Department of Electrical and Computer Engineering

University of Rochester, Rochester, NY

Ph.D. Public Defense

Wednesday, January 25, 2017
9:00 AM

Hopeman Building, Room 335

Vehicle Detection and Tracking in Wide Area Motion Imagery by Co-Registering and Exploiting Vector Roadmaps

Ahmed Elliethy

Supervised by
Professor Gaurav Sharma

Abstract

Wide Area Motion Imagery (WAMI) is an emerging technology that allow images to be captured for relatively large areas that span substantial parts of a city with high spatial detail typically at low temporal rates. WAMI captured imagery is feeding a corresponding thrust in large scale visual data analytics such as large scale vehicle tracking. The effective- ness of such analytics can be enhanced by combining the WAMI with alternative sources of rich geo-spatial information such as road maps. Toward this end, in this thesis, we propose methods for road map alignment and vehicle tracking using the aligned road map.

First, we propose a novel methodology for accurately registering a geo-referenced vector roadmap to WAMI by using locations of detected vehicles and determining a para- metric transform that aligns these locations with the network of roads in the roadmap. Specifically, the problem is formulated in a probabilistic framework, explicitly allowing for spurious detections that do not correspond to on-road vehicles. The registration is estimated via the EM algorithm as the planar homography that minimizes the sum of weighted squared distances between the homography-mapped detection locations and the corresponding closest point on the road network, where the weights are estimated posterior probabilities of detections being on-road vehicles. The weighted distance minimization is efficiently performed using the distance transform with the Levenberg Marquardt (LM) nonlinear least-squares minimization procedure and the fraction of spurious detections is estimated within the Expectation-Maximization (EM) framework. The proposed method effectively sidesteps the challenges of feature correspondence estimation, applies directly to different imaging modalities, is robust to spurious detections, and is also more appropriate than feature matching for a planar homography. Results over three WAMI datasets captured by both visual and infra-red sensors indicate the effectiveness of the proposed alignment methodology.

Next, we propose an acceleration of our WAMI to vector roadmaps registration framework via a parallelized and pipelined realization implemented on modern GPU architectures. Specifically, we map the LM calculations to the GPU by partitioning the vehicular detection locations and efficiently decomposing the expensive LM calculations across different GPU blocks. Additionally, we propose a novel pipelining of the LM algorithm that effectively exploits the low arithmetic latency of the GPU compared with its high global memory access latency. The pipelining is accomplished by combining the LM objective function evaluation with the gradient-Hessian calculations, which allows the gradient-Hessian calculations to be pre-computed relatively inexpensively (in time) by using data already fetched from GPU global memory to the per-thread memory. Our parallelized and pipelined realization achieves real time WAMI to vector roadmaps registration performance on our WAMI dataset.

Finally, we propose a computationally efficient tracking-by-detection approach that leverages the co-registered road network. The proposed approach iteratively estimates and combines the associations of vehicle detections over reasonably large temporal windows, which is beneficial for imposing motion models and obtaining a good estimate of the associations. To limit the exponential growth of association hypotheses with the large temporal windows, the proposed approach uses two innovative pruning strategies. First, guided by a pixel accurate co-registered road network, we disregard unlikely associations that do not agree with the road network. Second, through the iterations, we progressively enlarge the size of the temporal windows while limiting the related increase in association hypotheses by stochastically maintaining high confidence associations from previous iteration to the current one while disassociating low confidence associations. Thus, in iteratively improving estimated tracks, our approach revisits only the association choices made available by the stochastic disassociation from the prior iteration, benefiting from the enlarged temporal window at the current iteration. Vehicle tracking results obtained over test WAMI datasets indicate that our proposed approach provides significant performance improvements compared with three state of the art alternatives.