Twitter Thread by Santiago



Santiago
@svpino

The 8-step quick-start guide to learn Machine Learning.

■■

1■ Start with Python ■

Yes, you can do other languages, but Python is by far the most straightforward option.

■

2■ Get familiar with numpy, pandas, and matplotlib

These three libraries are probably the most common Python libraries you'll have to use every day.

(Even if you don't end up doing machine learning, these libraries are awesome and useful.)

■

3■ Start using notebooks

Look into Jupyter or Google Colab.

Notebooks are essential for data scientists and machine learning practitioners. Most of the code you'll read and write will be in notebooks.

■

4■ Find a problem (already solved)

In my opinion, the best way to start is by working through a problem —especially when you can learn from its solution.

Start with something simple. I usually recommend "Titanic" from Kaggle.

■

5■ Focus on the analysis and not the code

In the beginning, spend your time and energy analyzing the problem and its solution.

Code is not important at this stage. Code can come later.

■

6■ Start incorporating new algorithms

As you work through problems, start incorporating new algorithms into your toolset.

Here are a few great options to start:

1. Decision Trees
2. Linear regression
3. Logistic regression
4. Neural Networks
5. KNN

■

7■ Get familiar with a general process to approach problems

Here is a good start:

1. Define the problem
2. Prepare the data
3. Spot-ccheck algorithms
4. Improve the results
5. Present the results

■

8■ Pick a new problem and repeat

It shouldn't be surprising that the best way to improve is to practice and solve new problems.

If you don't have access to real-life problems, get familiar with Kaggle: everything you need will be there.

In the next coming weeks, I'll be posting a whole series of machine learning advice for people wanting to start.

Stay tuned!