Homework Assignment 6

15-463/663/862, Computational Photography, Fall 2017
Carnegie Mellon University

Due Friday, Dec. 8, at 11:59pm

The purpose of this assignment is to explore non-line-of-sight (NLOS) imaging. In particular, you will use a very straightforward and very intuitive *passive* technique for NLOS imaging termed corner cameras [1].

Unlike previous assignments, the focus of this (shorter) assignment will be on data acquisiton rather than coding. In particular, you are highly encouraged to use the publicly available implementation of the corner camera paper (though you are welcome to write your own implementations, partial or full, of the algorithm if you feel like doing so). The implementation can be easily used with newly captured data, with a few simple edits to replace filenames and data types. We will refer to this as the *reference implementation* throughout the homework.

In order to understand what the code is doing, we highly recommend that you watch this video and read the associated paper describing the technique. There is a "Hints and Information" section at the end of this document that is likely to help.

1. Reconstruction of motion trajectories around the corner in indoors scenes (100 points)

For the main part of this homework, you will need to reproduce results analogous to those shown in Figure 1, which is reproduced directly from [1]. Even though the figure shows an example with two people moving, it is up to to decide whether your capture videos and reconstructions will involve just one or two (or more if you want!) people.

To get full marks, you will need to reproduce motion trajectories from a corner video captured indoors (e.g., as shown in Figure 1). What follows gives an overview of the steps you will need to take for of the video.

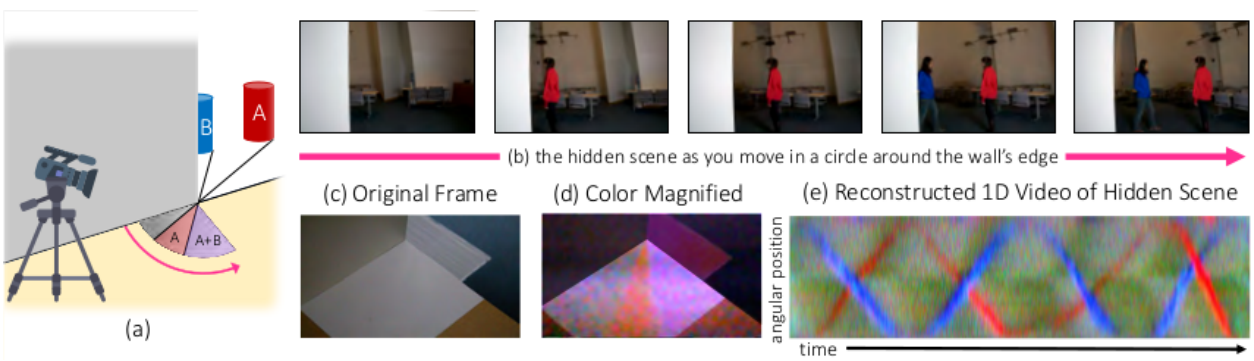

Figure 1: Reconstruction of motion projectories around the corner, by processing penumbrae. Reproduced from [1].

Set up the scene and your corner. Select the room where you will do the data capturing. We suggest that you use a room where the moving people will be well-lit relative to the rest of the environment.

Then, set up two cameras. One camera will be used to capture the penumbra, mounted as shown in Figure 1(a). We will call this the *corner camera*. We suggest that you use a high-sensitivity DSLR (such as the Nikon D3300 you picked up from the instructors) for this. You should also make sure to mount your camera on tripod, to avoid vibrations and so on.

The second camera should be used to capture video directly from the scene itself, like what is shown in Figure 1(b). We will call this the *direct-view camera*. This camera can be any video camera you have at

your disposal, e.g., your phone camera. You can hold this camera yourselves, or fix it on a book or other makeshift stand. The purpose of this second camera will become clear later.

Calibration for rectification. As discussed in Section 3.1.1 of [1], before performing motion trajectory reconstruction, it is necessary to rectify the captured frames so that they look as if they were captured with the camera looking at the floor head on. This can be done by a homography estimation (see also Lecture 19).

To compute the homography needed for rectification, place a square sheet of paper in the field of view of the corner camera, and capture a single image (or a few video frames if you prefer) of this scene. An example of this is shown in Figure 2. (Note that, even though the figure uses a checkerboard, a blank sheet of paper where all four corners are visible should suffice.) Then, use the four corners of this sheet to compute the rectification homography. The reference implementation includes all of the code needed for doing this computation, including a simple GUI for detecting the corners of the sheet of paper. Once you are done with calibration, remove the sheet of paper from the field of view of the corner camera.

The homography computation only needs to be done for the corner camera. Note that the corner camera *should not move at all* between capturing the video for rectification and the videos you will use for trajectory estimation. Therefore, we recommend that you first fix your corner camera in its final location, and then completely avoid touching it as you capture the various videos.

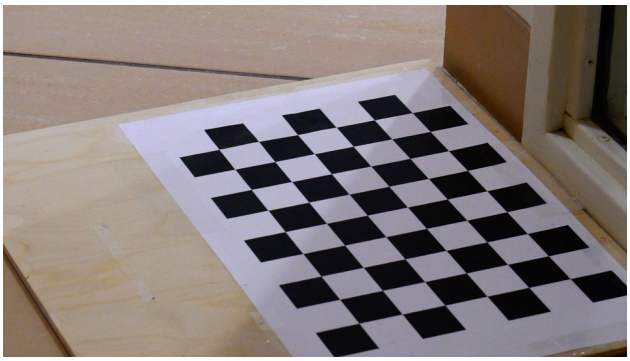

Figure 2: Frame captured by corner camera during calibration.

Penumbra measurements. After you have performed the above calibration, set both the corner and the direct-view cameras to start recording video. Try to have the two cameras start recording video as close to each other as possible: While perfect simultaneity is not required, try to have the start time of video capturing between the two cameras differ by no more than a second.

Then, have someone (either you or volunteers) start moving back and forth in the non-line-of-sight scene, similar to the movements shown in Figure 1.

At the end of this stage you will have three videos: 1) A short video (possibly even a single frame) captured by the corner camera for calibration. 2) A penumbra video captured by the corner camera while someone is moving in the non-line-of-sight scene (similar to Figure 1(c)). 3) A direct-view video captured by the direct-view camera while someone is moving in the non-line-of-sight scene (similar to Figure 1(b)).

Data processing and evaluation. Use the reference implementation to process the penumbra video you captured with the corner camera. This should result in a trajectory profile that looks similar to what is shown in Figure 1(e).

To evaluate how accurate the profile is, you can compare it with the direct-view video you captured with the direct-view camera. An example of this comparison is shown between minutes 1:42 and 2:06 of the companion video to [1]. If your videos correspond to only one moving person, then you can compare whether the reconstructed motion trajectory accurately captures the times at which the person turns around to move in a different direction. If your videos involved more than one moving people, then you can check whether the reconstructed motion trajectory accurately captures points where the people run into each other.

In your homework solution, make sure to show: 1) A photo of the scene you used, with the corner camera in place. 2) The reconstructed motion trajectories. 3) Frames from the direct-view video matching to key points in the reconstructed motion trajectories (changes in move direction, or more than one moving people running into each other).

2. Bonus: Reconstruction of motion trajectories around the corner in outdoors scenes (20 points)

Repeat problem 1, but this time for an outdoors scene, e.g., as shown in Figure 4 of [1]. We suggest that you use a location where the floor is tiled, and also where there are no moving shadows occuring nearby (i.e., no shadows from trees which may move due to the wind). You should also make sure to capture your outdoors videos during daylight.

In your submission, make sure to discuss any challenges you ran into, as well as differences you observed between the indoors and outdoors cases.

3. Bonus: Stereo around the corner (30 points)

Capture data and use it to reproduce a stereo localization result analogous to what is shown in Figure 3. The stereo localization procedure is described in Section 4.1 of [1], for where the previous figure is reproduced. The reference implementaion has most of the code you will need for handling this case, though you will need to write your own code for performing the triangulation needed for localization.

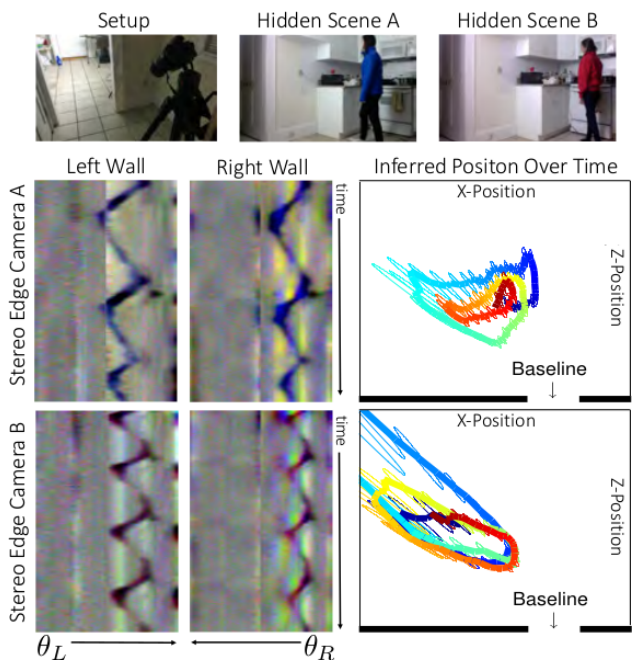

Figure 3: Stereo around the corner, by processing penumbrae from multiple corners. Reproduced from [1].

Deliverables

As described on the course website, solutions are submitted through Canvas. Your solution should be an archive (e.g., a ZIP file) that includes the following:

- A PDF report explaining what you did for each problem, including answers to all questions asked throughout problem 1. The report should include any figures and intermediate results that you think may help. Make sure to include explanations of any issues that you may have run into that prevented

you from fully solving the assignment, as this will help us determine partial credit. The report should also explain any additional image files you include in your solution (see below).

- For the indoors of problem 1, a photo of the scene you used, with the corner camera in place, the reconstructed motion trajectories, and frames from the direct-view video matching key points in the reconstructed motion trajectories You can also include additional image files for various experiments other than your final ones, if you think they show something important.
- If you do Bonus Problem 2: The same information as in Problem 1, but for the outdoors scene you captured.
- If you do Bonus Problem 3: Include a photo of the scene you used, the reconstructed motion trajectories for all corners you used, as well as visualization of localization results analogous to what is shown in Figure 3.

Please organize your solution submission using the following file structure:

```
.zip
├── .pdf ..................................... The PDF report.
└── data/ ..................................... Contains all image, video, and other data files.
```

Hints and Information

- To use a DSLR camera as the corner camera, you will need to operate it in video mode. For the Nikon D3300 DSLR that is provided by the instructors, you can find instructions on how to do this online.
- It will help if the people moving in the videos you capture wear some color that is very distinct from the surroundings (e.g., bright red or blue, as in most of the examples shown in [1]. If you choose to capture videos that have more than one people moving, then they should wear different colors from each other.
- Depending on how fast you move in your videos, and how long the total video sequences are, you may need to subsample the video sequences by some step size (e.g., every second or third frame), in order to be able to have them processed in a reasonable amount of time.
- The project website for [1] includes a lot of sample videos. You can take a look at those for inspiration and to see what your own videos should look like in order to be able to produce meaningful results.

References

[1] K. L. Bouman, V. Ye, A. B. Yedidia, F. Durand, G. W. Wornell, A. Torralba, and W. T. Freeman. Turning corners into cameras: Principles and methods. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.