

GPU Programming in Computer Vision: Day 2

Date: Tuesday, 18. September 2018

Exercise 6: Convolution - Shared Memory Speedup? (8P)

Implement the convolution $G_\sigma * u$ on the GPU, now using:

1. *Shared memory* for the image u , but still global memory for k . To solve this task, use shared memory in the following way:
 - (a) To compute the convolution in an output pixel (x, y) one needs values of the input image in (x, y) and also in the neighboring pixels. For each kernel block, load into shared memory all input image values needed to compute the convolution in every pixel of this block. Note that overall there are *more* input image values to load than the size of the block, see Figure 1 on the next page. Set the size of the shared memory array accordingly.
 - (b) When loading into shared memory, make the read accesses to global memory as coalesced as possible (for instance, wherever possible, neighboring threads should access neighboring memory regions).
 - (c) At image borders, use clamping.
 - (d) Don't forget to use `__syncthreads()` after you've finished loading the data into shared memory.
 - (e) The actual computation of the convolution may only use data from the shared memory (no global memory accesses allowed).
 - (f) Use dynamic allocation of shared memory. Make sure you allocate exactly the right amount of shared memory for your kernel, and not more than needed.
 - (g) For multi-channel images, apply the above procedure in a loop (within the kernel) separately for each channel. When you start processing each new channel, also synchronize *before* you begin writing to the shared memory.
2. *Texture memory* for the image u , but still global memory for k . (Optional)

Since we work with multi-channel images, but CUDA textures allow only one channel, define the texture as having width w and height $h \cdot n_c$. Here we use the fact that the channels are arranged in memory one after another. Therefore we can view a multi-channel image u as an n_c -times larger grayscale image u_1 , defined by $u(x, y, c) = u_1(x, y + h \cdot c)$.
3. Pick some memory variant for the input image u (global, shared, or texture), and use *constant memory* for the kernel k . To define the CUDA constant kernel array, assume a maximal kernel radius $r_{\max} = 20$.
4. For some fixed $\sigma > 0$, compare the run times for the different memory versions. Which combination is the fastest? How much faster is it compared to the CPU version?




Figure 1: Shared memory array for convolution

Exercise 7: Structure Tensor

(5P)

For an input image u , the smoothed version is defined as $S := G_\sigma * u$. The *structure tensor* T of u is defined at each pixel (x, y) as the smoothing

$$T := G_\sigma * M$$

of the matrix

$$M := \nabla S \cdot \nabla S^\top = \begin{pmatrix} (\partial_x S)^2 & (\partial_x S)(\partial_y S) \\ (\partial_x S)(\partial_y S) & (\partial_y S)^2 \end{pmatrix},$$

where $\sigma > 0$ is a scale parameter. The entries of M are scalar products over the n_c channels:

$$((\partial_x S)^2)(x, y) := \sum_{c=1}^{n_c} (\partial_x S_c)(x, y)^2, \quad ((\partial_y S)^2)(x, y) := \sum_{c=1}^{n_c} (\partial_y S_c)(x, y)^2,$$

and

$$((\partial_x S)(\partial_y S))(x, y) := \sum_{c=1}^{n_c} (\partial_x S_c)(x, y) \cdot (\partial_y S_c)(x, y).$$

Compute the structure tensor. Reuse your kernels for the convolution, don't write new kernels for this. You can use just the global memory for everything for convenience.

Implement this in several steps:

1. Compute $S = G_\sigma * u$.
2. Compute $v^1 := \partial_x S$ and $v^2 := \partial_y S$ using the *more rotationally symmetric* derivative discretizations $\partial_x^r, \partial_y^r$ as given in the lecture. Note that v^1, v^2 and S each have n_c channels.
3. Compute the matrix M at each pixel. The output should consist of three grayscale images m_{11}, m_{12}, m_{22} , corresponding to the three independent components of M at each pixel.
4. Compute $T = G_\sigma * M$ by convolving the three grayscale images m_{11}, m_{12}, m_{22} .
5. Visualize the grayscale images m_{11}, m_{12} and m_{22} . For this, you will need to define three new output images in the code framework.
Hint: You will need to scale up these images, otherwise they will appear too dark because m_{11}, m_{12}, m_{22} are usually very small. To multiply an OpenCV `cv::Mat` image `m` by a scalar factor `f`, use `m *= f`;

Exercise 8: Feature Detection

(3P)

Detect edges and corners in an image using the structure tensor T from the previous exercise:

1. Write a `__device__` function that computes the eigenvalues of a 2×2 matrix.¹
2. Compute at each pixel (x, y) the two eigenvalues λ_1, λ_2 of the structure tensor. We use the convention that $\lambda_1 \leq \lambda_2$.
 - If a pixel is on a corner ($\lambda_2 \geq \lambda_1 \geq \alpha$), mark it red.
 - If a pixel is on an edge ($\lambda_1 \leq \beta < \alpha \leq \lambda_2$), mark it yellow.
 - Otherwise, make the pixel darker by multiplying its components with 0.5.
3. The scale parameter $\sigma > 0$ and feature parameters $\alpha > \beta > 0$ depend on the input image u . Experiment with different values. As a start, choose $\sigma = 0.5$, $\alpha = 10^{-2}$ and $\beta = 10^{-3}$, for the input image `gaudi.png`. Test also on live webcam images.

¹http://www.math.harvard.edu/archive/21b_fall_04/exhibits/2dmatrices/index.html