Problem Set 1: A Simple Visual System

Problem 1 *A simple image formation model*

The goal of this first exercise is to take images with different settings of a camera to create pictures with perspective projection and with orthographic projection. Both pictures should cover the same piece of the scene. You can take pictures of real places (e.g., the street, a living room, ...) or you can also create your own simple world (e.g., you can print *simpleWorld.pdf* and create your own scenes. I recommend printing on mate paper).

To create pictures with orthographic projection you can do two things: 1) use the zoom of the camera, 2) crop the central part of a picture. You will have to play with the distance between the camera and the scene, and with the zoom (or amount of cropping) so that both images look as similar as possible only differing in the type of projection (similar to figure 1.4, in the lecture 1 notes). Submit the two pictures.

Problem 2 *Orthographic projection*

Prove the projection equations (eq. 1.2 and 1.3) that relate the coordinates of one point in the 3D world and the image coordinates of the projection of the point in the camera plane.

In the code, we assume $\alpha = 1$.

Problem 3 *Constraints*

In the lecture 1 notes, we have written all the derivative constraints for $Y(x, y)$. Write the constraints for $Z(x, y)$.

The code of the problem set will use constraints on Y.

Problem 4 *Approximation of derivatives*

Fill the missing kernels (lines 166 and 180) in the script: *simpleworldY.m*

Problem 5 *Run the code*

Select some of the images included with the code and show some new view points for them.

Optional: You can also try with new images taken by you if you decide to create your own simple world.

Problem 6 *Violating simple world assumptions*

Find one example from the images provided with the problem set or from one of your own images when the recovery of 3D information fails. Explain why it fail.

Research problem [optional] *The real world*

A research problem is a question for which we do not know the answer. In fact, there might not even be an answer. This question is optional and could be extended into a larger course project.

The goal of this problem is to test the 3D reconstruction code with real images. A number of the assumptions we have made will not work when the input are real images of more complex scence. For instance, the simple strategy of differentiating between foreground and background segmentation will not work with other scenes.

Try taking pictures of real world scenes (not the cubes) and propose modifications to the scheme proposed in this lecture so that you can get some better 3D reconstructions. The goal is not to build a general system, but to be able to handle a few more situations.