Exercise Sheet 8

Neural Computation
University of Birmingham

Week 8 2025

Exercise 1. Consider the following statements about DDPMs. Which of them are true/false? Explain your answer.

1. The forward process of the network is deterministic; meaning when the forward process is applied to an image, the same image at $t = 0$ will always lead to the same noise at $t = T$.
2. The backward process of the network is deterministic; meaning when it is applied to the same noise realisation at $t = T$ it will always lead to the same generated image at $t = 0$.
3. Similarly to a generator network in a GAN, the role of the UNet is to map noise to a clean image.
4. The UNet is deterministic; the same input will always produce the same output.
5. DDPMs are trained using backpropagation.
6. In DDPMs, the noise at $t = T$ corresponds to the latent code in an autoencoder or the noisy input of a GAN generator as it encodes the image in a compressed form.

Exercise 2. A diffusion model is trained to generate high-resolution images of faces. It uses a UNet with two levels (two pooling layers and two upsampling layers). The model frequently produces outputs like the one below, where the colour of the left and right eye does not match. While this can occur in humans it is quite rare. What could be the reason? Suggest architectural changes to address this.

