaluation of information retrieval for E-discovery. *Artificial Intelligence and Law* 18, 4 (2010), 347–386.
- [39] Frances A Pogacar, Amira Ghenai, Mark D Smucker, and Charles LA Clarke. 2017. The positive and negative influence of search results on people's decisions about the efficacy of medical treatments. In *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval*. 209–216.
- [40] Dan Primack. 2022. Global M&A just suffered its worst Q3 in a decade. <https://www.axios.com/2022/09/30/global-mergers-acquisition-market-q3>.
- [41] Suhas Ranganath, Suhang Wang, Xia Hu, Jiliang Tang, and Huan Liu. 2017. Facilitating Time Critical Information Seeking in Social Media. *IEEE Transactions on Knowledge and Data Engineering* (05 2017).
- [42] Adam Roegiest, Alexander K. Hudek, and Anne McNulty. 2018. A Dataset and an Examination of Identifying Passages for Due Diligence. In *Proc. SIGIR 2018*.
- [43] Adam Roegiest, Alexander K Hudek, and Anne McNulty. 2018. A Dataset and an Examination of Identifying Passages for Due Diligence. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 465–474.
- [44] Chirag Shah. 2013. Effects of awareness on coordination in collaborative information seeking. *Journal of the American Society for Information Science and Technology* 64, 6 (2013), 1122–1143.
- [45] Chirag Shah, Gary Marchionini, and Diane Kelly. 2009. Learning Design Principles for a Collaborative Information Seeking System. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems* (Boston, MA, USA) (CHI EA '09). Association for Computing Machinery, New York, NY, USA, 3419–3424. <https://doi.org/10.1145/1520340.1520496>
- [46] Anton Van der Vegt, Guido Zuccon, Bevan Koopman, and Anthony Deacon. 2020. How searching under time pressure impacts clinical decision making. *J. Medical Library Association* 108, 4 (2020).