Generating Credible Fake Images of White Blood Cells

Neutrophils are the most common type of white blood cell and the presence of abnormal variants can be an indication of health issues. Classification of neutrophils is therefore an important task. Some types of abnormalities are quite rare, which means there is not much data to train the classifier with. This is something that could lead to bad performance. In our thesis we seek to find the answers to the questions: Is it possible to generate synthetic images depicting white blood cells with abnormalities and can the inclusion of these in the training set improve the classification performance?

CellaVision is a company established in Lund that produces systems which find blood cells and give an initial classification. The user of the system can then examine the results and change classification if needed. Some serious diseases can only be diagnosed by looking at the distribution of cells in each class. Knowing the proportions of cells of each class is therefore very important. As previously mentioned, some types of cells are quite rare and there may not be a lot of data to train the classifier on. Finding more training data would demand a lot of lab samples as well as countless hours from a medical expert. In our thesis, we have therefore looked at the possibility of using deep learning methods to create synthetic images.

After generating images, their credibility was evaluated. A medical expert and an experienced CellaVision employee were given a set of 100 images, whereof 50 were generated with a diffusion model. They were given the task to classify which ones are real and which ones are not. The medical expert performed well when classifying real images, as 96% of these were classified as real. However, the expert classified less than a third of the generated images as fake and they can therefore be considered credible. The experienced employee classified less than half of the real images as real and less than a fourth of the generated images as fake.

The generated images were evaluated by comparing the performance of a classifier trained on only real images and a classifier trained on a data set where generated images had been added. The results of the tests were inconclusive. There were no obvious improvements but also not a clear decrease in performance. For some abnormalities, the classifier was able to correctly classify more images of abnormal cells, but this came at the cost of more images of normal cells being incorrectly classified. Examples of real cell images as well as our generated images can be seen in Figure 1 and 2.

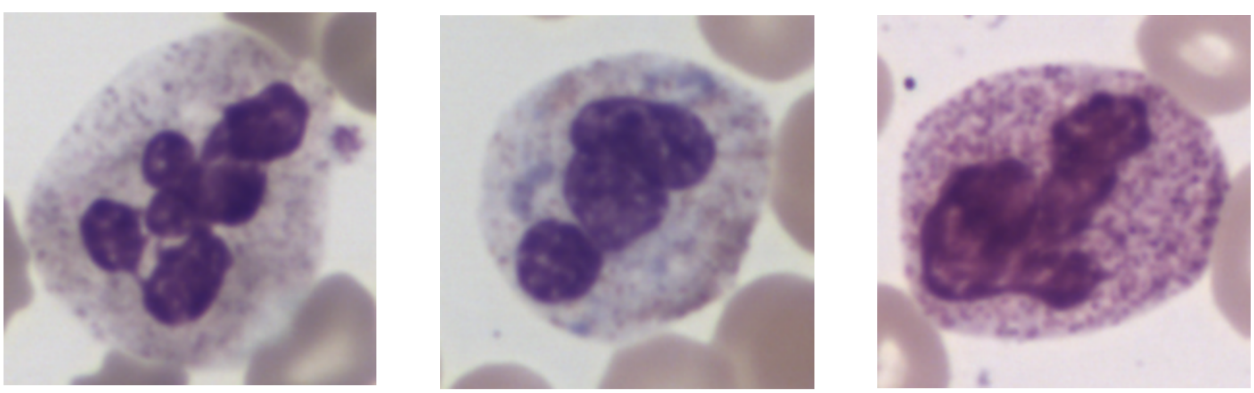

(a) Hypersegmented neutrophil. (b) Neutrophil with Döhle bodies. (c) Hypergranular neutrophil.

Figure 1: Examples of real images of cells with different abnormalities.

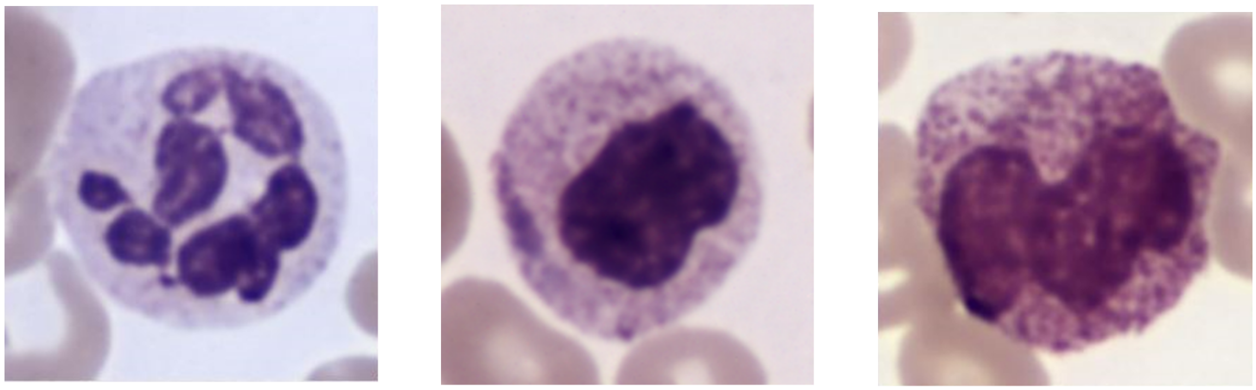

(a) Hypersegmented neutrophil. (b) Neutrophil with Döhle bodies. (c) Hypergranular neutrophil.

Figure 2: Examples of generated images of cells with different abnormalities.

Previously, a master thesis was done at CellaVision with the purpose of generating images. The images were supposed to depict white blood cells and were generated with a network type called GAN. The thesis students reached fairly successful results, but had some remaining problems with image quality. This thesis was done in 2020. Since then, a new type of generative model, the diffusion model, has been introduced to the world.

Diffusion models work by first adding more and more noise to the training images. The model then learns to remove the noise from the images. The intuition of diffusion models can be seen in Figure 3. When the model has been trained, images of the wanted size containing only noise is passed to the model. The model produces new synthetic images, by removing the noise from the input until it resembles the training data. Diffusion models have been shown to produce state of the art images and new improvements are published almost daily. Since the diffusion models seems to be the next big thing, we wanted to use this for our thesis.

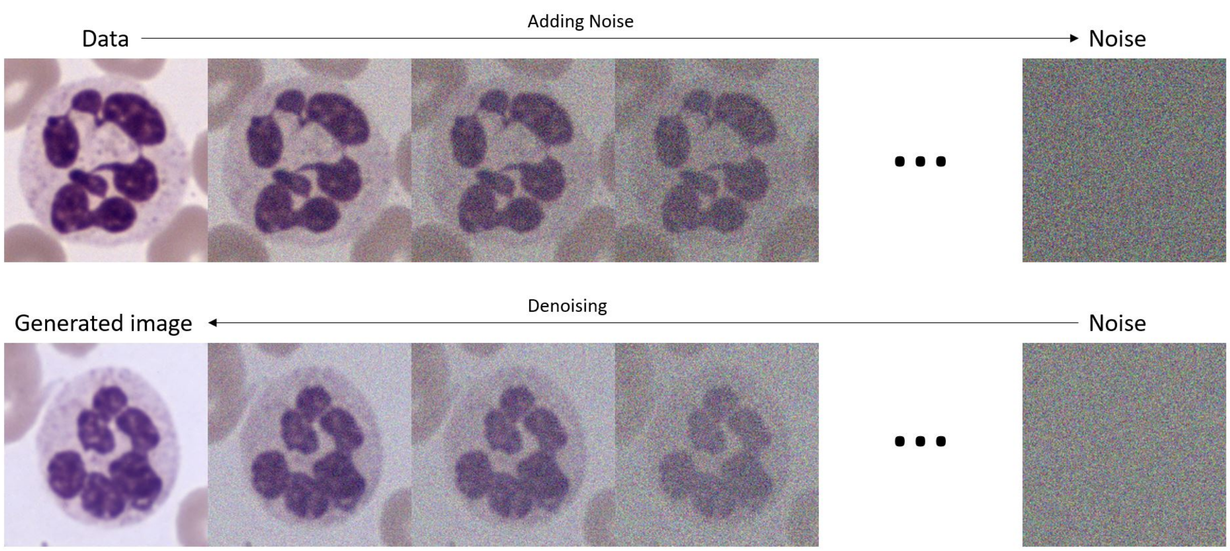

Figure 3: The intuition of diffusion models.