Who needs imagination? Exploring legal professionals lack of curiosity about e-discovery tools

EMMA LURIE, UC Berkeley School of Information
DEIRDRE K. MULLIGAN, UC Berkeley School of Information

In the legal profession, e-discovery platforms, which identify and classify documents relevant to a litigation matter, are increasingly relying on algorithmic and machine learning techniques. A series of interviews and review of training materials reveal that despite their positions of relative power and privilege, individual legal professionals–and the profession as a whole–do not demand that these tools be interpretable to them or even evaluated in a way that would give them a sense of how and how well they work. We identify three factors that contribute to this lack of curiosity 1) boundary work that is recasting parts of the discovery process as the work of new professionals; 2) market forces and institutional and professional choices that reduce incentives to identify weaknesses in algorithmic systems; and, 3) discovery rules that buffer legal professionals from externally cognizable failures TAR tools produce. This confluence of factors has produced a relatively uncritical adoption of a class of tools without meaningful evaluation.

ACM Reference Format:
Emma Lurie and Deirdre K. Mulligan. 2018. Who needs imagination? Exploring legal professionals lack of curiosity about e-discovery tools. In *CHI 2021: Designing Technological Systems with The Algorithmic Imaginations of Those Who Labor*. ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

In the legal profession, e-discovery platforms, which identify and classify documents relevant to a litigation matter, are increasingly relying on algorithmic and machine learning techniques. These systems, and the companies offering them, are reshaping relationships between lawyers and clients, introducing new kinds of professionals into legal practice, altering the discovery process, and shaping how lawyers construct knowledge about their cases and professional obligations. Given the ways in which these systems are shaping legal practice and the construction and presentation of legal knowledge, one might expect lawyers to be deeply curious about how they work. We conducted 20 interviews with U.S. legal professionals who are responsible for designing, using, or procuring these systems[1], and watched 15 hours of e-discovery software training videos.

Technology-assisted review (TAR), also called "predictive coding," systems for the discovery phase of litigation provide an interesting example of machine-learning based decision support systems entering a professional domain. TAR is used to help litigants sort through the enormous amounts of electronic data and identify relevant documents to provide the adversary. A 2012 RAND study found that using TAR reduced discovery costs by up to 70% [12]. The rise of TAR has created an ever-growing industry of e-discovery specialists, support staff, consultants, technology vendors,

[1]Methods further detailed in [10]

Manuscript submitted to ACM

and products. Predictive coding systems, under the umbrella of TAR, are marketed as tools to aid legal professionals in managing, classifying, and reviewing documents at a fraction of the cost.

Lawyers have resisted new technological tools, at least in part due to concerns about the logics and values they embed [7]. With respect to TAR tools we find that traditional resistance has been overwhelmed by a reconfiguration of market power as document management platforms build in discovery support, combined with client pressure for cost savings. The reliance on these tools has opened up space for new professionals in the legal field, and lawyers appear happy to defer to their expertise. Resistance has been further dissipated by court rulings permitting and at times legitimating the use of TAR, court deference to litigating parties' agreements about discovery processes, and discovery practices that buffer lawyers from the professional costs associated with the TAR failings most likely to be identified.

The result has produced a profession with little interest in understanding the inner workings of the systems they rely on to ply their trade. Interviews revealed both a lack of knowledge and a lack of interest in the algorithms and other technical details of these complex systems. Training materials designed to educate users of these tools focus on explaining the complex user interfaces with little description of the algorithms that underlie them. Professional materials focus on priming the legal community for adoption of TAR tools rather than promoting critical inquiry into how the logics they embed align with or connect to core legal concepts.

This case study provides an important point of contrast to other workers' curiosity about the algorithmic systems entangled in their work. The career opportunities of the legal professionals we interviewed are not directly controlled by the algorithms. Neither their productivity or their outputs are evaluated by TAR tools. This may reduce legal professionals' desire to understand the systems. Our interviewees were aware that these systems might automate legal work depressing future generations career opportunities, but they did not view their own opportunities as being limited. In fact, they consistently walled-off aspects of the work required to configure and operate TAR systems as technical work, despite its unavoidable connection to professional judgments for which they remain responsible. Unlike other workers who've personified algorithms [14] and in other ways sought to ascribe logic to their behavior [5], legal professionals showed limited interest in imagining, let alone understanding, how the algorithm systems they rely on reason. They do not engage in sensemaking about the algorithms that are necessary for their day-to-day work. Moreover, despite their positions of relative power and privilege, individual legal professionals–and the profession as a whole–do not demand that these tools be interpretable to them or even evaluated in a way that would give them a sense of how and how well they work. We identify three factors that contribute to this lack of curiosity 1) boundary work that is reallocating discovery work to new algorithmic professionals; 2) market forces and legal institutions that reduce incentives to identify weaknesses in algorithmic systems; and, 3) discovery rules that buffer legal professionals from the visible failures TAR tools produce. This confluence of factors has produced a relatively uncritical adoption of a class of tools without meaningful evaluation.

2 FINDING 1: IN TAR WE TRUST

In Christin's study of judges using risk assessment tools, she observes that "legal professionals have a deep lack of trust for algorithmic tools: they explain that they prefer to trust their own opinion rather than a number on a page" [7]. However, legal professionals interacting with TAR behave differently than the judges studied by Christin. As one TAR coding consultant explains that:

> As soon as it's software, nobody questions it and they think it's great. It's such a trap for attorneys and everyone. The more we get into AI, you just trust the computer, and you don't know what's making it tick.

A senior partner at a law firm that uses TAR explains that few people are questioning the TAR technology, instead: "They're just like, oh, it's AI. It's magic."

While many interviewees discussed their trust in the TAR tools as well as their colleague's trust in the tools many participants were reticent to speak about TAR in any depth, often claiming that we would be better off talking to someone else at their firm. This occurred despite the fact that participants had self-identified or been referred to as leaders in the use of TAR. As the interview progressed, many interviewees expanded on their thoughts about TAR, revealing their expertise or experience with the technology, but their initial resistance is telling.

Craig Ball, who creates TAR training materials, includes the following story in training materials: "A colleague once defended her ignorance of the technical fundamentals of electronically stored information by analogizing that "she didn't need to know how planes stay aloft to fly on one."

He continues to remark that "She had a point, but only for passengers. If you aspire to be a pilot or a rocket scientist—if you want to be at the controls or design the plane—you must understand the fundamentals of flight." Ball offers this example as indicative of the attitude of the legal profession which he perceives as uninterested in the underlying concepts and inner workings of the predictive coding algorithms.

The limited interest and attention to the workings of the algorithms is reflected in the training videos. In the industry leader's training videos (Relativity), approximately three quarters of primary training video focuses on navigating the interface of the software rather than any technical explanation of the software or warnings about its use and limitations.

3 FINDING 2: BOUNDARY DRAWING AND NEW PROFESSIONALS

Our findings suggest that other factors, beyond trust in the technology, including the emergence of new professionals–often tied to new entities who claim expertise in TAR–contribute to legal professions disinterest in the internal workings of TAR tools.

Kellog, Valentine, and Christin [9] identify three emerging "algorithmic occupations." Our interviews reveal lawyers' reliance on a new set of professional experts that relate to their concepts of "algorithmic brokerage", and "algorithmic articulation."

First, attorney's rely on what Kellog et al call "algorithmic brokers" to help firms make decisions about which TAR tools to acquire. Firms invest in more technically literate in-house litigation support expertise, they also heavily rely on the expertise of the TAR tool vendors to assist them in understanding and selecting tools. For example, one interviewee explained: "Attorneys simply don't have...the competencies to execute the very complex scientific tasks at this intersection of information retrieval, computer science, statistics, and linguistics…" An interviewee explains that the TAR work he manages "is a distinct professional domain, information retrieval. Finding information you want in very large datasets, especially in instruction datasets where you can't use a database, where you just look up, basically. It's truly a distinct professional field."

Training videos on best practices for procurement specify that both IT professionals as well as senior legal professionals should be involved in conversations with vendors. Given that IT professionals expect a specifications document, they recommend that a legal professional 'immerses' themselves in the area so that the e-discovery needs of the organization can be articulated to different stakeholders.

In addition, interviewees highlighted the distinct and important niche occupied by these new professionals in the legal process.Here we see professionals engaged in what Kellog et al call "algorithmic articulation." For example, interviewees identified the critical role they play in structuring discovery:

> "Discovery can be won or lost on how you negotiate the scope of discovery...Very often that hinges on, for example, what search terms do the parties agree to to reduce the data set or how is the technologists that review actually used and what you use to measure whether or not it's been accurate. That is very much a consulting kind of thing where we can sit with clients and say how can you at the lowest possible cost inflict the greatest amount of pain to the other side through the use of e-discovery in a smart way to help maximize the value to your case and put the other side in a bad place?"

Another explained all the ways she assisted the client from "interfac[ing] directly on the client's behalf... argu[ing] with the opposing side...sitting in the back of the room or whispering in my client's ear...educating my client...submitting testimony into the court...appear[ing] as an expert witness, a fact witness…" and noting that "Courts have something they refer to now as an eDiscovery liaison." where "We are sometimes introduced to the other side as a liaison where it's expected that we're just going to talk tech…"

Another explained the dire need for the expertise his firm provided particularly emphatically stating:

> "I can tell you that...we have nearly never seen a statistical measure in protocol that can withstand even rudimentary scrutiny. The reason for that is that even lawyers and judges today don't understand that your techie sitting in IT or discovery is not a statistician. You see in court – I mean, you literally see in court people debating statistics and, in fact, precision and recall measurement, the metrics, and no one in the room seems to understand that the main competence needed to even have this conversation is completely absent."

4 FINDING 3: MARKET FORCES

As we've described elsewhere [10], pressure to cut costs is driving the adoption of TAR tools. An additional driver is a seismic shift in organizational relationships: TAR vendors increasingly have their own direct, paid relationship with the law firm's corporate clients. This new configuration elevates the visibility of the vendors with the firms' clients, allows clients to influence lawyers' choice of discovery tools, and creates increased opportunity for cost savings to outweigh other considerations in decisions to use TAR.

While the switch to e-discovery is heavily marketed as a way to minimize costs, with the RAND study often cited in marketing materials, our consultant interviewees also commented on how they illustrate the value of their services:

> "$500,000 on a project. Is that a lot? Should you really have spent $400,000, and you made a mistake, or would you have spent 5 million, and you saved a ton of money? Those are the kinds of things that we try to show our client with reports."

Another interviewee remarked: We're seeing more and more that the general counsel has to show to their leadership and to their board that they have reduced costs by X percent or increased the use of technology. Then they're coming to us and saying how do we measure this? How do we show this? What do we do?

While cost-savings is one type of financial benefit, interviewees also discussed that TAR is viewed as the state of the art legal technology. One interviewee explained that "the primary motivators for law firms when purchasing this – when looking at these things is marketing. Can I use this to market to clients to show the PR? ...The cutting edge and we're using these new tools."

A TAR consultant explains that the direction to invest in AI comes from senior leadership at firms "Maybe someone went to a conference and was very motivated by some speaker talking about AI, whatever it is. There's an order from

the top that we need to be smarter, computer analytics can do that for us, and now go figure it out. There really is this almost weird, nervous energy among people who are supposed to implement it. What do I do with that?"

The idea that the push for AI e-discovery software comes from the top of firms highlights the role that organizational structure plays in the adoption of e-discovery technology, vendor-client relationships become essential. One interviewee explains " We have an array of products, and the client will tell us what review platform they want it to go up in." For others, e-discovery is part of a larger suite of electronically stored information tools. The training videos that discuss procurement emphasize the importance of making sure the technology clients (not firms) purchase fit into their existing technical infrastructure. Many of these systems are sold as part of a larger package, in which e-discovery is a single component.

While some vendors provide e-discovery software, others provide "e-discovery as a service." As one interviewee explains, the relationship between vendor-client-firm can be complex as the legal counsel is "not totally absolved of any duties. Their names are still going on the documents, and they're still the ones appearing for their client's in court. They should know what our process is, but there does have to be a line in separation of rules. We really need our client to help us make that clearer so that the firm doesn't just come in and do what they would normally do if they were not in that situation. I'm not going to tell the firm to go away. They're potentially a client too and often are our client too.

5 DISCUSSION

Why do we see such limited evidence of imaginaries [5]? We see multiple possibilities to explore further in future research.

The performance and outputs of the legal professionals we interviewed are not measured or evaluated by the TAR tools. We are certain that some of the lawyers and paralegals who directly interact with the systems are however evaluated on their ability to work well with them. One interviewee reflected on the way in which the use of TAR might push responsibility for a discovery failure up the chain to the senior attorney in charge hypothesizing that "...in a context where I'm not checking it document by document, I would be much more worried about the software than the human." because "if you trust the contract attorney, you can ultimately blame the contract attorney. If they're software, I would be blamed. I could not blame the software. I might be able to blame the vendor but unlikely."

This is an interesting, potentially unintuitive way in which the use of software shifts responsibility to higher paid workers and away from the more precarious workers who interact with the system–often, but not always, contract workers who work for the TAR provider. But this was an exception. As a general matter those we interviewed while they felt compelled to use TAR tools, at least to some extent, did not feel ruled by them. They were tools they used to practice, not tools used to control them.

Thus these tools while reconfiguring work life for the legal professionals we interviewed were not tools of subjugation or at least for the moment displacement. Without fear of being judged or evaluated by these tools, the legal professionals may have less interest in performing for them, being legible to them, or gaming them. They simply want some other experts to translate and bridge between them and the TAR tool so they can perform the discovery task for which they are professionally responsible without becoming an expert in information discovery, statistics, or machine learning. They rely on a new class of algorithmic workers who do the interpretive work, and therefore find little need to do it themselves.

Second, unlike the algorithmic systems low wage workers are subject too where they are typically kept in the dark about the metrics used to assign or evaluate their work, legal professionals directly input the examples that train the model and have the ability to view all of the predictions for all observations. So while the underlying model–the

decisional logic is opaque in every sense (proprietary, technical illiteracy, and uninterpretable models) [6]–they are not only privy to but literally control the inputs and outputs. They can tinker with inputs and outputs to interrogate the system. While this provides a very limited understanding of it, it is surely more than many low wage workers are afforded and in combination with the TAR tools' disconnection from performance evaluation, it may be enough to satisfy our interviewees curiosity about the systems. They need not imagine [5], because they think they understand how they work. This of course is far from true, but illusions of control can depress efforts to control and efforts to seek out information about what is going on in the black box.

Third, the legal system has accommodated and normalized the introduction of these tools in numerous ways. Boundary work is critical in professions. Professionals explicitly and implicitly defend their turf against outsiders and lower-status professionals within the field [13]. However, when new technologies entering a field or an organization, professional workers are influenced by the opinions of others, particularly those who claim special expertise [11]. Boundary work includes re-configuring professional identity in relation to both new technologies and the new workers that accompany their introduction into the field. This reconfiguration of socio-cultural beliefs and practices are important for legitimization and organizational change resulting from new technologies [8].

As described above, TAR is mediating the construction and reconstruction of lawyers' professional identity. Legal professionals appear to be drawing new boundaries between the work they do–managing the inputs and reviewing the outputs of TAR tools–and the work done by the new set of professionals who understand the back end of these algorithmic systems. They rely on "algorithmic brokers" who "seek to communicate the logic and value of the algorithmic systems to various groups" and "algorithmic articulators" to do so [9]. Firms have hired new professionals in their litigation support groups who help firms understand and assess the TAR tools, and others, emerging on the market. At the same time, both firm lawyers and lawyers working at TAR providers describe the role that vendors play in proselytizing and educating around these new tools.

But other actors are smoothing the way as well. Professional groups focused on TAR, such as the Sedona Conference[2] and the Bolch Judicial Institute[3] are part of a broader cross-institutional set of algorithmic brokers seeking to make the legal field amenable to TAR and other machine learning tools. In addition to the in-house litigation support experts and staff at the TAR providers who as "algorithmic articulators" are helping legal professionals "integrate and streamline algorithmic workflows", bridging among different professionals, and "addressing the failure" modes, hours of training videos provided by TAR vendors do additional articulation work helping legal professionals and those they supervise understand how to successfully work with these tools. Some have positioned themselves as key educators not just on their particular tool but on TAR tools more generally.

The incorporation of these new professionals is in part driven by client interest in automating as much of the document review process as possible, and the new market arrangements in which TAR tool providers frequently have their own relationship with a firms' clients. Earlier research documented the growing role of in-house attorneys during pretrial discovery. Other work has documented a trend for viewing legal representation as but one more service to let for bidding, and the increasing tendency for in-house attorneys to control the strategies and practices of outside counsel. The ability to directly contract with service providers who may host and manage documents as well as provide litigation support services, furthers this shift in power from firm lawyers to in-house attorneys and from firm lawyers to these new professionals who help rationalize and manage interactions with TAR tools. The tool and service providers, and the experts they employ, are playing an outsized role in interpreting how TAR tools relate to the legal requirements

[2] https://thesedonaconference.org/
[3] https://judicialstudies.duke.edu/

around discovery. Lacking appropriate expertise individual lawyers report great reliance on the representations of vendors, and to some degree evaluations by technical staff—although there was little clarity about how such evaluations occur.

Courts have also smoothed the path for TAR adoption. This is significant given the lack of standards for performance, testing and validation. While market forces may be driving lawyers to use TAR, the adoption requires the profession to at least acquiesce to their use. In an important moment for TAR tools, Judge Peck in Da Silva Moore v. Publicis found "that computer-assisted review is an acceptable way to search for relevant ESI in appropriate cases"[2]. As a general rule the party producing evidence has wide latitude to use the tools and methods they view fit. As a court nicely summed it up, "Plaintiff must request information, regardless of how or where it is maintained by Defendants, which Defendants must address as required by Rule 34. That is discovery: A party requests information and the burden is on the producing party to locate and produce it or object legitimately to production. To the extent Plaintiff is seeking to compel Defendants to conduct discovery as directed by Plaintiff, the Court declines to issue such an order." [4]. This places enormous weight on the parties' choices. If they agree to it, the court is unlikely to second guess it.

Our interviewees provide evidence of this deference to litigation parties. One said "courts have been entirely receptive to the use of predictive coding. Generally, what we see is courts just putting into an order whatever the parties have agreed to."

In addition our interviews suggested that the legal professionals trust one another to fulfill discovery obligations, and this includes making sound use of TAR tools. Interviewees were generally trusting of their peers efforts and not eager to meddle in their production choices. For example, one interviewee said, "if I were adjusting a dial, it wouldn't necessarily be so that I could peer into the TAR process of my adversary. It would be so that I could effectively trust and verify their discovery production."

This belief that the professional behavior of the parties will lead to a sound discovery process is captured in a recent district court case where the judge wrote:

> "In keeping with these principles, this Court is of the view that there is nothing so exceptional about [electronically stored information] production that should cause courts to insert themselves as super-managers of the parties' internal review processes, including training of TAR software, or to permit discovery about such process, in the absence of evidence of good cause such as a showing of gross negligence in the review and production process, the failure to produce relevant specific documents known to exist or that are likely to exist, or other malfeasance" [3].

Finally, lawyers use discovery rules to buffer themselves from the visible failures of TAR tools that might harm their clients, and themselves. In cases involving TAR, litigating parties generally negotiate for what are called "claw back" agreements under the Federal Rules 502(d) [1]. This allows a party who inappropriately discloses a privileged document during discovery to claw it back without waiving the privilege. Accidental disclosures of privileged documents are the most visible failure of TAR tools and the ones our interviewees were most interested in discussing. Our interviewees were generally uninterested and unconcerned with the possibility of TAR producing false negatives (material that ought to have been produced but was erroneously identified non-responsive). To the extent that they routinely reviewed material identified as responsive by the TAR tools it was almost exclusively the material identified for production. That material would be reviewed both to remove privileged documents and to correct false positives. Such secondary reviews might include complete human review on some subset of custodians (for example a company's top executives) or sampling of such a subset, or sampling of the entire corpus identified for production.

6 CONCLUSION

As we consider imaginaries of those who labor, it is useful to draw out the similarities and differences between lawyers and low-power workers. We highlight multiple factors that could contribute to the indifference to sense-making and lack of resistance to the machine learning systems despite how radically they are reshaping the legal profession. These include: the creation of new professionals, market incentives; the workers' performance is largely outside of the system's purview; the visibility front-line legal professionals have into the system outputs and predictions; and the fields acceptance of and accommodation to these tools.

REFERENCES

[1] 2011. Rule 502. Attorney-Client Privilege and Work Product; Limitations on Waiver. Federal Rules of Evidence.

[2] 2012. Da Silva Moore v. Publicis Groupe. 868 F. Supp. 2d 137 - Dist. Court, SD New York.

[3] 2017. Winfield et al v. City Of New York. Filing 217, Dist. Court, SD New York.

[4] 2021. DAVID HASTINGS, Plaintiff, v. FORD MOTOR COMPANY, et al., Defendants. United States District Court, S.D. California..

[5] Taina Bucher. 2017. The algorithmic imaginary: exploring the ordinary affects of Facebook algorithms. *Information, communication & society* 20, 1 (2017), 30–44.

[6] Jenna Burrell. 2016. How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society* 3, 1 (2016), 2053951715622512.

[7] Angèle Christin. 2017. Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society* 4, 2 (2017), 2053951717718855.

[8] Bob Hinings, Thomas Gegenhuber, and Royston Greenwood. 2018. Digital innovation and transformation: An institutional perspective. *Information and Organization* 28, 1 (2018), 52–61.

[9] Katherine C Kellogg, Melissa A Valentine, and Angèle Christin. 2020. Algorithms at work: The new contested terrain of control. *Academy of Management Annals* 14, 1 (2020), 366–410.

[10] Daniel N Kluttz and Deirdre K Mulligan. 2019. Automated decision support technologies and the Legal Profession. *Berkeley Tech. LJ* 34 (2019), 853.

[11] Wanda J Orlikowski. 2000. Using technology and constituting structures: A practice lens for studying technology in organizations. *Organization science* 11, 4 (2000), 404–428.

[12] Nicholas Michael Pace and Laura Zakaras. 2012. *Where the money goes: Understanding litigant expenditures for producing electronic discovery.* Rand.

[13] John Sutton. 2001. *Law/society: Origins, interactions, and change.* Vol. 474. Pine Forge Press.

[14] Eva Yiwei Wu, Emily Pedersen, and Niloufar Salehi. 2019. Agent, gatekeeper, drug dealer: How content creators craft algorithmic personas. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–27.