1.19.0 Digital Image Processing Using Neural Networks

Viktor Číhal, Tadeáš Fryčák, Jan Slíva, vedoucí stáže: Ing. Jan Schier, Ph.D.
Institute of Information Theory and Automation, Czech Academy of Sciences

Abstract
The mitotic spindle is a cellular structure essential during nuclear cell division (mitosis or meiosis), ensuring even distribution of chromosomes into daughter cells. However, errors can occur during chromosome division, which may result in serious developmental disorders (such as Down syndrome) or infertility. The process of cell division is recorded using a so-called light-sheet microscope, resulting in a time-lapse recording consisting of 117 3D images with a resolution of 31x750x750. The data set contains 70 positions in total, each representing a recording of an individual oocyte (file size of 3.6 GB). The data was provided by the Institute of Animal Physiology and Genetics of the Czech Academy of Sciences. The microscope used automatic centering, which kept the spindle mostly centered in the observed space. Due to time constraints, only every sixth frame was manually annotated.

Based on these requirements, we developed a web application called Spindle App, which allows users to automatically segment the recordings in bulk. We chose to implement multiple segmentation models primarily to provide users with options based on time and accuracy preferences. At the core of the application are:

Deterministic Model
The deterministic probabilistic model was developed mainly to set a baseline value for convolutional neural networks and to allow for fast segmentation of the spindle. To give an idea of its speed, the entire annotated dataset from Liběchov can be processed using a single CPU core in just a few seconds. This model achieves an average F1 score of over 92% for segmentation.

U-net
Another method we experimented with was the U-net convolutional architecture, which is the most commonly used technique for segmenting 2D images. To solve the problem of 3D image segmentation, we adapted the model to sequentially classify 2D slices and then combine the resulting segmentations. Since the spindle is often not present in most slices, we added an additional classification model to first determine whether a spindle is present in a given layer. This model has significantly fewer parameters than the U-net architecture, allowing for system-wide optimization.

Compared to the V-net architecture, it requires fewer parameters for quality segmentation, making it faster to train on input data. This also means it requires less memory and can be used for accurate segmentation on lower-performance devices. The average speed is approximately 1 second per 3D image. The average F1 score is 94%.

V-net
V-net is a convolutional neural network that takes 3D images as input. Its advantage lies in the ability to link information across multiple layers of the input image. For example, knowing that a spindle is located beneath a given pixel helps the model better infer that the pixel also belongs to the spindle. V-net can process one 31x750x750 cube on a CPU in 4 seconds. The average F1 score exceeds 95%.

After successful segmentation, the program calculates several biologically significant statistics, such as the spindle volume at a given time and the time of its formation or division.

Currently, the application has been handed over to the Institute of Animal Physiology and Genetics for more detailed testing on larger datasets.

Attachments

Spindle App, data selection window, experiment processing, method accuracy evaluation, segmentation example

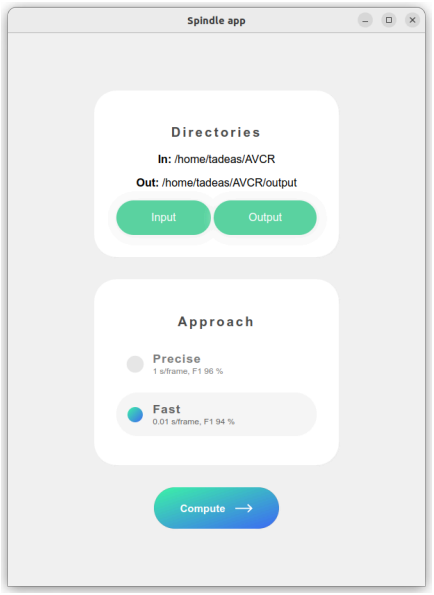

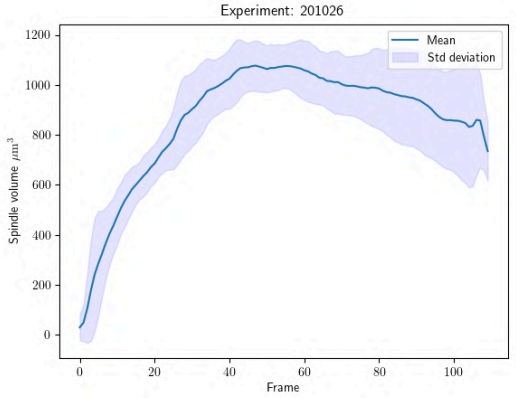

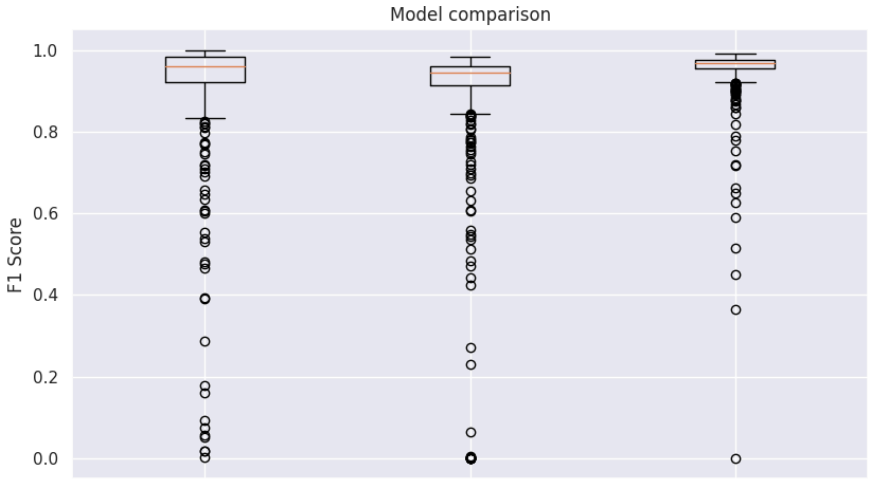

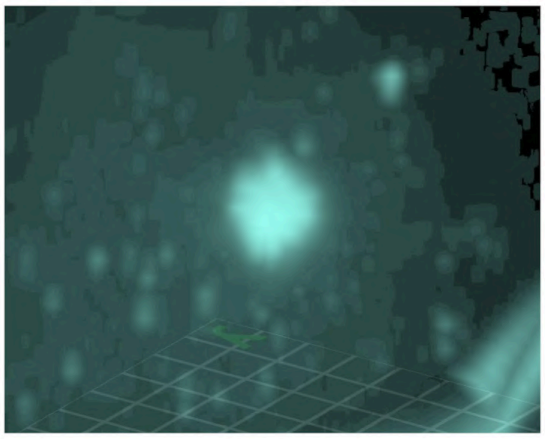

Vstup

Výstup