

ML Review Quiz: Big Recap

* Required

1. 3 main types of learning *

3 points

You were introduced to 3 different learning strategies utilized by machine learning algorithms: unsupervised learning, supervised learning, and reinforcement learning. Match each type of learning to one of the goals below:

Mark only one oval per row.

	unsupervised	supervised	reinforcement
Learn how to optimally behave in your environment	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Discover patterns in your data	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Predict Y from X	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

2. The main ML types also have subcategories... *

5 points

...match each subcategory to its main learning type.

Mark only one oval per row.

	Unsupervised Learning	Supervised Learning
Dimensionality Reduction	<input type="radio"/>	<input type="radio"/>
Clustering	<input type="radio"/>	<input type="radio"/>
Anomaly Detection	<input type="radio"/>	<input type="radio"/>
Regression	<input type="radio"/>	<input type="radio"/>
Classification	<input type="radio"/>	<input type="radio"/>

3. Assign each algorithm to one of these categories *

11 points

Mark only one oval per row.

	Dimensionality Reduction	Anomaly Detection	Clustering	Features-based Supervised Learning Models	Similarity-based Supervised Learning Models
PCA	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
t-SNE	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
k-means	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
DBSCAN	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
gamma-index	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
linear / logistic regression	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Decision Tree	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Random Forest	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Neural Network	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
k-nearest-neighbors	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
SVM (a kernel method)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

4. Data vs. Concept Drift *

2 points

We want to predict the price of a house from its size. Which of the scenarios results in either a data or a concept drift? Remember: Data Drift is when the input distribution changes, Concept Drift is when the input/output relation changes.

Mark only one oval per row.

	Data Drift	Concept Drift
Due to inflation, the same size house now costs 20% more than 5 years ago.	<input type="radio"/>	<input type="radio"/>
On average, people built larger houses in 2020 than in 2010.	<input type="radio"/>	<input type="radio"/>

5. Over- and Underfitting *

2 points

A model's performance on new data points can be bad for two reasons: underfitting (= the model has a high bias) or overfitting (= high variance). Which is which?

Mark only one oval per row.

	underfitting	overfitting
When you evaluate the model on the data it was trained on, the performance is close to that of a human, but on new data points it performs poorly.	<input type="radio"/>	<input type="radio"/>
No matter on what data (train or test) you evaluate the model, the performance is always far below that of a human.	<input type="radio"/>	<input type="radio"/>

6. Improving the performance *

5 points

Which of these actions can help if your model is either over- or underfitting?

Mark only one oval per row.

	underfitting	overfitting
try a more complex (non-linear) model	<input type="radio"/>	<input type="radio"/>
feature engineering	<input type="radio"/>	<input type="radio"/>
use regularization	<input type="radio"/>	<input type="radio"/>
get more data (samples)	<input type="radio"/>	<input type="radio"/>
feature selection	<input type="radio"/>	<input type="radio"/>

7. Machine learning is an “iterative” process, meaning that an AI team often has to try many ideas before arriving at a solution that’s good enough, rather than have the first thing they try work. *

1 point

Mark only one oval.

- True
 False

8. Which of these are reasons that it’s often unrealistic to expect an ML system to be 100% accurate? *

1 point

Mark only one oval.

- You might not have enough data
 Data can be mislabeled
 Data can be ambiguous
 All of the above