INTERVIEW: Andrew Obwocha with Kenneth Dwyer

Andrew Obwocha:

Hi Kenneth, thanks for joining us today. We're going to jump right into it — I've got four quick questions for you. Starting with this: I saw on your LinkedIn that you've built NLP pipelines to extract behavioural health insights from unstructured clinical notes. With the pace of AI development over recent years, and with more startups integrating AI into daily workflows, which industries do you feel are seeing the fastest or most interesting integration of AI today? I'm thinking of spaces like software engineering, healthcare, even education.

Kenneth Dwyer:

Yeah, that's a great question — though it's a pretty broad one. I'd say healthcare has traditionally lagged behind in adopting machine learning and AI. I remember when I joined my current company, my boss actually said that directly. But things are changing quickly. We're now using cutting-edge tools — GPT models, OpenAI APIs, and other state-of-the-art technologies. That's been accelerating the pace of adoption, at least in the kind of work I do, like information extraction from clinical text.

Outside of healthcare, it's harder for me to speak with authority. But I do come from a startup background — customer support, for instance, is an area where I saw clear potential, even back in 2020. I worked on moderation tools for children's games — the company was acquired by Microsoft, and the tech is now part of Xbox Live. I look back and think: wow, if we had access to something like GPT at the time, it would've been game-changing. A fluent, context-aware support agent would have outperformed our canned responses.

Andrew:

You also mentioned software engineering and education. Do you see anything notable happening in those areas?

Kenneth:

Absolutely. In software engineering, even for me personally, I've started using coding assistants. They speed things up in a major way. But there's a caveat — I've noticed junior team members sometimes rely too heavily on AI-generated code without fully understanding it. So we've had to re-emphasize testing, debugging, and knowing why things work.

Education is especially interesting. I have a three-year-old and a baby at home, so I've been thinking a lot about how AI will shape learning. Even now, my son enjoys generating images using Gemini — similar to DALL·E or other models — and then coloring them in. It's a light introduction, but it shows how intuitive and engaging AI can be. I think platforms like Khan Academy are doing some really exciting things here, too.

Andrew:

Just to circle back to healthcare for a moment — you mentioned earlier that it had lagged behind. Why do you think that's the case?

Kenneth:

Good question. Healthcare is such a broad space, but when it comes to working with electronic health records or patient data, the bottleneck is often regulation. Privacy and compliance requirements are strict — HIPAA in the U.S., and similar laws in Canada. At our company, for example, we go through multiple approval steps just to access data for research.

The process is improving, but it's still slow. There's actually a company we've looked into — Tonic.ai — that creates realistic synthetic datasets for testing and development. They keep the statistical properties but anonymize the content in smart ways. Solutions like that are helping us move faster without compromising privacy.

Andrew:

That's a great overview — now on to the second question. Your current work seems to blend technical skills, business impact, and communication. We hear a lot about those three pillars. But are there any lesser-known habits or personal workflows you've developed that have really made a difference in your career?

Kenneth:

Yeah, that's an interesting angle. One thing that's become essential for me is reproducibility — really keeping track of experiments and changes. In startups, your work often speaks for itself. But in a large company, clear communication and repeatability are crucial.

So if I'm running an NLP experiment — say, extracting breast cancer information from notes — I want to log every prompt tweak, every model version, every dataset. It's not just about version-controlling your code with Git — you also need to track your models, data versions, even prompts. Tools like MLflow are great for this.

Why? Because when leadership asks, "What made the model improve?" or "Can we ship this version?" — you need to have an exact answer. It also saves you from repeating the same work or guessing later. Honestly, I wasn't taught this in school, but it's one of the most useful habits I've picked up.

Andrew:

I really resonate with that. As a student, version control is mostly talked about in the context of code — but not so much models or datasets. That brings me to the third question. You spent

some time in a PhD program at the University of Alberta before moving into industry. Was there anything unexpected from that time that ended up shaping your career?

Kenneth:

Yeah… that's actually a bit of a sensitive topic for me. I didn't finish my PhD, and I still carry some regret about that. I had great mentors and learned a ton — but in hindsight, I think I lacked the drive to invent new things. I liked building things that worked.

What changed everything was this opportunity from a company in Kelowna — Two Hat Security. They were reaching out to Canadian universities for interns, and my advisor passed it along. I wasn't officially in an internship program, but I thought, "Why not?"

I loved it. I came back, tried balancing work and my PhD part-time, but eventually decided to leave and go full-time into industry. And I've never regretted that decision. That internship opened every door after. It was around 2014 or 2015, and I stayed with them until 2016 before moving to Silicon Valley.

Andrew:

And around that time, GPT was just starting to gain attention. What were your early thoughts on it?

Kenneth:

Honestly? I underestimated it. I thought it was mostly hype. I remember OpenAI warning about GPT-2 being dangerous, and I was skeptical — like, really? It just generates text…

But eventually, through work, I started using it more seriously. Even then, I debated between using these large models versus fine-tuning smaller ones. But now, with prices coming down and performance improving, it's hard to beat OpenAI's models for many tasks.

Even in my personal life — I recently tried Google's deep research tools and notebook-style LM apps. I was blown away. So yeah, I definitely came around.

Andrew:

That's a great arc — from skepticism to real-world use. So for our final question: If you had just 60 seconds to give advice to students aspiring to work in NLP or AI more broadly, what would you say?

Kenneth:

Okay… don't start the timer just yet — I need a second to think!

Alright. I'd say:

If you want to actually build or improve these models, load up on math. Optimization, linear algebra, calculus — that's your toolbox. You need it to understand how models are trained and how to push them further.

But if you're more interested in applying AI — using models, building products — then focus on problem-solving and engineering skills. Learn to identify real-world pain points and prototype quickly.

Across both paths, the most underrated skill is communication. If you can translate technical work into business value — in plain language — you'll stand out.

And lastly: keep experimenting. Build things. Break them. Learn fast. That's what the best folks in AI are doing every day.

Andrew:

That's fantastic advice, Kenneth. Thanks again for taking the time — this was a very thoughtful and inspiring conversation.

Kenneth:

Thanks, Andrew. Really enjoyed it.

Let me know if you'd like this formatted for publishing on LinkedIn, Medium, or the UAIS website.