

Tutorial 9

In this tutorial we will write a Cuda program that implements a 2D convolution. See Figure 1 for an example.

- 0.) Download the [skeleton code](#) and generate the build system using `cmake`.
- 1.) Understand the given host implementation.
- 2.) Develop a simple (i.e. unoptimized) Cuda implementation. Can we re-use some parts of the host implementation?
- 3.) How can the efficiency of your implementation be improved? Implement at least one improvement and measure the speed-up.

Please finish the implementation until the next tutorial (week of 08/01/2019).




Figure 1: Image before and after convolution with a Gaussian kernel.