

There is No AI without IA

In Conversation with Carol Smith

Sarah A. Rice and Andrea Resmini

rice@seneb.com

andrea.resmini@hh.se

Carol Smith is a Senior Research Scientist in Human-Machine Interaction at Carnegie Mellon University's Software Engineering Institute where she conducts research for AI and emerging technologies and teaches for CMU's Human-Computer Interaction Institute (HCII). Smith has a nearly twenty-year career as a user experience researcher and is an active community organizer. In 2018, she presented her work on the intersection of information architecture and artificial intelligence¹ at the 19th ASIS&T Information Architecture Summit in Chicago, USA.



¹ Smith, C. (2018) IA in the Age of AI. 19th ASIS&T Information Architecture, Chicago, USA. <https://www.slideshare.net/carologic/ia-in-the-age-of-ai-embracing-abstraction-and-change-at-ia-summit-2018>.

Q: How do you practice information architecture in your work at Carnegie Mellon University's Software Engineering Institute?

Sensemaking is an important part of information architecture, but there's an additional step we have to consider: now that we understand this content, this topic, or this body of knowledge, what are we going to do with it? Ideally, the answer to that question is "make it into something that really is helpful to other people". Applying it to a specific problem. Those particular individuals. Making sure that it not only makes sense but is also pragmatically useful, something that can be implemented.

This means organizing information and taking disparate data or what may be already semi organized and doing something with it. This part has an interaction side to it: how do we turn this information we have into something that supports someone's experience? How are we going to take this and really apply it to that particular problem?

In this sense, and in my specific professional environment, information architecture is different from design. Design here is intended to be leaning more towards the visual rendering of the information we work with. Design is the part that brings it to life. It's that necessary activity of really making something work for people and be beautiful or be functional or be whatever it needs to be at the end. I think of information architecture as that part of the work that really makes information first make sense, then be usable and then functional, all of which are necessary preliminary steps before you can get to design.

Q: How do you decide what is going to be the next problem you and the emerging technologies group work on? And what is the added value your participation brings to these projects?

Sometimes we get specific requests from customers such as the xView 2 Challenge². The Challenge was to explore how to assess disaster damage by using computer vision algorithms to analyze satellite imagery. In some cases, we can propose work that we're interested in. For example, this past year I worked on a framework to guide AI development and created a Checklist and Agreement³ to help teams implement ethics in their work. A project I'm working on is helping a customer figure out a better way to work with robotics and ways that they could be

² xView 2 Challenge. Software Engineering Institute. Carnegie Mellon University. https://www.sei.cmu.edu/research-capabilities/all-work/display.cfm?customel_datapageid_4050=295280.

³ Designing Ethical AI Experiences: Checklist and Agreement. Software Engineering Institute. Carnegie Mellon University. <https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=636620>.

explored more fully. We are figuring out the human teaming aspect of working with robots and figuring out how we can communicate with them. How can we make sure that we understand what a robot is doing, and when a machine or a robot needs help? How do we address that situation?

With another customer, and with a larger team, we are figuring out how to work with multiple robots, multiple people and more complex systems. We are figuring out a lot of the sticky environment-related problems: for example, how do I know that that particular robot sees the other robot, and who's dealing with that robot? That's more about the sense of place and about making sense, and how information flows relate to their physicality. The relationship with human beings in that space is another element we have to consider: lots and lots of information can be gathered from a robot, but what's really important if we want to help people to understand what's going on or what the robot will do next? Conversely, how do we get machines to focus on whatever the most important thing is at that particular moment?

When working on a particular project I'll usually start from the big picture so that I can figure out what aspects need to be worked on from that specific, pre-design, structural point of view of information architecture I mentioned earlier. Then I'll take a particular topic area or work with an existing interface and try to improve it. I'll take it step by step and try to understand what I have to work with, what information I think is relevant and how I could make this even more helpful or how can I implement some new functionality that we've been asked for when looking into this situation as part of the project objective.

If the project is more of an open space, one where we're not really even sure what we're going to end up with, more time is spent early on researching the problem space, understanding how people are thinking about it, and building up from there what the information architecture might be and what the space needs to solve the problems I'm seeing.

Another side of my job is to figure out how to explain AI engineering to people who are either new to AI in general or thinking about starting projects with AI. I work on answering questions like "what is the practice?" It's still early enough in artificial intelligence as a practice that people are really grasping at straws and have very different approaches, despite some agreed upon best practices. Trying to help think through and organize these processes is a big challenge, mentally. How do you even talk about it when there's just so many different types of technologies that encompass the area?

Q: Why is it important to focus on information, and the structure of information that information architecture provides, when doing artificial intelligence work?

Something is missing in the teaching of artificial intelligence development right now. People don't understand how important the organization of information itself is when we have to teach a new system how to understand that information. That's unfortunate. As a result, these systems are often not created in a way that is smart and that makes sense: the unspoken assumption is that it will all work out somehow. That the language will work itself out: 'it's a computer. It'll get figured out.' That's how people think, but that's not how it works.

There's a big knowledge gap. Unfortunately, it will be a while yet before people really realize that to make a smart system we have to do the hard work necessary to understand "smart" first. Right now, that's not happening in a systematic, intentional, routine way. Not that I can see.

Q: There is so much complexity around human expectations, communication, and the context in which we operate. What is important to make a system truly "smart"?

While some of these systems, the more specialized ones, can claim a measure of success, in a lot of cases there's a lot of work that needs to be done before we can even claim they work according to expectations. This is a really important point when we want to broaden them to more complex fields, such as law enforcement, or healthcare.

If we consider healthcare, for example, an information architecture approach has an immediate and very visible positive impact. With such disparate systems in healthcare, it is common to find that people use the same words to mean different things, or different words to mean the same things. Human communication is in itself a very complex system, and trying to break it down for software to understand, particularly words that mean the same thing, or words that mean different things but read the same, it's just really hard. The human side of things doesn't make it a simpler task. Education, or practical knowledge derived from previous encounters, has some weight: in describing symptoms or some past ailment, someone may not know the "right" word, the one used by a medical professional, for example. They may say 'I have a bruise', but medically that's a contusion. How do you help the individual figure out what information they need to provide or look up? And how do you know the difference? When they say 'bruise', do they actually mean a cut maybe, which is a very different thing?

These are relatively simple concepts: the overall complexity only keeps increasing if we consider, say, diagnosis. Imagine you have or are treating a particular illness: will the treatment be very aggressive or very conservative? Think about the differences there. One person may choose to do something very aggressive to treat cancer, for example, while someone else may choose to wait and see. How do you train an AI system to understand the difference between the patient's point of view and the physicians' points of view? Contextual sense-making, that information

architecture perspective, cannot be avoided. There's a lot of nuances that need to be properly understood for us to be able to teach them to a system so that in turn it can be able to present the right information or the right suggestions at the right time in the right way.

Contextualization and context-switching is what people are very good at. We find it fairly easy to switch between topics and we can actually have two different conversations going on at the same time and still keep track of them. You can be on the phone with somebody and watching a webinar and not necessarily lose contexts of either. At least for brief periods of time, anyway. We can recognize separate types of information (e.g. dinner options vs. results of a card sort), know what the differences are, and when these contexts are appropriate. We know when and why to switch between them. We understand the context: AI systems can't yet do that. And here you have one of those systemic issues we face constantly in the field. This is both an information architecture problem, centering on meaning, context, and the relationships between concepts and the environment; and a technological problem, because computing power is a finite quantity and scalability a partial answer.

Q: One thing that information architecture has focused on in the past fifteen years is bottom-up categorization, recommendation systems, and filtering. Could filtering be part of the solution for these systems?

I believe that helping artificially intelligent systems filter the content is definitely an interesting approach, especially since discriminating, finding patterns in data, is something that machines can do really well when they're properly trained. Once they have an understanding of deep-sea biology, for example, we can give them a thousand papers and ask them to single out the ones that discuss a specific topic of interest that is appropriate, such as articles on deep-sea worms. Figuring out what matches means also figuring out what doesn't match. So outliers, such as a paper on archeology, could be recognized as not fitting with the context. However, the system will not know what to do about the situation without being trained in that regard. AI systems only know what they are taught and, if left to determine their own meaning, will rarely do what a human would do because they do not have general intelligence.

Q: You touched on this briefly when you discussed the human angle of dealing with an AI system managing a healthcare service and the complex problems posed by language and context, but what's your position when it comes to the social and cultural implications of moving towards real artificial intelligence? In 2018, CNN Business ran Samantha Murphy Kelly's story of how "(t)he first four words (her) toddler understood were 'mom,' 'dad,' 'cat' and 'Alexa'". Kelly mentioned in her

⁴ Murphy Kelly, S. (2018). Growing up with Alexa: A child's relationship with Amazon's voice assistant. CNN Business.
<https://edition.cnn.com/2018/10/16/tech/alexa-child-development/index.html>.

piece how there's both a dark side to this, after all what Amazon wants is to sell you goods and services, as much as a bright side: while the research is not really there, preliminary studies seem to indicate that children exposed to voice assistants are more autonomous and develop more structured and clear speech patterns earlier. Should we be worried then or should we not?

I see this as a big problem and one we've already seen in many different systems that have been developed. For whatever reason, teams weren't speculative enough about the potential implications and unintended, or unwanted, consequences, or the danger that they're exposing people to. Data collection is an important aspect of this problem that we often discuss: how we are collecting it, what data we are collecting, what data we are not collecting. But there's a prior step, one we don't take often -thinking through how ethics should be discussed and implemented when any problem with technology is being approached.

We should take time to think through the worst-case scenarios: for systems that accumulate data and self-improve, those are often scenarios where humans misbehave. Assuming not everybody is going to be nice should be the default. When we're making a system that we're going to release to the public, or even within an organization, we really need to think about the potential misuse and abuse, and how we're going to prevent them, ideally, or at least mitigate those issues when they arise. Thinking it through ahead of time and not after the fact is key. It's not enough to claim "Oh, we didn't know". Figuring out those types of situations ahead of time and doing a little bit of extra homework can go a long way to protect people.

Q. How can information architecture help artificial intelligence research and implementation become more focused on the ethical side of the process? Is there an ethical benefit to combining information architecture with artificial intelligence?

There are signs these are becoming more important questions. We're seeing a lot more people and organizations develop sets of technical ethics which identify values to be upheld. That's very positive. We're also starting to discuss the original sources of harm, why we have racism, why sexism, and what we need to do to address these problems. These have organizational consequences in the way we recruit and reach out, in the way we assemble diverse teams.

In a podcast⁵ aired earlier this year, Ayana Howard, a professor at Georgia Tech and the Chair of the School of Interactive Computing, mentioned how it is possible to gamify ethical challenges by providing benefits to team members who find ethical bugs, similar to what we already do with software.

⁵ Fridman, L. (2020). Ayanna Howard: Human-Robot Interaction and Ethics of Safety-Critical Systems. Lex Fridman Podcast on Artificial Intelligence. <https://lexfridman.com/ayanna-howard/>.

And then I think we just need to have those difficult conversations, those where we ask each of us to take a hard look at our own biases. Being aware of bias is really the first step. Acknowledging that anyone can be racist, for example, and how we can work on being anti-racist, address the problem and move forward without bringing that stigma into the systems that we're building. Rather we are protecting people and creating a more equitable and fair landscape.

In a way, this is still about the human angle we mentioned at the beginning of this conversation. Humans bring in complexity in many different ways and mathematics and AI alone cannot solve all of the problems we encounter. There isn't an easy, algorithmic solution for a lot of this: it's messy problems with messy people.

artificial intelligence, 3, 4
complexity, 4, 6
data collection, 5
ethics, 6
information architecture
and design, 2
information organization, 3

placemaking, 2
sensemaking, 1, 2
smart systems, 3
structure, 3