Investigating Small Motions in Videos

Amanda Stowers: astowers@stanford.edu
Matt Estrada: estrada1@stanford.edu

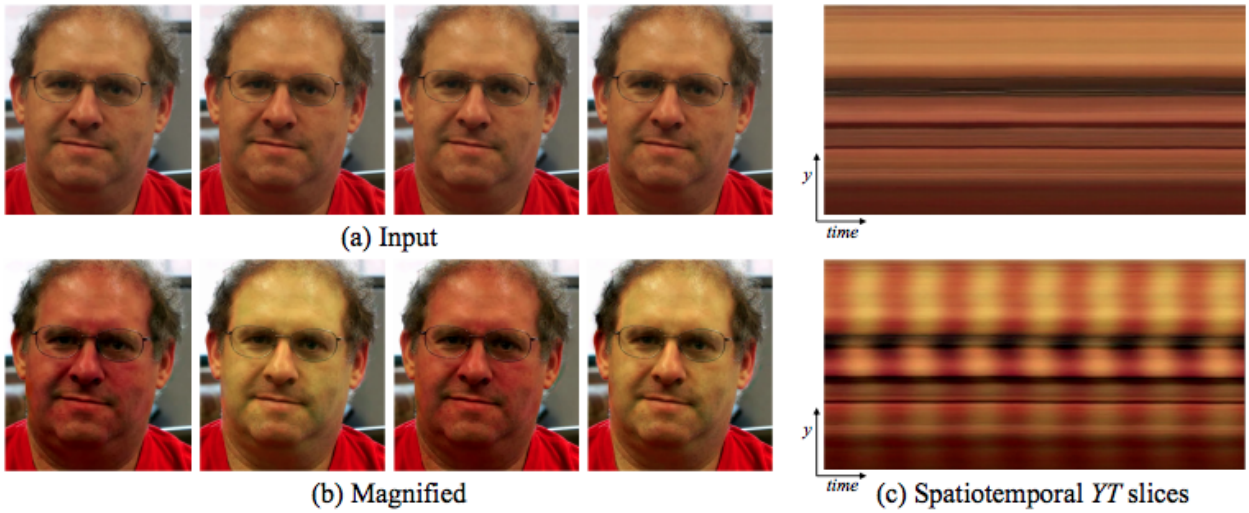

Example of amplifying facial color changes in time which enables heartbeat recognition. Figure from [1].

Goals and Work:
We plan to implement algorithms to enhance small (imperceptible to human) differences in videos in order to gain greater insight into everyday phenomenons. Many interesting phenomena occur with only very small temporal variations in terms of visible light [1]. For instance a person's heart rate causes subtle changes in their skin color over time which could indicate nervousness, lying in a game of cards, or impending medical crisis. A flower might adjust it's position in response to the sun's day-night cycle.

Our work will be based on recent research coming out of MIT's CSAIL, with Frédo Durand and William T. Freeman [1,2,3]. These methods encompass two main types of amplification in the time domain - color changes [1] and small motion [2,3]. We plan to focus on amplification of color changes, particularly in scenarios where there are multiple people in an image. For instance, if a group of people are sitting around a table playing poker, the goal would be to identify the humans and determine their heart rate. As a player, you could see if your opponents heart rate increased, indicating that they were perhaps lying or had an unusually good hand.

We plan to use relatively high quality, tripod mounted video captured from a DSLR camera. One of our challenges will be measuring fine signals such that the changing color of a human face. To facilitate this, we plan to use a steady light source and have people engaged in relatively low motion activities. This should minimize artifacts from incidental motion and varying light sources.

To process the video, we would identify where people were located using methods similar to those we learned in class. After identifying faces, we would then utilize the methods described in [1-3]. This would involve determining the time varying changes in color present in the video, seeing how these present in the frequency domain, amplifying the components found in our desired frequency range and then converting back to the time domain.

Original Contribution and Application

The specific application we will be looking at is heart rate detection in videos with multiple people. A practical use for this capability would be monitoring an ER waiting room and flagging patients with very high or low heart rates, so that they can be seen more quickly.

Currently, the available implementations for heart rate detection do so only in frames where a single person's face fills most of the frame. This requires a person to intentionally position themself immediately in front of the camera. Therefore, we plan to implement a method by which we identify the locations in the video with faces and then detect the heart rate for each of them, increasing the robustness to position and enabling multiple people to be monitored at once. One technical challenge will be the inherent loss of resolution of our signal by only using a portion of the video that corresponds to one of multiple individuals' faces.

References:

[1] Wu, H. et al, Eulerian Video Magnification for Revealing Subtle Changes in the World, ACM Trans. Graph., 2012

[2] Liu, C. et al, Motion Magnification, ACM Trans. Graph., 2005

[3] Neal Wadhwa, et al, Phase-based Video Motion Processing, ACM Trans on Graph., 2013

We do not plan to use a DROID phone.