Pioneers and Pathfinders: Katya Linossi
(This transcript was generated through AI technology.)

Kevin Young: Welcome to Pioneers and Pathfinders, Seyfarth's podcast about the people and ideas reshaping the legal industry. I'm Kevin Young, an employment litigator and counselor, and I've spent my career putting new tools and ideas to work in client service and helping colleagues and clients to do the same.

Zeynep Ersin: And I'm Zeynep Ersin. I lead innovation and design thinking at Seyfarth, focusing on reimagining how legal services are built, delivered, and experienced. Each episode, we sit down with someone on the front lines of legal innovation to unpack what they're doing and what it means for where our industry is headed. We'll keep it fresh, we'll keep it fun, and we'll keep it to 30 minutes. Let's jump in.

We're very excited to welcome Katya Linossi, co-founder of Atlas Fuse by Clear People and a recognized leader in knowledge management and legal innovation. From helping law firms tackle information overload to building AI-ready knowledge infrastructure, Katya has spent years exploring how organizations can turn fragmented information into trusted, actionable intelligence.

Kevin Young: In today's conversation, we'll discuss the evolving role of knowledge management, the foundation firms need for enterprise AI, and why knowledge could be the ultimate competitive advantage in the years to come.

Katya, thank you so much for joining us on the podcast. We're thrilled to have you. Just to kick things off, can you tell us and our listeners just a little bit about what you're focusing on day-to-day now at Clear People and also with Atlas Views?

Katya Linossi: I probably took one of the least conventional routes into tech and legal tech. When I moved to Denmark, I was very fortunate to be offered a job in software testing—and that was actually without a STEM background. This was my introduction to technology. I basically moved into a couple of other roles. And then when we moved to the UK, I started Clear People actually just three months after our first child was born, whilst I was also completing my MBA. So pretty intense. You could say that my introduction to legal was about 20 years ago when I started working with law firms. We were delivering accessible websites in response to the updated Disability Discrimination Act of 2004. And then I was pulled into a law firm knowledge management project. In fact, that inspired my MBA dissertation on the topic of law firm knowledge management. And I was drawing on some of Richard Susskind's work, so it was great to see that you've had him on the podcast as well.

But essentially, what we kept seeing was the same fundamental problems. Firms were not short of information. They were basically overwhelmed by it. And you could say that lawyers and business professionals have been spending far too much time searching for that information. So that led us to pivot in 2019, where we created Atlas Views. And so that's what actually made us think about—we keep rebuilding bespoke projects for our customers. And we thought, you know what, we need to build something that we can actually reuse over time. And so that's where Atlas Views was kind of born from. And if I just kind of give you a quick summary of what it is, Atlas Views is the enterprise knowledge infrastructure that rapidly transforms fragmented information into trusted AI-ready knowledge.

Kevin Young: That's excellent. And one thing you said, I like it so much, I just wrote it down. You said that firms are not short of information, that they're overwhelmed by it. Can you tell me just a little bit—I mean, from your perspective, what did or what does that look like?

Katya Linossi: When you think about it, we have got information in emails, in daily work, whether it's documents that we're creating, precedents that we're creating, whether it's in our teams—because, I don't know if you work with Teams or use Teams on a regular basis, but think about all those conversations that we're having. There's actually good nuggets of knowledge in those conversations we're having. For example, there's information that we have from third-party sources. If you think across a lot of the platforms that you're using on a daily basis as a lawyer, so you've got all this fragmented content all over the place. And it's how are you bringing that all together? You've got perhaps iManage even, or NetDocuments, documents sitting in there. And so you're constantly moving between all these different platforms, first of all. So there's that switching between everything. But there's also the daily alert. Everything is just coming at you from different places. And so how are we as humans trying to figure out what's important? What do I need to focus on? What do I need to do? What is my trusted source of truth? What is my authoritative source of truth? We talk a lot about that. And that can be very frustrating sometimes, overwhelming in that context.

Zeynep Ersin: And how widely accepted or known would you say that is amongst law firms or even corporate legal departments to understand what constitutes knowledge management, right? I think there are assumptions that typically it's just the documents we have in our DMS, but do you think that there's a wider scale appreciation for all that is involved when it comes to how you and others might today define KM?

Katya Linossi: Yeah, I mean, it's an interesting question because you could argue that knowledge management has always been important, but what is that definition of knowledge management? And so it hasn't always received board-level attention until AI fundamentally changed that. And I think we're seeing that across firms right now. For a lot of law firms, knowledge management is essentially maybe documents, precedents, playbooks, a variety of things like that. But for myself personally and for those in the knowledge management sphere, it's more than that. We've got explicit knowledge, tacit knowledge.

So we even have like a triangle around this where we speak specifically about—you've got your documents essentially over here at the bottom. And those are easy; those are things that you know can be turned into knowledge. Then you've got the bit in the middle, which is a little bit more where you need context to it. This might be conversations or something around that document, around what you're doing. Maybe even a playbook to some extent would also be that, because it's providing context to that. And then you've got the top of the triangle, which is your tacit knowledge. And that's really difficult to capture. That's because it's something that a lawyer would naturally, intuitively know how to do because it's based on their years of experience, years of doing something. And so knowledge is that combination of all those things. And so it is very difficult to define that, but at the same time, it's so important for AI.

And so that's why, as we move into using generative AI just as an assistant to actually agents, that challenge becomes even bigger. And I could imagine in the not-too-distant future, we're going to have firms have hundreds, potentially thousands of assistants and agents performing these different tasks. And every one of them needs a reliable knowledge foundation. So unlike an experienced lawyer, an AI agent cannot rely on these years of professional experience or intuition when information is ambiguous. It depends a lot on that—and what I talk about—the knowledge, the context that makes it available. And so to make AI work, you need trusted knowledge, governance, business context, permissions, explainability, and even now auditability, which a lot of clients are asking firms to provide. So these are not nice-to-haves, but they're actually critical. And so you could say that knowledge now is becoming that part of the operating system for enterprise AI. So that's in a way incredibly exciting to see the evolution of knowledge management and how it's actually becoming imperative for AI strategy, enterprise architecture, and you could argue competitive advantage. I mean, we've definitely seen that also across some of the things that I'm reading across the industry, understanding that knowledge is key to that.

Zeynep Ersin: So is it your sense, Katya, that the knowledge management problem in law firms—and opportunity in law firms—has been around for quite some time, but it's just become more of an imperative and maybe more prominent with the proliferation of AI?

Katya Linossi: Yeah, absolutely, exactly that. And I think law firms are realizing that you can put together a collection of documents for an AI tool, but how do you keep that up to date? What happens with context around that? And then every time you need to create a new AI tool, you suddenly create more documents, more playbooks for it. So that's why we talk about knowledge infrastructure a lot. And knowledge infrastructure—now unpack that phrase of knowledge infrastructure—it's not just another knowledge repository, because I think that's one of the things that maybe a lot of law firms fall into. They think, oh, we're just going to create this knowledge repository and great, that's going to work.

Knowledge infrastructure is about bringing together your information, your people, your processes, your business context, governance, and technology. And that's so that knowledge can be managed. So we talk about the knowledge lifecycle and also safely reused for both your people and AI at scale. And I think this is also the other big factor. If you think about it, a lot of law firms have piloted stuff and now realize, how do we make this at scale? And also how can we do this if we're using multiple AI tools, agents, things like that? That's what's also becoming imperative, because remember, an agent takes an action on your behalf, not just giving you an answer. You really need to provide it with a lot more context, a lot more meaning than you would just an AI assistant. So for me, that means the biggest difference is that firms that are AI-ready are treating this knowledge as infrastructure rather than just a simple collection of their documents.

Zeynep Ersin: So Katya, building off of what you said, from your perspective, in this moment, what does a successful KM strategy look like? What does it include? What does it take into consideration? In addition to the infrastructure, what are other components?

Katya Linossi: Well, I think that's why you've got to think about how are you connecting knowledge across systems? How are you then enriching it with metadata and business context? And I think that's one of the key things. How are you maintaining permissions, managing the knowledge lifecycle, preserving that provenance? I think that is critical. So a lot of times we'll question someone and say, well, hold on, you're building something for a very specific AI tool. What if that AI tool is replaced? And we can see the same situation like with a DMS. We built a lot of stuff, put a lot of stuff in our DMS, now we are stuck in there. So we talk about silos of knowledge management, silos of content.

You've got to basically make sure that you are no longer building silos, but having a consistent—and it's not just one, I'm not talking about having one consistent knowledge base—but having a knowledge layer. That's what we talk about, a knowledge layer on top of that, that's bringing things together. So you still have the content sitting in that source of truth, in that system of record essentially, but then you're building this knowledge layer on top of it. And that's what's bringing in the metadata, the business context, that's what's bringing in the governance so that you can reuse that knowledge—whether it's for search, whether it's for your AI system, whether it's for your agents of the future. And I think that's where some firms are already thinking about that.

I know, for example, I ran a webinar with Sarah Pullen from Baker McKenzie and she spoke very much about strengthening that knowledge layer, creating a foundation that made it easy when they rolled out their AI at enterprise rather than jurisdictional level. And so it opened up new opportunities to reuse and recombine that trusted knowledge for AI. There's also another important characteristic that I kind of leaned on, spoke about, but didn't speak about a lot—which was about reuse. So I think we need to think about how are we reusing knowledge? So not just having it in one source, but how we can reuse it for different things.

Kevin Young: In terms of building out the knowledge layer, how important is it in your view, and how hard is it in your view also, to ensure that the knowledge layer evolves? When I think about it, the issue is that our attorneys are attacking, the cases we defend, the deals that we work on—they're changing all the time. And what's great optimal advice on a given issue today could change a little bit a few months from now. We could be dealing with laws that don't even exist, we could be dealing with the type of case that we've not seen before. And so in your view, how important and also just how difficult is it to ensure that the knowledge foundation isn't just right out the gate once you really invest the time—that it's built to evolve as quickly as the firm and as quickly as the work does?

Katya Linossi: And that's a great summary, Kevin, because exactly to your point, knowledge gets outdated. So we talk about—you need to constantly feed that knowledge lifecycle: identify, create, and then curate essentially, manage it, organize, and then have it available. And then also monitor that to see what do we need to refigure and put back into the system again. It's not an easy thing to do, but it can be done. And it doesn't mean having this big, huge project to clean up all your data. It doesn't mean that you need to create the single repository, but it's about rethinking how you go about doing things in the flow of work.

So one of the things that we focus on, for example, is every time you create something, it's automatically tagged within our system. That means when you tag something, you're providing it context, you're providing it meaning. So if I tag something and I provide a meaning, it means that then I can use that, first of all, in different places—whether it's search, AI, or whatever the case may be. But also we have something that we call gold standard. So you can actually tag things as, this is my gold standard of a playbook or a precedent. So that means it's a signal to search—and the same principle applies for AI—that this is something authoritative.

We can also do things in the background like tagging: is this current? So for example, if you ask the question, can I please have the latest document? Because of the way we've tagged stuff and timestamped stuff, we can actually do in the background, show you only everything defined by your rules in the last year. So suddenly, instead of a person manually having to figure out, is this my latest? We can actually do a lot of that. And that's through that meaningful metadata taxonomy—common taxonomy, by the way, that's also very important. I don't know if I mentioned that. Having a common taxonomy across all your documents, across all your information, means then that if AI is asking a question and you might have something in iManage—and if that taxonomy is different to something like in SharePoint, it's not going to understand how to correlate those things. And so again, by having a common taxonomy, it knows and understands, okay, this jurisdiction means the same as this jurisdiction in SharePoint, for example. So you're bringing all those things together. So you're creating those meaningful connections, but you're also giving that auto-tagging that provides the metadata, which then instructs basically behind the scenes for search, for AI, what do I need to pull out as most relevant?

Kevin Young: I think you kind of answered this as you gave that answer, which is super insightful, but in your view—at least present-day optimal state, maybe the answer changes a little bit over time—but present day, an optimal KM state within a law firm, how involved are the practicing lawyers in that? I mean, you talked a little about auto-tagging and a lot of that happening behind the scenes, but optimal state today in 2026, is the level of involvement that an organization needs for sound KM from the practicing lawyers who are researching briefs—are they having to do much more? Should they be doing much more?

Katya Linossi: Yeah, this is a question we get quite regularly, because of course you could say knowledge management in the UK is maybe seen as a little bit more mature than it is in other parts of the world. But you could argue that yes, you need a KM team maybe to be there to just basically make sure that you've had a human in the loop to make sure something has been checked and it is considered as an authoritative form of knowledge. And that's why I talk a little bit about having knowledge available, but also managed through your flow of work.

So for a lawyer, for example, if they've just completed a transaction and they said, oh, I really like this document that I've produced, I want to submit it to the knowledge team to review and submit it as knowledge. So for example, in the way that we work, you could have it in iManage or within SharePoint or whatever, you could actually submit it into a workflow to say this is something that I think is actually really useful for the firm. And so you could actually have lawyers involved. I know their time is restricted, but you could have lawyers involved with simple things like that. So we also have a way of just simply dragging and dropping things like that to make it so much easier for the lawyer so that their time—we know and understand is restricted—but it also allows them to contribute in a way that's meaningful and time efficient. So I think there is an ability to do that.

Then there's more the knowledge team will need to be involved in the process. It'll go through a workflow where they'll verify, check it, say, okay, this all looks good, and then it gets submitted. But I think also what we're doing and where we're going, for example, is things like ambient knowledge, because often sometimes somebody might be creating stuff but not really shouting out about it. Through connections of what we understand by tagging stuff, we can actually see, hold on, there's some really great stuff here. So that could actually be submitted to the knowledge team to say, hey, we think this is really good knowledge. Do we think we should add this into—because we call knowledge collections—or should we add this to our dynamic knowledge collection and be part of what AI or my agent might be using for that? It's a complex thing, so that's why I'm trying to break it down into easier parts. But the whole idea is behind the scenes: if you are identifying things because you have a common taxonomy, a common metadata, because you have some of the tools that we do—for example, in Bold—this is already simpler than it should be, and will become.

Zeynep Ersin: So based on your work being in this industry and certainly focusing on legal—and I know there's a lot that is very unique to the legal industry as it relates to knowledge management—but are there any other industries that have made significant strides in terms of KM, whether it's other professional services industries or any lessons from how other industries have handled KM that might be relevant or transferable, or at least knowing about that could be pertinent for legal?

Katya Linossi: Yeah, so we work across legal and professional services. So I actually do see that as being very complementary to each other. There is no doubt there's some very specific ways that law firms work that do not translate to other industries. But then there's other things in terms of the whole knowledge lifecycle management—all of that you could say is the same for all industries.

If I look, for example—I'll give you an example. We work also sometimes with construction, and actually we got a great idea from them about creating a tool we called HoverPoint. So HoverPoint is an image that you can hover over and it gives you information about it. So for the context of the construction firms, they needed it for safety. But when you go on site, you can then hover over things and you can see, where do I need—where's my health and safety guidance or whatever? But interestingly enough, it's become one of our most popular tools for law firms, especially in the knowledge function, because it provides a very visual way for checklists, for workflows, a visual way of doing things. So sometimes one industry might give an idea for something for another industry. And I think borrowing those best practices from there and using them is really helpful.

When it comes to AI, I do think though that legal is leading. There's no doubt about that. So I think that one's a bit more tricky, because when I look across other industries, there's no other industry that has so many solutions, specifically AI solutions, compared to legal. So that's in some ways useful for us; we're learning from that and then applying it to other industries too.

Zeynep Ersin: Why do you think AI governance needs to be treated more today as a business issue than a technology issue?

Katya Linossi: Yeah, it's a timely topic. In fact, I recently co-authored a governance guide on Microsoft 365 and Copilot with my colleague, Steven Grant, because we're getting this question regularly, by the way. And I think there's two big shifts driving this. First, AI adoption is faster than basically any enterprise technology in history. So that's bringing a lot of complications and risks with that. Second, we're also getting a lot of regulatory and client expectations around governance, transparency, and auditability. I think I saw something on Gartner where only 23% of IT leaders actually express confidence in safely managing security and governance controls when deploying AI. So that's quite a low number. And that was some about six months ago. My guess is with agents and stuff like that, the agentic framework, that will almost probably be even lower than that. So it's complicated. We see knowledge governance and AI governance becoming inseparable. And so that's, I think, one of the things that we've got to think about, because as I mentioned earlier, knowledge is what provides those agents with guidance. And so therefore the same for AI—they need to work hand in hand. So I think those two are the main things for me.

Kevin Young: If you had to think out—look into the crystal ball and think out maybe, I don't know, five, six years from now—it seems like right now, investment and advancement in AI and probably in KM too is a competitive advantage that some firms are very far ahead, some are probably having them started somehow. But how long do you think that remains the case? Versus, I think about other innovations that weren't as innovative, but things that have come along in the last 20 years where it seems like every firm kind of has access to the same tech and that kind of thing. Do you think by six years from now, good AI and decent AI governance—is that still like a distinguishing factor among big firms, or is it something that's kind of just the cost of admission?

Katya Linossi: I would actually argue the other way around, that we can all have the same access to technology, but knowledge is what makes us different, isn't it? Expertise is what differentiates you, isn't it? That essentially is your USP—how you do things, how you provide advice, how you interact with clients. That is your USP. And it's also not easy necessarily to capture. So I think the component of an advantage is around your institutional knowledge. It is around your expertise. It's around your workflows, your client context, and your ability to operationalize that at scale. So it's the firm's ability to make—we call it—collective knowledge and judgment. And how do you make that reasonable through AI? So that's why I think that is the more important part. The tools will evolve, the tools will change, but your knowledge and your collective knowledge and judgment won't essentially.

Kevin Young: Yeah, and I should be clear—I mean, I agree. My question is mainly on the knowledge management function within the knowledge base. I totally agree. I feel like where we're at now and the things we can very safely dream of right now in terms of what's possible are things that we thought about all along, maybe not using terms like AI, but 10 years ago, 15 years ago when I joined the firm, you would have thoughts that we are at this great firm that we all believe in, with decades and decades of experience advising clients in the same area—and how can we leverage the power of what we've done in the past to be as quick and efficient and consistent as possible? And so the way to operationalize that, or the ability to operationalize that in the most efficient way we could have come up with, is something that's always, I think, kind of been part of the dream state. It just feels very attainable now.

Katya Linossi: Yeah, oh, absolutely. And I think that's why a lot of the concepts in knowledge management are the same, you could argue. We just haven't had the technology that's kept up with that. And so that's what makes this exciting. I think if we look back 10 years from now, we'll say, well, that's crazy. Why did we leave valuable knowledge trapped in documents and emails and in people's heads? We can do that so easily now. For example, we might even think how incredibly inefficient it was that we weren't using context and systems that could understand all of this. So absolutely, I think there will be things that technology will be able to do for us. But I think again, it's how are you doing that—and that goes back to the knowledge infrastructure. How are you doing that with people, processes, combined with the technology? And I think that's the important part. We've got to have all three of those working in harmony essentially.

Zeynep Ersin: I think we're at this crossroads too, where I'm happy to see and hear just how much knowledge management has become a huge hot topic, right? I think there seems to be a pretty wide acceptance of the importance of your knowledge and the role that it plays, in particular now with AI. I think it'll be interesting to see how firms think about their AI strategy coupled with their KM strategy moving forward, because so many have systems in place, processes in place, but without the benefit of generative AI to even organize what we were just talking about. But now it's, okay, well, how do we elevate this even further? Where can we go with this in ways that we hadn't even thought about before? I think that this opens up a lot of additional opportunity for innovation in terms of what is your KM strategy? How do you leverage it? How do you leverage the data for benchmarking, for predictive analysis? So much more than what I think historically was a reference point or a search function. And so I'm excited to see the evolution of where this goes next, in particular just now having the benefit of generative AI to be part of the conversation to move it to the next level.

Katya Linossi: Yeah, absolutely. We're thinking about that. Where is this evolving? So we talk about the trusted knowledge layer now. Where we see it going is the enterprise knowledge infrastructure. So this is also where we see things like skills, agents, plugins working together—and I talked about automatic knowledge creation, ambient knowledge capture, doing things like knowledge health analysis, things like knowledge graphs and ontology are going to be really key in that. But also you've got to think about things like AI harness, trust, and risk. So how are you bringing all that into play? And then down the line, the next level is we talk about collective intelligence. How are we creating this? Because the whole idea is if you're capturing knowledge, it's going to evolve over time and it needs to get better over time too. So you're going to have this, what we call an institutional memory, essentially. So that's where I think things are going to go. We're going to have things like organizational reasoning, predictive knowledge even. So we talk about ambient capture, but what about predictive knowledge in the future, this continuous knowledge? So that's the exciting part. That's the technology part. But again, I'm going to reiterate, we need to make sure that we've got the humans in the loop, we've got people, the processes, and we're getting this all working together. But that's where we see things going, in that whole enterprise collective intelligence.

Zeynep Ersin: I was going to ask—the next question I was going to ask is, what are you excited about and what developments are you looking forward to? And I wish our listeners could see Katya's excitement when she's talking about the potential for what's to come, because it's a very real excitement. Are we far off from that, from that being right within reach?

Katya Linossi: No, I mean, some of what I spoke about is next year. We're already putting this entire roadmap together—some of the things that we're doing, some of it's later this year, next year. But I'd say that collective intelligence, maybe that's off a couple of years. Some of that stuff is quite advanced. Let's just put it that way. And it's also about how you're going to do it in a way that's—and this is an important thing—safe and trusted and authoritative.

Zeynep Ersin: What are the types of KM roles in law firms that perhaps don't exist or aren't widely existing yet, in terms of those roles with the future of KM in a law firm environment?

Katya Linossi: Yeah, it's a good question. We're going to see your traditional roles still be available, but I think we're already seeing that hybrid between knowledge and AI roles. LinkedIn leaders have done that recently. So you're going to see more of that. So you're going to have to have people that understand not just knowledge, but also AI as a capability. I think you're also going to have to have someone who is going to, in some ways or some form, be in charge of governance and all of that too—in a maybe softer role, but just making sure that you've got somebody who is looking after maybe even the governance of all the AI tools that you have, whether they're in the knowledge function or the IT function, I'm not sure, but somebody has to be able to make sure that whatever's being outputted, there's transparency on this, what's happening.

I mean, we're even thinking about this within our tool, which we're doing right now—we're logging everything. At some point, somebody needs to go back if a client requests something to know, how did that output come about? I mean, in fact, you could argue AI is already one of our biggest consumers of knowledge, but it'll actually become our biggest creator. So we've got to understand how did AI come up with this answer, especially agents. How did it come up with this? What did it do? Somebody's going to have to audit that and check it and make sure that that's doing what it's meant to be doing. And then maybe in the future, there might be other roles around the tacit knowledge. So maybe capturing that in a way that might not be as manual as it is right now, but basically having ways to be able to capture that in a less manual way.

Kevin Young: On that note, any books or articles that you've read lately, seen lately, that our listeners should know about? I'd love for you to be able to plug.

Katya Linossi: Yeah, I'm constantly reading a lot of stuff. I had to punt a little bit from the book that I contributed to called *The Future of Legal Knowledge Management: Harnessing Artificial Intelligence*. Also Grant did a chapter as well with Sarah Pullen from Baker on this.

Zeynep Ersin: Katya, thank you so much for joining us today. It was an absolute pleasure having you on. Thank you for your time and your insights, and we'll be in touch.

Katya Linossi: Wonderful, and thank you for having me today.

Zeynep Ersin: Thanks for spending time with us today on Pioneers and Pathfinders. I'm Zeynep Ersin.

Kevin Young: And I'm Kevin Young. We'll see you next time.