

# Assignment: Video Stabilization

October 12, 2011

The goal of this assignment is to stabilize a video. This means that you will take a video which shows some jitter and process it so that it is smoother. I do not expect this to work perfectly, but you should be able to see it working on many of the frames.

You will turn in 4 things: your original video, 2 processed videos (see below), and your code. The steps are as follows.

1. Find a video on the web that has some jitter. Alternatively, you can take your own video. If you don't have a video camera, try to borrow one.
2. You will need to convert the video into a form that can be used in matlab. There are many ways to do this, but I want you to research it on your own. You are free to use any means at your disposal (other than have someone do it for you) to get your video converted into a bunch of images in matlab so that they will be easy to work with.
3. You may want to reduce the size of the video frames (now images) so that your code runs faster. Once you get your code working you can try running it on the full size video frames. I also recommend that you convert all of the images to grayscale in matlab. Again, I want you to research how to do this on your own.
4. Next, choose some portion of the first video frame. I suggest a rectangular set of pixels in the middle. Then, go to the second frame of video and try to find a portion of the second frame that matches that frame as well as possible. I want you to do this using a gradient descent method that only searches over translations. In your first implementation, I want you to use the L1 difference between the frames. Repeat this process for each frame, picking a new "target" from the middle of the current frame and trying to find it in the next frame using gradient descent. Create a sequence of images which aligns each frame to its successor to "stabilize" the video.
5. Once you have a version working with L1, I want you to try a version based on distribution fields. That is, you will make a distribution field of the patch in frame 1, and try to align it to the distribution field of frame 2, using an L1 difference between the distribution fields.

Turn in your original video, your L1 video, your distribution field video, and your code.

You will want to start this soon so you can get help if you need it.