

Production machine learning

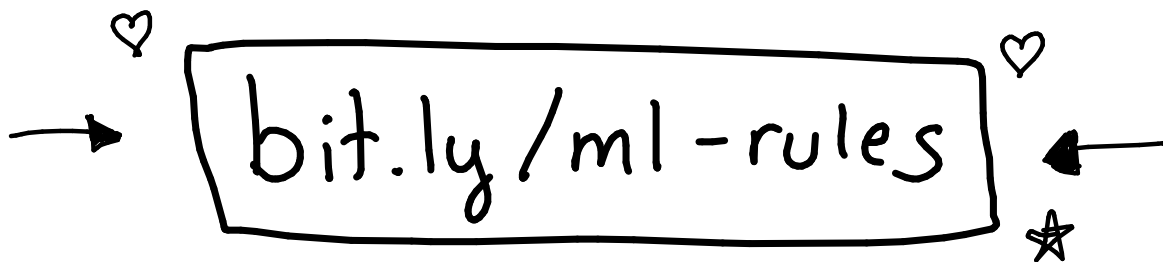


by Julia Evans
twitter: @b0rk
blog: jvns.ca

What's this?

I read "Rules of Machine Learning: Best Practices for ML Engineering" and a lot of it really resonated with my experience doing machine learning.

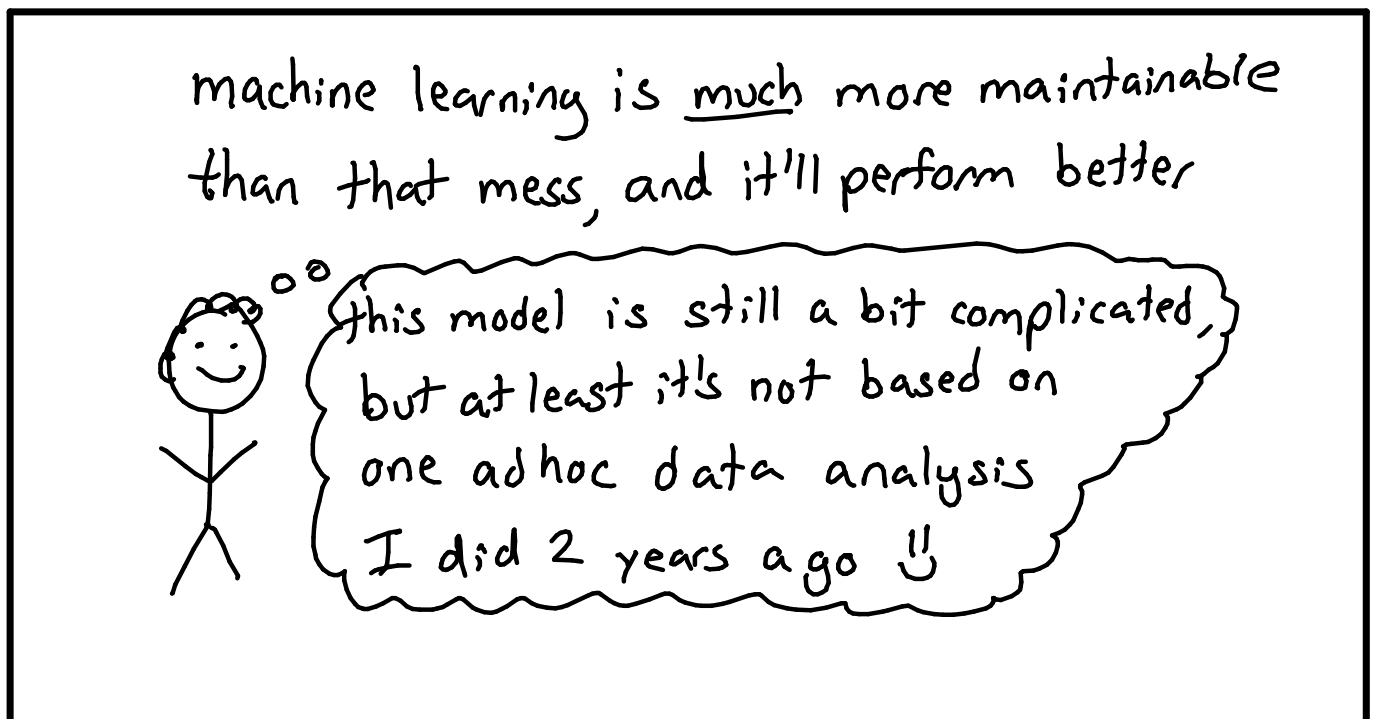
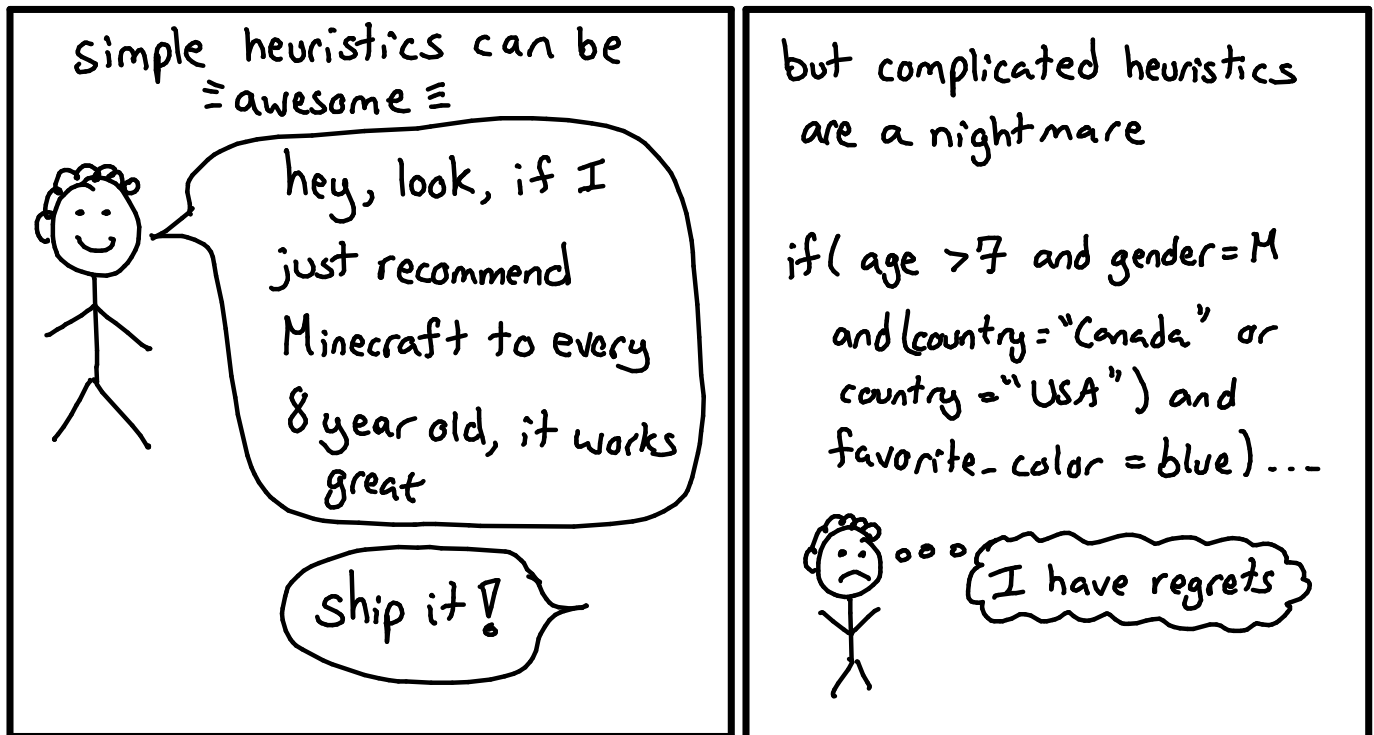
You can find that PDF at:

→  bit.ly/ml-rules

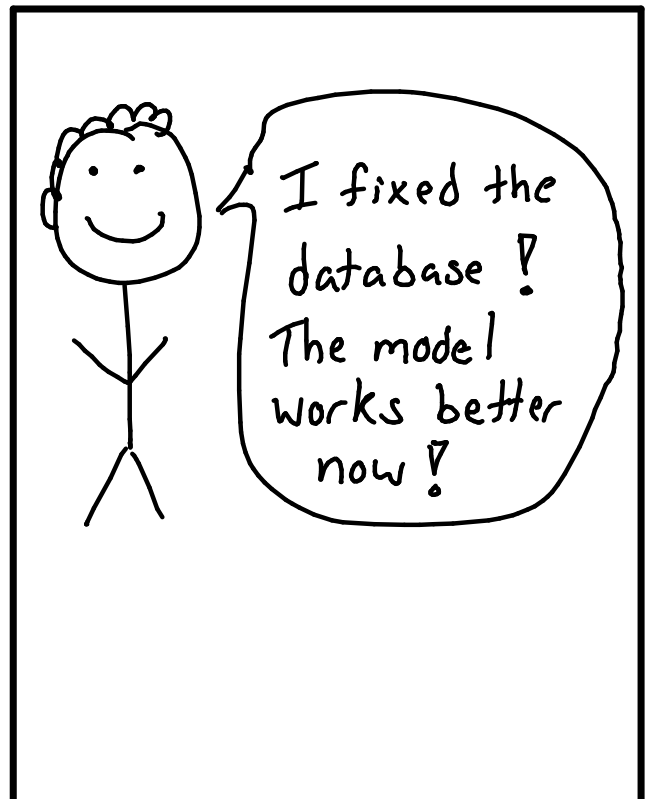
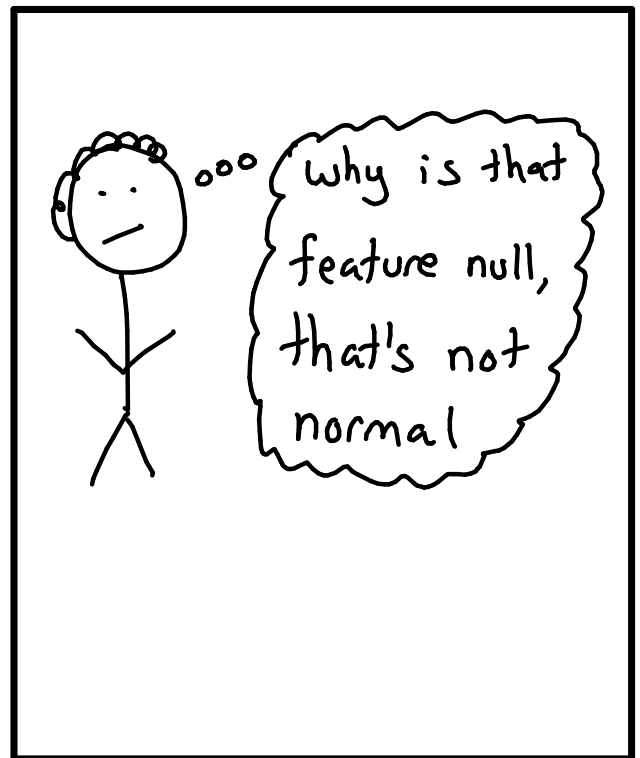
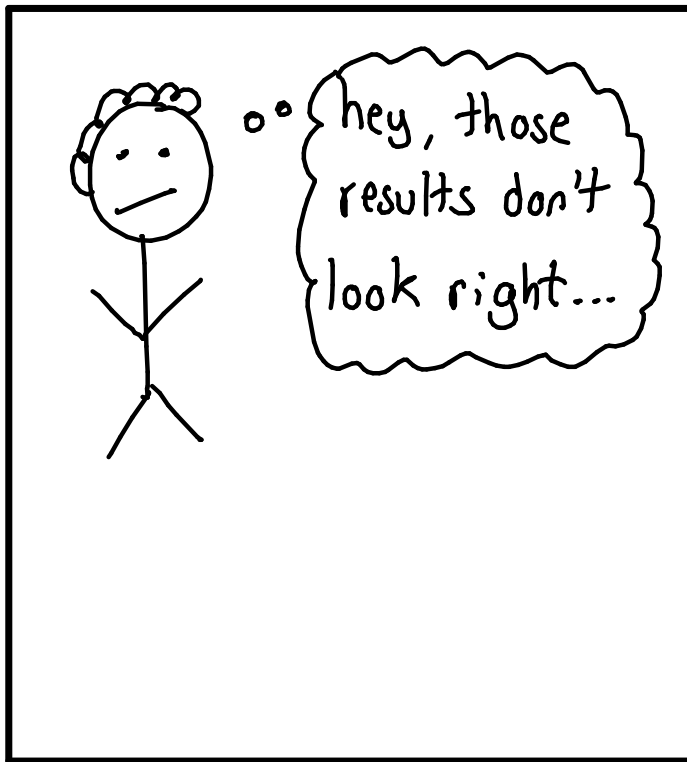
In this zine, I'll talk about what I've learned about machine learning! Basically I want to talk about 5 of my favourite rules from that PDF.

(there are 43 rules, it's great advice, you should read them all)

Rule #3: Use machine learning instead of complicated heuristics

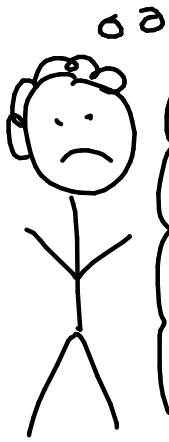


Rule # 10: watch for silent failures



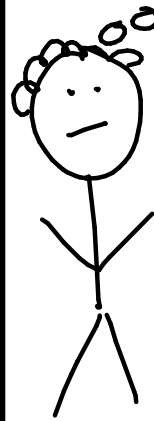
Rule #16 : keep your pipeline healthy

this is sad:



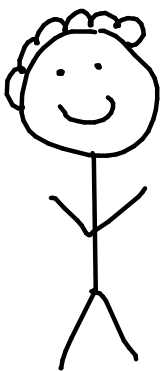
I can't generate the features I need to train a new model!

and this



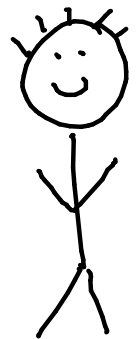
I'm worried there's a bug + I don't trust our trained models

what we want:

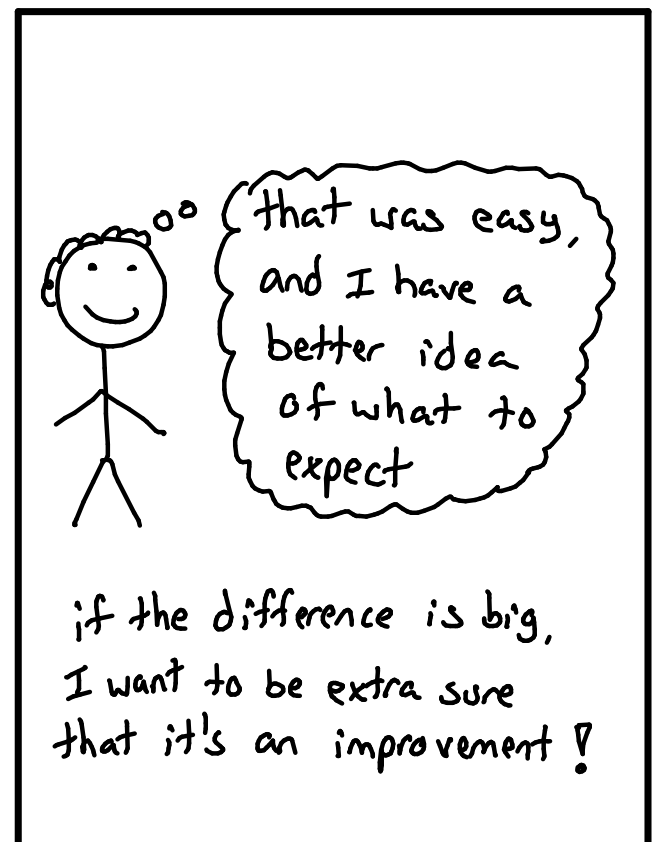
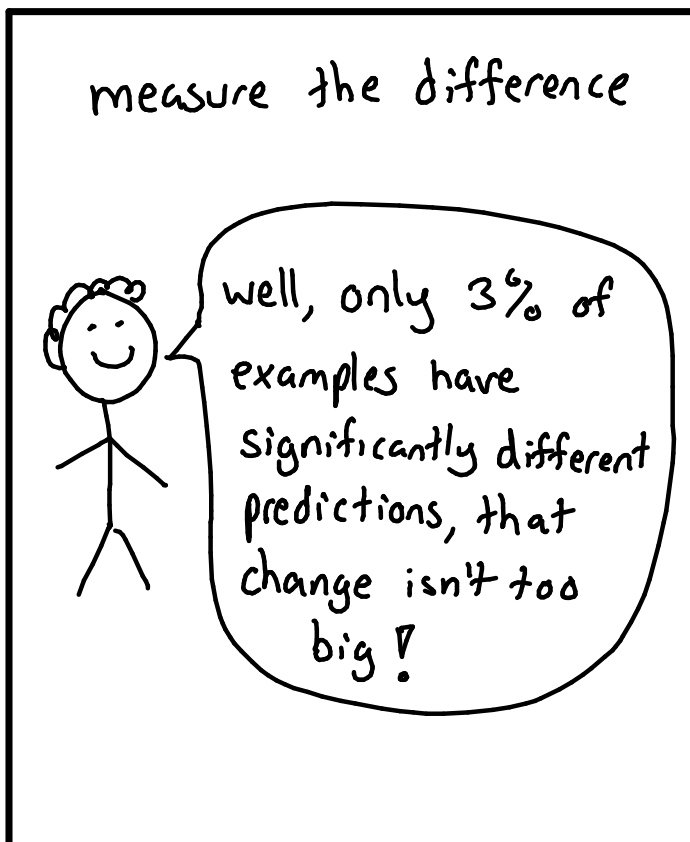
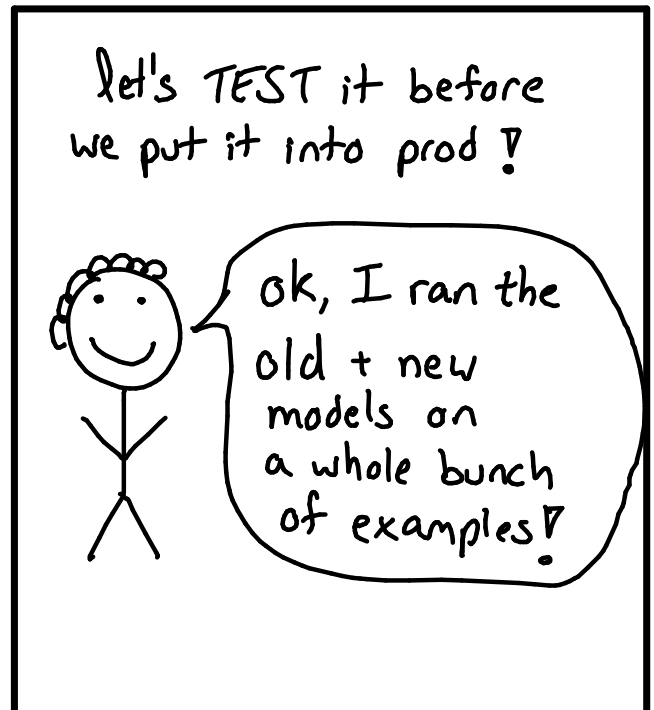
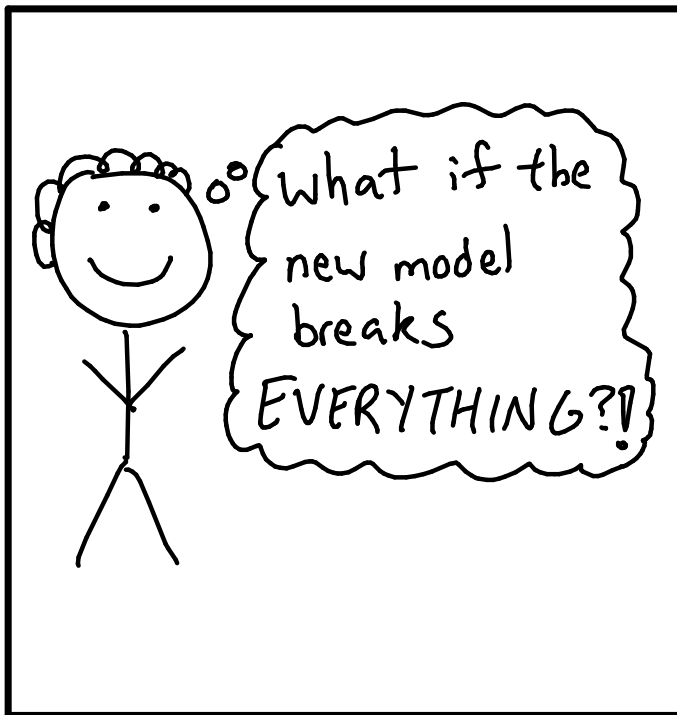


when I have new ideas / data, I can easily launch a new model!

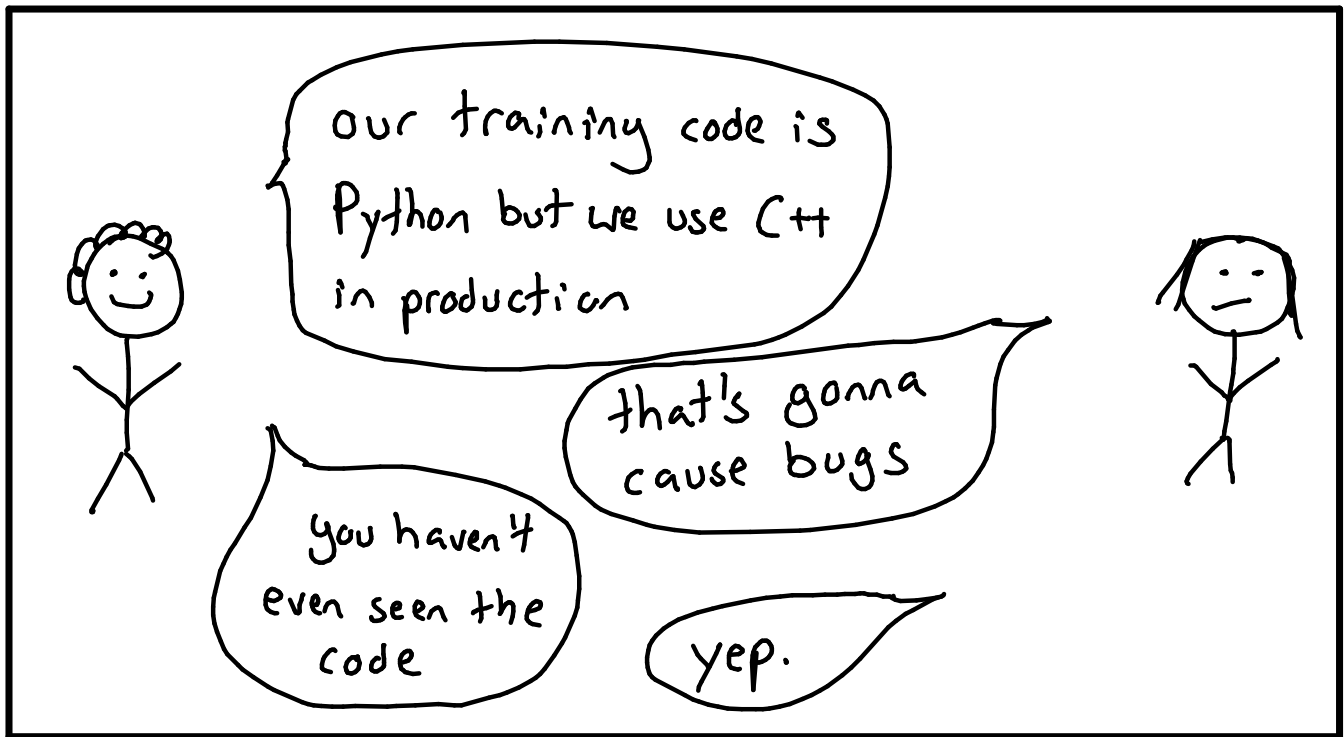
we're careful about adding complexity!



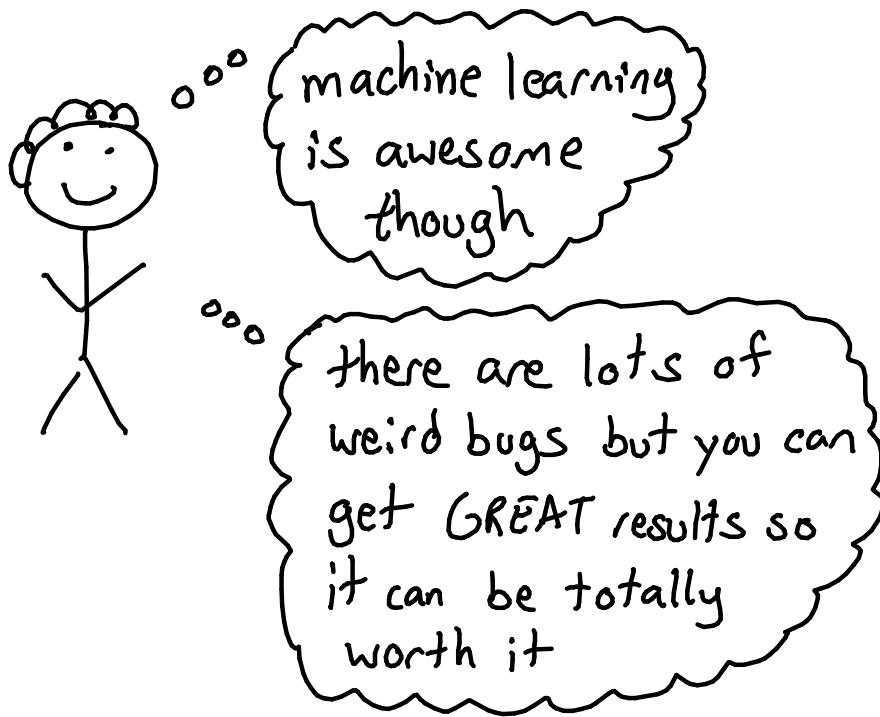
Rule #24: Measure the change between models



Rule # 32 : Reuse code between training + production when you can



- you're gonna have bugs related to training / production differences
- you probably can't make this perfect
- just try to monitor differences.



♡ this? There are more
of these at

→ jvns.ca/zines ←