Twitter Thread by Francesco Pochetti

Francesco Pochetti
@Fra_Pochetti

All ML projects which turned into a disaster in my career have a single common point:

■ I didn't understand the business context first, got over-excited about the tech, and jumped into coding too early.

When someone asks you for a model, always ask:

■ why do you need it?
■ what is your current solution (e.g. what is the baseline)?
■ who is going to use the predictions and how?
■ what is the impact of the model's downtime or mistakes?
■ which metrics do we care about?

Once you have your answers, back them up with a solid exploratory data analysis, and, when done, loop in the biz team again.

This is a critical moment as your results will translate into 3 potential outcomes:

■ "Really? This is weird. Well, in this case, the ML model doesn't make much sense anymore". You are off the hook ■
■ "Interesting. I guess we'll have to change requirements/scope then." Course-correct before moving forward ■
■ "This is what I expected. Let's go ahead".■

Might seem silly, but skip the above and you are all set for failure.
Trust me, I learned it the hard way ■

Also, always remember that the best model is no model.