GPU-based real-time trajectory estimation from videos of vehicle-mounted cameras

Balázs Piszkor[1], István Gergő Gál[1], Tekla Tóth[1], Levente Hajder[1]

[1] Department of Algorithms and their Applications, Faculty of Informatics, Eötvös Loránd University, Budapest, Hungary

Autonomous driving is one of the most popular research areas nowadays. Real-time sensing and control of vehicles require very high computational demand that cannot be possible without the usage of Graphics Processing Units (GPUs).

The aim of our presentation is to show that real-time camera-based visual odometry is possible. Visual odometry is a basic problem for both computer vision and robotics. The main idea behind the paper is that affine transformations between corresponding regions of stereo images give very useful information for the computation of camera motion. In the literature, only feature locations are used and spatial coordinates are computed by the so-called triangulation technique. We show that surface normals can also be computed from the affine transformations between corresponding image patches. Moreover, affine transformations depend on camera parameters as well, therefore camera extrinsic parameters, i.e. the relative pose between a stereo image pair, can also be retrieved.

For planar motion, which is common for vehicles in roads, only one affine transformation is required to estimate the extrinsic motion parameters. To the best of our knowledge, it is the first method in the literature that needs only a single correspondence to compute the relative camera pose between two images. Therefore, only local information has to be processed contrary to the mainstream pose estimators where at least two points correspondences are required even for planar motion.

The main challenge for a real trajectory estimator is that input data are highly contaminated by outliers, i.e. wrong point and/or affine correspondences. The well-know RANSAC algorithm is frequently applied to filter out outliers, however, CPU-implementations of RANSAC is very slow. In our case, histogram voting can also be used for robustification. Fortunately, both RANSAC and histogram voting can be efficiently parallelized, therefore the application of GPUs is preferred.

We have developed a CUDA implementation for the planar odometry problem. The most complex mathematical problem for the trajectory computation is to find the roots of a quartic polynomial. Closed-form solutions exist for this problem that can be implemented for GPU threads. The speed of the application is approximately 5FPS that is satisfactory for real-time vision of an autonomous vehicle.