SDS PODCAST EPISODE 924:

95% OF ENTERPRISE AI PROJECTS FAIL (PER MIT RESEARCH)

Jon Krohn:	00:00	This is episode number 924 on the MIT report that 95% of enterprise AI projects fail. Welcome back to the SuperDataScience Podcast. I'm your host Jon Krohn, a new report from an MIT lab called NANDA Networked AI Agents in decentralized architecture. NANDA really catchy that MIT lab made waves in recent weeks by declaring that 95% of enterprise AI projects fail to deliver measurable business impact. It's a headline that's stops you in your tracks. So is it true? And if so, why are so many projects falling short? The researchers at MIT call this phenomenon the Gen AI divide. On one side, almost every company is experimenting with tools like Chat GPT or Microsoft copilot. Adoption is high, but true transformation is rare. The report found that while 40% of firms deploy AI pilots, only 5% ever reach production at scale. Most projects stall because they don't integrate smoothly into workflows.
	01:09	And crucially, the models used in these projects don't learn. There's static. They can't retain context, adapt to feedback, or evolve with the business. Meanwhile, on the other hand, employees aren't waiting for official tools to be allocated to them. In over 90% of organizations workers use personal AI accounts to get their jobs done. A so-called Shadow AI economy that often delivers more ROI return on investment than the sanctioned corporate projects do. You might be familiar with this yourself in your own work, and it shows that there is hunger for ai, but companies aren't yet delivering solutions that fit how people actually work. Now, the 95% enterprise AI project failure rate has critics, for example, analysts at a tech research firm called futu. I've got a link to their rebuttal in the show notes, they argue that the number is inflated, that the 95% figure is inflated, and they point out that the report comes from an MIT lab with ties to vendors pushing agentic ai.

02:17   Still, even if the true figure is lower, the big picture holds most enterprise AI projects today aren't creating real impact. So assuming that that 95% figure is correct, what about the 5% that do succeed or the small number of projects that succeed regardless of what the exact number is? Well, these projects share a common playbook that I'm delighted to be able to share with you. The 5% of AI projects that succeed involve models that learn, that have tight integration into workflows and are ag agentic systems. Systems that remember, adapt and act within specified constraints instead of a demo bot that forgets yesterday's feedback. The tools in the 5% of projects that are successful improve with use and embed into daily processes. Winning teams also start small, prove value on a narrow but critical task and then scale from there. The takeaway, AI isn't failing us.

03:17   We failing to deploy it, right? If organizations shift from chasing hype to building adaptive integrated systems, the kind that learn and evolve more of them can move into that successful 5%. And if you don't know where to start on building a successful AI project, you can of course chat with your favorite conversational AI interface, be it chat, GPT, Claude or Gemini or whatever you use to give it the context on the kinds of solutions you're looking for. You could include the full MIT Nando report in your chat. I've got that report for you in the show notes. And of course, if you want support from human experts on how to prioritize, develop and or deploy AI solutions successfully, there are consulting firms out there that can help you. In the us. A particularly well-known and well-regarded AI consulting option is a firm called Tribe AI.

04:12   I've got a link to them, the show notes, and this isn't promotional on their part, I just happen to, yeah, they're really well regarded in general, and so that's just a organic suggestion. But I do have a shameless plug for

you because as of a few months ago, I also have my own consulting firm called Y Carrot, and we can also get you rolling. Feel free to DM me on LinkedIn or click the partner with us button in the top right corner of why carrot.com. Yes, there is your shameless plug of the week. Alright, that's it for today's episode. I'm Jon Krohn and you've been listening to this SuperDataScience podcast. If you enjoyed today's episode or know someone who might consider sharing this episode with them, leave a review of the show on your favorite podcasting platform. Tag me in a LinkedIn post with your thoughts, and if you aren't already, obviously subscribe to the show. The most important thing, however, is I just hope you'll keep on listening. Until next time, keep on rocking it out there, and I'm looking forward to enjoying another round of the SuperDataScience Podcast with you very soon.