

Real-time Image Segmentation: GPU Lab Project

Group 1: Ravikishore Kommajosyula, Miklos Homolya, Gaurav Kukreja

In this project, the problem of image segmentation using variational methods has been explored, with a primary objective to achieve real-time capabilities, i.e., to improve the efficiency. The basic algorithm for image segmentation of a binary image was implemented based on [1] and was extended to include non-binary, single channel images based on [2]. With this as an input, the update scheme was replaced by a primal dual approach based on [3] to improve the convergence. In GPU based codes, the costliest operations are the copying of data from CPU to GPU and vice-versa. To render an output using OpenCV, the data has to be copied onto CPU first and then rendered. Instead, we use OpenGL-CUDA interoperability [4] to render the data that has been generated by CUDA, directly from the GPU. Lastly, Qt was used to build a GUI to provide the user with options to interact with the segmentation by changing the parameters real-time and output the resulting frame-per-second (fps). On our machine, with a CC 2.1 GPU, we achieved a fps of 7.6 with the default settings.

Installation Instructions:

To run this project, you need to install CUDA, OpenCV, OpenGL, Qt and plug in a camera to the computer. All of these are pre-installed on the lab computers.

On the lab machine, the code has been tested to run out of the box. Issuing the command

`qmake && make`

will build the code and generate the binary `segment` and `./segment` (with the camera connected to the computer) will fire up the GUI and run the code.

On another machine, there are two things that should be configured in the `segment.pro` file -

1. The cuda installation directory - On the lab machine the installation directory is `/usr` and on our machines on which we installed cuda using default settings, it is `/usr/local/cuda`. This entry should be changed based on the target system
2. The NVIDIA driver directory - On the lab machine, the nvidia drivers are of version 331, however that could be different on other machines. The path should be updated to `/usr/lib/nvidia-XXX-updates` where XXX is the version of the driver.

By default, these settings have been set to `/usr` for the cuda installation directory and 331 for the nvidia driver version.

References:

- [1] Tony Chan, Selim Esedoglu, and Mila Nikolova, "Finding the global minimum for binary image restoration", ICIP 1, page 121-124. IEEE, (2005)
- [2] Tony Chan, Selim Esedoglu, and Mila Nikolova, "Algorithms for finding global minimizers of image segmentation and denoising models," Tech. Rep., CAM Report 04-54, 2004.
- [3] https://vision.in.tum.de/teaching/ws2013/vmcv2013/material/variational_methods9.pdf
- [4] <http://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html#graphics-interoperability>