

Employing AI for Retention & Disposition in Digital Information and Recordkeeping Systems

AI User Focus Groups



An ITrust^{AI} Report
April 19, 2024

Report Prepared by Jenny Bunn, Patricia C. Franks, Pat Moore, Katherine Hodgson.

Table of Contents

Introduction

About the Study.....	1
Methodology	1

Summary of Results

Part 1: Demographics	2
Part 2: Responses to Prompts.....	2
Part 3: Current Perspectives	3
Part 4: Opportunities	6

Conclusion	8
------------------	---

Introduction

This report outlines the results from user focus group discussions conducted as part of an ongoing study by InterPARES Trust AI. Between March 7th, 2023, and April 15th, 2023, an initial survey was sent to potential participants. Of the total respondents, 12 individuals who had expressed a willingness to be contacted for further questions were asked to participate in focus group discussions throughout September 2023. These discussions, facilitated by lead researchers, were a chance for records and information management (RIM) professionals to discuss their understanding of and experiences with artificial intelligence (AI) in their work and the implications that its use may have on the future of the profession.

Methodology

The goal of the focus groups was to promote dialogue on the role of AI in RIM and to establish a point of reference for where certain professions are in their integration of AI tools into daily operations. In contrast to our initial survey, which focused on acquiring quantitative results, the focus groups allowed for the collection of qualitative data regarding the present lived experience of RIM professionals navigating the widespread adoption of AI technology into their work.

Prior to the discussions, invitees were sent a list of high-level questions to reflect on pertaining to AI in recordkeeping. The focus groups took place virtually, over Zoom, and were structured as a natural conversation rather than a question-answer process. Facilitators were given the following guiding prompts for each group:

- What differentiates automation from AI?
- How much discussion is happening in your work about AI and its uses?
- How is AI used in your work, if at all?
- How will the use of AI change the role of the records manager, if at all?

Each prompt invited the participants to reflect on the current state of AI in their own work as well as to consider what the future may hold for it, be it an opportunity or threat. The following results come from the analysis of the recordings and transcriptions of the focus groups.

Summary of Results

Part 1: Demographics

Participants were all professionals specializing in information governance, records management, and information security. Nine different job industries were represented among them, with the most prevalent being Records and Information Management (25%) and Educational Services (17%). Other represented industries were Public Administration, Legal, Finance and Insurance, Utilities, Archives Management, Library and Information Science, and Information Governance. The majority (83%) of invitees worked within the United States and Canada, with the remaining participants operating within Spain and the United Kingdom.

Part 2: Responses to Prompts

While each focus group discussion offered unique insights, there were some recurring themes and patterns among the responses to the four guiding questions posed to each group.

What differentiates automation from AI?

Participants were split on how different automation and AI are, with some believing them to be so similar that they seem “married” in a sense while others set firmer boundaries between the two. The most common differentiating feature between the two was reported to be the system’s ability to learn. Once implemented, AI programs are expected to understand the data they are sorting through and adapt their processes based on what they find, meaning they “learn” without manual human intervention. In this way, their output resembles that of a human with room for interpretation and, as mentioned by a few participants, error. Automation, in contrast, provides consistent output and requires more human interaction by way of initiating processes and adjusting their capabilities. Automated processes were most heavily associated with rule-based decision making and working with known quantities.

How much discussion is happening in your work about AI and its uses?

The majority of participants indicated that there was at least some ongoing discussion within their workplace about AI systems, whether they had already been implemented or not. Most discussions centered around the development of policies and initial inquiries into what sorts of systems could be useful in their industry. Often these discussions did not include records managers or the discussion surrounding their use was not focused on managing information.

How is AI used in your work, if at all?

Throughout the focus group discussions, 42% of participants reported prior experience with AI via personal use or research. In terms of work, 25% reported that tools with AI capabilities were already in use to some extent in their place of employment. The majority of participants expressed an expectation that AI tools will be implemented in the future, either for the first time or in addition to pre-existing ones. The most commonly reported use for AI tools within their industries was for data analysis and extraction rather than for use in recordkeeping and management. Many reported that the state of AI in their workplace was still quite new, with it often being the first time anything of its kind had been implemented.

How will the use of AI change the role of the records manager, if at all?

Responses to this prompt showed an overall optimism that the role of record keeper will not be erased entirely by the implementation of AI, or at least a belief that these workers cannot be so easily replaced. Instead of replacing the human worker it is understood that new skill sets will have to be developed in order to work alongside these new systems and in order to do so, education is necessary to bring record keepers up to a level of understanding in which they feel more comfortable with AI as a whole. There is hope that these systems can handle some of the more time-intensive and often reiterative work performed by record managers, such as sorting and tagging large quantities of data (e.g., emails), so that workers can instead focus on the organization, proper storage, and lifecycle of the records.

Part 3: Current Perspectives

The nuances of AI technology and its implications for the future of recordkeeping provoked an understandably mixed response from participants in terms of opinions. Negative expressions often centered around the emotion these systems evoked, with repeated descriptors being “creepy” or “scary”, for example. In contrast, positive expressions tended to focus more in the realm of what AI is capable of, functionally, with it being most commonly referred to as an opportunity:

“I think we should embrace the technology because it offers tremendous possibilities.”

“I think [AI] is going to create more opportunity for us.”

A common opinion expressed by participants was that the current state of AI is not yet advanced or understood enough to be purchased and implemented, particularly for use in RIM. This is compounded by a perceived lack of transparency in how these systems interact with

data across all environments, leading to issues of safety, integrity, and ownership of records enacted upon or produced by AI.

“When you start digging in and asking the questions about AI, ... you’re not getting the true answers that you’re hoping [for].”

“So to even start to look at [implementation], ... we need to understand it more.”

“I don't even think [AI professionals] understand the background of how [AI] is grabbing the information and what it's using ...”

“We need to learn what it does well and what it doesn't do well. And then we need to validate the results to make sure we're getting what we expected.”

These sentiments go hand-in-hand with a generally expressed confusion or lack of applicable knowledge expressed by participants with regards to what AI is and what it truly does as a tool:

“I feel kind of like a new baby ... we’re just like learning [from the] very beginning”

“AI also, you know, as a concept, is still a little scary ... what does it mean?”

“[It’s] all very new to me still.”

“I feel like I'm a little behind in terms of fully understanding.”

This lack of understanding stems both from a deficit of experience with these systems for personal use and a lack of communication between AI, IT, and RIM professionals, leading to confusion and frustration at being “left in the dark.” For those with limited knowledge about AI as a whole, it is difficult to speculate what benefits it may have in RIM moving forward. Doubts also remain about the intent and motives of those developing these programs and whether they align with the goals of RIM professionals.

“I don't know if anybody in computing will really care about record management at this point.”

RIM has already shown itself to be capable of evolving as technology changes, evidenced by the introductions of the Internet and virtual contact systems. For the

professionals who experienced this shift, the idea of something new is not novel, no matter how challenging or daunting it may seem.

"It's not the first time we've seen something that's going to "change the way we do things"."

Knowing that RIM has already proven itself as capable of adapting, some professionals shared the opinion that the implementation of AI is unavoidable, particularly considering the rate with which it is currently being adopted in other fields. Despite any caution or hesitancy, participants recognized AI as something that they will have to account for sooner or later:

"It's something that we can't ignore and it's going to become absolutely massive..."

"We're kind of being pushed into where we're going with AI just because the products we use are embedded with A.I. in the next version."

"It's just another technological development and we'll have to figure out a way to deal with it."

Many participants reported a lack of clarity or timeliness when enacting retention and disposition schedules within their work. The instant nature of digital records compounds this issue as they are able to accumulate faster than a professional can sort through them, leading to backlogs and a lack of organization. Some professionals raised concerns that, at the current level of sophistication, AI can scarcely be expected to perform like a human worker when said individuals can still be unreliable or unpredictable in and of themselves. The system is only as smart as what it is trained on—and who trains it. If one does not have a full understanding of how a particular record must be treated within their system, how could an AI program learn to do so?

"Would it be helpful, or would it be like humans?"

"None of this will matter if we can't have comprehensive information governance within our organizations, right?"

"Yes, it does things faster than what a normal human can do. But does that mean that it's that much smarter?"

Overall, interest is evident, but there is hesitance present both in a lack of understanding of AI overall and trust in the way it will access and potentially share information. There is no streamlined, accessible source that RIM professionals can turn to on AI aside from

company and departmental policies, should they exist, which often leaves these individuals having to perform external research. AI is considered as new, inevitable, and difficult to trust, but full of possibilities should it be implemented responsibly.

Part 4: Opportunities

When defining the problem or task the participants envisioned (or perhaps hoped) AI could solve/undertake for them, many responses expressed this at a fairly general level in terms of identifying records/important information, and thereby categorizing or classifying it as requiring particular treatment, e.g., retention for a specific period.

“I'm hoping that they can solve that problem by just being able to grab what a record is of the company and help it with life cycle management, including the disposition.”

“Having something that maybe can help identify records, help [...] do some automatic assigning of retention to certain types of content.”

“Something that hopefully will take away some of the guesswork for my users when it comes to [...] doing the work to classify information.”

“If it's useful in a way that I [...] imagine that it could be [...] identifying what is the important information that needs [...] security [...] to be applied to it, and retention needs to be applied to it.”

“Millions of records and we do not have the manpower to go in and individually categorize that. [...] I would welcome it. AI could do that work. I would be very happy.”

At a slightly more granular, albeit more abstract level than these more general expressions, there were a few responses that described what participants would like AI to be able to do in terms of an ability to undertake inference or analysis.

“Auto apply retention labels and learn about the emails and give us a report that says [...] these are the types of emails we're seeing [...] and [...] give us an analysis so we can sort of say, yes, that's okay.”

“But also we were looking for something that might then make inferences, identifying, for example, if certain coded documents were clearly identified as potential attorney client privilege because of certain characteristics within them [...] Could the machine then take an [...] intellectual leap and say, hey, two of these participants are also involved in this document?”

Although potential records were often seen as the target of AI analysis, there were also some more specific examples given of analysis directed towards non-record material. The hope being that AI could carry out some of the so-called ‘drudge work’ of reading and analyzing non-record material, such as retention policies.

“I would love to say, hey, Hubert, whatever, get to name it right, ‘Please take all of these inputs and evaluate and reconcile these [...] for human review. [...] another thing where I would love to be able to point [...] saying, Hey, please evaluate applicable retention schedules and the roles within these different departments and come up with suggested base retentions for these areas.”

As can be seen from the previous statement, there remained a clear sense that records managers would not be relying on the results of such analysis without human review, but nonetheless the idea that AI could provide the first ‘rough draft’ of a retention schedule or reconciliation of multiple retention schedules was very attractive. One participant also spoke in a similar vein with respect to a non-records management output, the drafting of skeleton contracts.

“Our precedents are clause based [...] we're curating our long-term precedents and then using them as examples so that we can specify [...] standards and also familiar clauses or standard gold templates, clauses that are used in comparing different clauses from different documents. We think of it like building the skeleton of a contract and then obviously someone would go through as a professional.”

At the most abstract level, one participant spoke of how “Maybe A.I. could help us to be more consistent,” but there were also some more narrowly defined tasks discussed, for example:

- automated redaction—identifying and ‘taking out’ information such as social security numbers.
- carrying out an initial pass on email enquiries to ‘understand’ them enough to direct them to the right person
- to detect mistakes in the classification
- self-description of [...] images
- the transcription of the audio and description of the audio visual
- personalizing searches

Email was commonly mentioned as the record type to which AI could most usefully be applied, particularly perhaps in the sense that an email inbox offered a site where users could be brought to work with AI (rather than the records manager directly), and hence to work out for themselves and to refine a shared understanding of the ‘right’ classifications to be applied.

“I wish there was a way that we could have something, identify emails and sort of classify it and then allow the user to say, no, that's the wrong classification, and that we wouldn't have to feed a lot into it to make that happen.”

“I do think email is the right place to start and I think it will be the place that we start and hopefully with some quick wins. But if they can have a naming tool that picks out the personal stuff and then they can be happy that everything else is operational content...”

The participants did therefore see many opportunities in the form of possible applications to which AI could be directed, but one person also spoke of how the new forms of content and data being generated by AI also offered a different form of opportunity for records professionals.

“I think that as content gets created, what from these systems or as spatial technology becomes a thing [...]. I think that's all new content and new information that we have to understand to preserve and retain.”

Conclusion

Amidst the mixture of confusion, anxiety, and optimism that AI brings to RIM, our focus group discussions discovered that an underlying theme in nearly every conversation about these technologies appears to be people. Whether positive or negative, many of the shared perspectives directly or indirectly discussed the roles, responsibilities, and intentions behind those who develop, utilize, and implement AI. Open communication may be one key way to address issues of trust and transparency within the realm of AI that cause hesitance in the minds of recordkeepers, from addressing the black box problem to having ongoing discussions with peers about the use of AI. Despite the evolution of AI being considered a development in technology, this shows that it is still viewed as a tool that can be wielded and defined by human intent, as demonstrated in the below quote:

“I'm not afraid of [AI]. I'm afraid of some of the ways some people might use it.”

As with previous introductions of new technologies into RIM, AI has more than ever emphasized the shifting nature of the record and the difficulties that can arise in identifying it in an ever-evolving landscape. However, AI, as a set of technologies, does not just cause issues within the processing and identification of records. It is also involved in their overall conception, creation, and evaluation. Its introduction into the field is not linear or measured and neither are its widespread impacts. It is clear to RIM professionals that this change cannot be approached in the same way other technologies were introduced and it must be understood that there is no

one seamless way to integrate AI into current operations due to how entangled it may be with the overall lifecycle of the records in a system.

The need to clearly define the tasks that RIM professionals undertake and the rationale behind their actions is not a new one. However, in the face of this massive shift there is incentive and possibility for self-reflection that opens an opportunity for professionals to clarify tasks that may seem second nature to them and clearly identify the reasons behind them. Without doing so, it would be difficult to train an AI program to do aspects of the job that are considered intuitive, thus limiting how useful they could be in this field. This, in turn, can and should be viewed as a chance to redefine the modern records manager and to more clearly determine what could be done by AI and what must continue to be done by humans.