

10 Things Everyone Should Know About Machine Learning

By Daniel Tunkelang

As someone who often finds himself explaining machine learning to non-experts, I offer the following list as a public service announcement.

1. **Machine learning means learning from data; AI is a buzzword.** Machine learning lives up to the hype: there are an incredible number of problems that you can solve by providing the right training data to the right learning algorithms. Call it AI if that helps you sell it, but know that AI, at least as used outside of academia, is often a buzzword that can mean whatever people want it to mean.
2. **Machine learning is about data and algorithms, but mostly data.** There's a lot of excitement about advances in machine learning algorithms, and particularly about deep learning. But data is the key ingredient that makes machine learning possible. You can have machine learning without sophisticated algorithms, but not without good data.
3. **Unless you have a lot of data, you should stick to simple models.** Machine learning trains a model from patterns in your data, exploring a space of possible models defined by parameters. If your parameter space is too big, you'll overfit to your training data and train a model that

doesn't generalize beyond it. A detailed explanation requires more math, but as a rule you should keep your models as simple as possible.

4. **Machine learning can only be as good as the data you use to train it.** The phrase “garbage in, garbage out” predates machine learning, but it aptly characterizes a key limitation of machine learning. Machine learning can only discover patterns that are present in your training data. For supervised machine learning tasks like classification, you'll need a robust collection of correctly labeled, richly featured training data.
5. **Machine learning only works if your training data is representative.** Just as a fund prospectus warns that “past performance is no guarantee of future results”, machine learning should warn that it's only guaranteed to work for data generated by the same distribution that generated its training data. Be vigilant of skews between training data and production data, and retrain your models frequently so they don't become stale.
6. **Most of the hard work for machine learning is data transformation.** From reading the hype about new machine learning techniques, you might think that machine learning is mostly about selecting and tuning algorithms. The reality is more prosaic: most of your time and effort goes into data cleansing and feature engineering — that is, transforming raw features into features that better represent the signal in

your data.

7. **Deep learning is a revolutionary advance, but it isn't a magic bullet.** Deep learning has earned its hype by delivering advances across a broad range of machine learning application areas. Moreover, deep learning automates some of the work traditionally performed through feature engineering, especially for image and video data. But deep learning isn't a silver bullet. You can't just use it out of the box, and you'll still need to invest significant effort in data cleansing and transformation.
8. **Machine learning systems are highly vulnerable to operator error.** With apologies to the NRA, "Machine learning algorithms don't kill people; people kill people." When machine learning systems fail, it's rarely because of problems with the machine learning algorithm. More likely, you've introduced human error into the training data, creating bias or some other systematic error. Always be skeptical, and approach machine learning with the discipline you apply to software engineering.
9. **Machine learning can inadvertently create a self-fulfilling prophecy.** In many applications of machine learning, the decisions you make today affect the training data you collect tomorrow. Once your machine learning system embeds biases into its model, it can continue generating new training data that reinforces those biases. And some biases can ruin people's lives. Be responsible: don't create self-fulfilling prophecies.

10. **AI is not going to become self-aware, rise up, and destroy humanity.** A surprising number of people (cough) seem to be getting their ideas about artificial intelligence from science fiction movies. We should be inspired by science fiction, but not so credulous that we mistake it for reality. There are enough real and present dangers to worry about, from consciously evil human beings to unconsciously biased machine learning models. So you can stop worrying about SkyNet and “superintelligence”.

There’s far more to machine learning than I can explain in a top-10 list. But hopefully this serves as a useful introduction for non-experts.