

Machine Learning Engineers Will Not Exist In 10 Years.

The landscape is evolving quickly.



Luke Posey



Note: this is an opinion piece, feel free to share your own opinion so we can continue to move our field in the right direction.

Machine Learning will transition to a commonplace part of every Software Engineer's toolkit.

In every field we get specialized roles in the early days, replaced by the commonplace role over time. It seems like this is another case of just that. Let's unpack.

Machine Learning Engineer as a role is a consequence of the massive hype fueling buzzwords like AI and Data Science in the enterprise. In the early days of Machine Learning, it was a very necessary role. And it commanded a nice little pay bump for many! But Machine Learning Engineer has taken on many different personalities depending on who you ask.

The purists among us say a Machine Learning Engineer is someone who takes models out of the lab and into production. They scale Machine Learning systems, turn reference implementations into production-ready software, and oftentimes cross over into Data Engineering. **They're typically strong programmers who also have some fundamental knowledge of the models they work with.**

But this sounds a lot like a normal software engineer.

Ask some of the top tech companies what Machine Learning Engineer means to them and you might get 10 different answers from 10 survey participants. This should be unsurprising. This is a relatively young role and the folks posting these jobs are managers, oftentimes of many decades who don't have the time (or will) to understand the space.

Here are a few requirements from job listings from some of the top tech companies, notice how vastly they differ:

This first one is spicy. Are you sure this isn't a researcher? How is this a Machine Learning Engineer?

- PhD in Math, Stats, Operations Research. Knowledge of R, SQL, and modern Machine Learning techniques.

This next one's more on-brand. And it comes from the top so it shouldn't be a surprise.

- BS or MS in Computer Science. 1–5 years work or academic experience in software development. Exposure to Computer Vision, NLP, etc a plus.

And finally drilling down on your stereotypical ML Engineer posting.

- BS/MS in Computer Science. 3 or more years building production Machine Learning systems and efficient code. Experience with Big Data a plus.

Some companies have started a new approach and I think most will follow. The approach is to list a Software Engineering role with exposure to Machine Learning as a core requirement + a few years of experience as a preferred qualification. Employers will take a preference to engineers with experience building and scaling systems, regardless of whether it was based on Machine Learning or some other technology.

The Machine Learning Engineer is necessary as long as Machine Learning understanding is rare and has a high barrier to entry.

It's my earnest belief that the role of Machine Learning Engineer will be taken over entirely by the common software engineer. It will transition to a standard engineering role where the engineer will get a spec or reference implementation from someone upstream, turn it into production code, and ship and scale applications.

For now, much of many Machine Learning roles exist in this weird space where we're attacking problems with ML that just haven't been attacked before. By consequence, ML Engineers are in many cases half researcher, half engineer. I've come across my fair share of Machine Learning Engineers who play across the entire stack. I've come across others who have a more narrow skillset but spend more time reading new research papers and turning them into usable code.

We're at a weird crossroads where we're defining where the members of our teams fit into the puzzle.



research → spec/reference implementation/PoC → production

By consequence of the way we work, we tend to shove ourselves into discussions and sit in meetings regardless of whether it's core to our expertise. We accept any and every meeting invite... It's my opinion the Machine Learning Engineer belongs at the tail end of building a reference implementation and then owns turning any of that into production code.

Not long from now, most enterprises will have little need for research efforts to get their projects to the finish line. Only niche use-cases and deep technical efforts will require a special skillset. Engineers will consume APIs and the world will move on; Machine Learning becoming a commonplace tool in every new engineer's toolkit. We're already seeing this as more and more

exposure to Machine Learning trickles into universities. Go to a Machine Learning course at a university and it's packed to the brim. Almost every graduate will leave university with some exposure to the field.

We can draw an analogy to Blockchain where the Distributed Systems Engineer became **hot**. The vast majority of Blockchain projects since Nakamoto's white paper have been spending their efforts on building the fundamental technology and infrastructure. To do so you had to have incredibly strong engineering skills, most-often described as a Distributed Systems Engineer. You're finally seeing a shift where things are getting abstracted, enterprises are starting to find use-cases, and the everyday engineer can now build novel use-cases using blockchain. We're seeing the same general shift in AI/ML.

Some Valid Counter Points

- It's possible that the Silicon Valley theme of "One API to rule them all" is a bunch of bogus and Machine Learning will always require some degree of customization at the infrastructure level. It's my opinion that what HuggingFace is to NLP will happen to every other domain. We'll be able to conquer the majority of use-cases with a simple API.
- "It's just a title dude. Machine Learning Engineer just means someone with a heavier background in Math and Stats than your average CS graduate." Totally agree. It's just a title. But if that role is no longer necessary will the title exist? But you're right, it's just a title.
- "In my organization that's not what Machine Learning Engineer means at all." Let me know what it means to your organization so I can learn. I'm constantly surveying the field to understand where things are at and where they're headed. I would love to hear your outlook.
- "It's just a title. Who cares?" You're right, but it's fun to consider anyways.
- "Machine Learning is a nascent field with new use-cases and research being constantly realized; to think this will slow down in the next decade is naive." Very possible!

One of my favorite responses to the article, from Varii on Twitter:

"Like you said, it's a title. Most employers expect you to have overlapping skillsets. I feel like in the end it's not about who gets wiped out, it's about who is versatile enough to constantly adapt to the ever changing industry."

Tons of great input from the broader community that I'm learning from. But my opinion will never shift on one thing: if you're passionate about something it doesn't matter what happens to a title, a field, or a trend, there will always be a place for you to pursue your passion and build cool things.

Stay safe and build on!

[I started a \(free\) analytics group called Dataset Daily where we share a dataset every Monday and code throughout the week.](#)

Let's continue the conversation on [Twitter](#).