Advance Python Programming

A Machine Learning-Focused Approach

Dionisios N. Sotiropoulos

November 10, 2021

University of Piraeus
Department of Informatics

Table of contents

A Toy Example of Model-Based Learning

Can money buy happiness?

Model-based Learning:

- A typical way to **generalize** from a set of examples is to build a **model** of these examples and then use that model to make predictions. This is called **model-based learning**.
- For example, suppose you want to know if money makes people happy, so you download the **Better Life Index** data from the OECD's website and stats about gross domestic product (**GDP**) per capita from the IMF's website. Then you join the tables and sort by GDP per capita.

Country	GDP per Capita (USD)	Life Satisfaction
Hungary	12,240	4.9
Korea	27,195	5.8
France	37,675	6.5
Australia	50,962	07.3
United States	55,805	7.2

Table 1: Does money make people happier

Can money buy happiness?

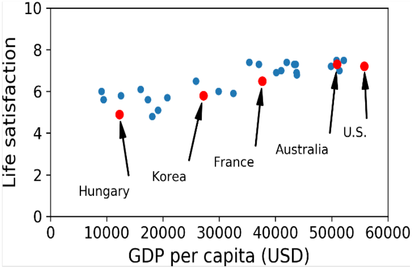

Figure 1: Do you see a trend here?

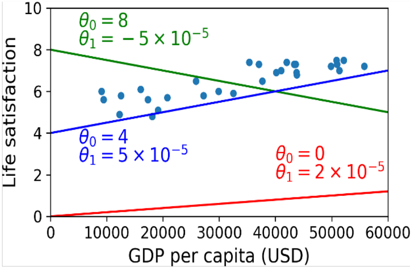

Figure 2: A few possible linear models.

Model Selection:
There does seem to be a trend here! Although the data is noisy (i.e., partly random), it looks like life satisfaction goes up more or less linearly as the country's GDP per capita increases. So we may decide to model life satisfaction as a **linear function** of GDP per capita.

$$life\_satisfaction = \theta_0 + \theta_1 \cdot GDP\_per\_capita \tag{1}$$

Can money buy happiness?

This model has two model parameters, $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$. By tweaking these parameters, you can make your model represent any linear function as shown in Fig.2.

- What is the optimal choice for the parameters θ_0 and θ_1?
- Answering this question requires the definition of an appropriate **performance measure**.
- You can either define a **utility function** (or **fitness function**) that measures how good your model is, or you can define a **cost function** that measures how bad it is.
- For Linear Regression problems, people typically use a cost function that measures the distance between the linear model's predictions and the training examples; the **objective** is to **minimize** this distance.

Can money buy happiness?

Linear Regression Algorithm:
This is where the Linear Regression algorithm comes in: you feed it your training examples, and it estimates the parameters that make the linear model fit best to your data. This is called **training** the model.

- In our case, the algorithm finds that the **optimal** parameter values are $\theta_0^* = 4.85$ and $\theta_1^* = 4.91 \cdot 10^{-5}$.

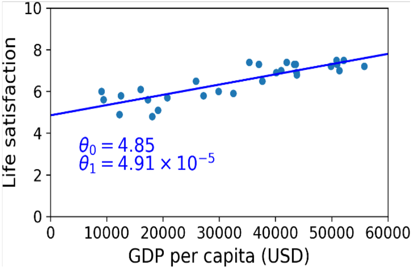

Figure 3: The linear model that fits the training data best.

Can money buy happiness?

We are finally ready to run the model to make predictions. For example, say you want to know how happy Cypriots are, and the OECD data does not have the answer.

- Fortunately, you can use your model to make a good prediction: you look up Cyprus's GDP per capita, find $\$22,587$, and then apply your model and find that life satisfaction is likely to be somewhere around $4.85 + 22,587 \cdot 4.9110^{-5} = 5.96$.