Deadline: We. October 30th , 16:00 Drop your printed or legible handwritten submissions into the boxes at Samelsonplatz. Alternatively upload a `.pdf` file via LearnWeb. (e.g. exported Jupyter notebook)

1. Machine Learning Basics (10 points)

Which kind of machine learning problem are the following tasks? Use the scheme from the "Fundamental Machine Learning Problems" slide (slide 33 in the 2018 script).

1. Predicting tomorrows temperature through meteorological data.
2. Predicting whether a hard drive is going to fail within the next year using usage statistics.
3. Grouping songs into genres based on their musical similarity.
4. A bot learning chess by playing against itself.
5. Generating photo-realistic looking faces.
6. Recommending movies to users in a video streaming platform.
7. Detecting whether a bank transfer is suspicious.
8. Automatically identifying street signs with camera.
9. Anti-aliasing computer graphics.
10. Replacing one person's face with another's in a video clip ("deep fakes").

2. Linear Regression (10 points)

bonus points	0	0.5	1	1.5	2	2.5	3	3.5	4
num. students	18	6	8	2	4	3	8	5	17
num. passed	3	5	3	2	3	3	8	5	16

Table 1: results from last years final exam

Table 1 shows how many students passed in last years final exam, given how many bonus points they obtained.

A. [5p] Fit a Linear Regression model to predict the chance of a student passing given the number of bonus points they obtained. (compute $\hat{\beta}_0$ and $\hat{\beta}_1$)

B. [2p] Make a scatter plot of the data and add the fitted regression curve.

C. [1p] According to the regression model, how likely is a student to pass if they obtained 2 bonus points compared to 0?

D. [2p] What is a potential problem of the linear model in this example?

3⋆ Bonus Questions (10 points)

A. [5p] Propose a better model for problem 2, and explain why it is better. Repeat parts A-C with your proposed model. You are allowed to use `scipy.optimize.minimize` to find optimal model parameters for your model.

B. [5p] Prove that the formulas

$$\hat{\beta}_0 = \bar{y} - \hat{\beta}_1 \bar{x} \qquad \hat{\beta}_1 = \frac{\frac{1}{N} x^\mathsf{T} y - \bar{x} \cdot \bar{y}}{\frac{1}{N} x^\mathsf{T} x - \bar{x}^2}$$

Indeed yield the optimal parameters for a 1-dimensional linear regression model with regards to the least-squares loss.