5.When a machine learning model gives good accuracy on the train set and bad accuracy on the test set, we say that the model is:

(4 Points)

- Over-fitting
- Under-fitting
- Needs more data
- Needs more representative data
- None of the other options

6.What is the output of the following Python code?

```
L = []
for i in range(10):
        L.append(i%2)
print (sum(L)//2)
```

(4 Points)

- 2
- 2.5
- 3
- 5
- None of the other options

7.What is the sum of Array A elements after executing the following Python Code:

```
import numpy as np
A = np.ones((4,4))
A *= 2
A[0, :] *= 0
```

(4 Points)

- 0
- 2

- 16
- 32
- None of the other options

8.For the following Python code, what is the sum of all elements of Array A? from sklearn.preprocessing import StandardScaler

A = [[0], [1], [2], [3], [4]]

scaler = StandardScaler()

scaler.fit(A) A = scaler.transform(A)

(4 Points)

- 0.0
- 1.0
- 5.0
- 10.0
- None of the other options

9.Given the following confusion matrix, what is the prediction accuracy?

	C0	C1	C2
C0	50	5	5
C1	2	55	3
C2	7	8	45

(4 Points)

- 83%
- 90%
- 85%
- 75%
- None of the other options

10.After executing the following code, the user notices that the accuracy on the train set is 100% and 60% on the test set. Which one of the following lines could work as a solution for this problem?

tree_clf = DecisionTreeClassifier(max_depth=6)

tree_clf.fit(X, y)

(4 Points)

- tree_clf = DecisionTreeClassifier(max_depth=None)
- tree_clf = DecisionTreeClassifier(min_samples_split=2, max_depth=6)
- tree_clf = DecisionTreeClassifier(max_depth=4)
- tree_clf = DecisionTreeClassifier(max_depth=8)
- None of the other options

11.How many trainable parameters does the following model has?

model = keras.models.Sequential()

model.add(keras.layers.Dense(100, activation="relu", input_shape=[10]))

model.add(keras.layers.Dense(100, activation="relu"))

model.add(keras.layers.Dense(10, activation="softmax"))

(4 Points)

- 12,210
- 12,000
- 12,200
- 12,100
- None of the other options

12.When training a model using the following code, the training stops at the end of Epoch 95. Which epoch has the best validation loss?

checkpoint_cb = keras.callbacks.ModelCheckpoint("my_keras_model.h5", save_best_only=True)

early_stopping_cb = keras.callbacks.EarlyStopping(patience=5, restore_best_weights=True)

history = model.fit(X_train, y_train, epochs=100, validation_data=(X_valid, y_valid), callbacks=[checkpoint_cb, early_stopping_cb])

(4 Points)

- 90

- 95
- 100
- 105
- None of the other options

13.How many training iterations are conducting in the following tuning experiment?

```
from sklearn import svm, datasets
from sklearn.model_selection import GridSearchCV
iris = datasets.load_iris()
parameters = {'kernel':('linear', 'rbf'), 'C':[1, 5, 10]}
svc = svm.SVC() clf = GridSearchCV(svc, parameters, cv=3)
clf.fit(iris.data, iris.target)
```

(4 Points)

- 18
- 6
- 5
- 15
- None of the other options

14.How many trainable parameters are included in the following three layers?

```
[keras.layers.Conv2D(64, 7, activation="relu", padding="same", input_shape=[28, 28, 1]),
keras.layers.MaxPooling2D(2),
keras.layers.Conv2D(128, 3, activation="relu", padding="same")]
```

(4 Points)

- 77,056
- 76,864
- 77,058
- 76,866
- None of the other options