DTU Compute
Department of Applied Mathematics and Computer Science

Measuring particle statistics using a CNN-based segmentation

Anders B. L. Larsen   Anders B. Dahl   Rasmus Larsen

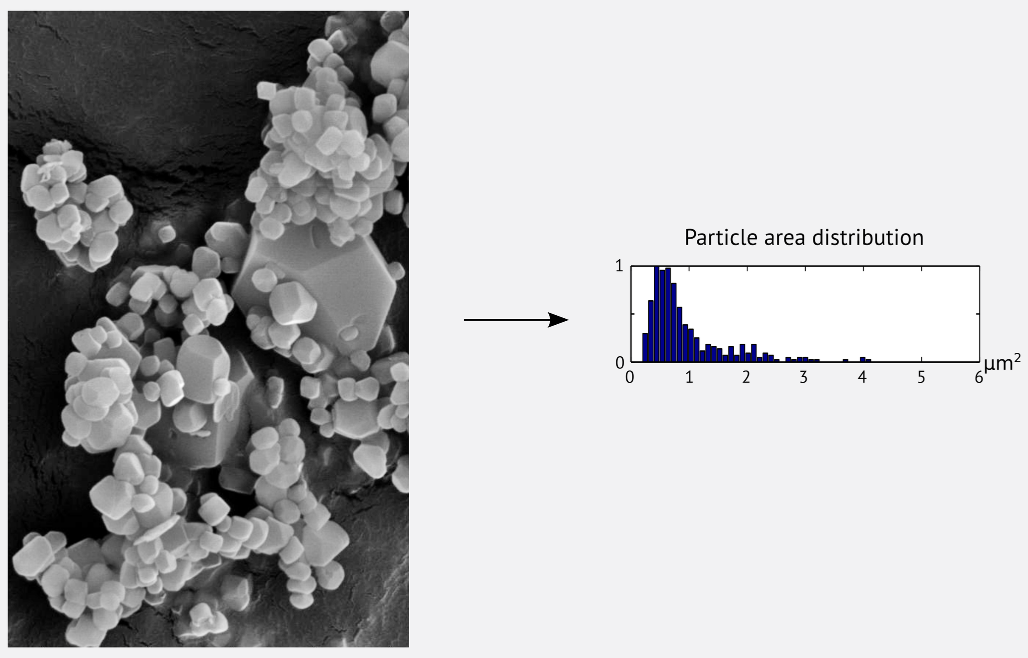

Motivation

We wish to analyze the behavior of chemical compounds over time in order to e.g. estimate shelf life of pharmaceutical drugs. Many chemical compounds have the ability to *crystallize* leading to changes in the physical properties of a drug. One way to characterize the crystallization process is by monitoring the size of the individual particles/crystals over time using *scanning electron microscopy* (SEM) images. This process requires manual inspection and is cumbersome and labor intensive. We show that this task can be automated using a convolutional neural network (CNN) to segment the particles.

Dataset

The dataset consists of a series of SEM images taken over time showing particles of a chemical compound.

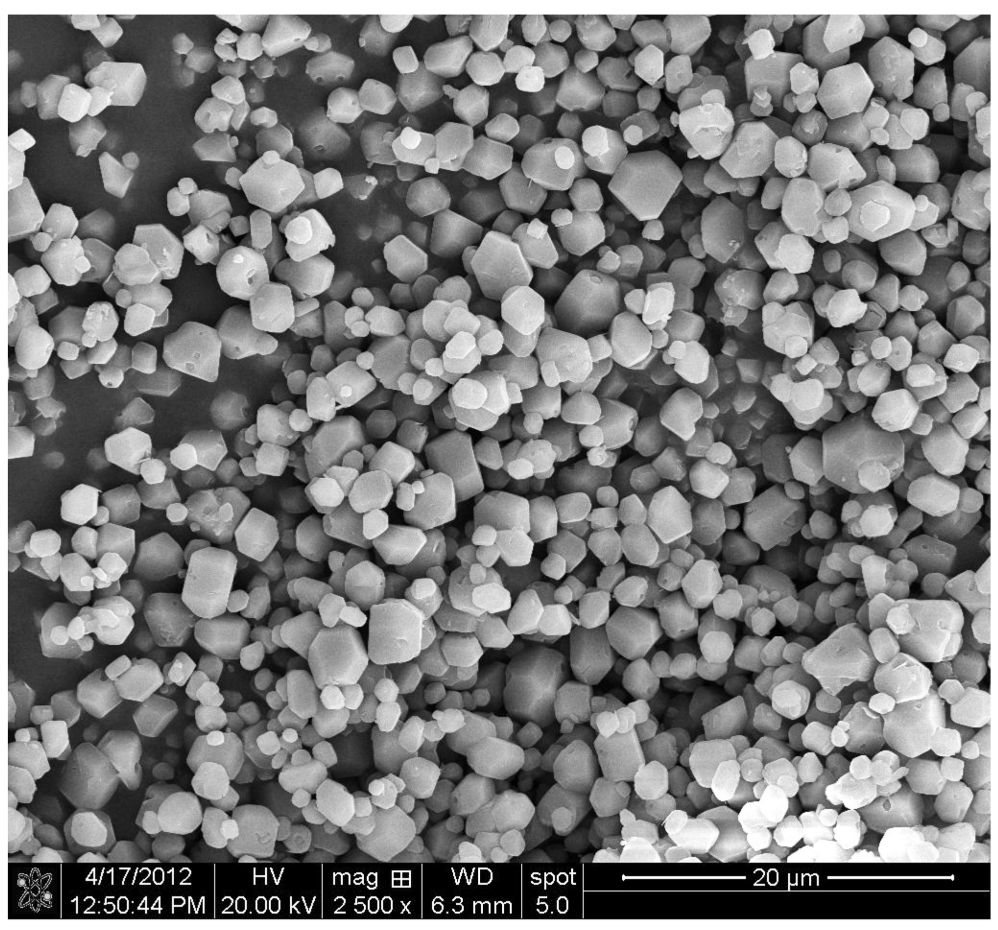

SEM image example.

Annotation scheme

We wish to determine particle boundaries in order to separate the particles. Therefore, we manually annotate a small region of an image into background (blue), background borders (red), and inter-particle borders (green).

NB: We do not handle overlapping particles in a clever way. The current manual scheme for measuring particle sizes has similar (and other) limitations.

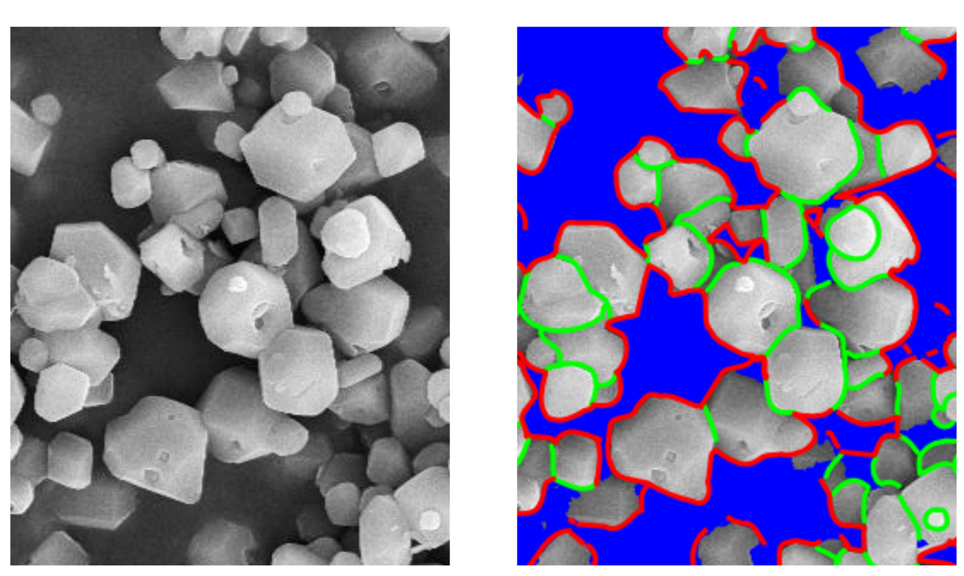

Annotation example.

Method

An image is processed by extracting overlapping patches and assigning class probabilities to each patch. The pixel-wise class probabilities are then used as segmentation.

Preprocessing: Local contrast normalization to compensate for exposure and contrast perturbations.

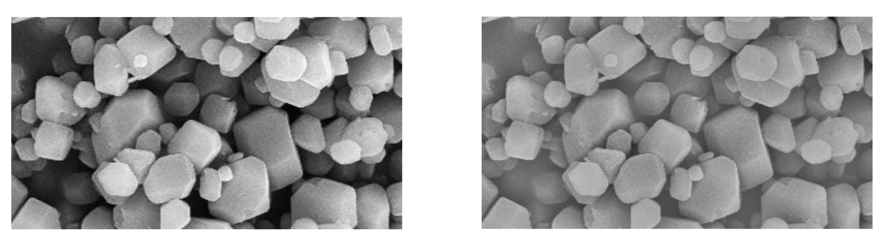

Before (left) and after (right) local contrast normalization.

Convolutional neural network

We employ a feed-forward convolutional neural network with the following implementation details.

Mini-batch SGD with momentum For more stable gradient updates and faster convergence.

Dropout We use dropout regularization in the fully connected layer to prevent overfitting.

Rectified linear units For faster convergence.

Multi-target, multinomial logistic regression We assign class labels to a 3×3 neighborhood around each patch center. The label overlap yields a smoother segmentation.

FFT-based convolutions Allow for a GPU architecture agnostic implementation.

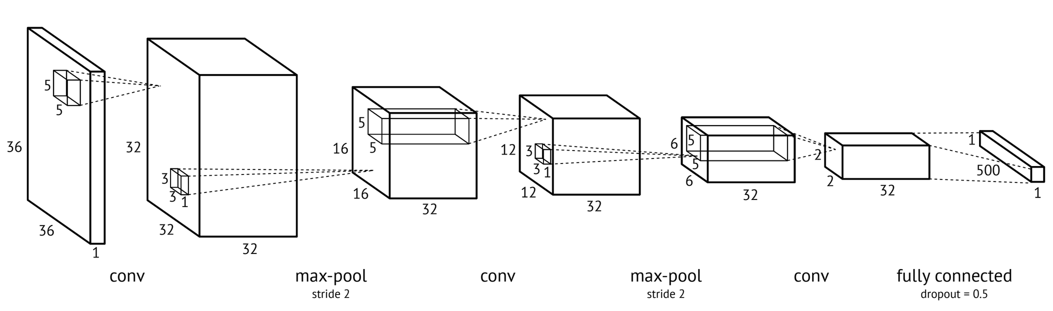

The convolutional neural network architecture.

Measuring particle sizes

With the segmentation image, we can extract particles using a *watershed* transformation. The area of each particle is then measured by counting pixels (the physical size of a pixel is known).

Results

We are able to measure particle size distributions similar to what has been measured manually.

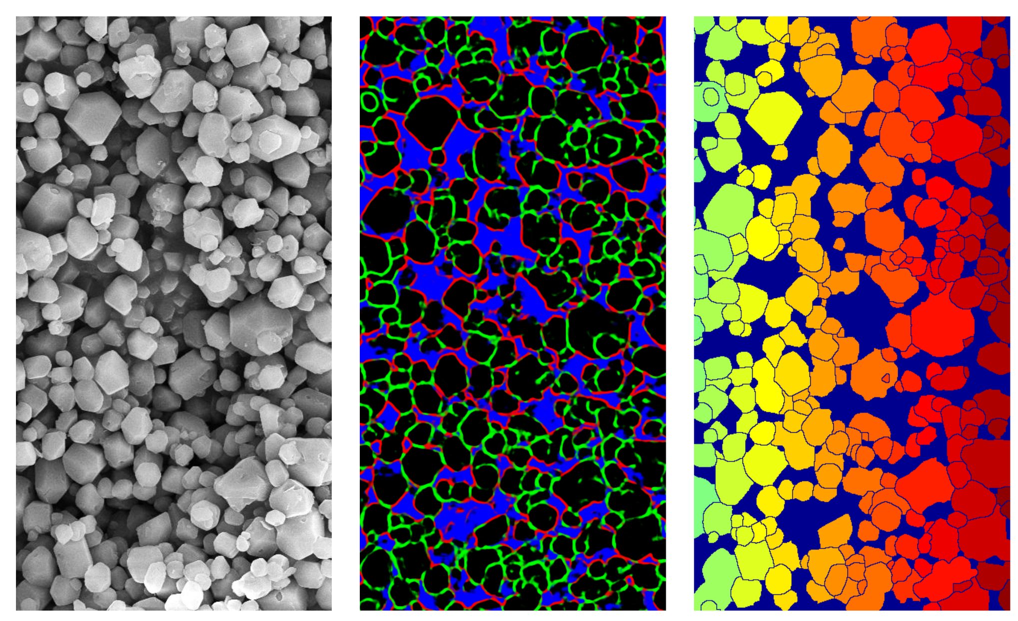

Left: Input image. Middle: CNN segmentation. Right: Extracted particles.

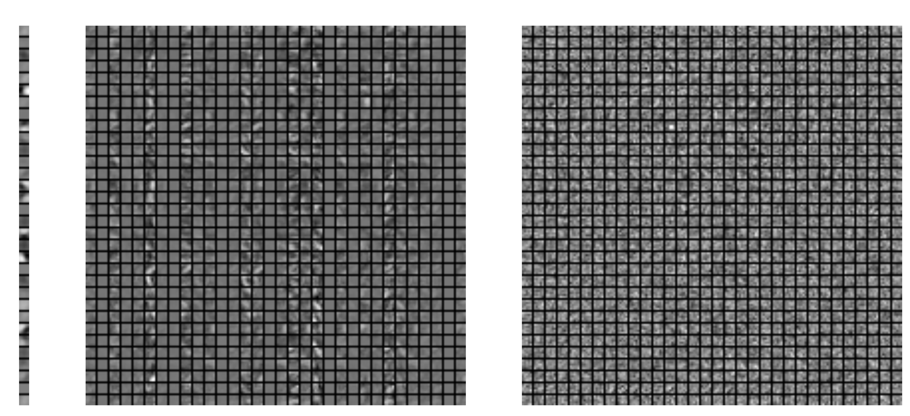

Learned CNN filter banks for the 3 convolutional layers.

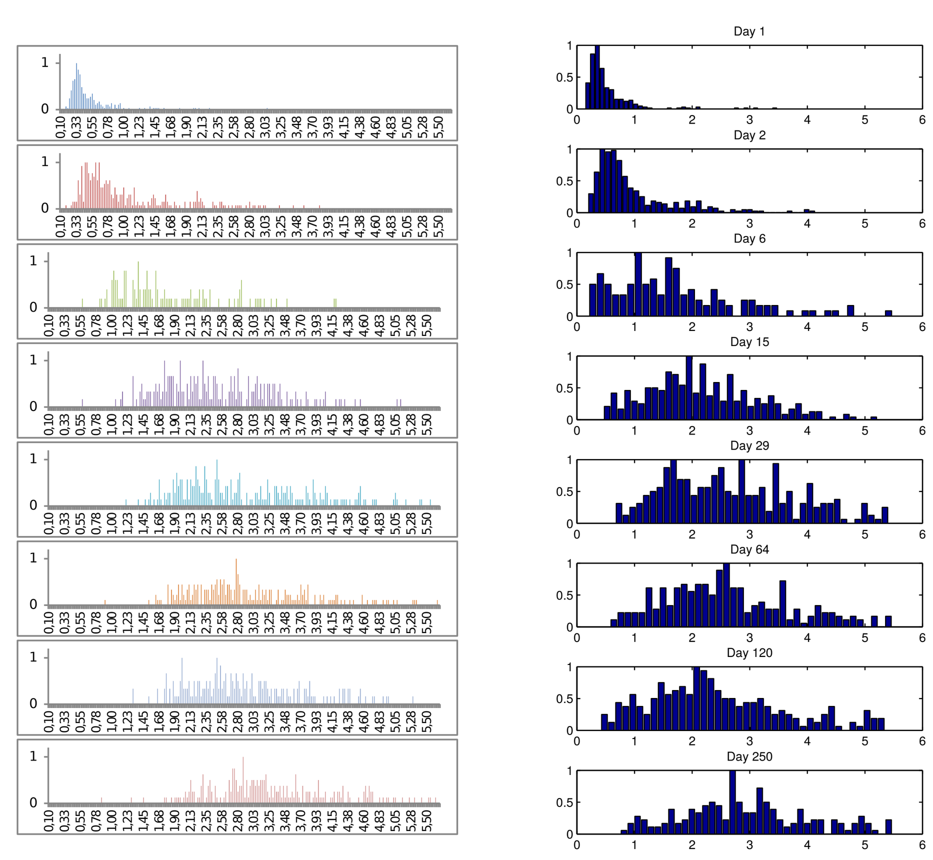

Particle area distributions over 1 to 250 days. Manual (left) vs. automatic (right) measurements.