

Machine Learning for Data Journalism

Hi, my name is Dale Markowitz, and I'm applied A.I. engineer at Google. Among other things, I work as a machine learning advisor to the Google News Lab, where I help newsrooms understand how they can use this type of technology. And there are lots of ways that machine learning or M.L. can help journalists.

And every year I see more and more examples of reporting that takes advantage of this technology to write stories that wouldn't have been possible otherwise. Part of the reason for this is that the tools that we have for building machine learning models are getting much, much easier to use and more approachable.

To run a machine learning powered investigation today, you probably have to work with some sort of an engineer or a programmer, or maybe you're a programmer yourself. But you no longer need the specialized training you would have had to have had in the past to do this sort of stuff. And we'll talk more about tools to do machine learning later.

But that said, I also read a lot of things that I wish I hadn't – articles that get what machine learning is wrong or that use it incorrectly or misleadingly. So in this week's lesson, I'll try my hardest to make sure that you're never the byline of one of those stories.

I'll start off by introducing you to machine learning at a very high level. And then I'll show you all of the different types of problems that it can solve, so you know what tools you have in your toolbox.

We'll also talk about some of my favorite projects and stories powered by machine learning. We'll also spend some time discussing when to use machine learning, and how to explain that you've used M.L. to your readers, and especially when not to use it. This part is really important. So if you're going to skip anything, make sure it's not this bit. At the end, I'll show you some of my favorite tools for building models and the ones that I think are the easiest to use. Thanks for listening and let's get started.