Machine Learning Course Syllabus

The course is separated into 14 blocks, that should take about 2-3 hours each to complete (depending on your Python proficiency; *case study* about 5h). First read the respective sections in the book, then complete the associated exercises. Feel free to schedule an individual coaching session anytime to ask questions or get feedback on your solutions!

Book: https://franziskahorn.de/mlbook/
Exercises: https://github.com/cod3licious/ml_exercises

Part 1:	**Introduction: Solving Problems with ML**	
	Why should you care about ML?	
	How to create value with ML	
	Different ML algorithms & their use cases	[Quiz 1]
	ML with Python	Python tutorial / installation
	Data & Preprocessing	
	What constitutes 1 data point?	
	Feature Extraction	
	Feature Engineering & Transformations	
	Computing similarities	
	Garbage in, garbage out!	[Quiz 2]
	ML Algorithms: Unsupervised & Supervised Learning	
Part 2:	UL: Dimensionality Reduction	[NB 1] visualize text
	UL: Outlier / Anomaly Detection	
	UL: Clustering	[NB 2] image quantization
	Supervised Learning: Overview	[Quiz 3]
Part 3:		[NB 3] supervised comparison:
	SL: Linear Models	- linear models
	SL: Decision Trees	- decision trees
	SL: Ensemble Methods	- random forest
	SL: k-nearest neighbors (kNN)	- kNN
	SL: Kernel Methods	- SVM
	Information Retrieval (Similarity Search)	[NB 4] information retrieval
	SL: Neural Networks	[NB 5] MNIST with torch (or keras)
	Time Series Forecasting	
	Recommender Systems (Pairwise Data)	
Part 4:	**Avoiding Common Pitfalls**	
	Interpolation: Does the model generalize?	
	Extrapolation: Correlation vs. Causation	
	Explainability & Interpretable ML	[Quiz 4]
		[NB 6] analyze toy dataset
		[NB 7] *predict hard drive failures*
Part 5:	**ML Algorithms: Reinforcement Learning**	[NB 8] RL gridmove
	Conclusion: Using ML in Practice	[Quiz 5]
	AI Transformation of a Company	
	Additional Resources	[Exercise] Your next ML Project