AI LEADERS PODCAST
EP 85 DEMYSTIFYING CONTEXT: WHY YOUR AI AGENTS NEED A SEMANTIC LAYER
TRANSCRIPT

Juan Sequeda [00:00:00] Your knowledge, your semantics, your context, like that is truly the DNA of your business.

Teresa Tung [00:00:13] Welcome to Accenture AI Leaders podcast. My name is Teresa Tung; I'm global data lead at Accenture and your host for this episode on demystifying context and why a semantic layer is needed for AI. We're going to talk about why do agents and this age of AI, generative AI needs context. What is a semantic player? And did you even know you needed one? And I'm thrilled to be joined by two experts in this space. So, I'm going to start with Juan. Juan, can you introduce yourself?

Juan Sequeda [00:00:47] Thank you, Teresa. So, my name is Juan Sequeda. I am the principal data strategist at ServiceNow. And I joined ServiceNow about a year ago when they acquired our previous company, Data.World, which is a data catalog and governance platform all built on semantic knowledge graph technologies. And I've been working in the space of semantics, ontologies, knowledge graph for about 20 years. So, I carry my academic scientific hat but also been working with very large organizations for a long time, so. Really excited to be here and thanks for the opportunity.

Teresa Tung [00:01:19] And thank you, Juan. And many of you know Juan will give no BS. That's literally the name of one of his podcasts. So, I'm very happy to get Juan's perspective here. And finally, Accenture's very own Ora. Ora, can you also introduce yourself?

Ora Lassila [00:01:35] Thank you, Teresa. My name is Ora Lassila. I've fairly recently joined Accenture, but I've been doing knowledge graphs and ontologies and semantics for a long time. Back in the late 90s, I worked when I was at MIT, got to work on RDF and then the whole semantic web activity started, and I have basically done nothing but that ever since.

Teresa Tung [00:02:03] So lucky to have Ora join Accenture. He's being so modest, but he is the founder of the semantic web concept and is really ahead of its time. So, we're going to get into that, into this session. So, let's start by setting the scene about why context matters. And so, Juan, I'm going to start with you. We keep hearing that context is everything in AI. And so, what does context actually mean when an AI agent is trying to answer a question against. Enterprise data, and why is it so hard to get right?

Juan Sequeda [00:02:36] Let me zoom out for a second around this, and I want to talk about what I'm seeing is kind of like two types of semantics that's happening right now. Now, obviously, everybody's talking about semantics and context all right now with respect to agents. And actually, our research lab a couple of years ago was kind of the first to actually start pushing out the research and actually numbers and benchmarks when it shows that if you were going to do question answering over enterprise data, if you would do it over a knowledge graph, versus not one, like over directly relational database, we showed that your accuracy would increase like three times.

And this was back now three years ago, right? So, the models have changed and everything is, but at that time, it was really, it was kind of that first evidence that people were kind of hoping to have around how the semantics can actually increase your accuracy. That's kind of what kicked off the whole movement, I would argue. Now, going back to the two parts of semantics, right now people are focusing on semantics to get the context for your AI agents. That's one, and that's what we need. But if that's your ultimate focus, I'm already seeing that kind of be a little loose semantics in a way, saying like, oh, maybe it's just like some documentation, or we'll just put into some YAML syntax so I can just process, or some skills, and that it. But then there's also the semantics for the interoperability of all types of systems that you have in an organization. And guess what? These AI agents. Is another system that needs to interoperate with the rest of the data you have. So, when you start thinking about these two types of contexts, the types of semantics, if one is only for the agents, the context for that stuff, all right, great. But when we really think about it from a foundational perspective, you really want to start to find that semantics for the interoperability of all those systems. And I'm not saying there are two different things. They actually should be the same. But what happens is that if people are starting to go think about the context just for agents. They are not going to be doing that foundational work for the semantics of inoperability of systems. And guess what? The complaints that people have had always about, I don't understand my data, or why do I get different answers from the different answers for the same question from different dashboards, we're going to have that same thing. We're already seeing it. Is I'm asking a question, these agents are giving me different answers. It's because the semantics, you're thinking about it very loosely instead of really defining that foundation for the Semantics of Inoperability of Systems, which also needs to be governed.

Teresa Tung [00:04:58] We've jumped right into semantics, right, and as part of context. So, semantics just means meaning, right? So, I think, or I'm not sure if you wanted to add upon that, but how are semantics related to context here?

Ora Lassila [00:05:09] First of all, thank you Juan for bringing up interoperability because that's the whole reason why we started the semantic web activity originally. I also wanted to say that we build semantics also for humans because thinking carefully about semantics forces us to really think about what our data is about. So now we can talk about context.

Juan Sequeda [00:05:36] Well, so I think context is going to be a mix of multiple aspects. And I think, number one, I would say you want to be able to understand the meaning of the data, meaning that's going to be one part. But when I look at how we're actually looking at context is you have users, people have context around that to understand who is that employee, and what department do they work in, so forth. What are their particular tasks and goals? Then you actually have context around security, for example, and identity. Like do they have access to some data or so forth like that's extra context. Then you also have the context of your decision traces in your workflows about how a particular business process is done. So if you start combining all those contexts together the typical example I say is like okay somebody has a question about a customer well okay what do you mean by customer I have three definitions of customer oh I know that the person asking that is in the marketing department so they probably care about the marketing definition of customer do they actually have access to the data for that stuff. And, by the way, I probably know what marketing people do with customer data. So, I can probably make a suggestion of what, or probably predict, what are you trying to do with that customer data? And probably no other people have done things. All of that is the context that comes in. And I think semantics plays a key role in that stuff. But again, we're, we, we were using this word context so much. And I think the other aspect is when I look inside the industry and if we look at talk, talk to folks who are just in the data analytics world.

They may be thinking of context just in the aspects of, again, data analytics, but when you start combining analytics with the rest of the operational world, everything I just described, the users, your decision traces, security, that transcends the data analytics side. That's more on the operational side. That's why it's really important to start kind of combining the operational analytical today.

Ora Lassila [00:07:18] As far as users are concerned, good that you mentioned that, but there's also kind of like the situational context is that not only who the user is, but where is the user? What's happening around him or her? What was happening just before and things like that, which we often might forget, but those can be important factors when actually trying to solve problems for the user.

Teresa Tung [00:07:45] And I think the good news through all of this is that, while the term context as it's used for AI, specifically agents, is new, that the technology and the treatment of context and specifically semantics, we know how to do this as an industry, right? So anytime any of us have done either, I'm sure internet shopping, watching a video on an online streaming service or a TikTok video or just internet search, all of us have experienced the context or that you and Juan were just talking through, right? Based off of what I have purchased in the past, I actually recommended certain things, right, because of that very dynamic notion of context. When I do an internet search and I search for Accenture, you're going to know that Accenture, one of the things is we're a company and company has places and locations. A good way to see that is on a map. And so, I'm going to see the location office closest to where I'm sitting right now, right? So, all of that context is something that, again, or I think this is some of your seminal work. So, the role of semantics with this knowledge graph. So, you have co-authored the original semantic web vision with Tim Berners-Lee back in 2001. So, 20 plus years later, AI agents have arrived and suddenly the industry needs pretty much exactly what you were describing. Does it feel like vindication or is this a different problem?

Ora Lassila [00:09:16] It's not a different problem. In some ways, maybe the scope and the setting have changed a little bit, but it definitely feels like Vindication have been talking about semantics and ontologies and agents even for a really long time and all of a sudden everybody's excited about that and they're no longer running away when I'm talking about it. So, this is, this is all very good. I think what has changed is that originally, we thought that the problem sort of encompasses the entire web and we solve it for the whole web. Now it's maybe more sort of an enterprise context. We are solving problems within an enterprise using enterprise data. In some ways that simplifies the problem a little bit. But I think that fundamentally we're still very excited about the idea that we have data that can be understood, we have agents that do the understanding, the agents apply context, which the way I think about it is everything we know is context. And giving proper formal meaning to these things will make the agents function better. So, we're living in very exciting times.

Teresa Tung [00:10:44] I think as an industry, we've been really good at describing context maybe as data as it's born, right? So, a lot of the practice around data products was more describing some of the technical operational and increasing the business metadata for these data products for people to better understand. I think what we've not done yet at scale that's needed now is to describe the context of the data, how it's used. I think Juan, even in your example about customer data. I'm going to use the same customer data for many use cases, prospective sales, customer support and service, right? And that description of how I find and use and apply the right data, I think that's something we haven't done as an industry. So, I think maybe with your experience at ServiceNow, connecting enterprise data and data services with Knowledge Graphs, Juan. When an agent hits a data source without a semantic layer in this context, what can go wrong? Can you give some concrete examples?

Juan Sequeda [00:11:50] Yeah, so I think, let's start off, first of all, with the accuracy, right? I think one of the clear things that we know that we've already learned, I think everybody's learned this, they've hit the wall themselves, is that without the semantics around this, you don't know what something means. And I think at some point, that's why you have the whole hallucinations and stuff, because then these things will return back something that you don't know if it's correct. So, I think that's the first thing that we see that without semantics, around here, without your semantic layer, your ontology, and so forth, you're going to have reduced accuracy completely. So that's number one around that. But number two is the governance around it. Because I also want to have the explainability, going back to my answer, if I ask a question about customer, well, you want to know that there are multiple definitions of customer. And I'm giving you an answer with respect to the definition of what is a customer for the marketing department. And as a user, you may want to note, but why did you give me that one? Who said that was the right thing? When was the last time that was governed and audited to know about that? All of that is the governance and the traces that you want to keep track of. And then, because at the end of the day, that depends on your industry, regulatory type of stuff, you need to keep of all this lineage, right? I think now, I always say that governance is getting much more complicated. There's the surface area is greater, but it's also becoming non-negotiable. Everybody's acknowledging right now that LLMs because of their non-deterministic nature and you're going to have agents that are going to come up with things that if you ask it again It may come up a different plan. You want to be able to say, okay this part actually needs to get assigned to a Quota skill, which is really a very deterministic Workflow or it's going to be something that is well-governed that I know what it is because I there's no guessing allowed here right and I think that's something that we're seeing a lot right now, especially because when it comes to agents. Pareto principle, I think, applies here. 80-20 rule, 80% are probably going to be things that is just these processes that we know how it gets done. I just want to fully, not just automate it, like have this fully autonomous. And the humans will be there just to be notified, hey, we're doing this, or when an exception comes up. And at the end of the day, that goes back into workflows. And connecting to ServiceNow, like that I think is the originally fascinating, the bread and butter of ServiceNow being workflows, that your time workflows. With semantics altogether, this is like that, I think that perfect marriage that you're seeing. And what I'm observing is that a lot of people come to this conclusion on their own, which is great. And then they're somewhat also reinventing the wheel. And I think this goes back to your earlier point, Teresa's like, we've known how to go do this. And there are organizations, there are people who've had experience around this. And I think it's time to really kind of come together to make sure that. Let's not waste some time, and let's not reinvent the wheel. Let's reuse what we know, and let's go focus on the truly unknown things we have to go figure out going forward.

Teresa Tung [00:14:52] I mean, to that point, as you were talking about what could go wrong with agents and the hallucinations and getting them not knowing what data they're using when they're talking about customers and is it marketing versus is it financial, right? In your example, it reminds me again about that internet search example. Many of us are old enough that we remember internet search before knowledge graphs where you would again search for Accenture and there would just be all these links and then we would have to like to click through and search through all the links and find Accenture.com.

You find maybe the locations of the office, and then you kick open like a map quest or something, and you map it, because that's what you want to do. All of that, some of us remember that. And I guess that ability to have something not just easier to do, but repeatable, accurate, you talked about explainable workflows, right? You would think that a lot of the use case people want agents to do for enterprise are just that, right. It's precise, it's not different every time, it is explainable. So, makes sense that we know how to do it. Why will we not do something like this? I see another side of the argument though. So, there's a version of the argument that says where LLMs are good enough at inference, right? They have enough inference capabilities that they'll figure out the schema and this ontology from the metadata. Is this wrong? Or is it sometimes right, or can you use an LLM to generate an ontology?

Ora Lassila [00:16:22] You can, and sometimes they do a very good job, but it's only a starting point. Ontologies, the way we traditionally crafted them, represent a great deal of very careful thought. And we can use an LLM, we can use an agent to get started, get you the sort of skeletal frame of an ontology. But after that comes all the things that the LLN probably does not see or won't consider. Humans are very important for this. So, I would not accept an ontological created purely by an LLLM. But if it speeds up the work that ontologists, these people who actually understand how to build ontologies, how they create ontologies. If it helps them, then I'm all for it. Then I've seen good results in that regard.

Teresa Tung [00:17:27] I think, Juan, you were talking about governance, right, and I think that probably goes hand in hand with Ora's point. Can you talk through why governance is needed in the ontology?

Juan Sequeda [00:17:37] We were talking about this the other day is like you could you definitely could you definitely can and actually should kind of use LLM's to kind of kick off your building your ontologies But that is still a means to win end at the end of the day vibe coding can help but you're not going to vibe code your key algorithms, right? You're going to go back and look at it and you want to make sure this is the correct thing and there are things that may have nuances that you want you want To be able to confirm with other people and so forth and I think that whole human aspect right there that that is the governance because governance at the end of the day is to provide trust and we want to be able to provide trust so we can do things at scale so people can go they can go use their systems and they know that they can run with it and because they trust it is working well so governance should be seen as an enabler governance people traditionally think about it as like the brakes in a car where you slows you down which yeah you have to make you want to make sure bad things don't happen but governance must also be seen as brakes in a car that enable you to drive fast safely. So, make sure that great things happen around that. And I think when we start doing this, defining your ontologies, governing everything you're doing, that's so later on systems can say, oh, I can just go run with this because I know that I have my guard rails kind of installed around that, so I think that's what is really important is that it's this combination and to know that you're in a way quote, you're slowing down to speed up time of approach. And again, that's how we're actually going to go scale right now.

Teresa Tung [00:19:06] I think there's two takeaways that I'm getting from this. So, I think one is we should use vibe coding for ontologies, right? So, so definitely do that. That's going to be helpful for everybody to build this layer at the rigor that has to be done, but it's missing the mark.

You both said that we're only documenting what is documented, right. This is like the easy stuff. If I have a report or if I have pipeline on it, most of, and I'll use Accenture, like our best data is I think our finance, our sales, our IT, right, our, our HR. I hope that's not what anybody thinks about Accenture as what makes us really special, right? I hope it's our deep industry, technology functional knowledge. And a lot of that is it, if you're trying to use LLMs to kind of learn from or reason from that's generic outside knowledge, or maybe it's some of our best data, but it's not what makes it special. Our best data is not even governed in a place that is. You know, what does Teresa think about when she's making some of these decisions? You're saying that needs to be actually governed and then somehow extracted.

Juan Sequeda [00:20:11] Yeah, and I think what's actually really exciting right now with LLMs is that we can really scale the function of knowledge engineering and these roles with knowledge engineers, which we were doing this in the 80s and 90s, and it quote, quote, fail because we just couldn't scale it. But the practices were needed. Now we're realizing that we need to go do this. And guess we've got great technology now that can actually help the scale. So, it is a way now to kind of extract what Teresa thinks about this concept and so forth and be able to kind of do that scale. And one thing to add to this here is also we need to be comfortable with somewhat of, call it chaos, maybe it's too part of a word, but agreeing that we disagree, acknowledging our disagreements, is an agreement. And I think that is important to understand. So, if you're going back to like, well, what is a customer? We can't expect to have just one definition of customer. It's fine if we have multiple definitions of customers. At least let's make them explicit, such that when an agent comes in and is asking about something about customer, it will know that there are different definitions. And it may have some extra context to say, okay, the one that I want is this one, the marketing one. Or for some reason, if it still doesn't know.

Because the context doesn't help it to disambiguate, maybe that's a reason why you want to go back to the user saying, hey, we need your help here. And by the way, you'll learn from that.

Ora Lassila [00:21:36] This is an important point because ontologies, in some sense, represent an agreement. They are almost like a contract. It's like we agreed what certain things mean. And I think just for this reason, it's important to have humans involved in this. And I think the governance that Juan's talking about is exactly this, humans agreeing on things.

Teresa Tung [00:22:06] And so the AI makes it easier to document and capture what the humans agreed or explicitly disagreed, which is an agreement in itself, right, on and now the AI knows what to do. So, I think, you know, it still gets harder, right? So now that we know AI solves some of that part, right. Humans still have to agree and that remains true, but the AI is going to make it easier to document and get that, but. There's a difference, is there a difference between a semantic layer that's genuinely AI ready versus one that just has good documentation? Hey Juan, like what has to be true about how meaning is represented?

Juan Sequeda [00:22:49] So this goes back to my earlier point, like, there's these two types of semantics around this and, and look, we need to acknowledge that this this whole topic of semantics, right? I've been working on this for 20 years or have been doing it for double that. So, but anyways, the point is that this has always been an area of work, but we have never been invested into it. So, one I think there's, we've lacked incentives. But also, humans want always the magic wand. They want something super easy and automated, but in reality, want to be able to lower the barrier. And to be very frank, the barrier has kind of always been high when it comes to dealing with managing semantics. That has been the hard thing to do.

Heck, you need somebody with a PhD to do this stuff. So, what is changing now is that we have new tickles, technologies, vendors, kind of providing more automated ways of doing this, right? You've seen this from the latest announcements from the list, from the latest kind of big, other platform companies. But, and that's a great start. That's a good start to go do. But I will say that a lot of these are coming in from, again, also your data and analytics world where your starting point may be your effect and dimension tables, right. It's an important start, but again. Your enterprise is much more than your fact and dimension tables. But that's a good point to go do. So now you can probably start from there and then build out your semantic layer for your BI, and then you can go edit things or other. That's a great start for sure. But then if we're really thinking about kind of enterprise-wide meaning around things and goes back into my other point of, we need to tie analytics with operational world. That's when we start getting things on a larger scale, which is much more than just your data lakes and stuff that we've seen that you start getting into, wait, if not even humans will agree on this, why do we think an agent is going to be the right thing? I'm like it's going to provide another perspective that could be used. So, I think that's where we start seeing that it gets much harder when you're thinking about this at a true enterprise. Kind of interoperability of systems within your enterprise and the way to go deal with this is to really you got to start you got a start small think big type of approach but what I tell folks always is that you want to start you want make this investment because you are going to generate some compounding value so the most important thing is if I'm going to think about this work it is going to be work we're not going to we're not we're going to sugarcoat this you have to do work but I want that unit of work that I could go for the next use case, reuse 80% of it. And that's the compounding value that comes over and over again. Then suddenly you're like, well, I did all this work and then all these use cases are starting to happen more and more without my involvement in it. That's a compounding of value and that's showing how semantics starts to help you scale.

Teresa Tung [00:25:51] I'll say the other side to it it's just inaccurate without it I think we started at it but I had a case where we started using a Conversational AI this AI for BI type of use case where they had data products and they were actually well documented Right that the data as it's born and it was in a sales function there were only 11 data products, but you can imagine these are pretty big Data products right it has a lot of features to it. We can argue if that's a good practice, but that's what they had, right? And there's just too much similarity, right? Too much similarity. You weren't able to understand what you were getting. And so, you had to go back and build in semantics. Should also go and refactor the data products, but it's thinking through how I actually make it so that I can get value from the data. And then if you're doing that and you're fixing the data product, you're building out how it's used, as new data comes along, you don't have to be explicit anymore. Where you could actually find and incorporate that. That is more easily within your agent workflows.

Ora Lassila [00:26:54] Some of the foundational technologies that we use for this actually make it easier for us to take individual data products or like data from these proverbial silos even, and then ask questions, write queries, what have you across all of these. And now adding new data products, adding new sources of data is no longer an engineering problem the way it used to be.

Teresa Tung [00:27:21] I'm going to give you a critique, though, that I have heard, Ora, since you founded the Knowledge Graphs, right? The other critique I hear is, I've done Knowledge Graph before, and it didn't work. What's your answer for somebody who's criticizing it like that?

Ora Lassila [00:27:38] Rather than saying you did it wrong, I'd say that, well, first of all, we have the AI to help now, and it is going to make things easier. So first of, starting to build an ontology is a lot easier. But I think it's also that AI is kind of the incentive to do this now. We want something that the agents can consume. And so now you have a higher incentive to do this, and the barrier is lower and the AI is going to help you. So, I think we're in a good place.

Juan Sequeda [00:28:17] Yeah, and another thing to add, that I've seen this a lot too, and that was one of the mistakes I made at the beginning too, is that we do things, we're bowing the ocean, right? Which means that we're trying to go so much and we don't have a very clear kind of path on like, what are the use cases and what is the value that we are delivering these use cases? So, I mean, a lot of these are platitudes, I'll acknowledge, but I see this over and over again, that like, well, it's that you didn't have a lot like the- the basic foundations from organizational support around that, which is not a fault of the technology. It is not understanding the bigger picture and connecting it all together. So, I think that is now. And then the other one is the incentives. I think AI is the biggest incentive right now. But also, what I'm cautioning, folks, is going back to my other point. There're two types of semantics. You're going to take the easy, quick route. And then we're going to have the situations like you just discussed, Teresa. It's like, oh, I have this stuff. It has semantics in there. But I'm still getting wrong data, because you really didn't do it right. You did it quickly and stuff, right? So that's the stuff we're going to go see. But I am seeing other organizations and people doing this, quote, the foundational right way. And what we're going to be seeing is like two organizations saying, hey, I'm doing Symantec, it's not working. I'm like, I am doing it, and it's working great. And it's actually finally, and then we'll see, I mean, we're already starting to see more and more of that. And I think that's what's going to kind of be a wake-up call for industry.

Ora Lassila [00:29:45] One of the interesting aspects of this is that in many ways, technology is not the real challenge here. We are dealing with human problems and organizational problems. In fact, these are problems we've dealt with for a long time. I think technology is going to help us now, but we mustn't forget that some things haven't changed.

Teresa Tung [00:30:08] Yeah, so, I mean, who owns that problem today in most organizations, right? And why is the ownership so messy, building on Ora's point? It's a human problem. Juan, what are you seeing from, why is it messy from an organizational problem?

Juan Sequeda [00:30:26] Call this the politics of knowledge, I mean, first of all, this is all spread around within an organization, and I think it goes back into what are the incentives to go do this. And if we look at, when it comes to interoperability of data, data is a means to an end, right? Governance is a mean to an ending. All of this is a means of an end where the real goal here is to get some work done, right? There are some business functions that need to occur, which is going to increase revenue, reduce costs, or whatever, and that's what needs to happen. Now, we need data for that. And we probably need to go bring in data from multiple sources to be able to accomplish that particular work goal. But because of our incentive to do something so quickly, fast, we just do a quick point-to-point solution and we get it done. We finish that particular piece of work and we're done. So, we were incentivized for short-term success and we did it. And then the next task comes along and I need to go integrate other data which may have been data already integrated before but either I didn't know or I'm just going to do it quickly my way another point-to-point solution happens so all this is what we've been doing because we've incentivized for short-term and then later on people come in saying well I need to go do another thing and then well I should be able to go reuse this data but I can't make sense of this or I get this is all

I have all this debt around so now you are really kicking the can down the road, so I think what is happening, what needs to happen is that there needs to be kind of this incentive change. And this is a really a leadership perspective that if we're really going to abuse AI as a transformation to really reinvent our business, because this is the moment that we're in, we really need to think about incentive structures at the leadership level and saying, look, we want to be able to do things for short-term, but also for long-term investments. That I want to make sure that whatever I'm investing today, going back to my previous point without compounding value. We should incentivize for compounding value, for example. And I think that's the change we're making. So, companies who are making that investment now are going to be leapfrogging. They're finding ways of doing both. I'm not saying only to do long-term investment. I'll do short-term. This goes back into identifying what are the use cases, how to start small, think big, and type of stuff. I mean, frankly, heck, this is why ServiceNow even acquired our company Data World has been. From the beginning, a foundation built on semantic knowledge, graph technologies, and data as a service now was also heavily invested in doing the same type of approaches. And when they started their new data analytics business unit, they said we need to have a catalog. Went out to the market. They saw that we were the best one because not just of our market, of what we were doing with our customers, but also of our technological innovations, which align with what they're doing. So, I think we're just going to need to see that. And we're starting to see this, right? People making investments. Of finding that balance of short-term but also thinking about long-term and incentivizing those long terms.

Ora Lassila [00:33:20] I think interoperability is one of those long-term things. Interoperability was never an easy sell when people were looking for these quick point-to-point solutions. But interoperability really starts to matter when you're playing the long game and you're thinking about what happens in the long term where you are building some kind of a system where you just can keep adding new sources of information, new knowledge.

Teresa Tung [00:33:47] Yeah, I mean, I don't think enterprise data is ever going to be centralized in one platform, and it probably shouldn't be, right, from a sovereignty and just a cost and an access perspective. But if we think about being able to have agents, much less humans, be able to find what to use, to use it, to make the right decision to learn from and to come back to it, that sort of bet starts with, I think you're saying some of this good practice, right? It has semantics in place. So that you can make sense of all of this. And if you kick the can down the road, you might have hard coded it within the agent skill itself or within the prompt. As soon as something changes or some of that data source changes, then you have to change all of that. Even if I were able to find what to change would be so hard. But if you started right now, while you're saying ServiceNow is doing, we're saying this is a strategic sort of competitive advantage for ServiceNow, it sounds like, right? That's and companies should think about what it means for them. And to be able to model the semantics in that domain sounds like a really important moat for a company to think about doing right now.

Juan Sequeda [00:34:58] 100%, and I think this is why we're seeing now that, and we're using these terms interchangeably, but that your knowledge, your semantics, your context, like that is truly the DNA of your business. And that's something that you need to be able to manage. And the management part is not just creating it, right? But it's like, what is the life cycle around this? And then again, it goes back into like, you need to know what does this aspect mean? And has this definition of our metric been revisited?

IHow often like all this is the governance that keeps our stuff our knowledge alive and that's why governance is going to be such a bigger problem that's what we need to automate more of it right in that's where workflows are going to play a key role of being able to go automate the governance of all this knowledge that we need to go maintain and that is why I am personally excited about how everything is coming together at our side of the world here at ServiceNow.

Teresa Tung [00:35:52] I want to ask you each for a practical takeaway, right? So hopefully we made the case why semantics and knowledge graphs are needed for context for AI. But for a leader listening to this right now, where do they start? What is the minimum viable semantic layer, and or I'll start with you.

Ora Lassila [00:36:09] Yeah, semantically or the way I sort of look at it is that it's, it's a knowledge graph between your legacy data and the applications and agents that you build. It insulates the applications and agents from all the mess that your legacy database. But it also provides this. What I, what I mentioned, you can keep adding things. Interoperability comes into play. And the kind of the marginal cost of adding new things is lower and lower. And I said this in a talk once that I gave that if you're trying to understand the short point-to-point solutions versus the long-term thing, these are now technologies that essentially are perfect for the use case you have not yet articulated. They are kind of like the means to future-proof your investment.

Juan Sequeda [00:37:04] A hundred percent that goes back to the incentives that I've been kind of argue for, my side the practical takeaway is so first of all, we have to be pragmatic and part of is we need to make sure that we need a start with the use case that is connected to a top-level value. So, I think about this as like the analogy is the elevator, right? This is Gregor Hopes architect's elevator.

The elevator I should be able to go up to the penthouse, go to the right, understand what is being discussed strategically. What are the top-level objectives we're trying to accomplish in the organization? And then figure out what is the data and the knowledge that is needed to go answer that. So, then you want to be able to take your elevator all the way down to the engine room where the work gets done to be able to figure out what is data and figure out kind of the different pieces of knowledge for there. But what is important is to start thinking, making the roadmap of your use cases. That can leverage the previous work that you've done. And this is, so the big advice is, think about compounding value. What you really want to do is say, I'm focusing on one use case, and I know that there's a next use case on my roadmap that is going to leverage 80%, right? 80-20 Pareto Principle, 80% of the work that I've done for that first use case. And that's how you start showing the value of semantics and you start scaling this out. And then the other thing is, think about semantics. As the context for AI and the interoperability of systems. It is not just for the context for AI because if that's all you're going to be doing, you are going to take shortcuts. You're going have a bunch of POCs and then your POCs are going to show low accuracy or people have different POCs which are supposedly on the same day but getting different answers. And that's how people are going to lose trust. While the folks who are actually investing for inoperability of systems, where the AI agent is just another system that needs inoperable data, that's where we're going to see the wins there.

Teresa Tung [00:39:05] Okay, so we're banging the audience. Invest in semantics. Do it right. If you do it right now, you're going to be able to unlock the questions you've yet to articulate rather than just the questions that you have right in front of you and then you're going to build upon that compounding value that is your moat, right? So, thank you, Ora and Juan.

For those who have not read Ora's main paper, Semantic Web, I recommend you search for it on Google or Bing, and then take a read because it talks about this, you know, I think you know like a father that a brother and sister are trying to coordinate his doctor's appointment, and it has the agents for the brother and the sister and the dad all working to schedule this appointment knowing the different schedules, the insurance provider, the in-network, out-of-networks services, and talks about the world of today. And that's in 2001. So, it's just you're like a prognosticator.

Ora Lassila [00:40:08] I am very excited about the prospect that the Age of Agents has finally arrived.

Teresa Tung [00:40:12] That you were right. Okay, well thank you guys and thank you all for listening. Take care.