Final Project - Algorithms in Machine Learning and their Application

Group D: Valerio Cini, Marco La Vecchia

July 11, 2019

CIFAR-10 (1/3)

The dataset CIFAR-10 consists of 60000 images of dimension (32, 32) with three color channels. The dataset is split into:

- training set: 50000 samples
- test set: 10000 samples

CIFAR-10 (2/3)

Each image belongs to one of the following classes :

0) airplane,
1) automobile,
2) bird,
3) cat,
4) deer,
5) dog,
6) frog,
7) horse,
8) ship,
9) truck.

CIFAR-10 3/3

We can see some examples:

Settings

For all our models we used :

- **ReLU** as activation function,
- **SoftMax** for the last layer,
- **Adam** optimizer,
- **NB_EPOCH** $= 15$ (CNN) or 50 (perceptron),
- **BATCH_SIZE** $= 128$.

Our Favorite Models : Model 1

Layer (type)	Output Shape	Param #
dense_56 (Dense)	(None, 512)	1573376
dense_57 (Dense)	(None, 256)	131328
dense_58 (Dense)	(None, 256)	65792
dense_59 (Dense)	(None, 256)	65792
dense_60 (Dense)	(None, 512)	131584
dense_61 (Dense)	(None, 128)	65664
dense_62 (Dense)	(None, 64)	8256
dense_63 (Dense)	(None, 32)	2080
dense_64 (Dense)	(None, 16)	528
dense_65 (Dense)	(None, 10)	170

Total params: 2,044,570
Trainable params: 2,044,570
Non-trainable params: 0

Our Favorite Models : Model 1

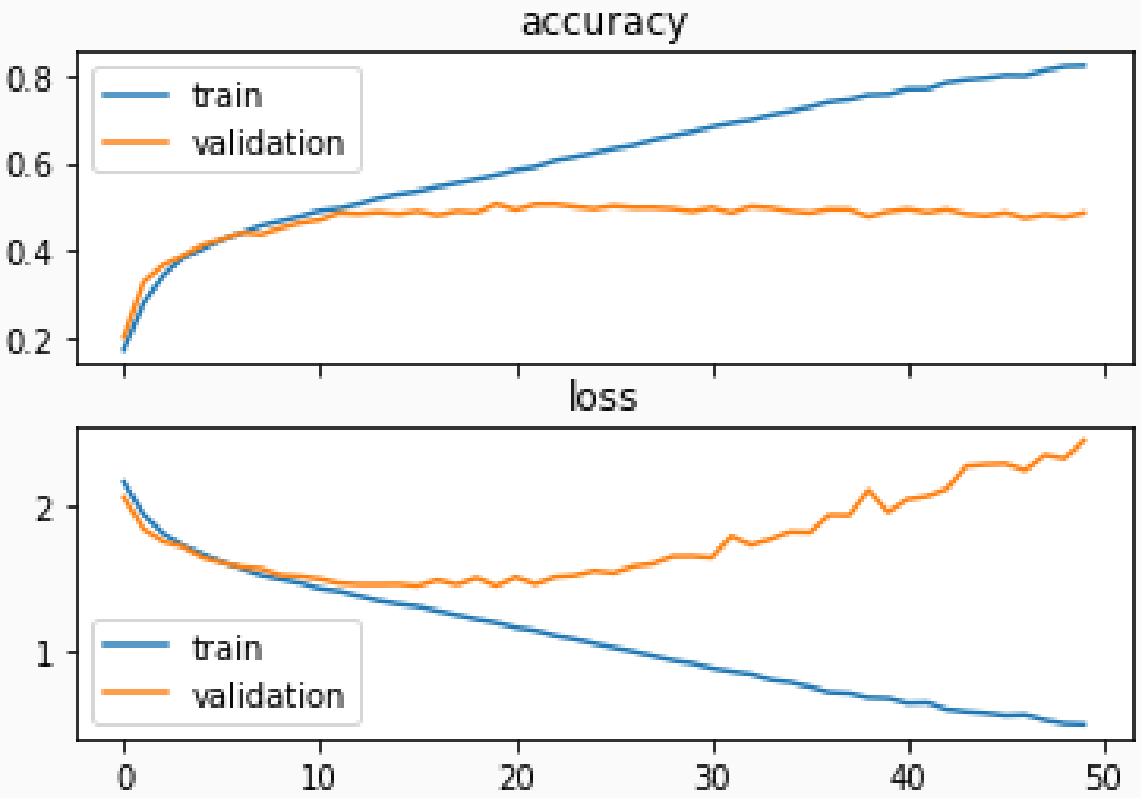

- The final accuracy of model 1 on the test set is 48.44 %,
- Overfitting after 10 epochs: the accuracy on the training set is 82.18 %!

Our Favorite Models : Model 1 + PCA

Layer (type)	Output Shape	Param #
dense_46 (Dense)	(None, 512)	337408
dense_47 (Dense)	(None, 256)	131328
dense_48 (Dense)	(None, 256)	65792
dense_49 (Dense)	(None, 256)	65792
dense_50 (Dense)	(None, 512)	131584
dense_51 (Dense)	(None, 128)	65664
dense_52 (Dense)	(None, 64)	8256
dense_53 (Dense)	(None, 32)	2080
dense_54 (Dense)	(None, 16)	528
dense_55 (Dense)	(None, 10)	170

Total params: 808,602
Trainable params: 808,602
Non-trainable params: 0

Our Favorite Models : Model 1 + PCA

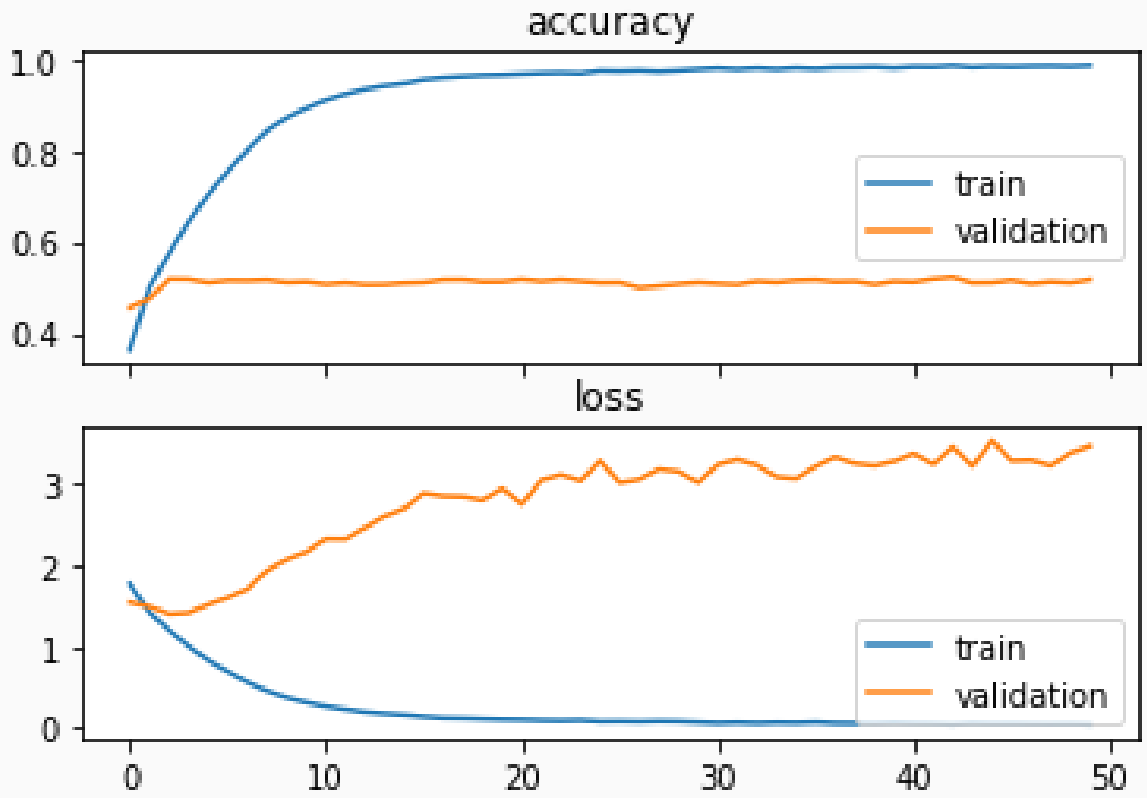

- The final accuracy of model 2 on the test set is 51.67 %
- Overfitting after just 2 epochs: the accuracy on the training set goes up to 98.54 %!

Our Favorite Models : Model 3

Layer (type)	Output Shape	Param #
conv2d_5 (Conv2D)	(None, 32, 32, 32)	896
conv2d_6 (Conv2D)	(None, 32, 32, 32)	9248
max_pooling2d_3 (MaxPooling2	(None, 16, 16, 32)	0
dropout_5 (Dropout)	(None, 16, 16, 32)	0
conv2d_7 (Conv2D)	(None, 16, 16, 64)	18496
conv2d_8 (Conv2D)	(None, 16, 16, 64)	36928
max_pooling2d_4 (MaxPooling2	(None, 8, 8, 64)	0
dropout_6 (Dropout)	(None, 8, 8, 64)	0
conv2d_9 (Conv2D)	(None, 8, 8, 128)	73856
conv2d_10 (Conv2D)	(None, 8, 8, 128)	147584
max_pooling2d_5 (MaxPooling2	(None, 4, 4, 128)	0
dropout_7 (Dropout)	(None, 4, 4, 128)	0
flatten_3 (Flatten)	(None, 2048)	0
dense_5 (Dense)	(None, 128)	262272
dropout_8 (Dropout)	(None, 128)	0
dense_6 (Dense)	(None, 10)	1290

Total params: 550,570
Trainable params: 550,570
Non-trainable params: 0

Our Favorite Models : Model 3

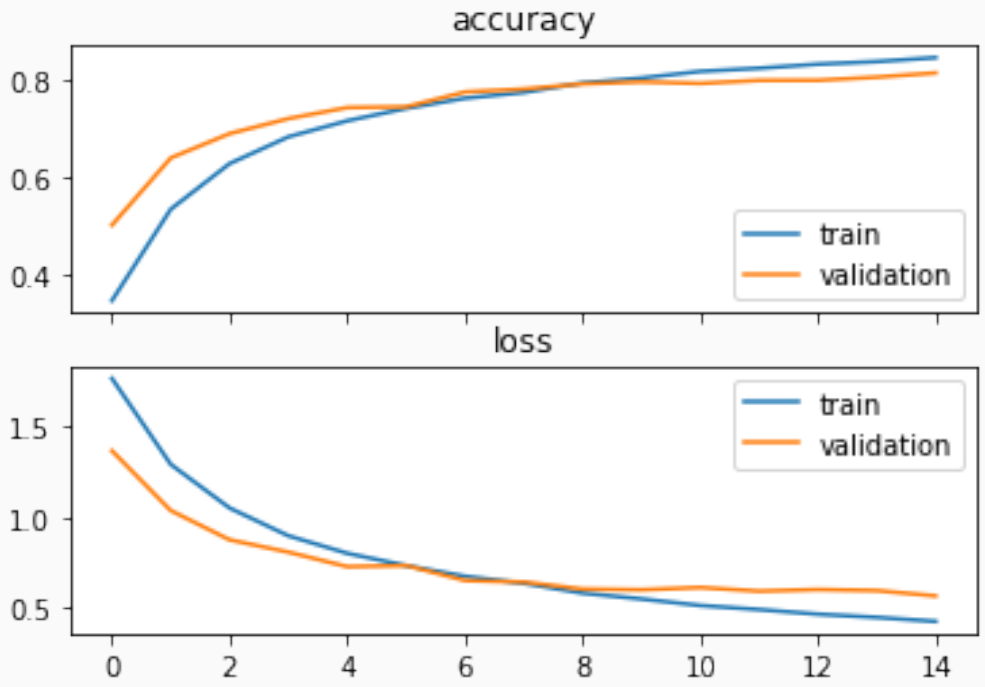

- The final accuracy of model 3 on the test set is 81.580 % !!,
- The overfitting phenomenon is very mild: just 3 percentage points difference between training and test set.

Into the Wild

We now test Model 3 on images of our pets.

Well-classified: Argo 1/2

Well-classified: Argo 2/2

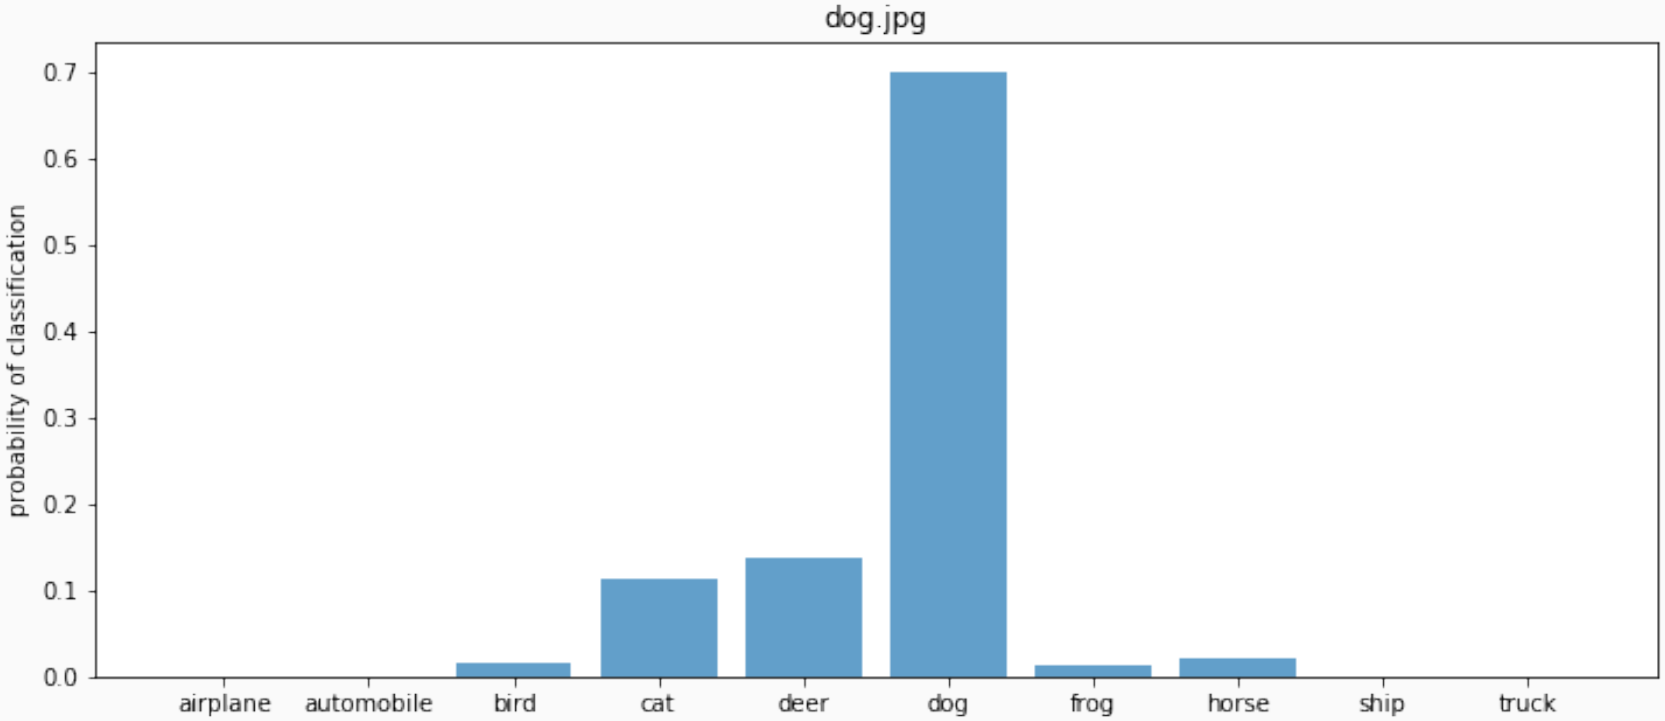

Mis-classified: Bijoux 1/4

Mis-classified: Bijoux 2/4

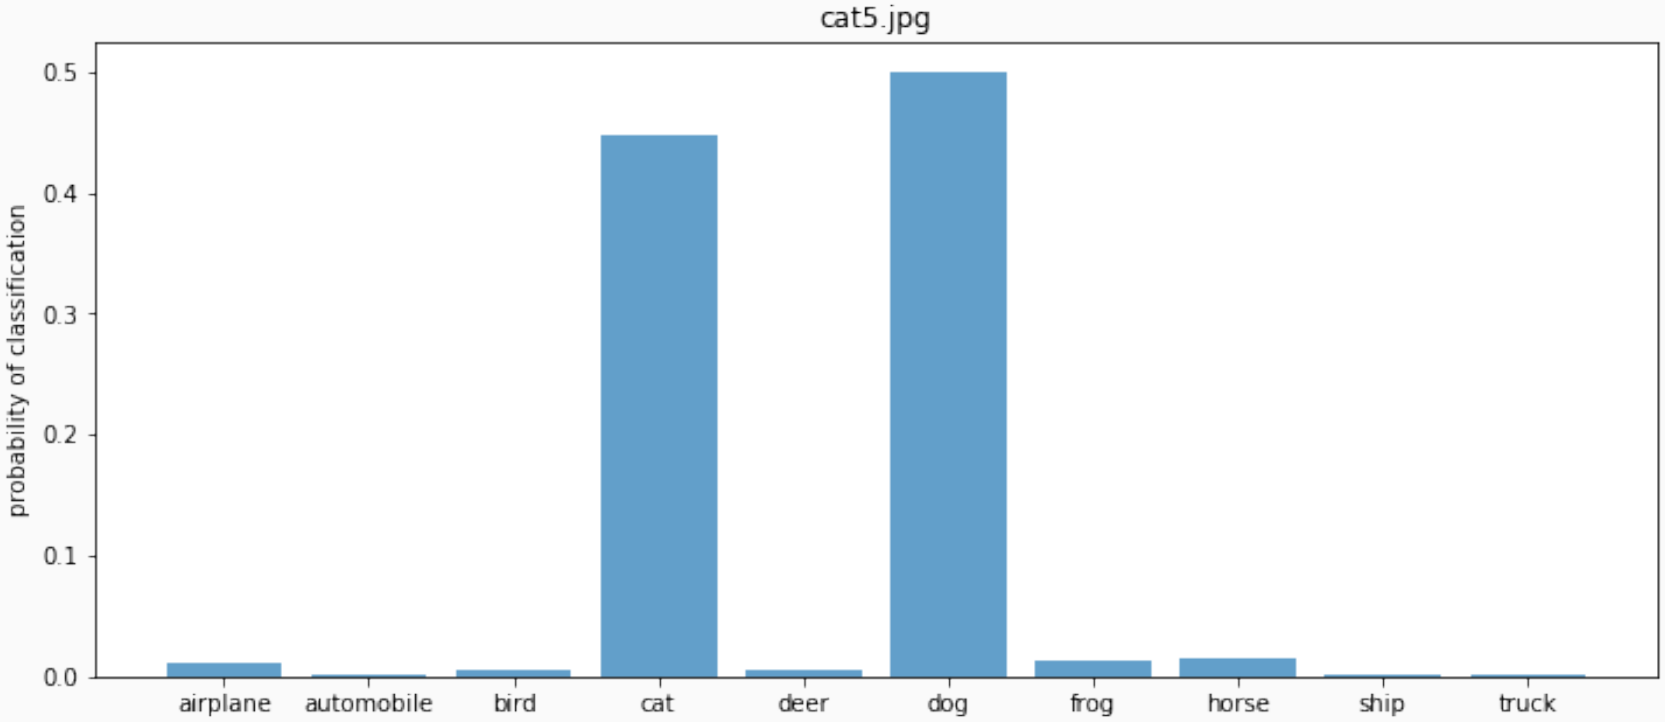

Mis-classified: Zoe 3/4

Mis-classified: Zoe 4/4

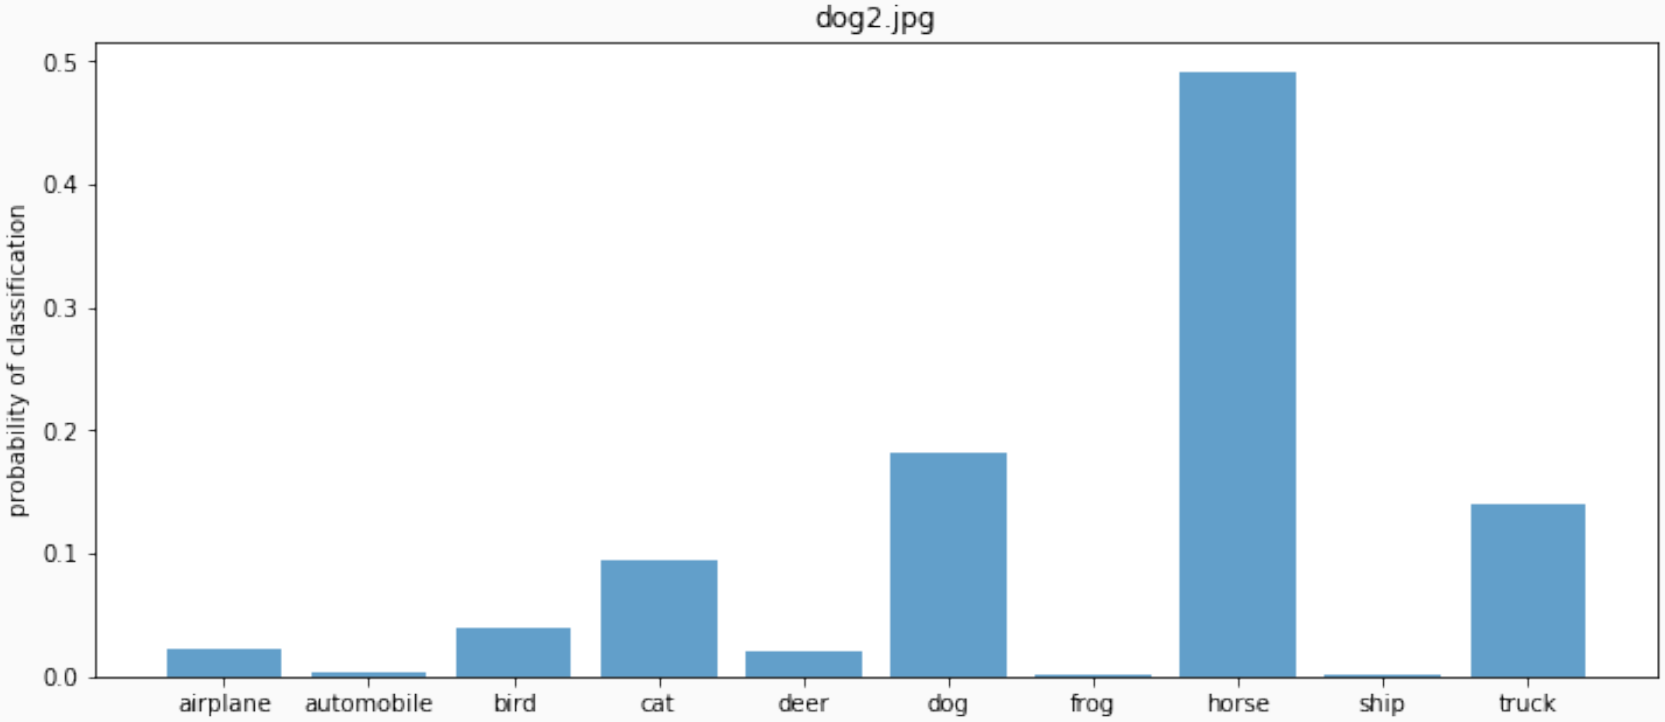

Thanks for your attention!