



Computational Video: Post-processing Methods for Stabilization, Retargeting and Segmentation

Irfan Essa
Georgia Institute of Technology, USA

Abstract

We address a variety of challenges for analysis and enhancement of Computational Video. We present novel post-processing methods to bridge the difference between professional and casually shot videos mostly seen on online sites. Our research presents solutions to three well-defined problems: (1) Video stabilization and rolling shutter removal in casually-shot, uncalibrated videos; (2) Content-aware video retargeting; and (3) spatio-temporal video segmentation to enable efficient video annotation. We showcase several real-world applications building on these techniques.

We start by proposing a novel algorithm for video stabilization that generates stabilized videos by employing L1-optimal camera paths to remove undesirable motions. We compute camera paths that are optimally partitioned into constant, linear and parabolic segments mimicking the camera motions employed by professional cinematographers. To achieve this, we propose a linear programming framework to minimize the first, second, and third derivatives of the resulting camera path. Our method allows for video stabilization beyond conventional filtering, that only suppresses high frequency jitter. An additional challenge in videos shot from mobile phones are rolling shutter distortions. Modern CMOS cameras capture the frame one scanline at a time, which results in non-rigid image distortions such as shear and wobble. We propose a solution based on a novel mixture model of homographies parametrized by scanline blocks to correct these rolling shutter distortions. Our method does not rely on a-priori knowledge of the readout time nor requires prior camera calibration. Our novel video stabilization and calibration free rolling shutter removal have been deployed on YouTube where they have successfully stabilized millions of videos. We also discuss several extensions to the stabilization algorithm and present technical details behind the widely used YouTube Video Stabilizer.

We address the challenge of changing the aspect ratio of videos, by proposing algorithms that retarget videos to fit the form factor of a given device without stretching or letter-boxing. Our approaches use all of the screens pixels, while striving to deliver as much video-content of the original as possible. First, we introduce a new algorithm that uses discontinuous seam-carving in both space and time for resizing videos. Our algorithm relies on a novel appearance-based temporal coherence formulation that allows for frame-by-frame processing and results in temporally discontinuous seams, as opposed to geometrically smooth and continuous seams.



Second, we present a technique, that builds on the above mentioned video stabilization approach. We effectively automate classical pan and scan techniques by smoothly guiding a virtual crop window via saliency constraints.

Finally, we introduce an efficient and scalable technique for spatio-temporal segmentation of long video sequences using a hierarchical graph-based algorithm. We begin by over-segmenting a volumetric video graph into space-time regions grouped by appearance. We then construct a region graph over the obtained segmentation and iteratively repeat this process over multiple levels to create a tree of spatio-temporal segmentations. This hierarchical approach generates high quality segmentations, and allows subsequent applications to choose from varying levels of granularity. We demonstrate the use of spatio-temporal segmentation as users interact with the video, enabling efficient annotation of objects within the video.

Part of this talks will will expose attendees to use the Video Stabilizer on YouTube and the video segmentation system at videosegmentation.com. Please find appropriate videos to test the systems.

Part of the work described above was done at Google, where Matthias Grundmann, Vivek Kwatra and Mei Han are, and Professor Essa is working as a Consultant. Part of the work were efforts of research by Matthias Grundmann and Daniel Castro, as part of their research efforts as students at GA Tech.

Keywords: Post-processing, Stabilization, Retargeting, Segmentation