Interview with Shibhansh Dohare

Andrew Obwocha: I went over your LinkedIn and saw that you've been researching for your doctorate in continual learning, deep learning, and reinforcement learning. What really stood out to you about those three fields over other areas of AI like NLP?

Shibhansh Dohare: The main focus for me is continual learning. And continual learning isn't really separate from reinforcement learning or NLP or vision—it's orthogonal. What that means is, instead of training systems once on big datasets and deploying them, continual learning is about systems that keep learning from new data as it comes in. That's really important, because the world is always changing. Like, you might train GPT on data up to a certain point, say Spider-Man 2, but when a new Spider-Man movie comes out, the system needs to learn about it. So continual learning is essential for that. It's also, in my view, an essential part of intelligence itself. And because it's missing from a lot of current literature, that's where I've chosen to focus.

Andrew Obwocha: Would you say there's also a deeper level of interest that drew you to it? Because continual learning is quite specialized. I get that it's orthogonal and incorporates other fields, but why did you go down this particular line instead of pushing frontiers in reinforcement learning?

Shibhansh Dohare: My main goal is long-term—solving intelligence. And that means identifying computational principles that give rise to it. I think continual learning is a key component of that journey. There are definitely important open problems in reinforcement learning, deep learning, and all the others, but in the long run, I believe continual learning will be one of the most critical areas. So it's not just an interest; it's aligned with that bigger objective.

Andrew Obwocha: Right now, AI is everywhere—LLMs, image generation, you name it. But you're currently a researcher at UofA, right?

Shibhansh Dohare: Yeah, that's right.

Andrew Obwocha: So at what point did you decide you were going to pursue research instead of going straight into industry? Was it during your undergrad or later in postgrad?

Shibhansh Dohare: That's actually a really good question, because I remember in my second or third year of undergrad, during the summer, I was talking to my mom and I was just really confused. I was like, why would anyone choose to do a PhD? You could just go into industry, make money, do real-world work. But everything changed when I met Rich—Richard Sutton, one of my advisors.

He had this big-picture view of AI. Not just that it's an important problem, but potentially *the* defining problem of our generation. Maybe even the last invention humanity makes on its own. It becomes a way of understanding ourselves—building beings that are just as smart as us. That perspective made me feel like this is the most impactful thing I could be doing. The impact might not be visible tomorrow or even in five years, but on a 50-year or 100-year timescale, this matters the most. Once I saw it that way, everything else felt smaller in comparison.

Andrew Obwocha: That's powerful. So Rich really inspired you?

Shibhansh Dohare: Absolutely.

Andrew Obwocha: During your undergrad or postgrad, I'm sure you had events you attended, interactions with professors or peers, maybe collaborations with companies. Was there any one moment or experience that really shaped how you see the field or impacted your research?

Shibhansh Dohare: Honestly, it ties back to the same thing. I did an internship with Rich, and in the first or second meeting, he asked me a question that really shook me. At that point, I'd taken a lot of machine learning courses, worked on computer vision, deep learning, lots of stuff—and it was all fun. But he asked: *What do you think are the main algorithmic ideas we need to actually make intelligence?* And I was stunned. No one had ever posed the problem that way before.

Instead of chasing benchmark performance, he was thinking about the foundational ideas—like gradient descent, TD learning—deep, high-level abstractions. And that stuck with me. It felt like the right way to think about intelligence. Almost like how physicists, like Newton, tried to understand the universe at a fundamental level. That mindset had a big impact on me.

Andrew Obwocha: So not every student gets the opportunity to meet someone like Richard Sutton, let alone be mentored by him. If you were to give advice to an aspiring AI undergrad today—someone who doesn't have that kind of direct access—what would you say to them?

Shibhansh Dohare: That's a good point. And I've heard this echoed by Rich, Yoshua, even Yann LeCun: The field will always have its hypes—whatever is trending. But if you feel like everyone is missing something, and you have a strong intuition about it, you should follow it.

You should have logical reasons, of course, but even if you're going against the grain, that's often where the breakthroughs happen. Science tends to evolve not by following the hype but by going into the unknown. That's how deep learning and reinforcement learning made it—they were fringe ideas at one point. So if you think you're right and everyone else is missing it, even if it feels arrogant, you should still pursue it. It might be wrong 90% of the time, but that 1% can make all the difference.

Andrew Obwocha: So trust your intuition—even when it contradicts the mainstream?

Shibhansh Dohare: Exactly. Even if it feels isolating or unlikely, it's worth the pursuit.

Andrew Obwocha: I want to pivot a bit. With LLMs, Copilot, Claude, and everything else evolving rapidly—what AI tools or resources do you personally use? How do you stay updated, considering there's something new every week?

Shibhansh Dohare: Yeah, it's tough. I'm not sure I do a great job of staying up to date, honestly. As you said, there's just so much coming out. But my current feeling is that the best models—ChatGPT, Claude, Gemini—they're all about the same right now. Maybe small differences in performance, but nothing drastic. That could obviously change if someone releases a breakthrough, but for now, I don't stress too much over the details.

Personally, I use Claude. It's helpful when I want to learn something outside my area—like if I'm reading neuroscience papers and need something explained, Claude usually does a better job than Google. If I'm debugging code or there's a bug I can't figure out, sometimes Claude can catch it faster than I can. And for things like LaTeX—which is tedious—it's great. I can just describe what I want, and it handles it. Even converting handwritten notes to LaTeX works really well.

Andrew Obwocha: And in terms of coding—how reliant are you on Claude or other AI assistants? Some people use it for everything, others just for debugging.

Shibhansh Dohare: I'm not fully reliant yet, but I'm moving in that direction. I use it for generating code and debugging—it definitely speeds things up. But in my research, the bigger challenge is coming up with ideas and designing experiments. That takes time no matter what. Claude helps, but it's not the bottleneck. In software engineering, maybe it's a bigger productivity boost, but in research, you still have to go through the full experiment cycle.

Andrew Obwocha: So AI tools might shorten the dev time, but experimentation still dominates the timeline.

Shibhansh Dohare: Exactly. The cycle is a bit shorter than it used to be, which is nice, but experimentation was always the bigger chunk of the timeline.

Andrew Obwocha: Got it. Thank you so much, Shibhansh. This has been incredibly insightful. We'll be publishing the interview, and once it's up, I'll send you the link on LinkedIn—and can forward it to your email as well if you'd like.

Shibhansh Dohare: Thanks for having me. I hope it's helpful.

Andrew Obwocha: Absolutely. Take care, and have a good night.

Shibhansh Dohare: You too.