Foundations of Deep Learning
Davide Giardini – mat. 897473

Flowers Classification

with

Convolutional Neural Networks

Data Preparation

Original dataset

Daisy: 799

Dandelion: 699

Roses: 641

Sunflowers: 898

Tulips: 633

Batch size: 32

20% test set

80%

20% validation set

80% train set

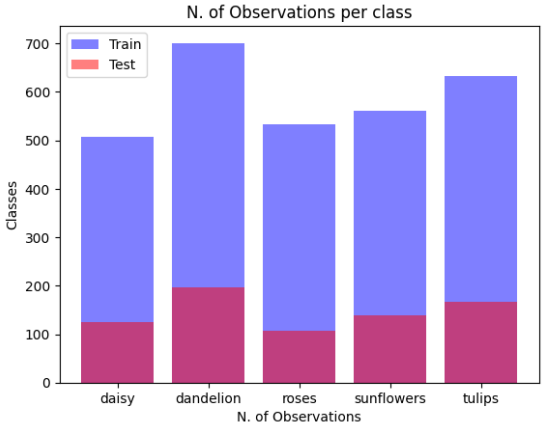

Final:

- 64% train
- 16% validation
- 20% test

First (naive) Model

naive

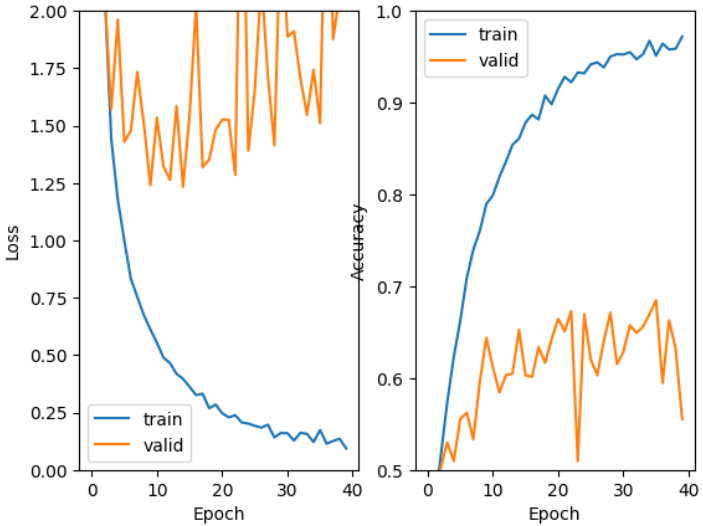

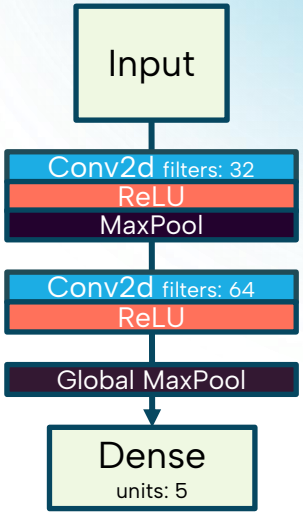

Optimizer: RMSprop
LR: 0,001
Epochs: 40

Train: 80%
Val: 20%

Model v2 (BatchNorm)

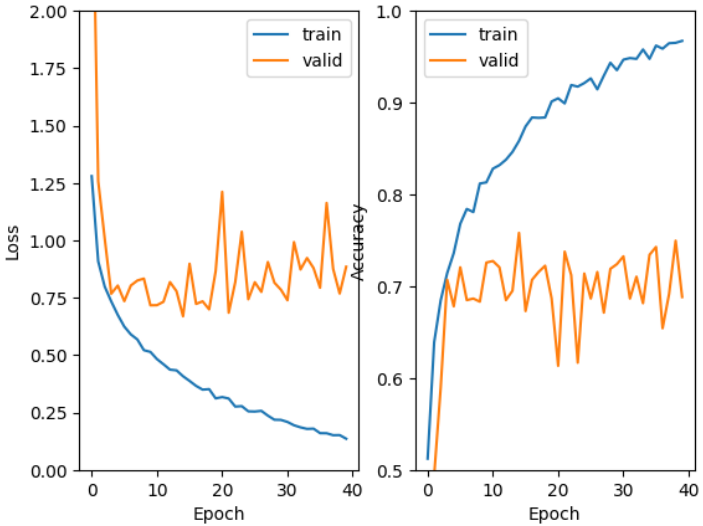

Optimizer: RMSprop
LR: 0,001
Epochs: 40

Train: 80%
Val: 20%

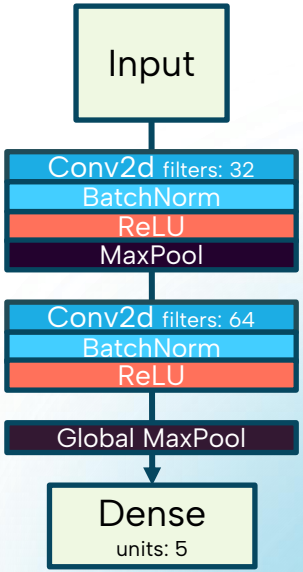

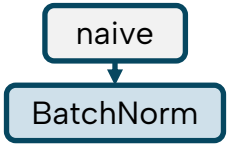

Model v3 (Data Augmentation)

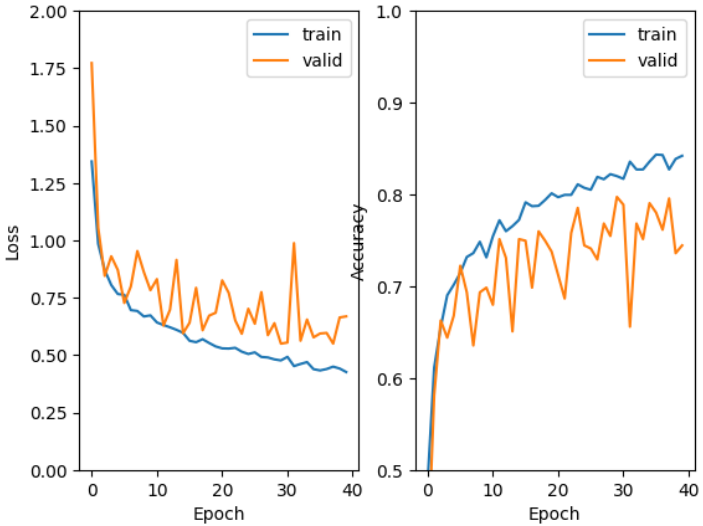

Optimizer: RMSprop
LR: 0,001
Epochs: 40

Train: 80%
Val: 20%

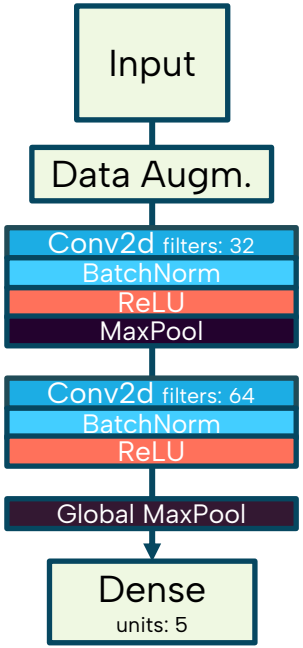

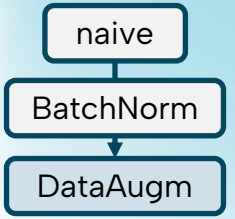

Model v3 (Lower LR)

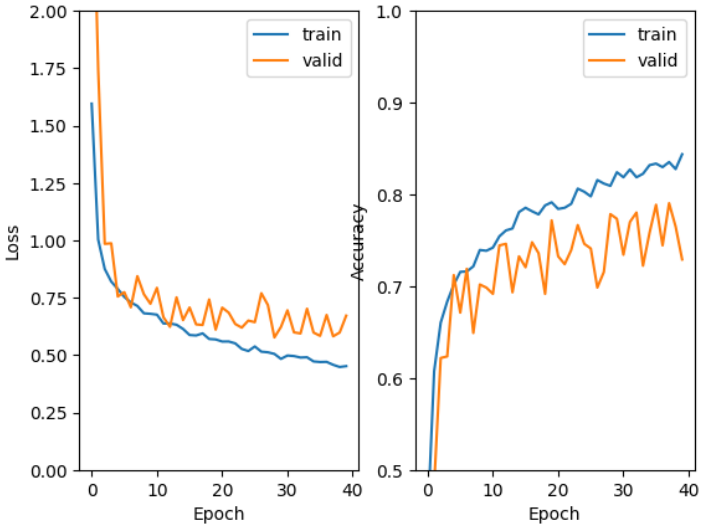

Optimizer: RMSprop
LR: 0,0005
Epochs: 40

Train: 80%
Val: 20%

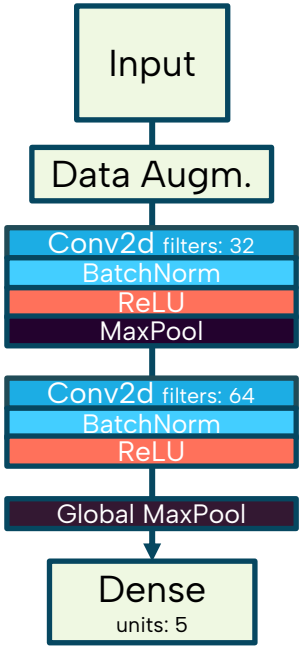

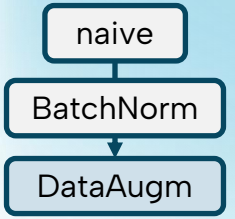

Model v3 (30% val set)

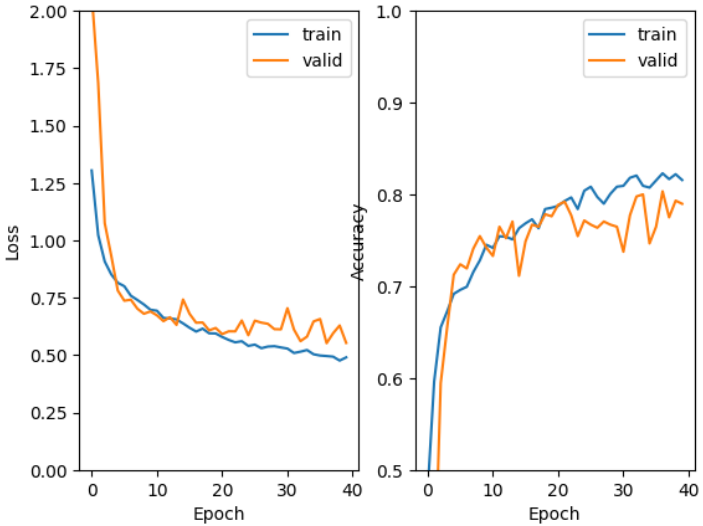

Optimizer: RMSprop
LR: 0,0005
Epochs: 40

Train: 70%
Val: 30%

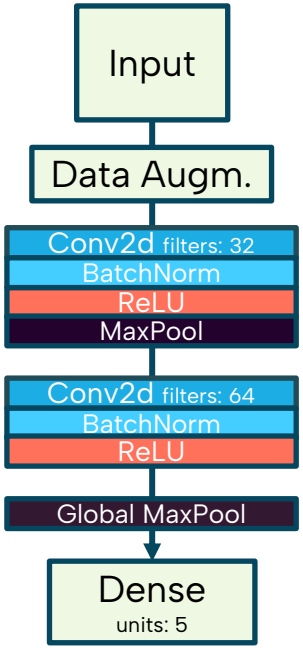

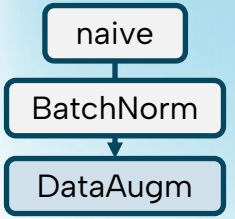

Model v4 (deeper)

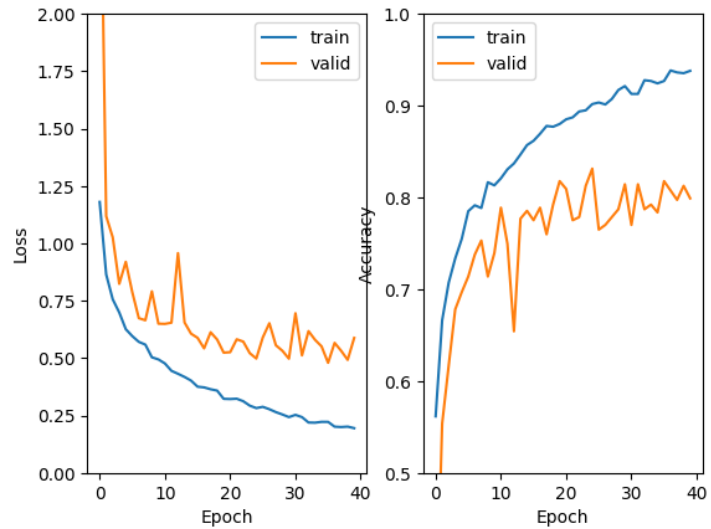

Optimizer: RMSprop
LR: 0,0005
Epochs: 40

Train: 70%
Val: 30%

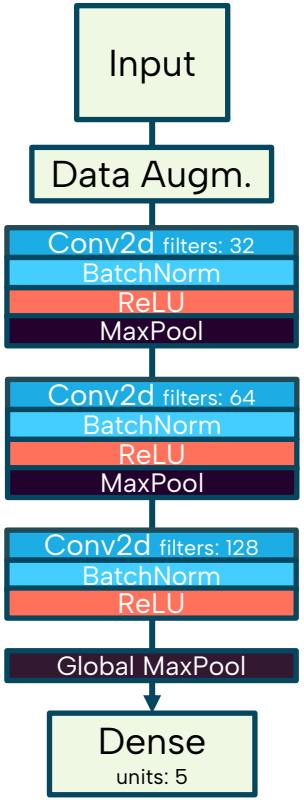

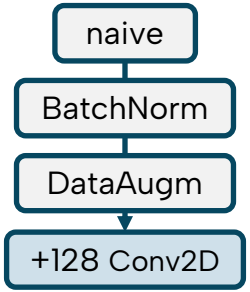

Model v4 (TimeDecay)

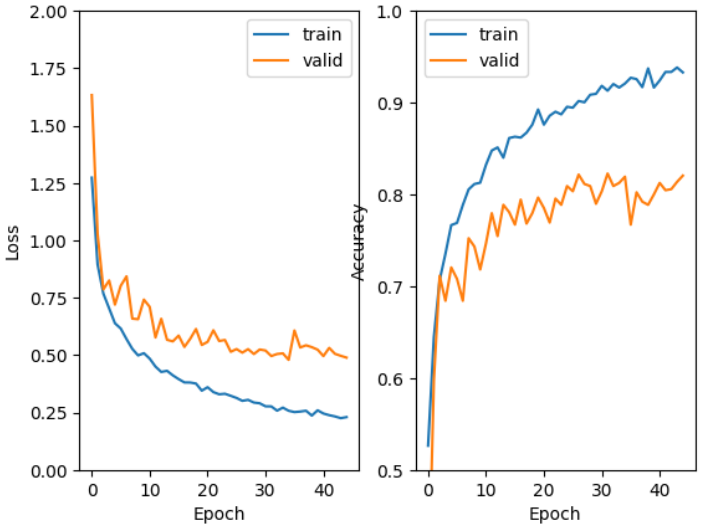

Optimizer: RMSprop
LR: TD (0,0005, 65, 0,08)
Epochs: 40 **Callback**: 10

Train: 70%
Val: 30%

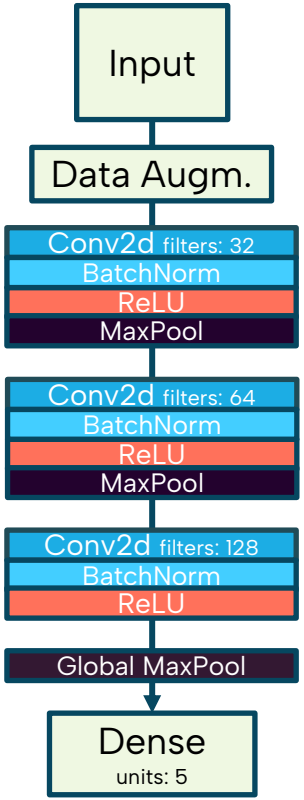

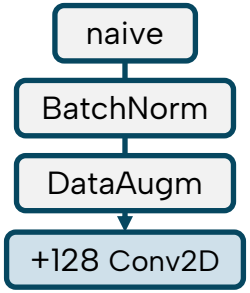

Model v4 (Adam)

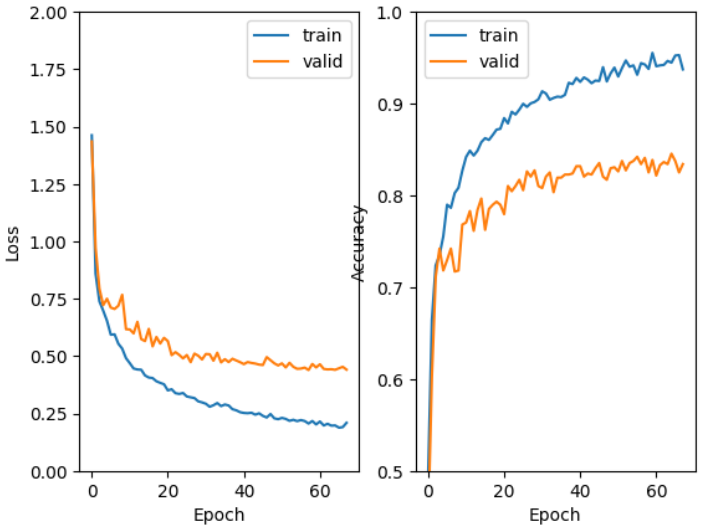

Optimizer: Adam
LR: TD (0,0005, 65, 0,08)
Epochs: 40 **Callback**: 10

Train: 70%
Val: 30%

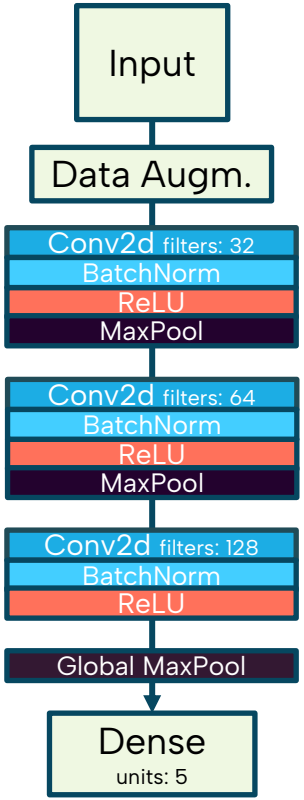

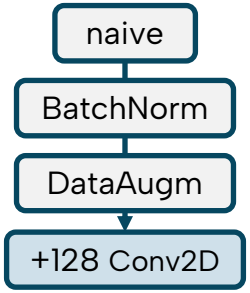

Model v5 (DropOut)

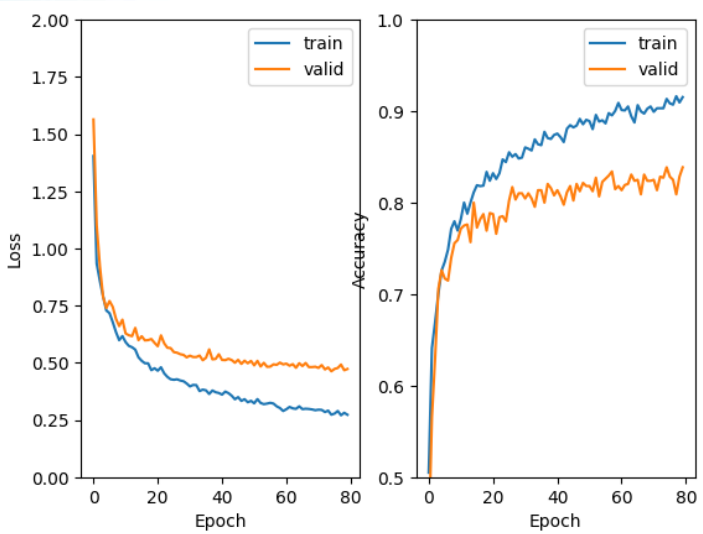

Optimizer: Adam
LR: TD (0,0005, 65, 0,08)
Epochs: 40 **Callback**: 10

Train: 70%
Val: 30%

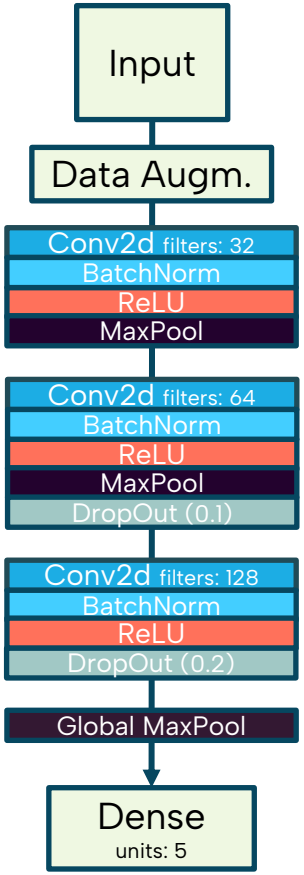

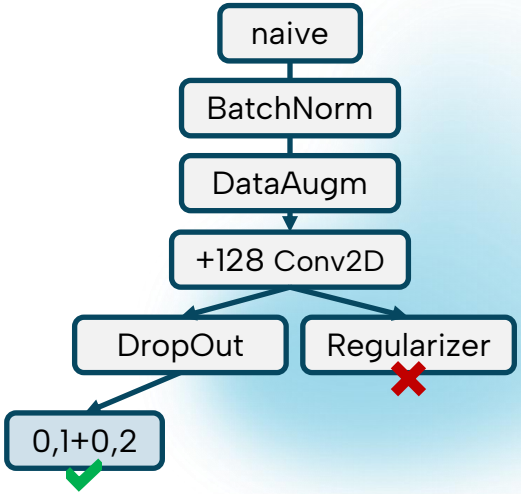

Model v6 (DropOut)

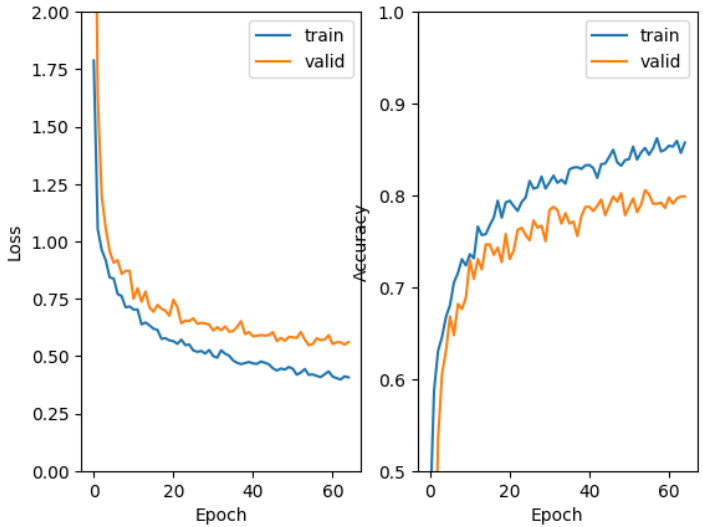

Optimizer: Adam
LR: TD (0,0005, 65, 0,08)
Epochs: 40 **Callback**: 10

Train: 70%
Val: 30%

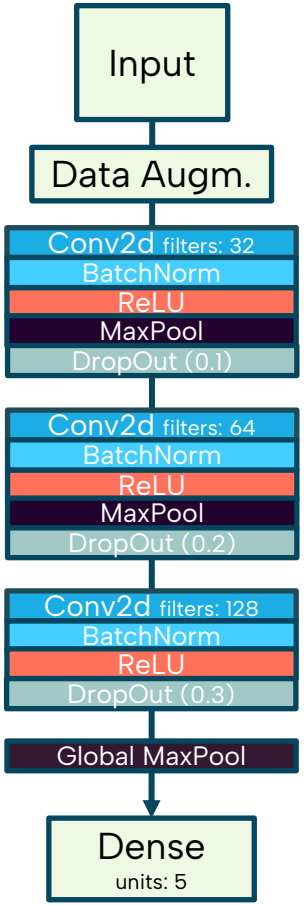

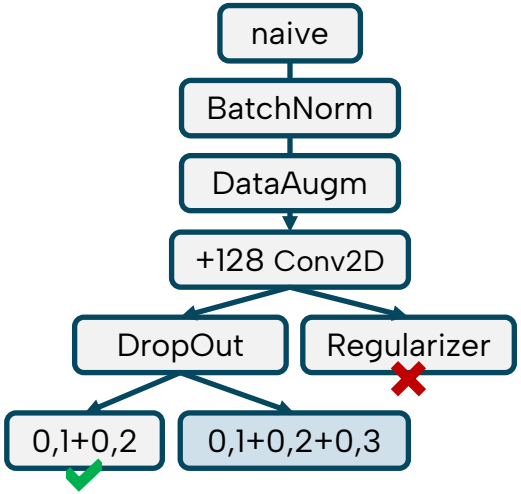

Model v7 (Deeper)

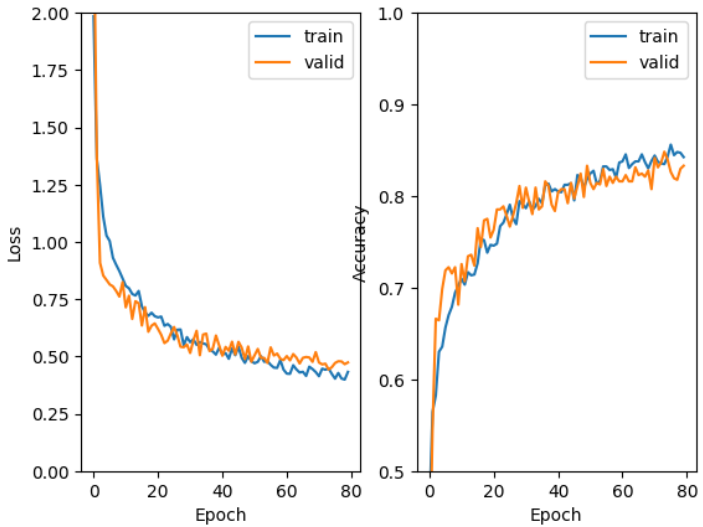

Optimizer: Adam
LR: TD (0,0005, 65, 0,08)
Epochs: 40 **Callback**: 10

Train: 70%
Val: 30%

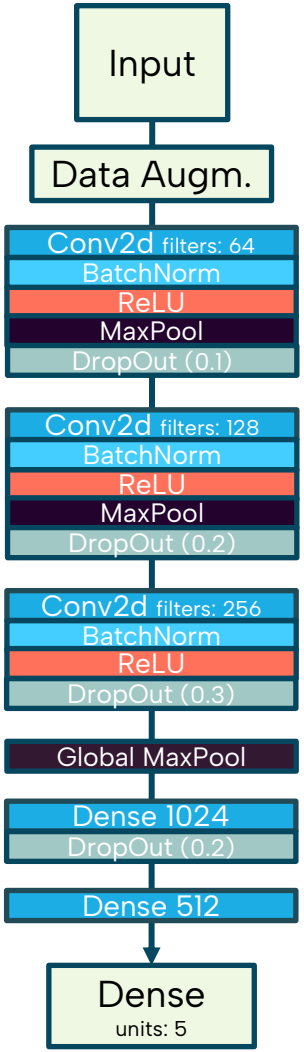

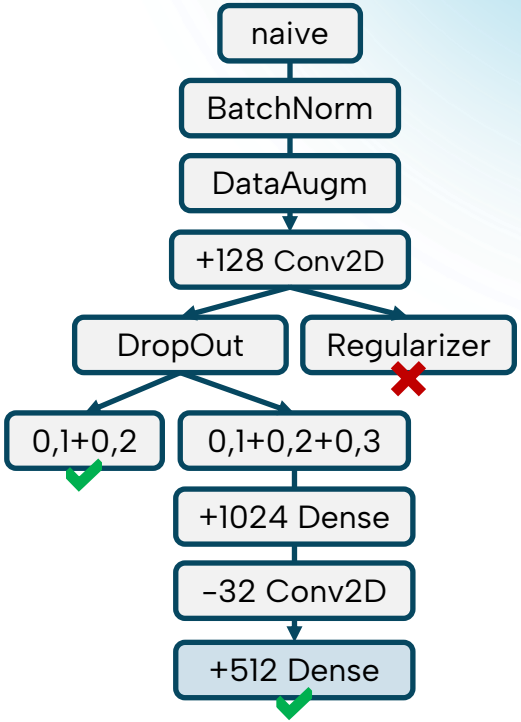

Most of the literature on flower recognition with CNN (Y Liu, 2016; I Gogul, 2017; T Nguyen) cites background as one of the biggest problem in flower classification. Let's try to tackle this issue using segmentation.

Segmentation

Lang SAM

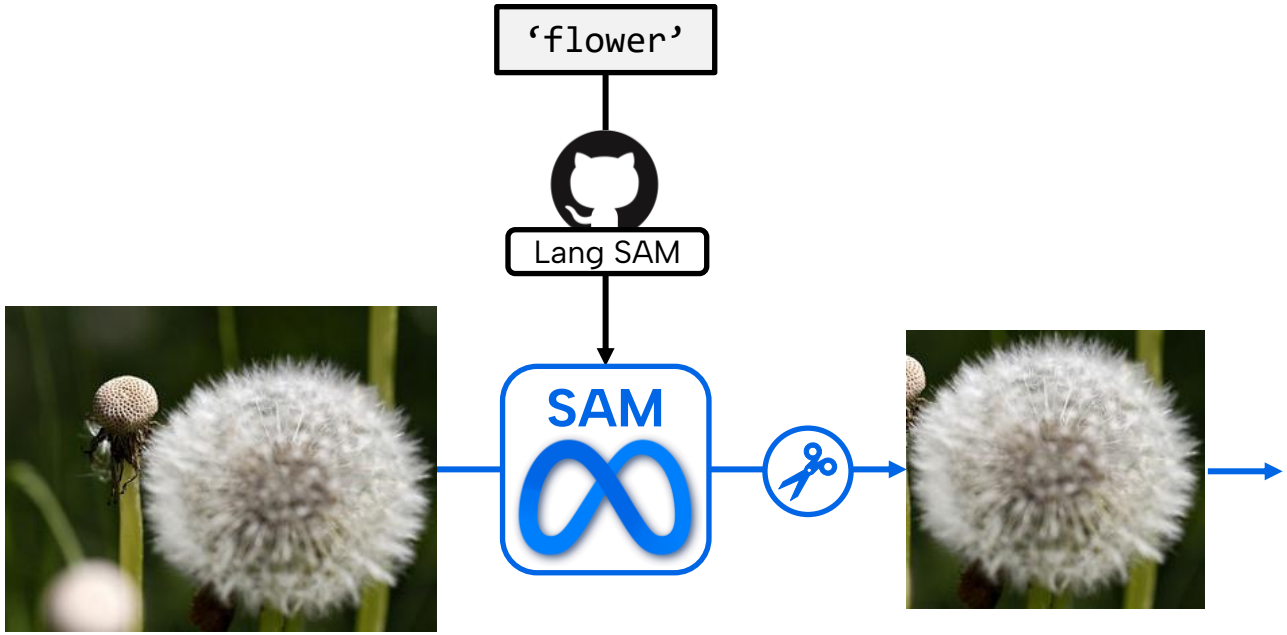
'flower'
Lang SAM
SAM

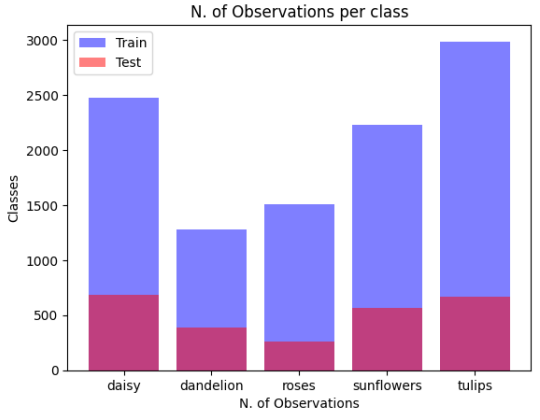
N. of Observations per class
Train
Test
3000
2500
2000
1500
1000
500
0
Classes
daisy
dandelion
roses
sunflowers
tulips
N. of Observations

Balancing
the data

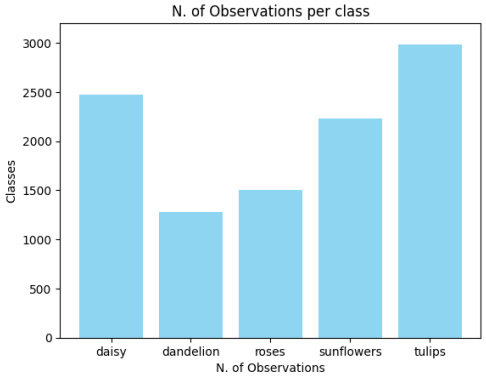

Subset all the classes to the dimension of the smallest class (dandelion: 1275)

The subset is made so that for each of the original photos there is at least one segmented photo within the new set

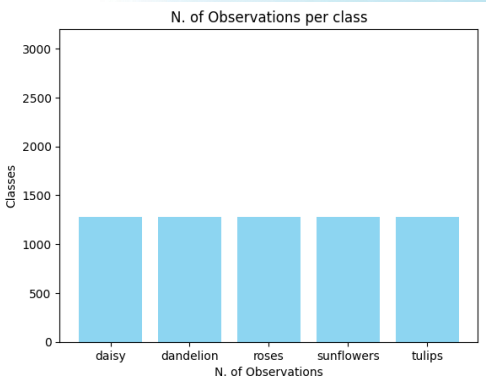

Test Sets

Test

The original test set with uncropped images. It is going to be used to evaluate the model performances to generalized even to the raw, unsegmented, original images.

Test_s

The unbalances test set with **s**egmented images. It is going to be used to evaluate the model performances on all the flowers that the segmented model has detected.

Problem: some original images (those with more flowers in them) will appear more times than other in the test set.

Test_sb

A **s**egmented and **b**alanced test set. Composed of one segmented image for each original picture. This resolves the problem of the unbalanced test, but it does so by selecting one random segmented image from each original picture, and this adds randomness.

Model v5 (before)

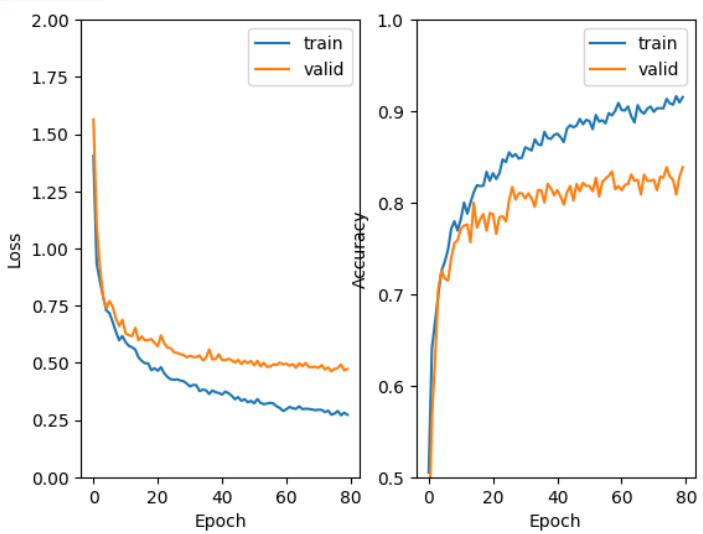

Accuracy:

Test 1 : 83,65%
Test 2: - %
Test 3: - %

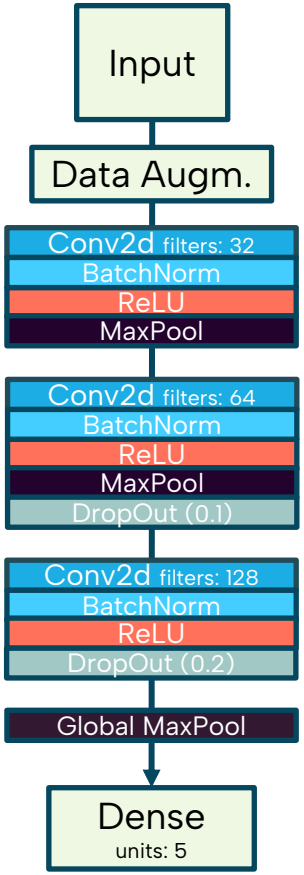

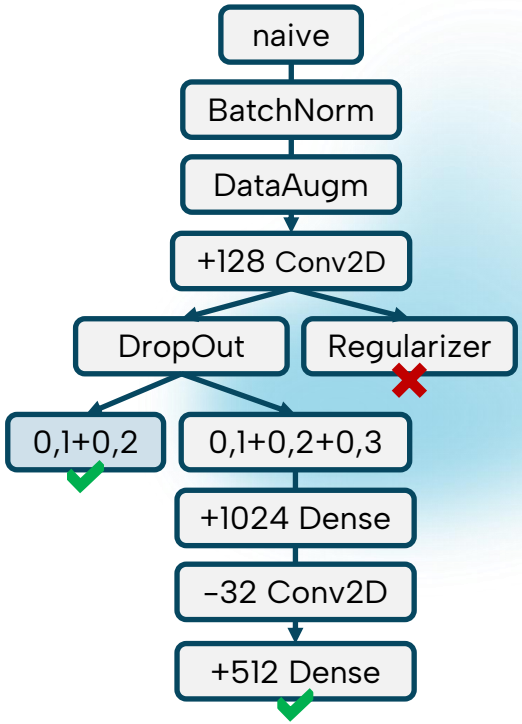

Segmented
unbalanced data

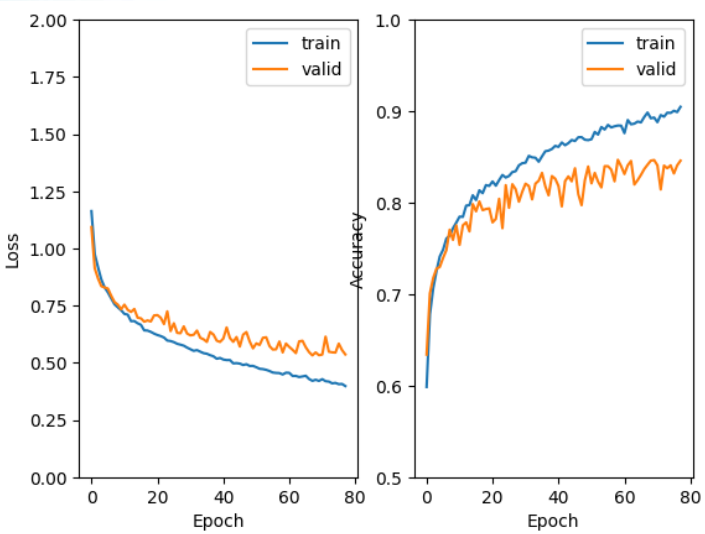

Accuracy:

Test 1 : 79,56%
Test 2: 80,85%
Test 3: 83,33%

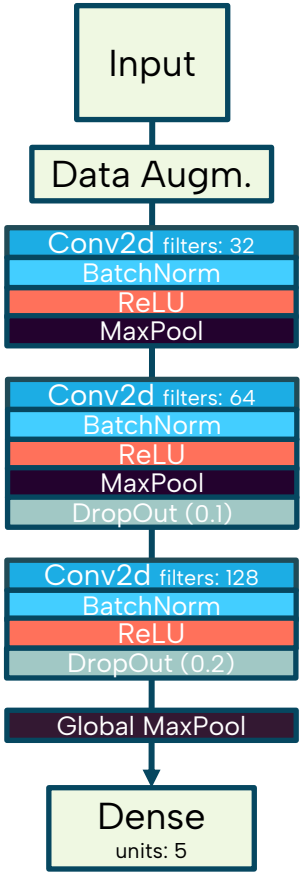

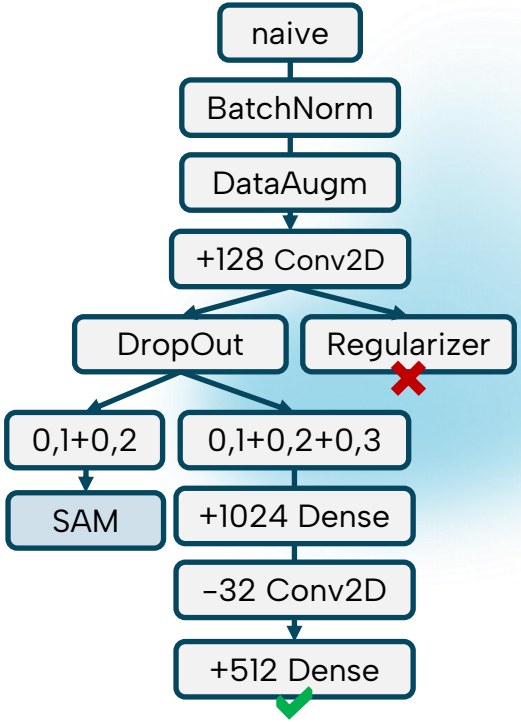

Segmented
balanced data

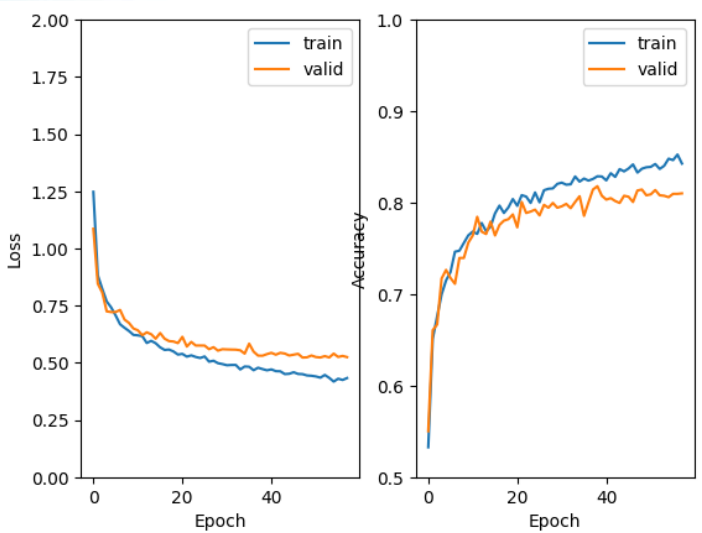

Accuracy:

Test 1 : 80,38%
Test 2: 78,48%
Test 3: 82,79%

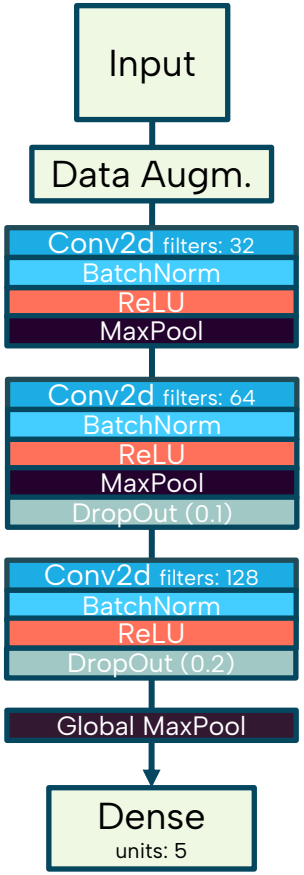

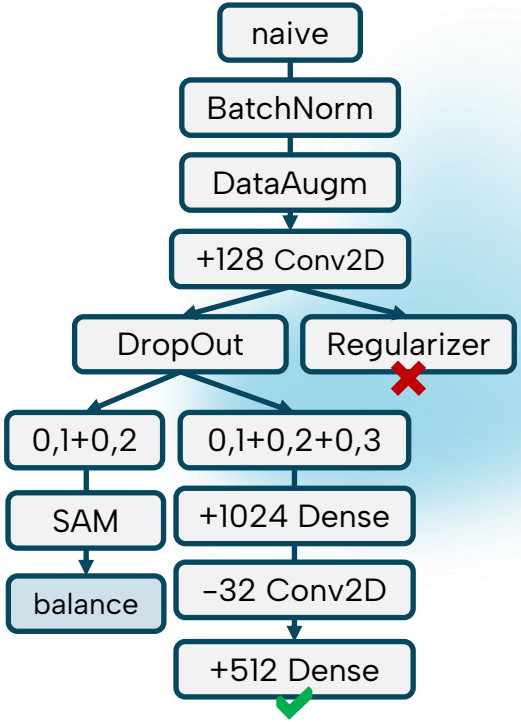

Model v7 (Deeper)

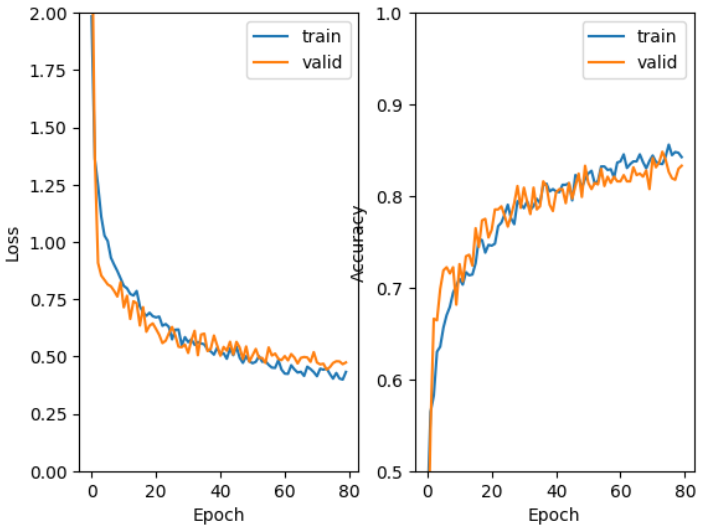

Accuracy:

Test 1 : 84,74%
Test 2: - %
Test 3: - %

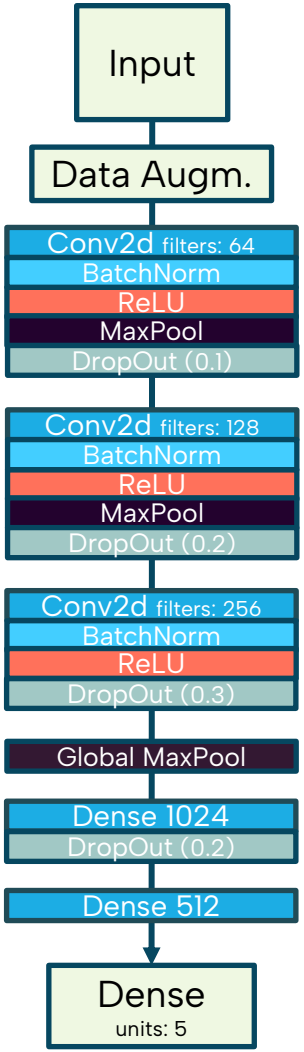

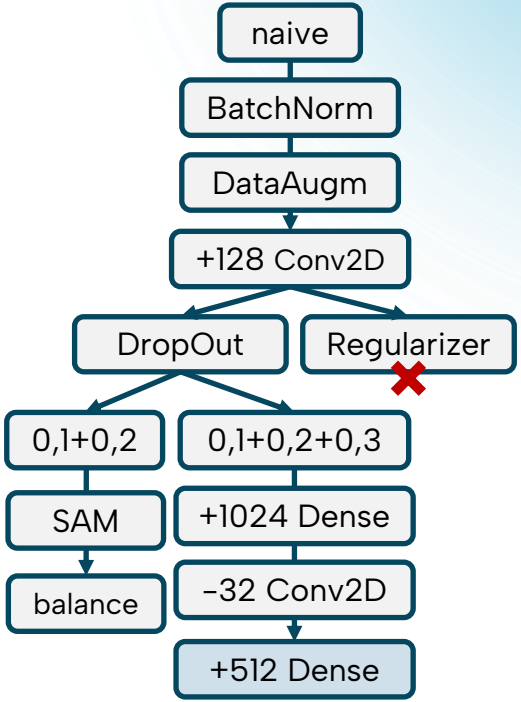

Segmented
unbalanced data

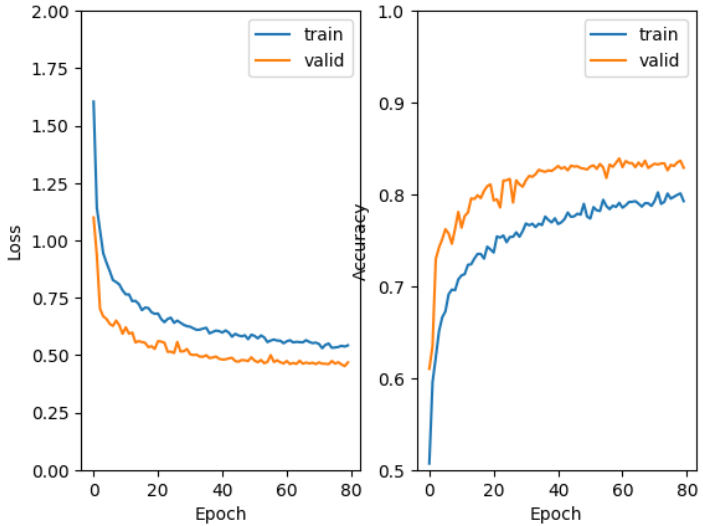

Accuracy:

Test 1 : 81,20%
Test 2: 81,08%
Test 3: 84,70%

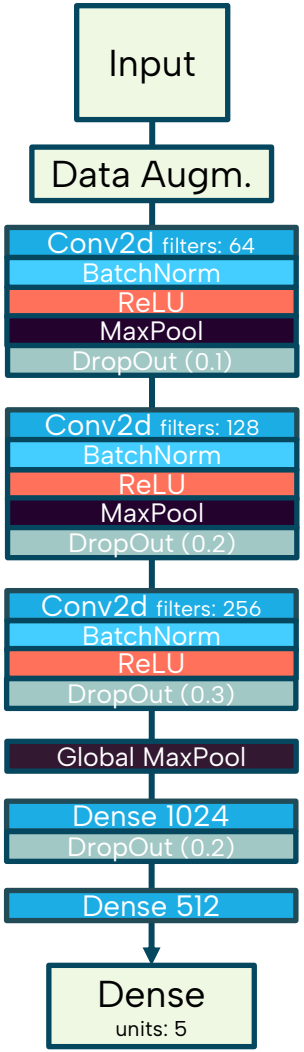

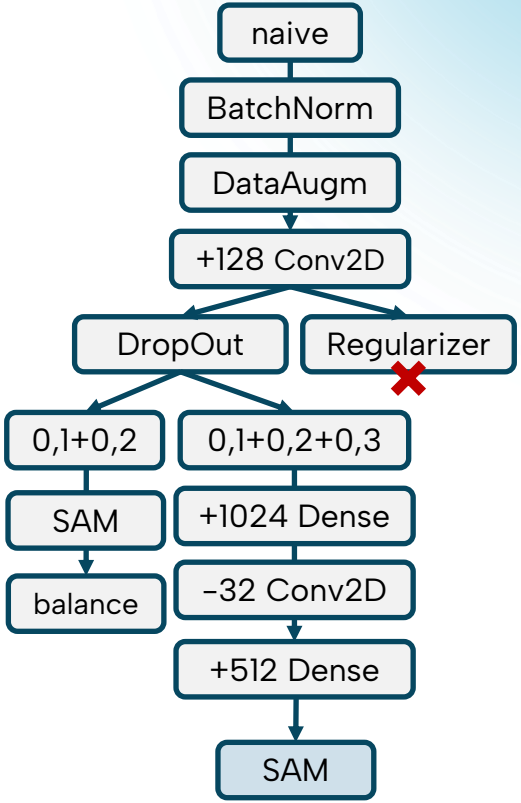

Segmented
unbalanced data

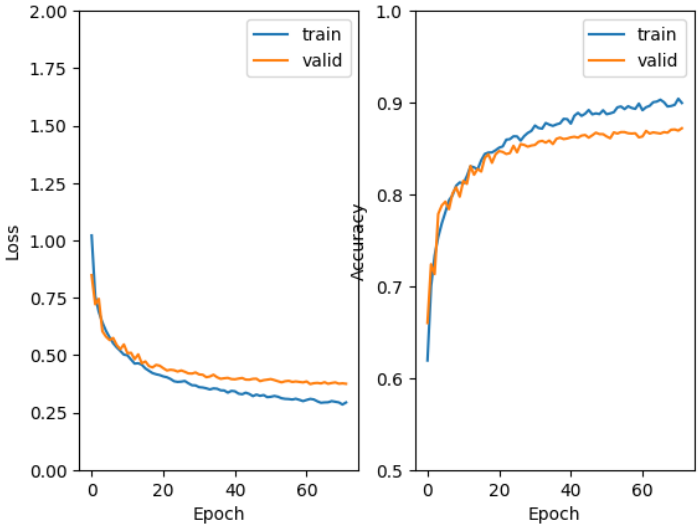

Accuracy:

Test 1 : 82,56%
Test 2: 83,02%
Test 3: 84,84%

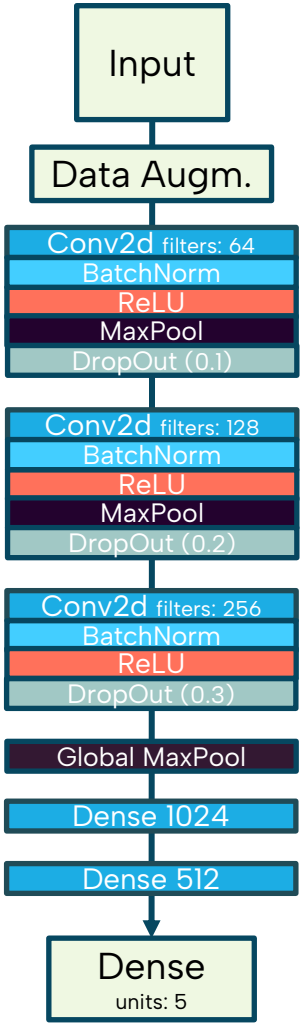

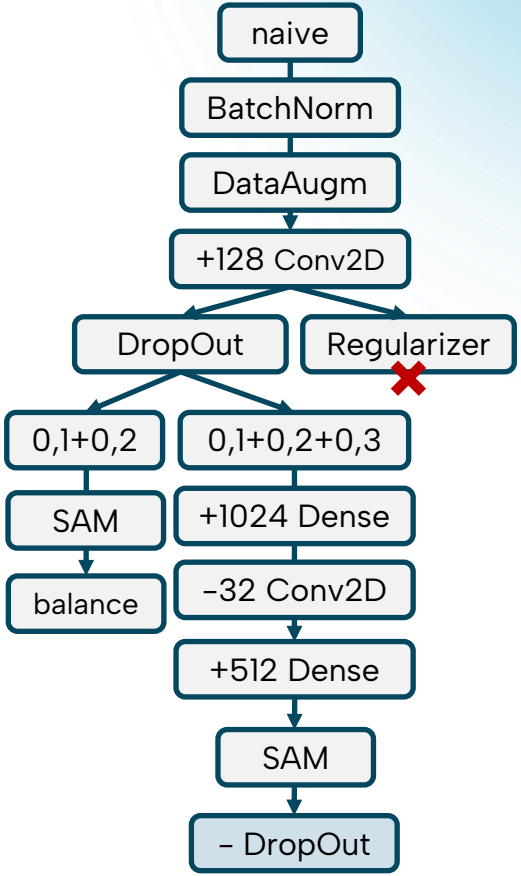

Segmented
balanced data

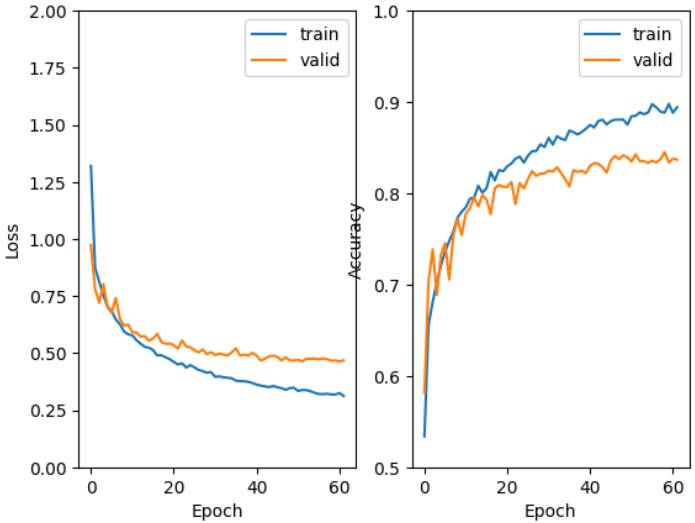

Accuracy:

Test 1 : 82,56%
Test 2: 83,02%
Test 3: 84,84%

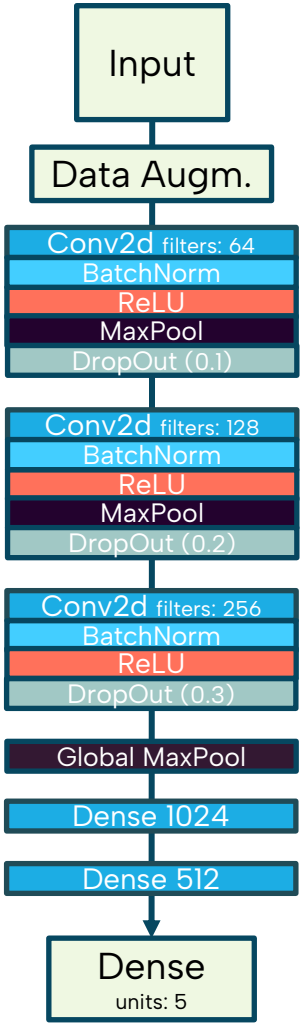

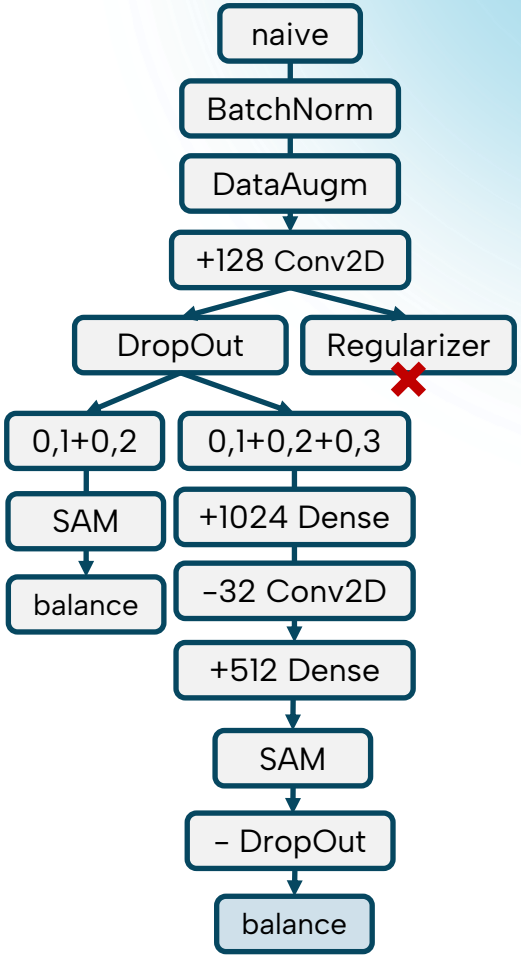

Thank You

Foundations of Deep Learning

Davide Giardini – mat. 897473