

Deep Learning

Pre-requisites:

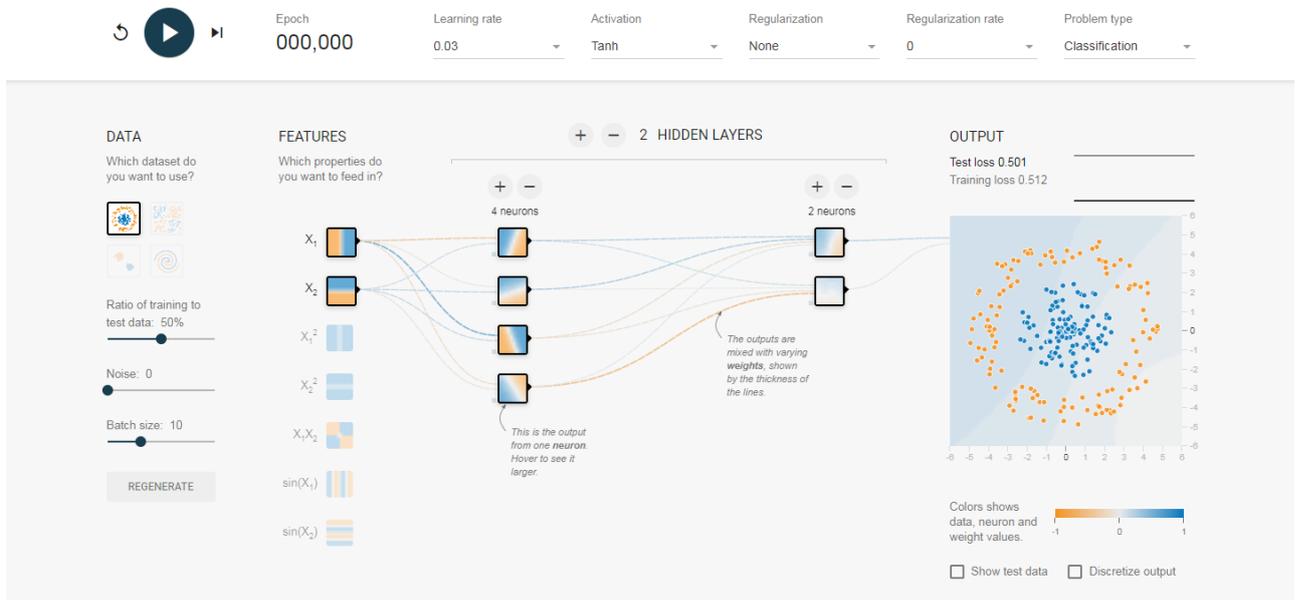
- What is Deep Learning?
- How are Deep Neural Networks different from the Neural Networks discussed in the previous labs?
- Define: learning rate, activation, regularization.

Problem definition:

Classify points based on their location in a 2D space, using Deep Neural Networks

Tool:

As exemplification, we use the [TensorFlow Playground](#), an online platform where one can train a Neural Network from a web browser, without any additional resources or the need to write code.



Task 1. Simple classification – 2 sets of points (Gaussian), in clear, easily distinguishable clusters

Set the following values:

Learning rate: 0.03

Activation: ReLU

Regularization: None

Regularization rate: 0

Problem type: Classification

Ratio of training to test data: 50%

Noise: 0

Batch size: 10

Hidden layers: 1

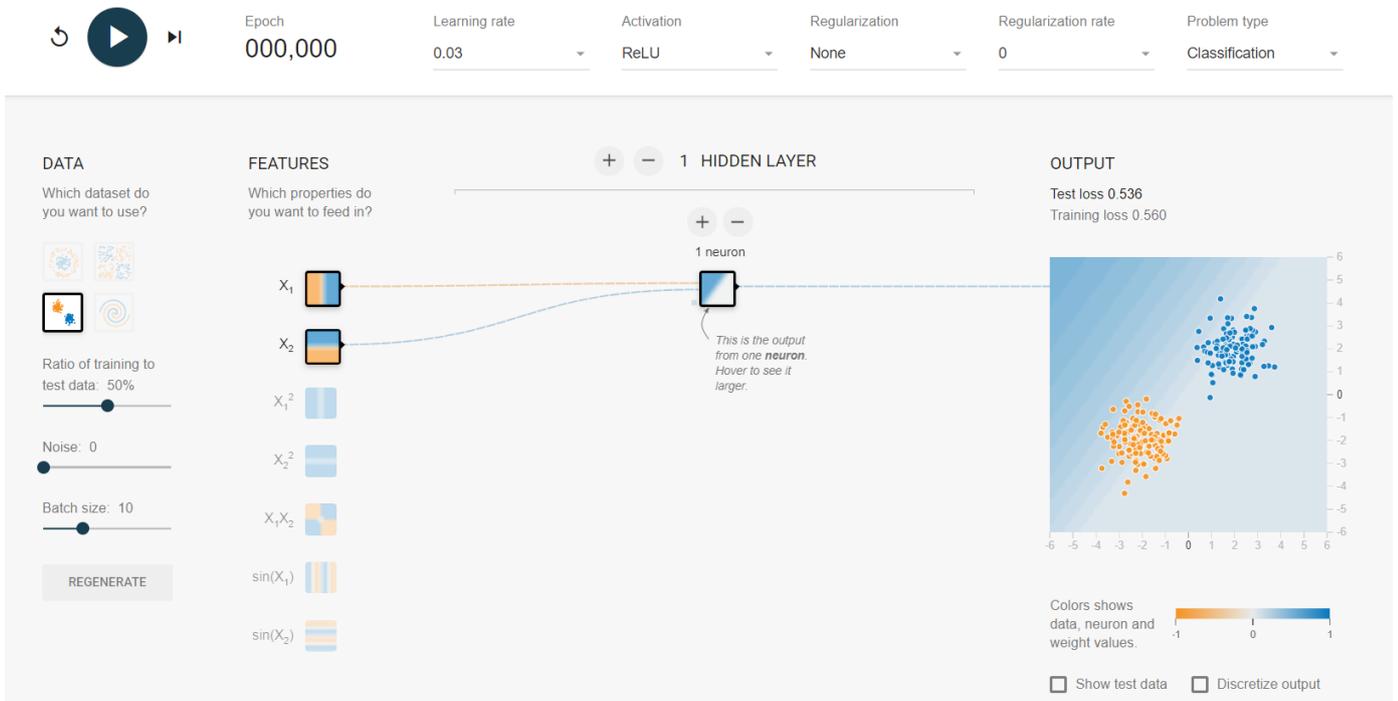
No. of neurons per hidden layer: 1

Start the classification. How long does it take to make the classification?

Further study: Change the activation function. Does it influence the results?

Design challenge:

Add some noise in the data set. Restart. Does the classification take longer? Is the network still able to classify the data?



Task 2. Non-linear classification – 2 sets of points, orange points surround blue points (Circle)

Set the following values:

Learning rate: 0.03

Activation: ReLU

Regularization: None

Regularization rate: 0

Problem type: Classification

Ratio of training to test data: 50%

Noise: 0

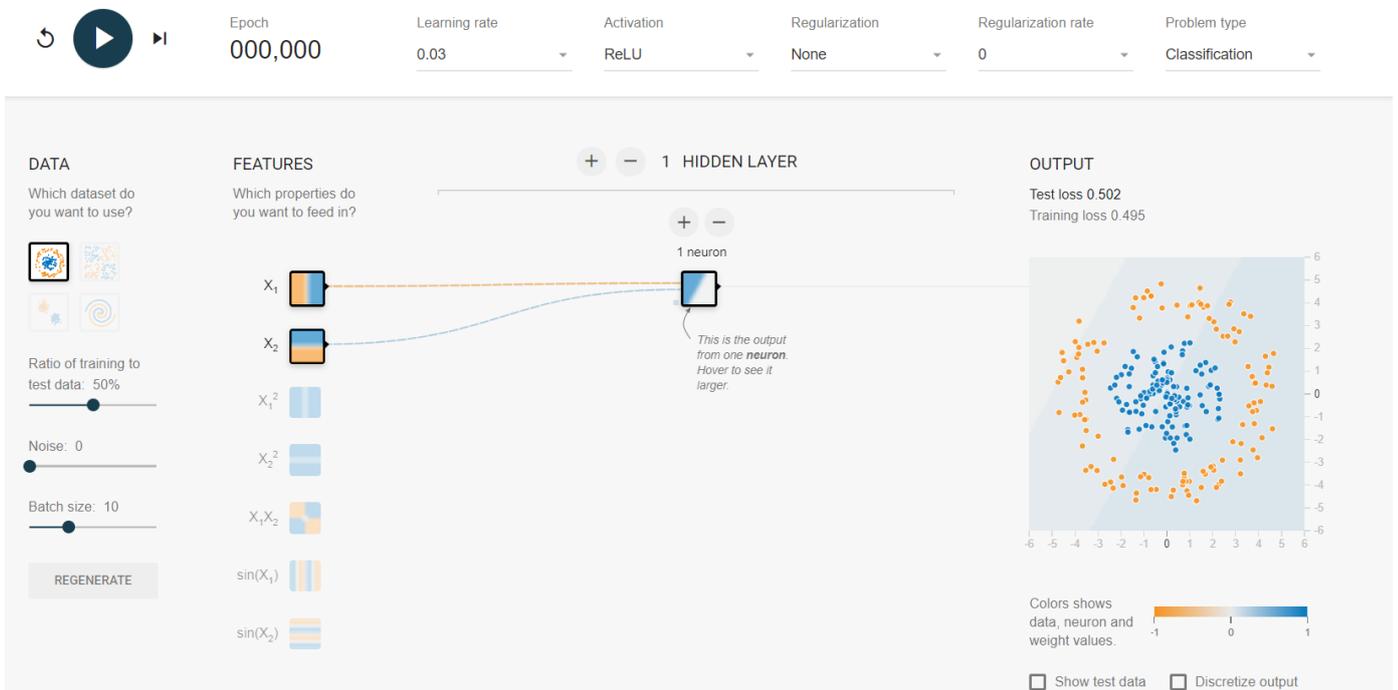
Batch size: 10

Hidden layers: 1

No. of neurons per hidden layer: 1

Start the classification. How long does it take to make the classification? Is 1 neuron enough?

Write down the values for training and test loss, for 1, 2 and 3 neurons respectively, after no more than 200 epochs. Compare the values. Is it worth adding more neurons? Will this make the classification faster?



Task 3. Non-linear classification – 4 clusters of points (Exclusive OR)

Set the following values:

Learning rate: 0.03

Activation: ReLU

Regularization: None

Regularization rate: 0

Problem type: Classification

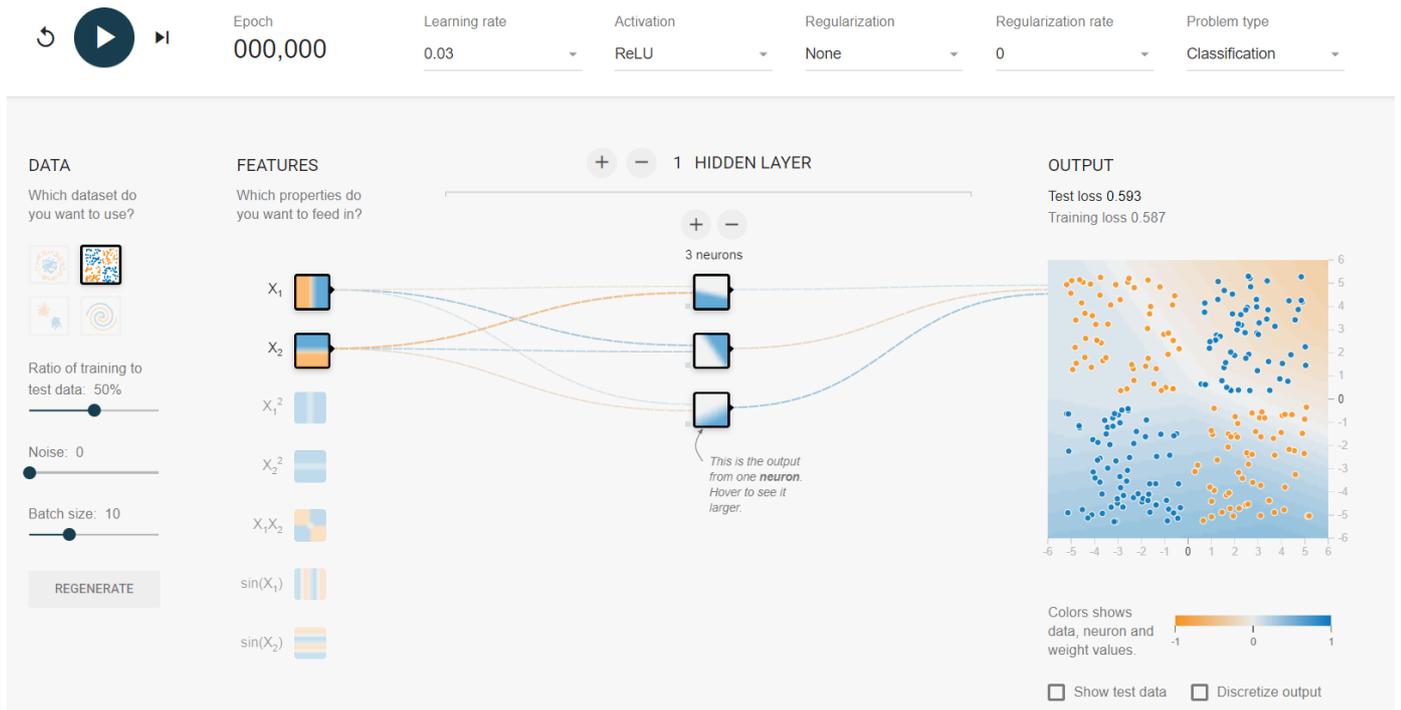
Ratio of training to test data: 50%

Noise: 0

Batch size: 10

Hidden layers: 1

No. of neurons per hidden layer: 3



Run the classification multiple times. Why are the results different?

Design challenge:

What is the minimum number of neurons that provides reliable results?

Task 4. Non-linear classification – points in a spiral (Spiral)

Set the following values:

Learning rate: 0.03

Activation: ReLU

Regularization: None

Regularization rate: 0

Problem type: Classification

Ratio of training to test data: 50%

Noise: 0

Batch size: 10

Hidden layers: 1

No. of neurons per hidden layer: 3

↶ ▶ ▶ Epoch: 000,000 Learning rate: 0.03 Activation: ReLU Regularization: None Regularization rate: 0 Problem type: Classification

DATA

Which dataset do you want to use?



Ratio of training to test data: 50%

Noise: 0

Batch size: 10

REGENERATE

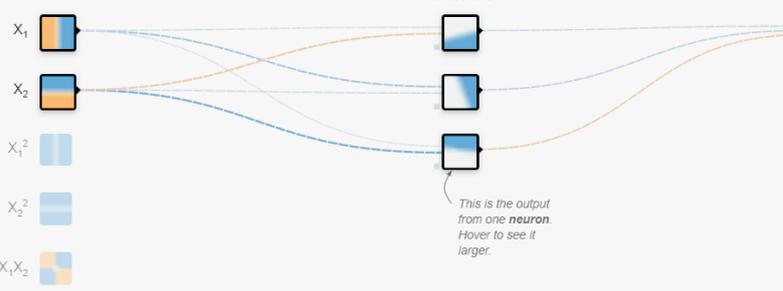
FEATURES

Which properties do you want to feed in?

- X_1
- X_2
- X_1^2
- X_2^2
- X_1X_2
- $\sin(X_1)$
- $\sin(X_2)$

1 HIDDEN LAYER

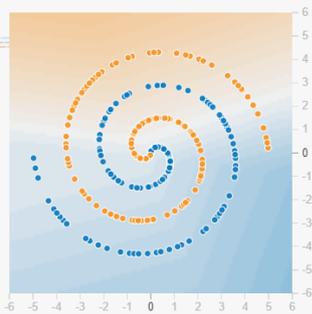
3 neurons



This is the output from one neuron. Hover to see it larger.

OUTPUT

Test loss 0.491
Training loss 0.464



Colors shows data, neuron and weight values.

Show test data Discretize output

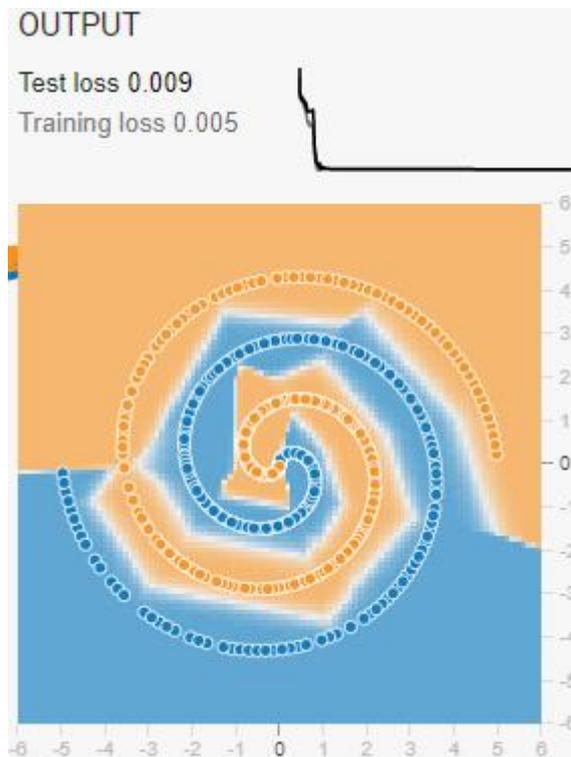
Start the classification. Add more neurons to the hidden layer. Does it work?

Add another layer. Are the results any better?

Gradually increase the number of hidden layers and neurons.

Design challenge:

1. What is the minimum number of hidden layers and neurons that can achieve a Test loss ≤ 0.05 and a Training loss ≤ 0.02 ?
2. Comment on the result below.



Online resources:

1. <https://ru.coursera.org/lecture/deep-learning-business/6-1-introduction-to-tensorflow-playground-ArfBs>
2. <https://towardsdatascience.com/a-visual-introduction-to-neural-networks-68586b0b733b>
3. <https://www.youtube.com/watch?v=ru9dXF04iSE>