AI and human-centered design

Host: Hanish Patel, *User Friendly* host and digital transformation leader, Deloitte Consulting LLP

Guests: Dr. Mahesh Saptharishi, executive vice president and chief technology officer, Motorola Solutions

Nitin Mittal, global AI business leader, Deloitte Consulting LLP

Hanish Patel:
Recent advances in artificial intelligence, more commonly referred to simply as AI, bring machines tantalizingly close to demonstrating intelligent behavior equivalent to humans. If interacting with a machine becomes indistinguishable from interacting with a human, what does human-centered design mean in this case?

If machines can be taught to interact, could one argue are designers in the future a combination of today's educators and psychologists?

In this episode of *User Friendly*, I'm joined today by Dr. Mahesh Saptharishi, executive vice president and chief technology officer at Motorola Solutions, and Nitin Mittal, Deloitte's global AI business leader.

Today, we'll discuss the intersections that these questions address, implications of the advent of generative AI, and why it's critical to rethink design in the age of AI. Mahesh, Nitin, welcome to the show.

Mahesh Saptharishi:
Thanks, Hanish.

Nitin Mittal:
 Thank you, and happy to be here.

Hanish Patel:
Alright, so to start us off, whether you're reading, watching, or listening to any form of news, you can't avoid the topic of generative AI; it's dominating headlines. Innovations in machine learning, and the cloud tech stack, coupled with the viral popularity of publicly released applications such as ChatGPT have propelled generative AI into the zeitgeist.

So, given that, Nitin and Mahesh, I'd like to open up with a definition of generative AI for our listeners from the two of you.

Nitin Mittal:
I can get us started. Generative AI is certainly one of those seminal moments that we're experiencing today. And frankly, if we think about what is this field? It is all about the ability of AI models to generate human cognitive-like capabilities.

In fact, two centuries back, the invention of the steam engine, in a sense, essentially automated and mechanized human energy production. Generative AI is, today, automating and mechanizing human cognitive production. It's the same analogy that we had two centuries back.

And when we talk of human cognitive production, it is the ability to write, it is the ability to make a speech, it is the ability to use logic, it is the ability to be creative, and it is the ability to use reasoning. Those are all innate human traits.

With generative AI, models are now being built which essentially are able to autonomously write, autonomously generate speech, autonomously be creative, autonomously use logic, and—at some point of time—autonomously use reasoning. That is the realm of generative AI: the ability to mechanize human cognitive production.

Mahesh Saptharishi:
So, Nitin, maybe just to add to what you said—and I think that was very eloquently put.

One of the things that really sort of excites me about generative AI is what is at the heart of generative AI. And that's really the capacity to take lots of data, lots of information, and condense it into knowledge, into a set of skills. And that knowledge and the set of skills are really allowing you to do things that typically has been attributed to what humans do.

And also, I think the other thing that really excites me about generative AI is the notion that it actually represents that collective experience.

And I think probably the last comment I'd make is the notion of databases. We take that for granted, and database technology has advanced in so many different ways over the

past many, many years. There's sophisticated and very interesting ways to access data, very large databases, perhaps, analytics and AI algorithms that can tap into that database and give you insights, etc.

The internet itself you could think of really as a giant database. We're moving into this realm where we're actually thinking about knowledge bases. And what I mean by that is it's not just that information and data stored as is, but it is knowledge stored in a way where it is actually very directly applicable to a task and helps you accelerate the completion of a task. So, that notion of a knowledge base is that evolution that we, I think, see enabled by generative AI.

Hanish Patel:
So, if I think about what the both of you have described there as a part of that definition for generative AI, it's probably an understatement here to say that the potential is immense, absolutely huge. But on the same token, when I think about other disruptive technologies, there's also got to be some element of risks and limitations that also come with that. Could you briefly touch upon some of those as well?

Mahesh Saptharishi:
I talked about this notion of generative AI, basically condensing this data and information that's representative of a community's experience. It is a collective experience that is reflected onto this AI, and that term *reflection* is key because the AI that learns from that community could intrinsically be biased as a reflection of that community, could not be as accurate, because it could be the community is not composed of a set of experts.

So, the AI effectively is going to inherit some of the elements of what is not good or perhaps not optimal in the community that it learned from. And today, the risk there really is that large language models have such a huge data set that sometimes it's really difficult to say whether or not you are learning from high-quality information.

And I think we're just starting to learn where it's applicable, what kind of test measures are
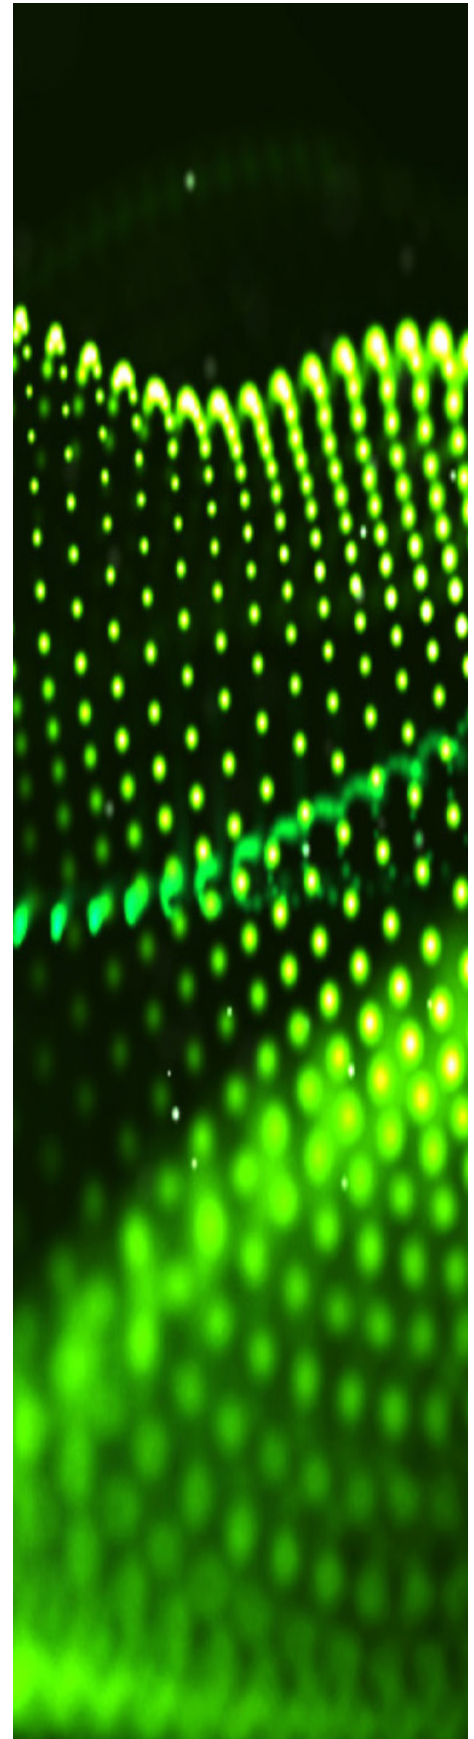
necessary, etc. Testability becomes a big deal in this context, when it comes to generative AI in particular—large language models and multimodal language models as well—that I think there's a level of risk that we need to be able to deal with there.

The second is, again, related to this notion of a community of users actually teaching this AI to have an understanding of a concept, to be able to acquire a set of skills. The other part of it is that sometimes members of this community may not necessarily know that they're providing information that this AI is learning from.

The second block of risk is really around trust. We tend to anthropomorphize AI; we've done it in science fiction, and now we're naturally doing it here. And the fact that generative AI is really allowing AI to kind of talk almost as if it were human, that it was a being of some sort.

There's a level of trust that is implied in something that is a machine, that we as users of this solution have to recalibrate ourselves.

And if the accuracy is not there, or if the AI is not communicating effectively a sense of confidence in the answer it's giving you or the suggestions that it's giving you, then at that moment of need, you may trust something that really should not be trusted, unless you carefully engineer the trust mechanisms involved here.

That is the other important element where it could be a risk, where today, we tend to ascribe trust to something that perhaps does not necessarily deserve that trust. It may lead you down a path where you're believing that it is actually very trustworthy, but based upon its tone, based upon its language, you need to be able to understand how exactly to use that information appropriately in any given situation.

The last block is really misuse, and I would say because of the ease with which you can build solutions here, sometimes you may choose to use AI for applications where an AI perhaps is not needed, and maybe there's something else other than AI that could

actually do a much better job of it. Then, of course, there's this notion of malicious use. So, we are all familiar with things like deepfakes and I think misinformation of a variety of kinds. And I think there's a deep risk there where malicious use of AI, not just by folks who are well-intentioned, but perhaps by bad actors, where it could actually end up being something that harms the society as a whole. So, at a high level, those are the three big blocks of risks that I see.

Nitin Mittal:
That's very, very well said Mahesh. In fact, I will say this, that in a lot of the work that Deloitte does with the clients globally, we absolutely stress on something that Mahesh indicated, which is trustworthy AI. The concept and notion of trustworthiness actually includes a number of dimensions.

It includes privacy, it includes ethical usage of AI, it includes safe and responsible AI. It includes right to privacy, and it also kind of includes the security of the information that is generated inclusive now of the content that is actually generated.

How is that content generated? How can it be explained? Does it infringe on the IP of what may have preceded it? And consequently, how do you make sure that it is something that you have the ability to generate in the first place? All of that is part and parcel of the trustworthiness of these systems.

And as a society at large, and businesses in particular, we have to be thoughtful of the guardrails that we actually put in place for all the reasons that Mahesh actually articulated.

It's true, the guardrails, the governance systems, the regulatory debate that is inevitably going to happen—as well as essentially some of the safety mechanisms that we actually put in place as businesses and society—that we would actually be able to fully harness the promise of generative AI and kind of see it elevating the level of productivity and elevating the disruptive aspects of it as opposed to just going with what is in the public's consciousness right now.

Hanish Patel:
Thank you both for one, the upfront definition, and for talking through some of the risks and limitations of what you see. And I'd like to really anchor in on what you introduced around businesses there, Nitin, because much of the media coverage to date has certainly focused on the consumer use cases.

But I would say that there's a whole wide spread of opportunities when we look at businesses, the enterprise view, so to speak. And when we think about the advent of generative AI, really kind of laser focusing in, what would you say would be the impact on things like software design and engineering as enterprises are looking at those components in relation to generative AI?

Nitin Mittal:
Yeah, we all recognize that at the moment that we are living in right now, generative AI is very much in the public's consciousness.

And frankly, as much as we read about it, we also have to be thoughtful around what are some of the more practical business-oriented applications of generative AI that are actually emerging. Frankly, our hypothesis is that the greatest benefit of generative AI is going to be in the B2B context. The greatest curiosity element is in the B2C context, but the greatest, frankly societal benefit and business benefit is going to be with enterprises in a B2B context.

This is where we are going to essentially see a movement where there are perhaps going to be purpose-built large language models that underpin generative AI that are going to be trained on enterprise data sets in a secure manner within the actual enterprise for generating very specific content, tasks, interactions, experiences, and the type of modalities that I was articulating earlier.

That is where businesses are going to see a logarithmic increase in productivity. That is where businesses are going to see a significant cost savings because the marginal cost of generating the necessary modality drops to almost zero.

And that is where businesses are going to see the disruption that is going to be fostered in their particular industries as it relates to new products and new services that they could be introducing in the marketplace at speed with a lower cost structure and with the ability to be a lot more insightful.

That's why the whole B2B enterprise space is absolutely going to blossom beyond just the new cycle, and frankly, it's going to define what it means to be a thriving business that has the ability to compete in today's digital economy as it relates to the industries that they are in.

Mahesh Saptharishi:
Hanish, I think you asked about what does this also do to just software engineering and design as we think about things in the future, and Nitin mentioned productivity. Very broadly speaking, I actually "bucketize" the impact of generative AI from a company standpoint, even if I speak from a perspective of Motorola Solutions. There are three real buckets that we think about generative AI really making a difference in.

And that's productivity being bucket one, the second being user experience—and that's really the interaction between users and the product. And the third being customer experience—and this is where our customers interact with us as a company.

In all those three areas is where I think generative AI stands to make some pretty big functional differences in how we do our jobs day to day.

The productivity increase is actually quite a significant one where we are still not ready to take ChatGPT as a care programming peer of ours when we do software development, but I think that day will come where we are able to do that.

And this notion that from an agile development standpoint, I can start with perhaps a user story and that user story— rather than having to go through a lengthy process of software design, a bunch of developers working through multiple sprint cycles, trying to write code, implement it, and test it—maybe one where that user story can

go directly into code where an application can be built almost automatically.

And then the burden of agile development is really making sure you are refining those user stories appropriately and rapidly testing it in the most appropriate way. That may be what the future of software development looks like. And with that increase in productivity, you are talking about iteration being something that is perhaps an order of magnitude if not more than what it is today.

So, if I just jump a little bit then to the user experience aspect of it, one of the things that I think, again, from a design standpoint, is going to be quite interesting where generative AI, and AI in general, is almost used as a design material.

We think of user experience, we think about UIs [user interfaces] perhaps, today, and we think about what kind of aesthetics lie behind it, how many mouse clicks do you need to do something, etc. And it also ends up being something that's largely a static quantity, for the most part.

But start thinking about it a bit more broadly from a generative AI standpoint, the type of interfaces you can engineer, whether that is a chat-based interface that is just a lot more powerful than a traditional chatbot, a voice-based interface, something that can generate visuals for you that you can consume in new and novel ways.

All of that may actually quite profoundly accelerate the ability to do your work faster, enhance that user experience in some very powerful ways. But then also think about it a bit more: Why does the user experience have to be static when you have an AI that potentially can also learn from its interactions with you?

So, now you're starting to think about software engineering and software design in terms of something where that piece of software that you have put forward with the interfaces that connect it with the user, being something that as that user skill goes up, the user interface also more appropriately adapts the needs of that user. It is one where it's no longer static.

So, perhaps in one case, you're talking about something that is more of a basic user interface, but as that user gets more proficient, more advanced capabilities come to the forefront. As a very simple example of this sort of basic user versus advanced user example that I can give you, is that when you think about a 911 call taker, somebody who answers the phone when you call 911, well that person oftentimes is using a fairly complex piece of software to do multiple different things.

You can give them a lot of capabilities in terms of a fancy UI, buttons to click, nice interfaces to do multiple different things. And that's what helps them perhaps when they're getting started, do their job quickly, learn how to do their job more quickly.

But then as they get more and more skilled, and as they sort of get more experience under their belt, many of them actually opt for some something that's more of a command line interface where they know they have a notion of a structured language, they can execute a lot of things very quickly in the command line, and that actually ends up being the faster approach for them.

Now, if there's a user interface that really evolves with their needs to give them the right sort of interface in that time of need, and also, be able to calibrate itself to the level of expertise of the user, personalize itself to that user, that is incredibly powerful. But that's also a completely different way of how we think about software.

It is completely a different way of how we think about user interface design and UX [user experience] design and human-centered design. And I think that is going to very dramatically change human-centered design thinking as a whole in this industry, and I think that's quite profound.

So, those, you know roughly, those three areas is I think where a lot of our attention's going to get focused on in the coming years.

Hanish Patel:
You mentioned it a couple of times there in terms of interfaces. So, if we are thinking

about from a human or a machine interface, what are the big design elements that should be considered? I'd love to sort of dig in a bit deeper there with you, Mahesh and Nitin, in terms of what your perspectives are there.

Mahesh Saptharishi:
Sure, Hanish, this is a deep topic. And if I were to just touch on a few areas, it really starts with a deep understanding of the job to be done, how our users are going to use our solutions, what is the job that they're trying to do, and how does this, as a tool, help them do their job effectively.

In computer or human-centric design, there's this notion of a journey map, being able to journey map each user in terms of what they are trying to do on a daily basis, and what are the things that typically add some level of friction to what they're doing, where some set of tools or some augmentation comes in handy.

It is very important for us to understand that user journey, it is very important for us to understand the job to be done very carefully. That's step number one.

Step number two is really trying to understand, OK, where does AI potentially apply? What can it do to really accelerate a particular task?

So, knowing where the AI can apply, where it can be productive, where it can actually have a positive impact in the outcome, that's the second part of this journey.

The third part of this journey is when you think about things like generative AI, as an example, you have this community that is potentially giving you the information that allows for this model to be this large language model. It gives the necessary information for that module to be either trained or fine-tuned either way, and there's a certain set of assumptions, biases, etc. or a set of demographics that that community really represents.

And you have to match that up against the demographics and the needs of the user, and also the demographics of who that user affects.

Understanding those three things and making sure that those three are actually matched and carefully calibrated is a very necessary condition for us to make sure that we're not unintentionally creating an issue with bias, with accuracy, with trust, and we are, in fact, very good practitioners of responsible AI in this process.

Then it's really an element of trust engineering. In a moment of need, you are going to be inclined to trust this thing that sounds very human, sounds very confident, even if potentially it is not giving you the right sorts of suggestions. So, being able to do that trust engineering, with a very clear idea of what the mental state of the user at any given moment in that workflow is, becomes a very important factor. This is essentially human-centered design as we see it.

The user experience can actually adapt to the needs of a particular user in a particular circumstance.

And that really requires a rethinking of what we have traditionally thought of as the design task.

So, that notion of an evolutionary UX, where the design itself evolves automatically with the changing needs of the user, or the changing environment in which the user is, that becomes a very powerful design element.

Hanish Patel:
Nitin, I'd love to get your perspective on that as well.

Nitin Mittal:
There is actually something I'd like to explore, given all the aspects that Mahesh laid out over here. We're actually seeing a fascinating trend that is very likely going to emerge at this particular moment. And this trend is where the power of intelligence is transcending beyond just the human species.

We're now going to see the advent of so-called intelligent machines that have the same cognitive capabilities that we associate with humans. What does that actually mean? That essentially means that we could be having coworkers that are not necessarily

bipedal and are neither carbon-based the same way that the homo sapiens race is.

And it essentially means that our coworker, our work environment, and the type of work that is actually conducted is going to be more and more with these intelligent machines that are non-bipedal and non-carbon-based. Which goes to the very essence of the question that frankly was asked, Hanish by yourself, which is, what does this do to the human machine interface and human machine interactivity?

That is some of the beauty of what Mahesh outlined as it relates to the human-centered design thinking that permeates beyond just how we as humans work, but more and more is in the realm of how we as humans essentially interact with these intelligent machines that happen to be our coworkers of tomorrow.

If we start thinking of human and machine interactivity and human and machine interface in that, let's say, framing and with that notion in mind, it leads to frankly a completely different paradigm as it relates to how are we going to be interacting? How are we going to be conversing? How are we going to be working? And ultimately, how are IT systems of tomorrow going to be kept?

Hanish Patel:
So, I want to stick on this train of tomorrow—and I'm going to ask the both of you to somewhat get your crystal ball out now. If we look at the next maybe five, maybe even 10, years from now, how do you see that evolution of software development in that time frame?

Nitin Mittal:
I think five years is a window that we can reasonably assume how this is going to progress. Ten years certainly kind of stretches the limit of our imagination given just the fast-paced revolutionary nature of this technology and the disruption it introduces.

Specific to software development, I think it is reasonably safe to say that we are going to be in the full-fledged era of autonomous engineering. Wherein today we interact with

computers using computer language and programs that define the software field.

And the output of it is software programs, and applications and systems that frankly many of the organizations have implemented. Now, we are going to be, interacting with the computers using human language. The medium of interaction is not necessarily going to be for example, Java code or Python as a software programming language that we are going to be using.

The medium of interaction with a computer is going to be speech or the ability to write in any of the languages that we have grown up with and happen to be our mother tongue, so the system of interacting changes. The output is going to be such where computers are going to be writing and generating computer code, and then that feeds into other computer systems that are able to execute tasks. That is the very essence of autonomous engineering wherein everyone becomes a programmer, the task and the activity of programming gets democratized.

You may not necessarily need specialized training on the art and science of programming, but rather, you would need the necessary ability to prompt your models. The ability to ask the right questions, the ability to think through the user stories, and feed those user stories in natural language that happens to be your mother tongue and the computer or the AI model that you're feeding to it, will generate the applications, the systems, the workflows, and the computer programs that essentially businesses use today.

Not only do that, they can feed it to another set of computers, who then take the tasks forward and go along that chain so that you end up with autonomous systems, autonomous workflows, autonomous tasks, and autonomous processes.

Mahesh Saptharishi:
I go back to what I said originally about this notion of us going from databases as being sort of this primary concept of storing information, to these knowledge bases, where really it's a combination of the

condensation of that data into knowledge and also a set of skills that are learned in the process as well.

And when you think about it from that standpoint, I think there's an element of this where the future is not going to be necessarily about us thinking about what a software developer lines a code.

You're not thinking about that as much as you're thinking about now as this system that is going to accept this user story and convert it into an application that potentially can be productized very quickly, making sure that the engine that is taking that user story and converting it to that application is able to do it in a way where it represents the end outcome that we desire.

You're talking about a very different modality here. Prompt engineering, I think that's a problem of the near future.

I think prompt engineering is likely going to go away, but the clarity of you being able to communicate that user story in an effective way is not going to go away. You are going to be effectively the teacher. You need to be an educator of the system to be able to say, "This is what we want the users to accomplish and by the way, you the AI, need to be mindful of these other constraints as appropriate to make sure that we're serving that user well."

So, that notion of being a psychologist; the notion of being an educator, a teacher; the notion of being an ethnographer, being someone who really understands the community that we serve; that is going to change software development in quite a profound way.

I think today, when we think about human-centered design, we do think about psychology, we do think about cognition, we do think about the mental state of our users, but the level to which we think about it today is going to be amplified quite significantly going into the future because that is going to be the focus.

So, the nature of what a software developer is, is going to very much change as we go

into the future. And that change is going to be one where it is not about programming languages, Nitin is absolutely right. But it is going to be one of understanding the human condition better, the state of mind better. And I think that's kind of an exciting place to be, to be honest.

Hanish Patel:
And I'm just processing and imagining a world that the both of you have described. So, I want to, with that, maybe close with this question—as that as a backdrop of what you've both highlighted, what the future could look like. So, if we think about what you've both said about where generative AI could be in the coming years, a direct question, so to speak, in the sense of do you think that's actually going to deepen our knowledge and expertise or will the notion of a human subject-matter expert or adviser be antiquated?

Nitin Mittal:
I personally think that it's going to deepen because the best way of thinking about generative AI as an aid to society and as an aid to business is augmentation. The ability for generative AI to augment human endeavor, human insights, human experience, and human tasks—that is where the greatest potential is.

Yes, there's always going to be the element of anxiety and a degree of fear and apprehension, which is inevitable with every technological revolution and has always been the case, whether it was the PC revolution, the internet revolution, the mobile computing revolution, the cloud revolution, etc.

But as we take a look at the history of technology and what has transpired, over a period of time and as we get more comfortable with the technological revolutions that take place, we learn to embrace it, we learn to co-op with it, and the more we think of generative AI as a means to augment what we do and elevate not only the level of productivity that could manifest itself, but also the level of creativity, it will actually unleash a completely, let's say, different part of the economy that we had never imagined, the same way that the steam engine was the

harbinger to the Industrial Revolution that changed civilization and society for good.

Mahesh Saptharishi:
If I think about it in very simple terms, I think of subject-matter experts really doing two things. One, helping me understand a topic that I don't have any expertise in but I need to understand just enough so that I can do my job. And then the second thing that I think often when I think of subject-matter experts is really discovery of new information.

Going into new territory, being able to create new knowledge, really, and I think—Nitin alluded to this really nicely—I think the shift is going to be where subject-matter experts are not necessarily the explainers, but the subject-matter experts really become the discoverers, the creators of new knowledge.

And that new knowledge, that creation process, is very much going to be aided by the capabilities that generative AI also provides.

I think what generative AI is going to allow us to do is for experts to actually broaden what it means by *expertise*. I think there's going to be a notion of what we today call *cross-disciplinary expertise.* I think that'll become the norm of tomorrow.

Hanish Patel:
Brilliant, I think a lot of listeners probably had a deep sigh of relief there, based on what both of you described. And I mean, there's a lot to unpack in terms of what we've discussed here, but I'll try and close with the fact that, yes, we're at the early stages.

It's hard to imagine to a certain degree where this would be—as both of you described the pace of change—where things will be in five years, never mind where we look 10 years from now. But what is certainly clear is the impact it's going to have on the future of work; where we need to consider, when it comes to trust, the importance of that interaction between the human and the machine.

And it's abundantly clear with all of that, that there's just huge potential for this disruptive technology to permeate in all aspects of our lives.

And with that, what I do want to do is thank the both of you, Nitin and Mahesh, for joining me today and helping us further understand the path forward and the future of what AI might look like. And with that, until next time, happy listening.

Explore more episodes of
User Friendly at:

userfriendly.deloitte.com