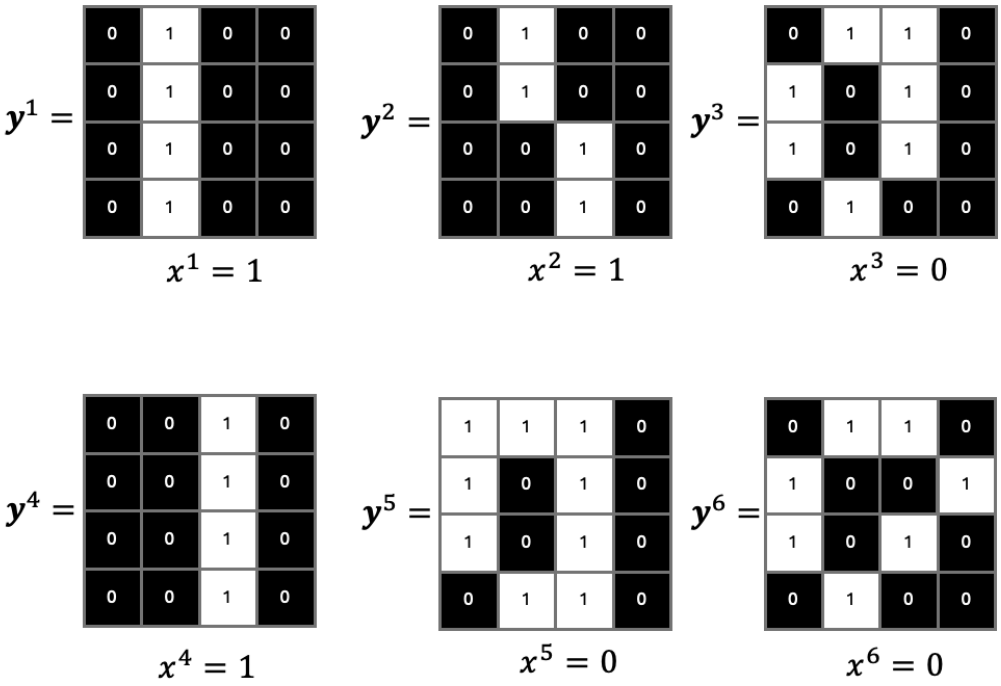
Exercise Sheet 7

Neural Computation
University of Birmingham

Week 7 2025

Exercise 1. You want to train a regression network to take a single digit number x as input and the produce a 4×4 pixel hand drawn image $\hat{\mathbf{y}}$ of that digit. You are using gradient descent and a mean square error loss function to achieve this; averaging the quadratic error between all pixels in the predicted image $\hat{\mathbf{y}}$ and the corresponding target image $\mathbf{y}$. Consider the following set of training pairs.



After training the network for a time, it is able to produce the optimal output (minimising the loss function). What are the outputs of the network for inputs $x = 1$ and $x = 0$ after the network has been trained. Draw the pixel grid and fill in numbers for each of the pixels. Do you see any problems with this approach?

Exercise 2. You are training a GAN on a dataset to generate 256×256 pixel colour images. Assume you are using a batch size of 4. What would be the size of the input tensor and output tensor of the discriminator network?

Exercise 3. Which of the following statements about GANs are true? Explain your answers.

1. The discriminator is only required for training the generator.
2. The generator in a GAN is trained to maximise the discriminator's ability to distinguish real from fake data.
3. The discriminator in a GAN is trained to minimise the probability of correctly classifying real data as real.
4. When using back propagation to train the GAN, to compute the derivatives for the discriminator's parameters we first have to compute the derivatives for the generator's parameters.