

**iLLUminate Blog Transcript: Ramayya Krishnan on AI and the Future of Work**

Recorded February 12, 2025. Listen to it [here](#).

- ANNOUNCER: 00:01 [music] This podcast is brought to you by iLLUminate, the Lehigh Business Blog. To learn more, please visit us at [business.lehigh.edu/news](https://business.lehigh.edu/news).
- JACK CROFT: 00:13 Welcome. I'm Jack Croft, host of the iLLUminate podcast for Lehigh University's College of Business. Today is February 12th, 2025, and we're talking with Ramayya Krishnan, Dean of the Heinz College of Information Systems and Public Policy at Carnegie Mellon University. Dr. Krishnan is on campus today to deliver the Lehigh College of Business Year of Learning Lecture on this year's topic, Generative AI and Business: Opportunities and Challenges. He is currently Faculty Director of the Block Center for Technology and Society, a university-wide center at Carnegie Mellon examining the societal consequences of technological change. And he is an expert in data and decision analytics and digital transformation. Ramayya, thanks for joining us on the iLLUminate podcast today.
- RAMAYYA KRISHNAN: 01:06 Thank you, Jack. Looking forward to it.
- CROFT: 01:08 Yeah, I'd like to start by looking at AI within the context of academic settings. Obviously, you're a dean at Carnegie Mellon. You're going to be talking to Lehigh students and faculty. And for students in particular, I'm wondering, how have you seen AI actually enhance education for students to this point?
- KRISHNAN: 01:31 Well, that's a great question. But let me first begin by saying thank you for this opportunity to chat with you. I should also note, given my public policy roles, that I'm here as a professor at Carnegie Mellon University and I'm not representing the National AI Advisory Committee or any other public policy role that I'm in. So this question that you pose actually is a really important one. There's been rapid development of AI over the last couple of years in particular since the advent of ChatGPT towards the end of 2022. And what we are seeing is a combination of things. I think one is students trying to use these technologies in ways that they think might benefit them. So, for instance, using ChatGPT for research purposes. And I mean ChatGPT broadly to refer to either Claude or Google Gemini or any of these class of AI assistants that you see out there. Using them either for research purposes or using them to write some of their reports. I think it's really essential in these cases that students and we, as instructors or teachers, ensure that students really understand and learn a core set of learning outcomes. One, I think to be good critical thinkers, to learn how to formulate problems, scope problems out, learn how to construct arguments, learn how to both verbally and to, in written form, be able to construct and convey arguments, learn to work well with others and teams.
- KRISHNAN: 03:21 These are all really important skills to go along with data and analytic and other skills that are needed in today's world. And this requires, I think, both a liberal arts education as well as a STEM type of education, some of which is certainly going to benefit from these advances in AI, but we shouldn't lose sight of what we think the core learning outcomes should be independent of the technology. The technology can

be a big productivity enhancer. It can be a creativity enhancer. But the core learning outcomes will have to be delivered.

CROFT: 03:58

And what are some of the potential pitfalls that students may be falling into already in the way that they use AI?

KRISHNAN: 04:11

So, for instance, let me give you some concrete examples. So suppose you were my instructor and you gave me a writing assignment and I used one of these AI assistants to produce a report and I turn it in. Have I checked the box on submitting a solution that is expected of me? Perhaps yes. But did I achieve the learning outcome that you as an instructor wanted me to attain? I think not. Because I think you want for the student to engage in the brainstorming, the ideation, the rough drafting, the editing, and learning the craft of actually creating the report. So it would be ideal if they would actually go through that process without the AI, meaning the brainstorming, the rough drafting, the editing, and then begin to use the AI to see how and where it could actually help them so that they understand the process of actually creating a well-reasoned report. So they've learned and absorbed the core ideas of how to create an argument and write a report, and then see the AI as a productivity enhancer. That's what I think would be a good outcome to have. And a pitfall would be where I short circuit the learning experience and then use this just as a means of reducing the time and effort to produce a report. And you can see why students may - they're busy, they think that this is something that they've got a lot of competing priorities, I think. But this is really important for students to understand that there are these learning outcomes to be achieved.

KRISHNAN: 05:57

And so we, as instructors, I think, have to set up the appropriate mechanisms. So I would think there would be opportunities, I think, in the classroom where just with a paper and pencil, you require students to go through these exercises without the technology so that they learn those aspects, and then have them use the technology so that they can get the benefits of both.

CROFT: 06:18

Now, the other side of this is, when they leave Carnegie Mellon or Lehigh, and particularly with business students and going into the business world, what will the expectation be in, we're talking in some cases one year, at most, probably four or five, by employers for what students really should know about AI when they get that first job?

KRISHNAN: 06:47

Right. And I think this is really essential to understand. It doesn't matter whether it's the business world or whether you go to work for government or whether you work in any organization, not for profit, because I know Lehigh also has arts organizations that you send students out into. I think all these organizations are wanting students to be good critical thinkers and problem solvers, but also be able to use these technologies coming when they graduate. That is that they're savvy, that they have some familiarity with these technologies. Imagine a world when Excel came out. How people now pretty much expect, if you've gone to college, you sort of have some familiarity with Excel. I think that's where we are going to be in a couple of years, that people will expect you to know and have used these tools. But we don't want them to be uninformed users of these tools because these AI tools are tools that can commit errors. So you need to have oversight of these tools so that when you use them, you're using them in an informed manner so that when it makes an error, these are called hallucinations as it's become part of our vocabulary now, you are able to catch that and fix it. And then when it actually helps you get the benefit or the productivity

benefit that comes with it, but use it in an informed manner. And I think increasingly, organizations economy-wide are going to expect people to be familiar with these technologies and to be informed consumers of these technologies for all the benefits they bring, but also be aware of the risks they represent. But also be critical thinkers and good problem solvers.

CROFT: 08:43

Right. I wonder, I've talked to people on both sides of this issue, some particularly faculty and business people who are concerned that there could be a diminishing of critical thinking skills, that if people go down that path of just putting the prompts in and kind of letting the machine do the thinking for them, that we've lost something really valuable. And then there's others on the other side who are more optimistic and think that, if anything, critical thinking skills could be even more important as we integrate AI into the work processes. So where do you fall on that range?

KRISHNAN: 09:31

I think I've always been somebody who has thought about how to sort of be creative as new technology develops, how to sort of deploy it in education in ways that achieve learning outcomes while yielding productivity benefits. So the keyword that you use, how do you integrate it into education, I think the integration piece is critical. Because if you sort of just give people ChatGPT or Google Gemini or any of these AI assistants, and they use it in the manner of, "For this task, I'm just going to put in this prompt, and I get a response, and I'm going to submit that response," then we haven't achieved the attainment of the learning outcomes. So this whole process begins, I think, with having clarity about what the learning outcomes are that we're seeking to achieve. Having that clarity in mind, then I think we could say what kind of instruction and experiences do we need to offer, some of which may be without the technology, like I said, be it in writing, be it in coding, be it in constructing arguments, where you may want students to do things with paper and pencil. And this might require more effort from the instructor to grade these things. But the advantage is that you've got them going through this process of developing ideas, editing ideas, and constructing an argument, or when you're writing code, thinking about these kinds of principles of what it takes to write good software.

KRISHNAN: 11:16

And having done those steps without the technology, then you could use the technology to see how it might speed up some of these activities or be a good completer. By completer, I mean where the technology may give you a first draft, you do an edit, and you then improve it, or like when you write code, it recommends some code to you. You edit it, it improves it. So then the combination of you working with the AI yields something that's better than you alone or the AI alone. And in fact, you're beginning to see this with there's a recent paper that looked at the use of AI for science and AI for scientific discovery. And they found that with AI generating the hypotheses, somebody who's skilled and who's knowledgeable could decide which hypotheses should be pursued and which should not be. That requires critical thinking. But if you did that right, you got an 80% increase in productivity. But if you're unskilled, it actually resulted in you chasing too many poor, bad hypotheses, and you ended up being less productive. So that, I think, is an example where you needed to have independent capacity to think and be able to evaluate these hypotheses, but the AI made you much more productive by expanding the set of hypotheses that you could consider.

CROFT: 12:47

In order for students to learn, obviously, faculty are going to have to know AI and be able to show them what it can do, what it can't do. So what have been some of the

effective ways you've seen that faculty have incorporated AI into their own research and into teaching?

KRISHNAN: 13:08

So I think there's a broad waterfront of ways in which faculty integrate AI into teaching and research. In some cases, research is AI itself. My own work is in AI. So I tend to be much more engaged in it because that happens to be the work I do. That's what I invest a lot of time in it. But you could imagine faculty members whose work is not in AI, but who seek to integrate AI in the manner that we've been talking about. So at Carnegie Mellon, I think the provost office and as dean, both at the school level as well as at the university level, there are investments being made to give faculty opportunities to experiment and try out new ideas about how they may integrate AI into their coursework. We also have begun a process of helping faculty standardize on how their core syllabi should provide guidance to students about how AI will be used and when it can be used, when should it not be used. We have these tiered levels of tier one, tier two, tier three, where these levels are based on the kind of course you want to teach and based on the learning outcomes you're seeking to have students attain in these courses. You may choose a particular tier, and we are trying to support faculty as they migrate from where they're at right now into the appropriate tier. Doesn't mean tier three is the best for a given course or tier one is the best. It's an appropriate mapping of which tier is best for the kind of course that they're teaching.

KRISHNAN: 14:51

But these kinds of efforts require an investment in faculty and giving them the time and the space to acquire knowledge about AI, be able to upskill themselves to figure out how best they can use AI to enhance the educational experience of the student that they're instructing. So if I'm teaching English and history versus if I'm teaching statistics, if I'm teaching AI itself versus I'm teaching software engineering, there's a range of different subjects and I haven't even spanned everything. There's arts and music and other things that go on at the university. So I think there's a whole range of effort, different kinds of effort that faculty teaching these different subjects will have to invest. At the college that I lead, which is the College of Information Systems and Public Policy, they've begun the discussion and there are always going to be early adopters who are inherently interested in figuring out how this technology is going to help in their courses and others that will need more support. So working together with the university and the provost office and the college, we are seeking to invest in our faculty, both in terms of supporting their upskilling, but also in terms of the processes they need for actually delivering a class. And the compute infrastructure is another important point as well, which I did not mention earlier. Because to deliver on AI, you need computational infrastructure as well.

CROFT: 16:24

Now looking at it in the broader context of the business environment as it is today and is rapidly changing, what have you seen that make you optimistic and hopeful that AI really can be a powerful productivity enhancer for companies?

KRISHNAN: 16:50

I mean, just a few examples, right? I mean, I think beyond productivity, it can increase access and opportunity. If you take [Khan Academy](#), which is AI-powered education, the number of kids who can get access to AP-level courses, that not every school has instructors who can deliver AP-level courses, you get access to really high-quality instruction using Khan Academy. And initially, it was all just videos that Sal Khan produced. But now that Khan Academy program has been-- there's been an integration of AI into Khan Academy's content so that you have tutors that are AI tutors that help students learn, be it math, be it history, be it English. That's a great

example of increasing access and opportunity. A second example is in health. So you have physicians engage in consults with patients. So if you are my physician, there's a startup from CMU called [Abridge](#) that actually listens in on the consult. And at the conclusion of the consult, creates a summary, a script of what the physician and the patient discussed. And the physician has to review it to make sure that it correctly summarized the content of the consult and then uses that summary to fill in what the electronic health record system needs and fill in what the claim system needs. And that supposedly returns between six to eight hours of what they call pajama time to a physician.

KRISHNAN: 18:41

This is the time that the physician at the conclusion of their workday spends in actually filling in and providing summaries of the various consults that they've had during the course of the day. Six to eight hours per week is a lot of time for physicians. So these two examples, one in education, one in health. And you can think of numerous other examples of AI yielding productivity benefits. I think the key issue is that as you think about responsible deployment of AI, you have to think carefully about how this is deployed in ways that one is cognizant of potential intended or unintended kinds of consequences that might come about from this deployment. So there's a risk management piece that needs to be thought about as well. So there is productivity, but you need to also consider how to assess risks and manage the potential risks of deploying AI. The simple example is the example of the hallucination example I gave you, making sure that you don't accept everything it says because especially generative AI is what's called a stochastic system, meaning it makes predictions. It can have errors in them. While the number of errors is reducing, it can still commit errors. And some of them could be serious, especially in high-stakes settings. So assessing these systems is still an ongoing effort at figuring out how to actually reliably get assurance on these systems. It's still an area where we are still understanding how to do it in a way that would allow us to gain complete confidence that these AIs that are being deployed are being deployed in ways that are responsible.

KRISHNAN: 20:32

So in truly high-stakes situations, so I think most of the deployments that you're seeing today are in deployments where the benefit cost assessments in terms of productivity gains and the potential risks are such that that's where you're seeing most of the applications of the sort that I was describing. You're not seeing deployments of it in areas say like loan granting where explainability is an important criteria and therefore these types of models that are hard to understand are not being deployed. Or in diagnosis, there are proofs of concept underway, but they're not being used in production yet because we are not fully confident yet about how to assure these systems in high-stake settings.

CROFT: 21:22

Now, one of the main concerns, particularly that workers have with AI is, "Oh my God, the machine's going to take my job." So first of all, how much of a threat is that? And what are the kind of categories of jobs that will probably be most vulnerable to having either whole or part of the work being taken over by AI? And how do we as a society balance that against reskilling, retraining, the other things so that people aren't kind of falling by the wayside?

KRISHNAN: 22:09

So this is a huge question for policymakers, right? And I think it boils down-- I think let me first start by saying it's really hard to predict how this will play out. First question. Let me give you two examples. We've both seen the advent of automated teller

machines [ATMs]. And when they first came out in the '70s, you'll remember that everybody said there are going to be no more tellers and banks. And if you look at the number of tellers and banks today, we have more tellers today than when ATMs came out in the '70s. Are they doing the same thing that tellers in the '70s did? No. They do more than debit and credit. They do marketing and selling of products to customers. So the job, even though the title has remained the same, the set of tasks and skills that they require has evolved. That's potentially one kind of future that might emerge. Another is on the turnpike, since we're both in Pennsylvania, I can talk about the Pennsylvania Turnpike, the toll booth operators, you remember they used to hand out those tickets and when you get out and you enter and when you exit, you would pay. Are those folks in those booths anymore? No, they aren't. That job has been entirely automated. So you have one example of automation where the technology is reliable enough to be able to perform that task in its entirety, whatever that role called for in its entirety, at a skill level and at a reliability level that is acceptable, in which case that role then disappeared.

KRISHNAN: 24:01

But on the other hand, with ATMs, it's actually you have more tellers, like I just told you, right? So I think you're going to see certainly some roles that are going to be like the toll collector roles. In fact, OpenAI had this article published by a colleague of mine at Penn, [Daniel Rock](#), that said 80% of the U.S. workforce could have at least 10% of their tasks affected. Affected. 10% of the tasks affected doesn't mean that 80% of the workforce is going to be displaced, but 80% of the workforce is going to have 10% of their tasks affected, and 19% of the workers may see 50% of their tasks impacted. So the question is, what does that mean? What does affected and impacted mean, right? And does impacted and affected mean augmentation? Does it mean automation? So I think these studies certainly result in people thinking, "Hey, the machine's going to take my job," which is exactly what you started out with. But I think it's hard to predict what exactly may come about. I think you will see certainly some jobs go the way of the tollbooth operator. I think some new jobs may come about that you and I are not even contemplating right now. And then you might have other jobs that get modified in ways that will require new tasks, new skills, like the automated teller machine and tellers. Under that state of the world, my view is that we as a society need to monitor instead of predicting because predictions can go wrong here, we need to monitor so that we can have quick response. That monitoring has to be of both public and private data systems, by which I mean the Bureau of Labor Statistics, which the Department of Labor runs, certainly is one source of data.

KRISHNAN: 26:00

But the LinkedIn, Indeed, and ADP, which does payrolls, these systems, they have a lot of data. Burning Glass, these are private entities that have a lot of data about the labor market. And the private sector data sources are not representative, nor are they complete, but they're very high frequency, meaning that you get very quick feedback about what's going on in the market. The public sector data sources tend to be representative and complete, but they tend to be slow. So I think by blending the two and monitoring that data, I think we'll have a much quicker sense of what's happening in the labor market. And that will allow, I think, policymakers and firms to ask the question, "What kinds of programs are best going to be able to help people acquire the new skills required by the new economy?" And we'll have to combine that with some kind of a safety net that will be required for people who are older, who might not have that much of a runway ahead of them, and therefore upskilling and reskilling may be more challenging. So we need a combination of wage insurance type of public

policy to allow for people to get the wherewithal, the funding to be able to acquire these new skills because it costs money to acquire new skills. So it's a combination of giving people the time and money to acquire new skills, and at the same time, monitoring the labor market to see what's going on. Because we are, I suspect, going to go through changes in the labor market. And we should be proactively monitoring it both as business decision leaders and policymakers to best support workers manage this transition.

CROFT: 27:55

Well, unfortunately, our time is up. I'd like to thank our guest today, Ramayya Krishnan. Thank you for a most informative and enlightening discussion. And I wish you the best with the lecture this afternoon.

KRISHNAN: 28:11

Thank you, Jack. I appreciate the opportunity to speak with you.

CROFT: 28:15

This podcast is brought to you by iLLUminate, the Lehigh Business Blog. To hear more podcasts featuring Lehigh Business thought leaders or to follow us on social media, please visit us at [business.lehigh.edu/news](https://business.lehigh.edu/news). I'm Jack Croft, host of the iLLUminate podcast. Thanks for listening. [music]