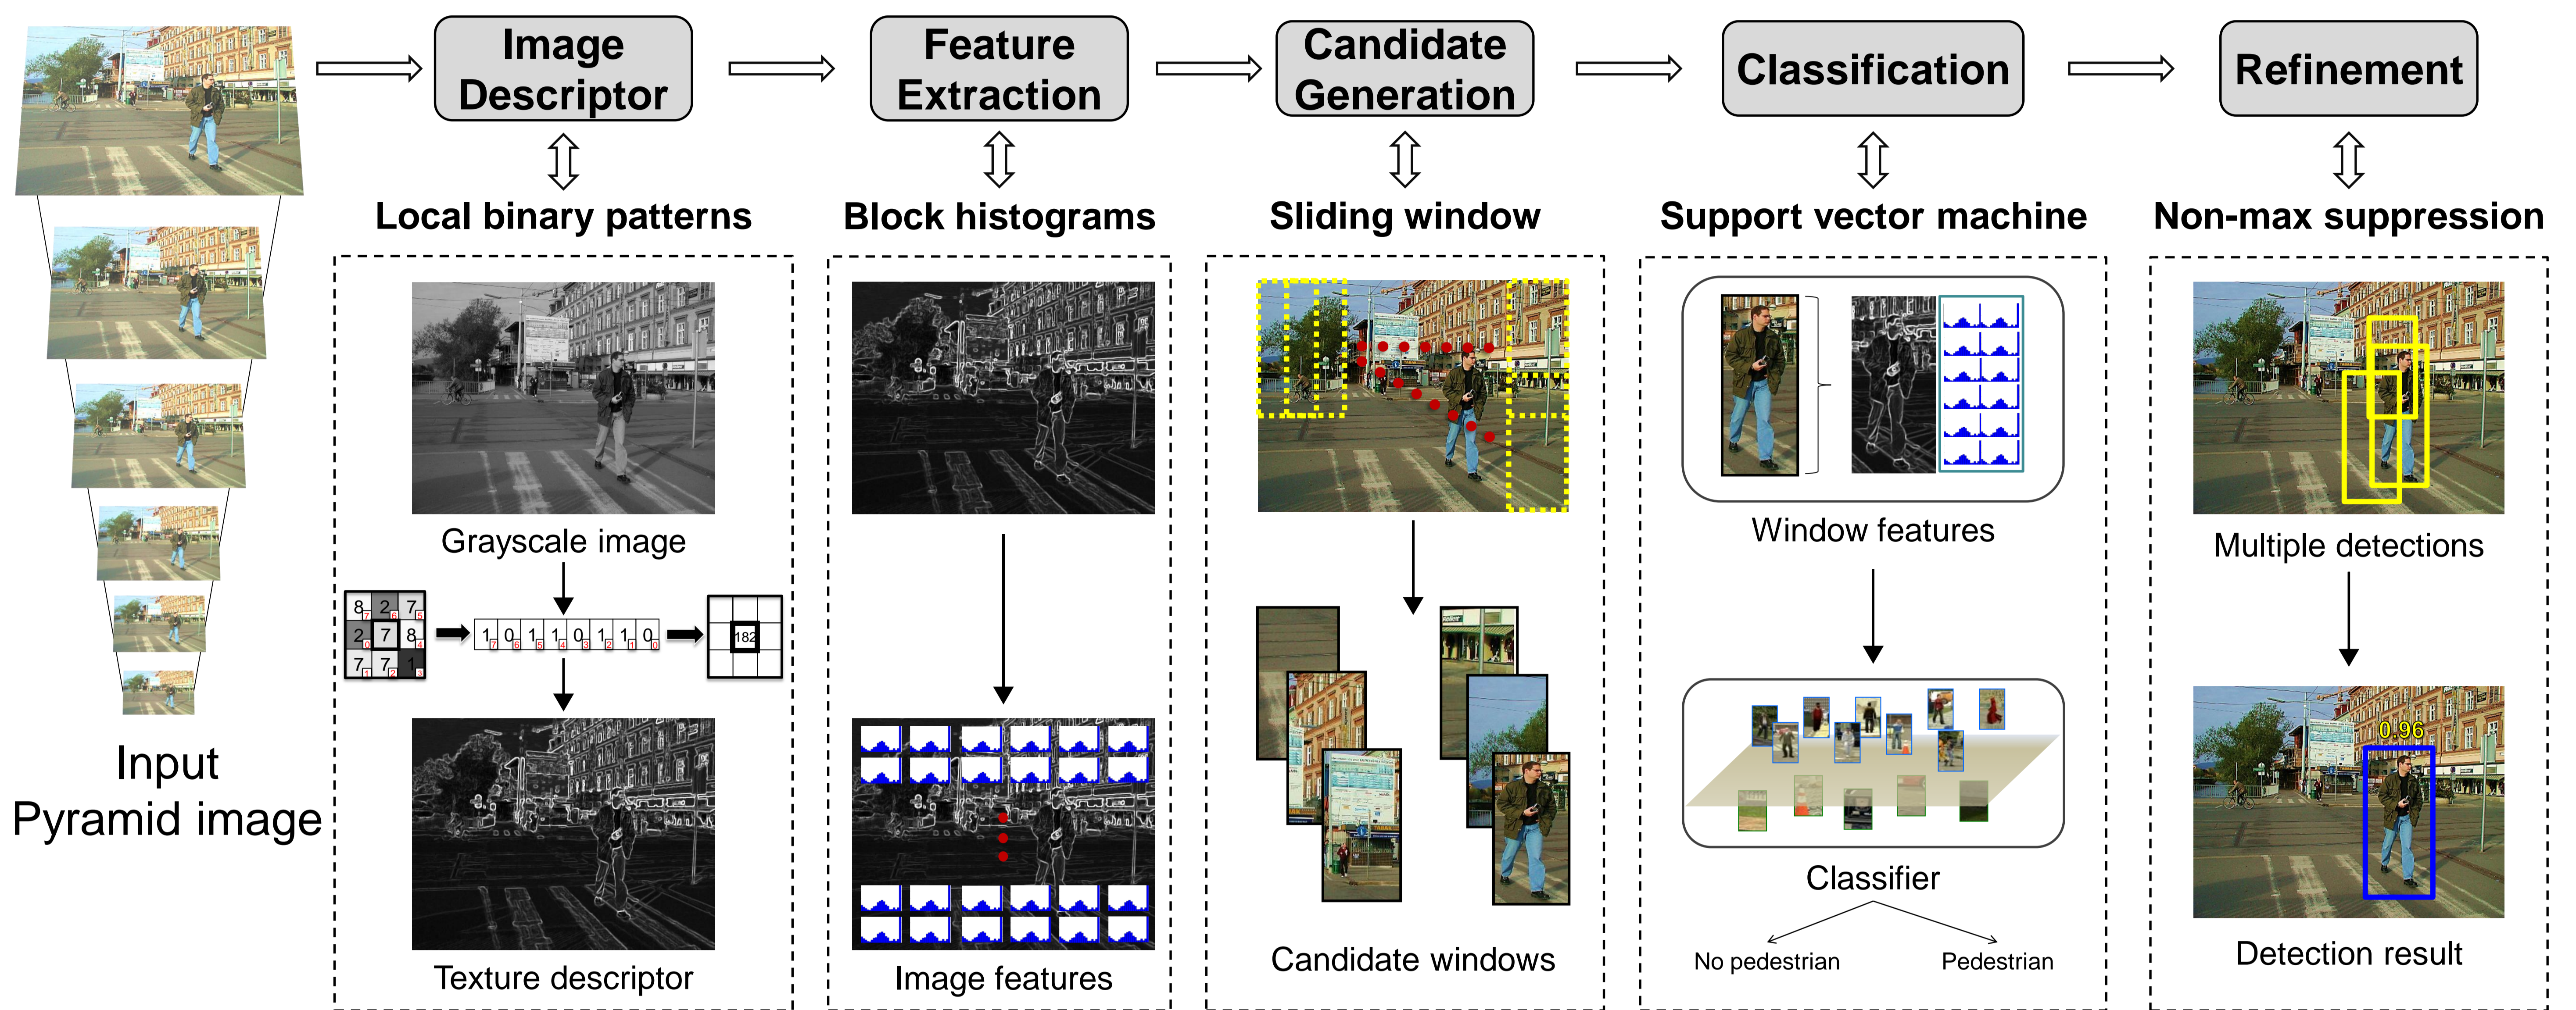


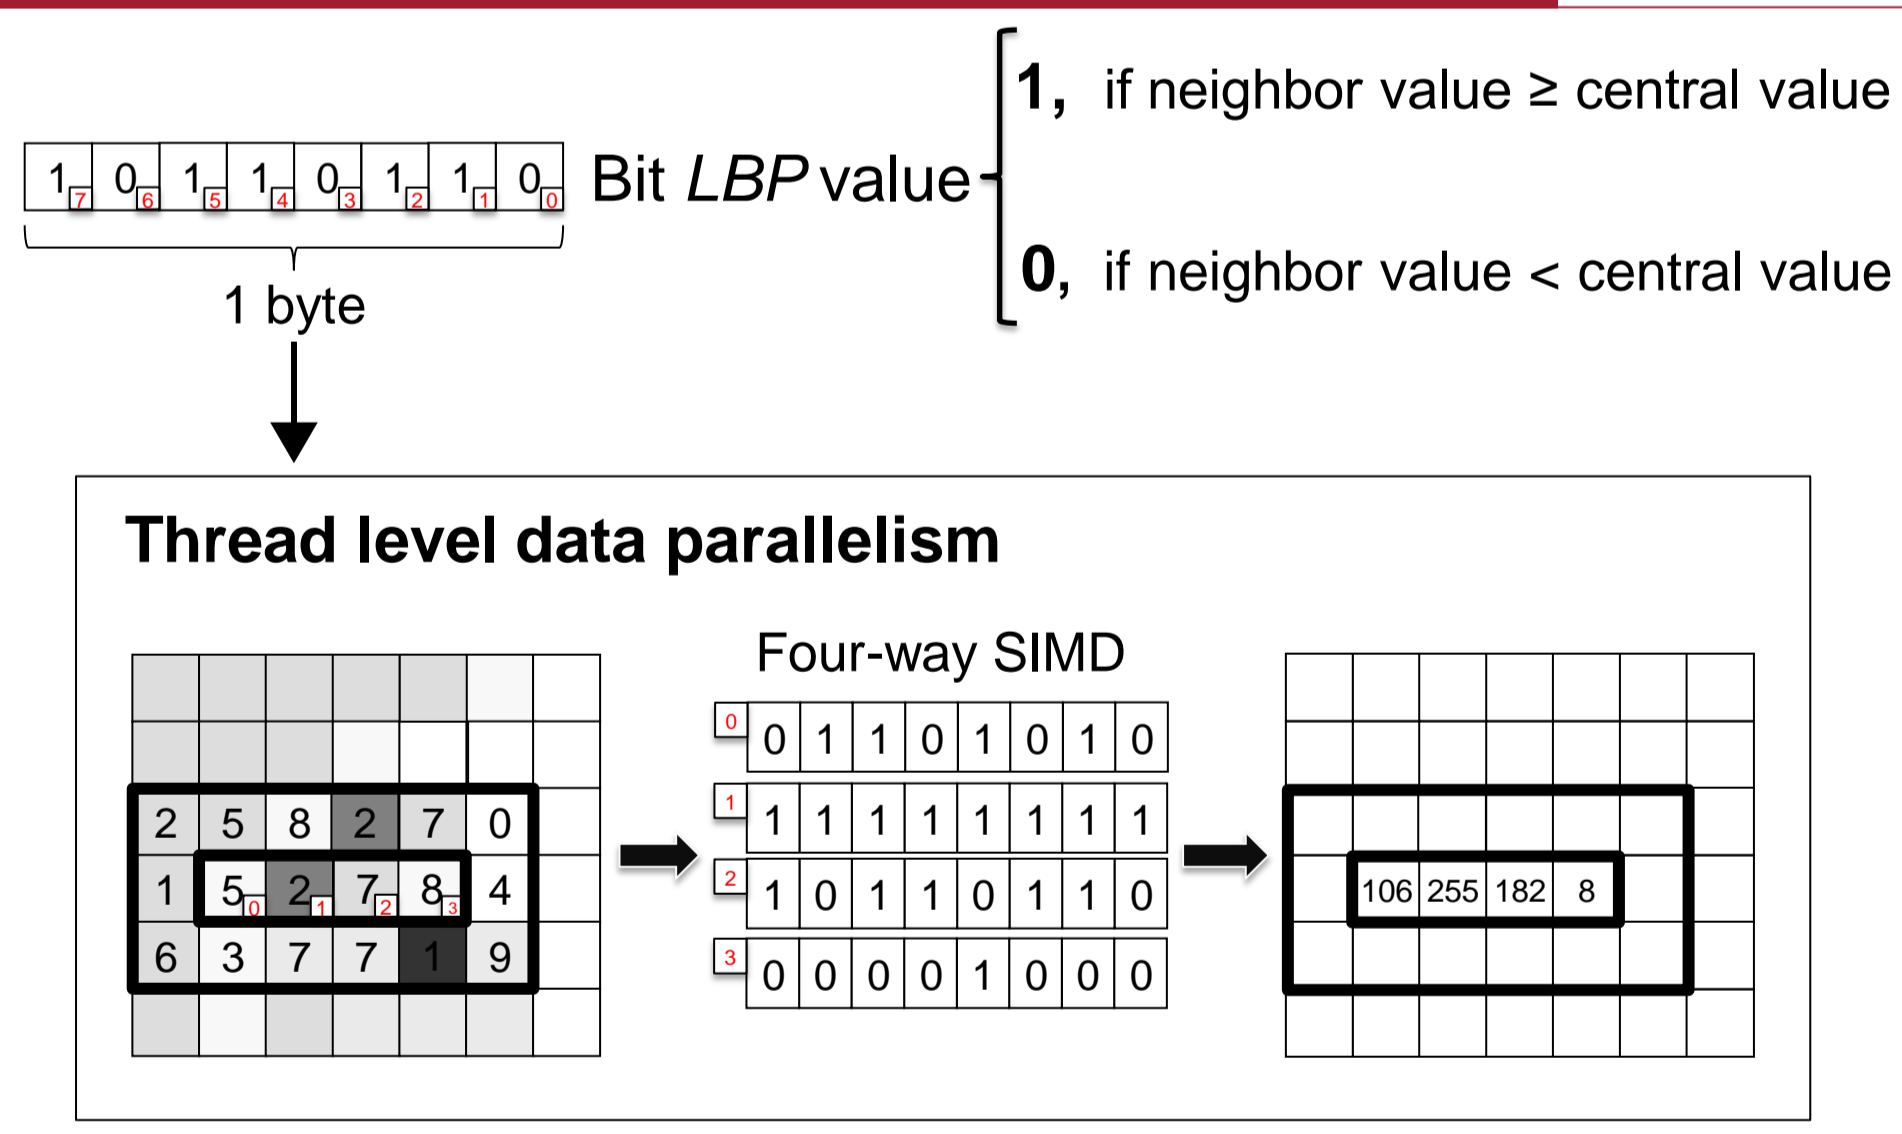
## Abstract

Pedestrian detection for autonomous driving has gained a lot of prominence during the last few years. Besides the fact that it is one of the hardest tasks within computer vision, it involves huge computational costs. The real-time constraints in the field are tight, and regular processors are not able to handle the workload obtaining an acceptable ratio of frames per second (fps). Moreover, multiple cameras are required to obtain accurate results, so the need to speed up the process is even higher. Taking the work in [1] as our baseline, we propose a CUDA implementation of a pedestrian detection system. Further, we introduce significant algorithmic adjustments and optimizations to adapt the problem to the GPU architecture. The aim is to provide a system capable of running in real-time obtaining reliable results.

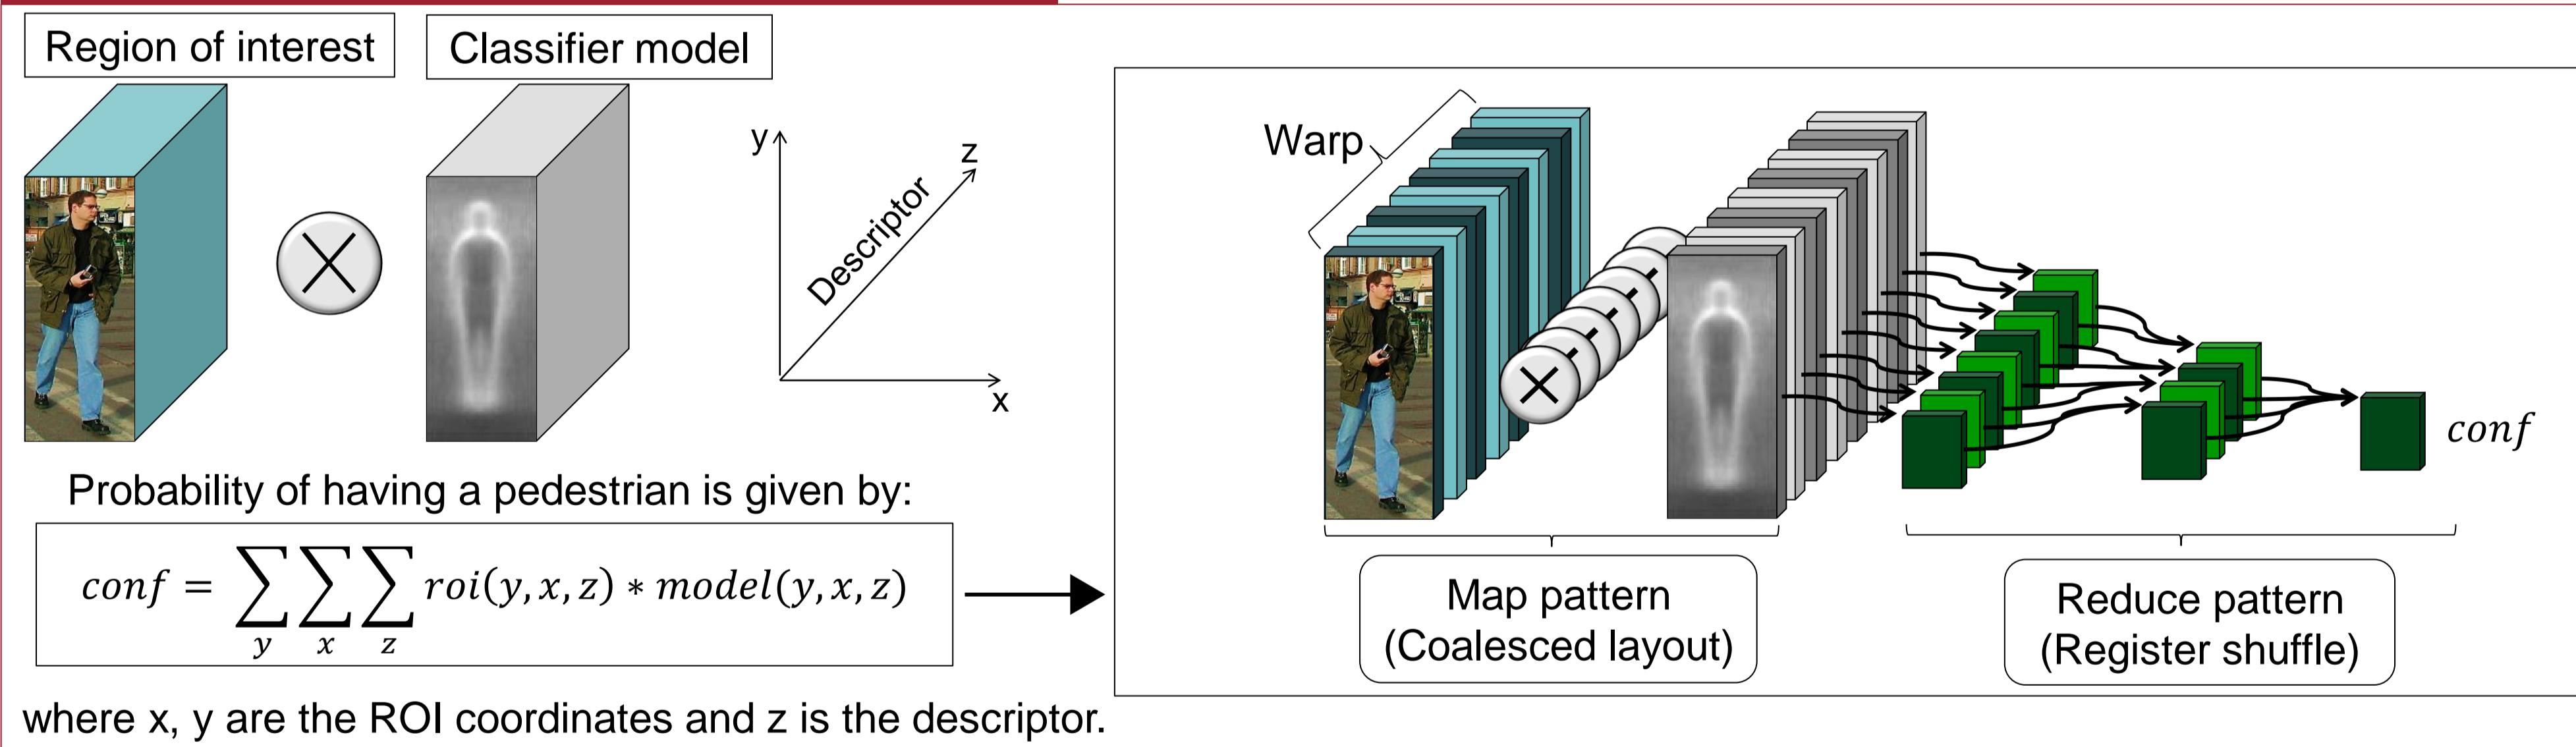
## Object Detection Pipeline



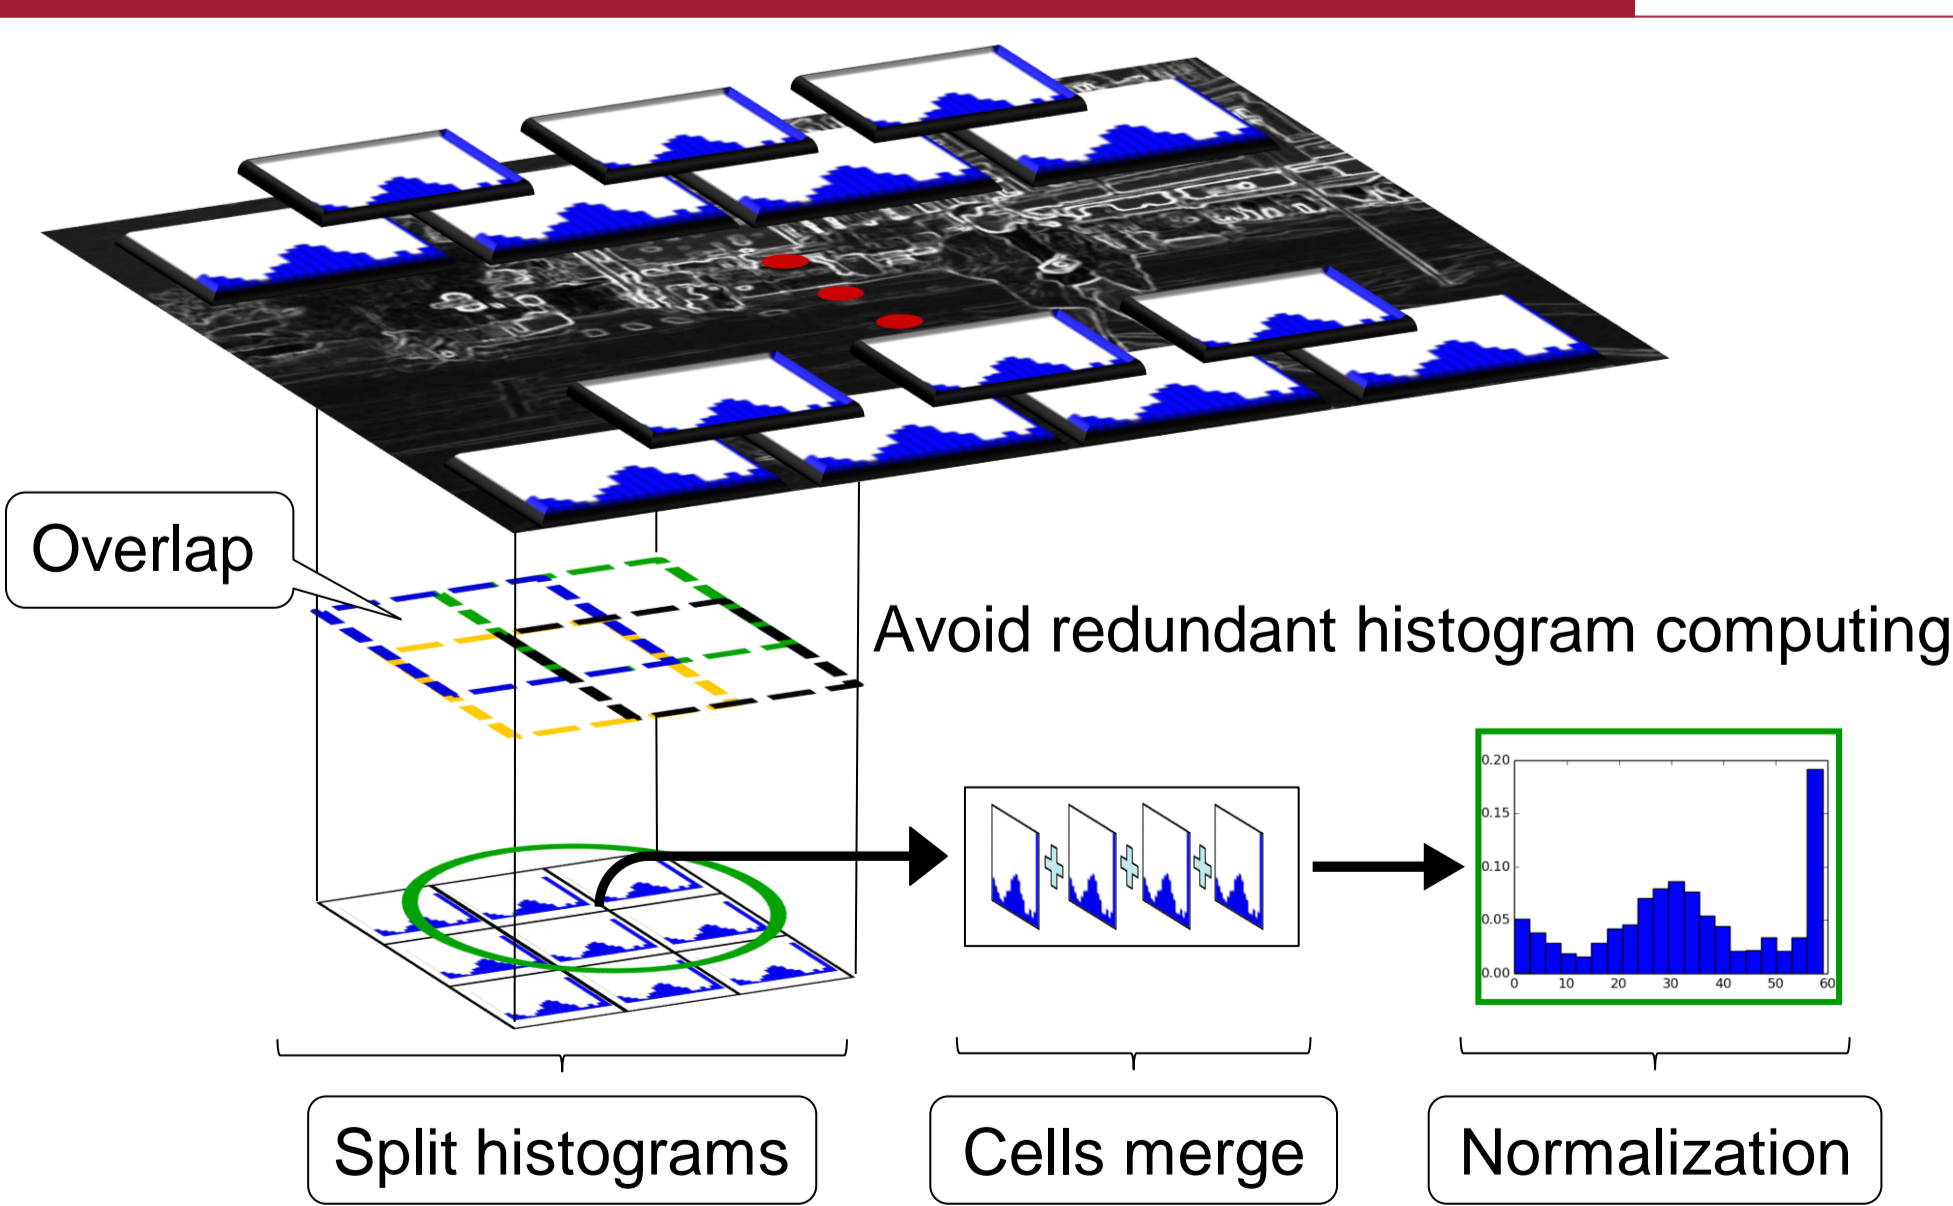
## Local Binary Patterns



## Support Vector Machine



## Block Histograms



## Results

