

Abstract

Pedestrian detection for autonomous driving is one of the hardest tasks within computer vision, and involves huge computational costs. Obtaining acceptable real-time performance, measured in frames per second (fps), for the most advanced algorithms is nowadays a hard challenge. Taking the work in [1] as our baseline, we propose a CUDA implementation of a pedestrian detection system that includes LBP and HOG as feature descriptors and SVM and Random forest as classifiers. We introduce significant algorithmic adjustments and optimizations to adapt the problem to the NVIDIA GPU architecture. The aim is to deploy a real-time system providing reliable results.

Object Detection Pipeline

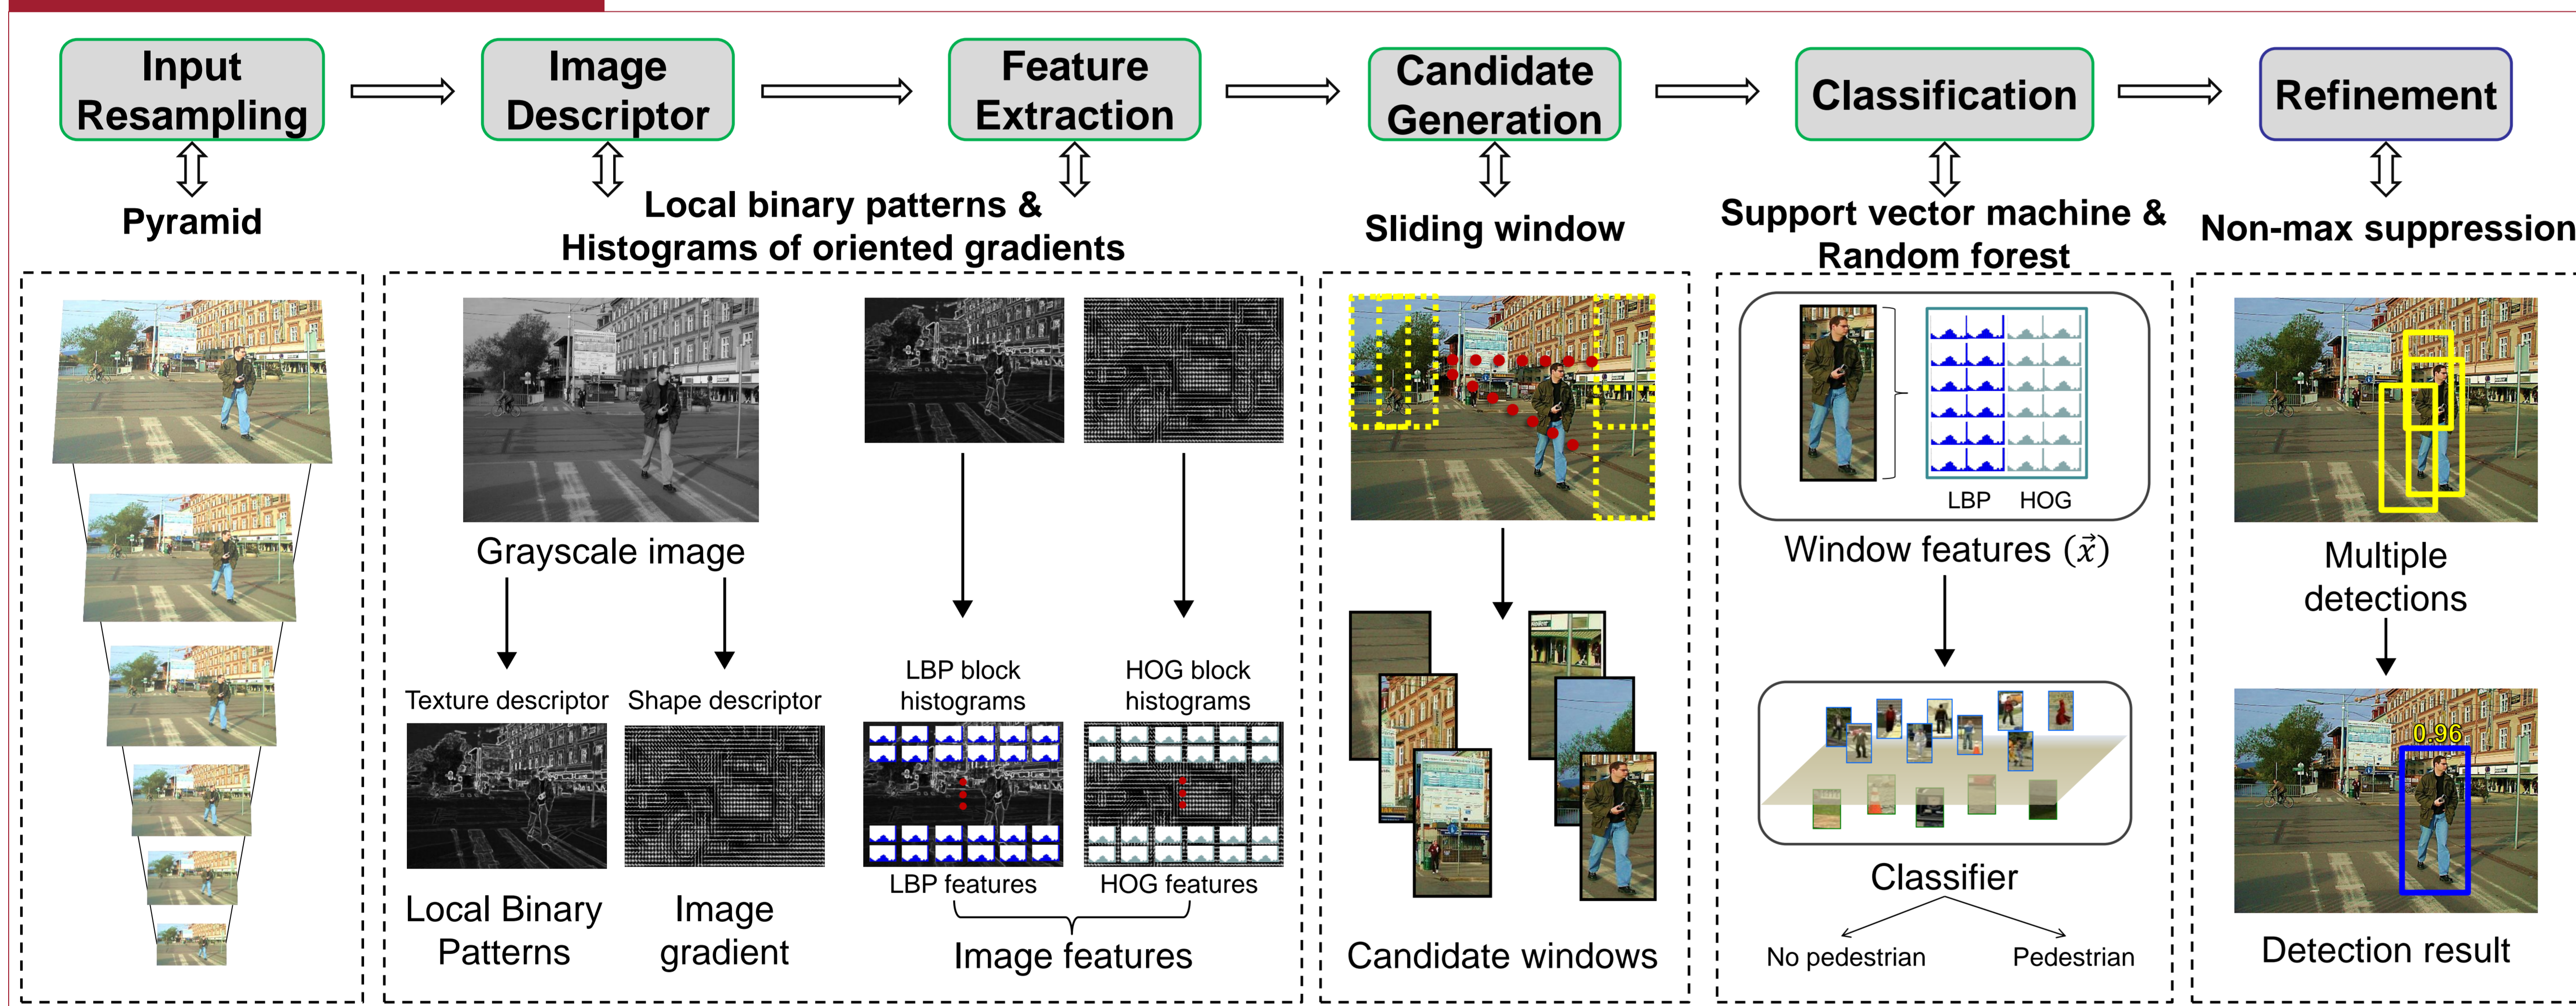
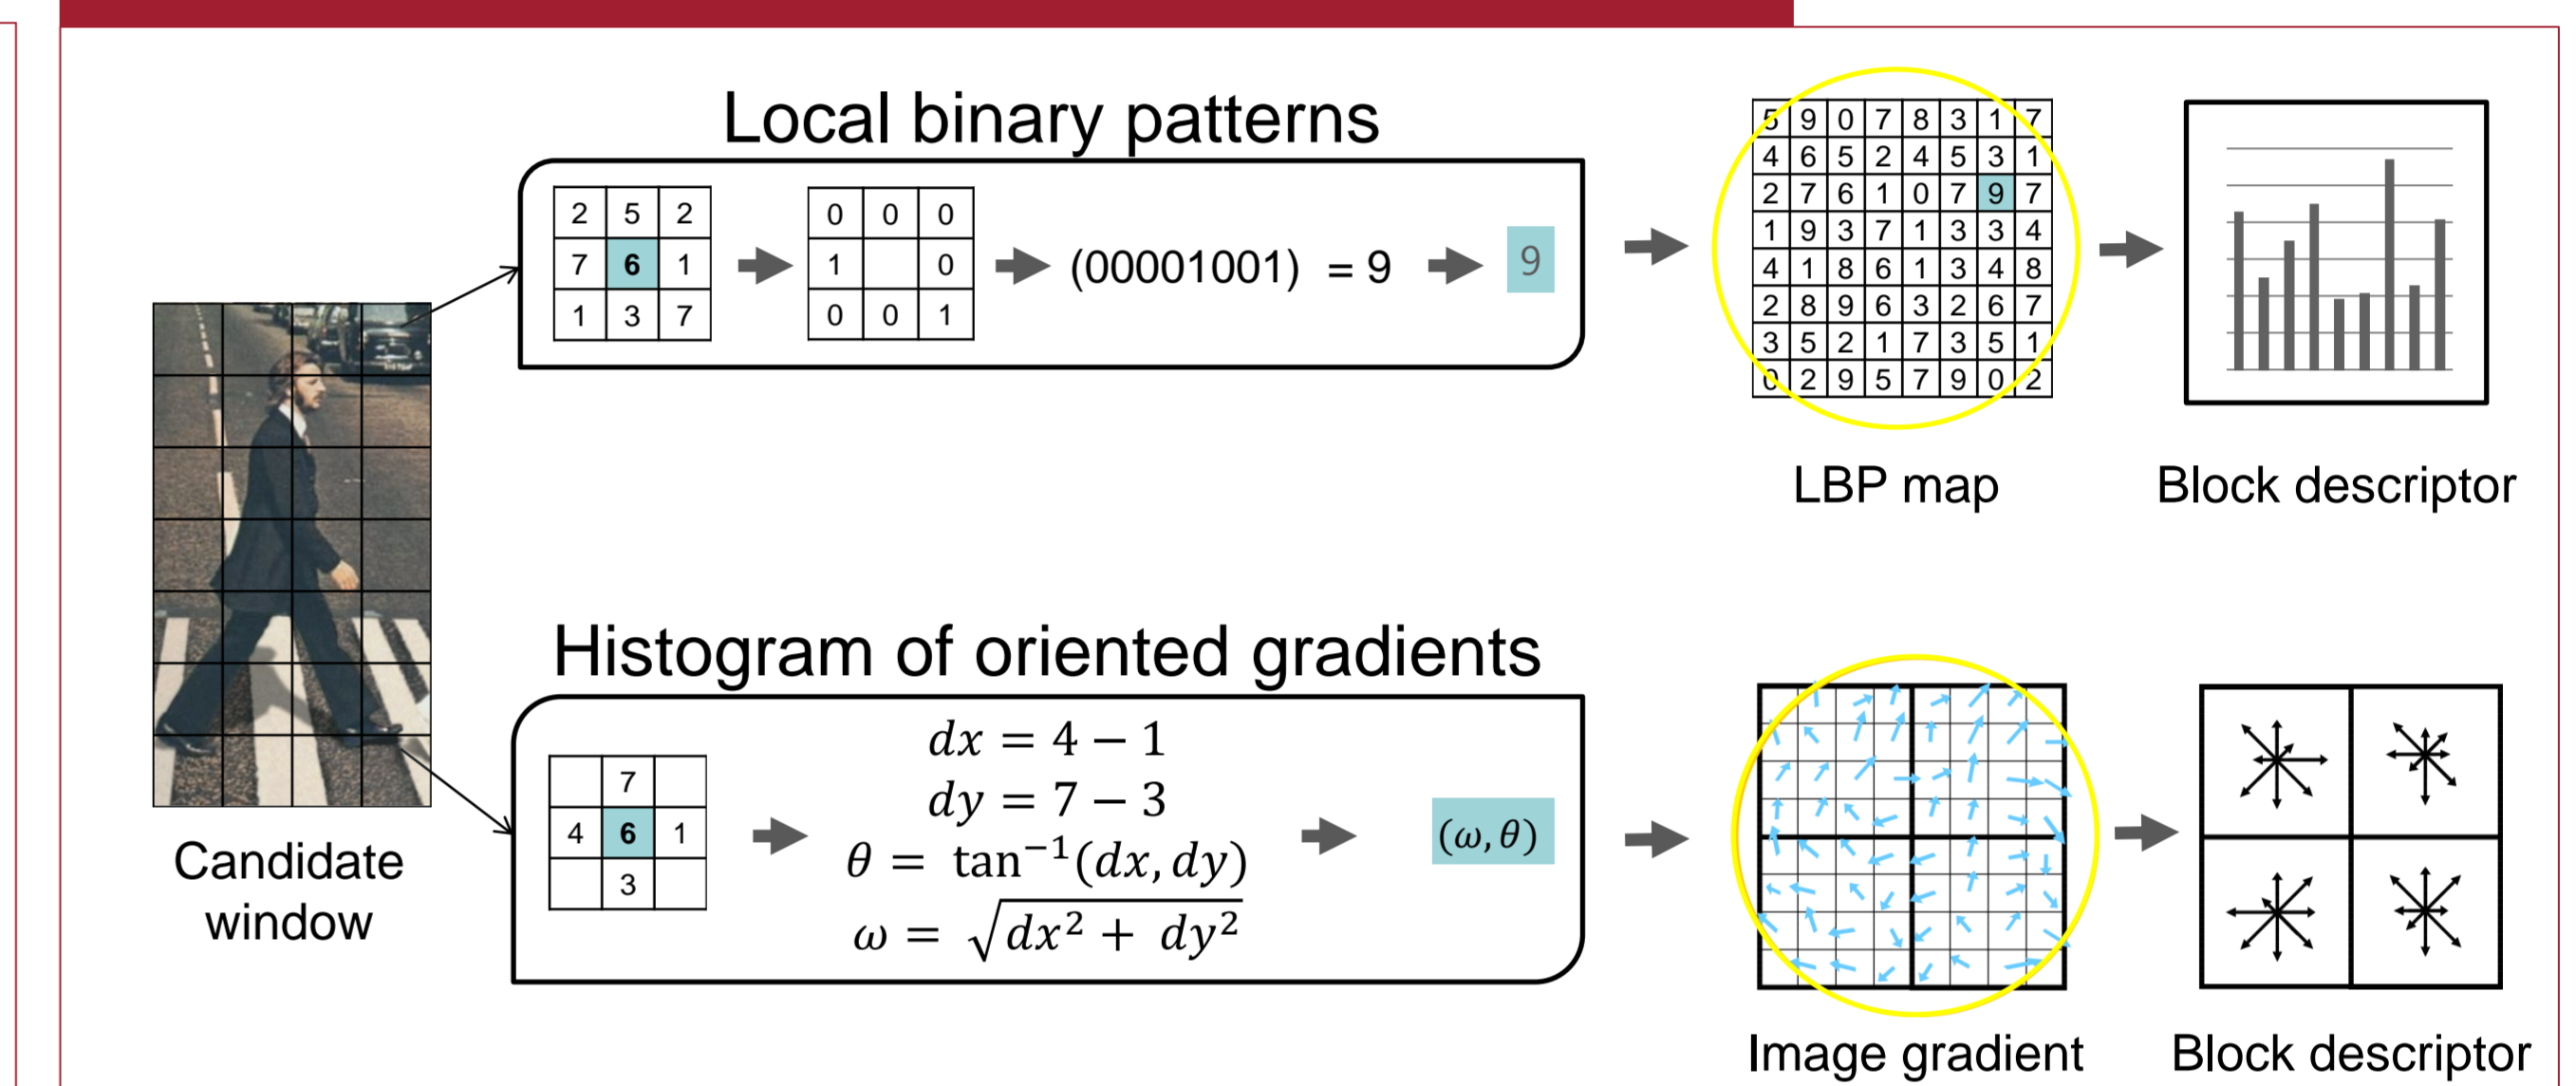
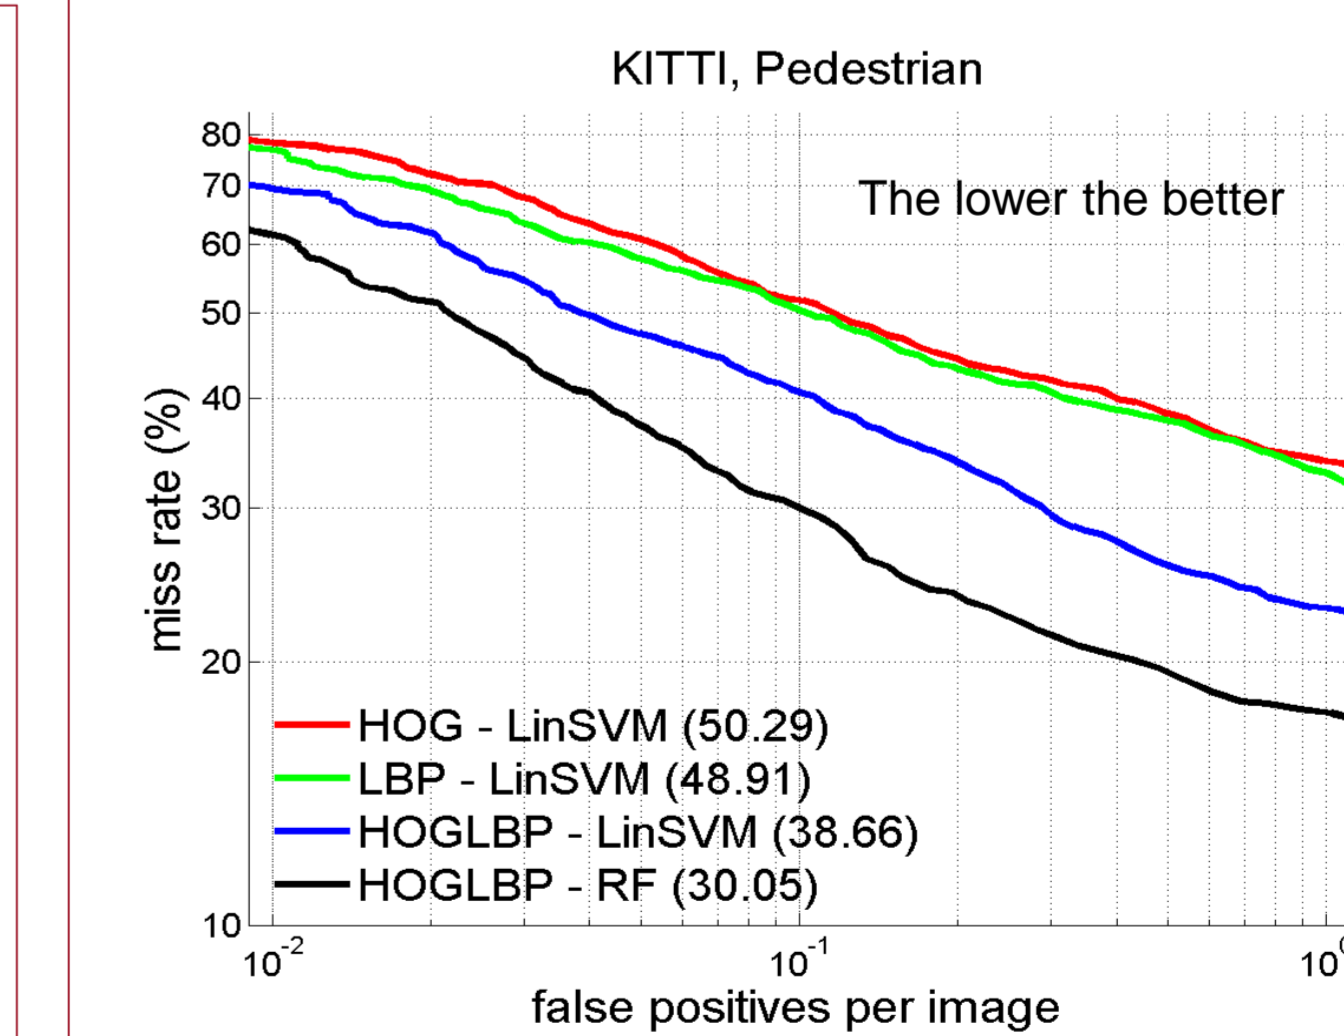



Image Descriptor & Feature Extraction



Results



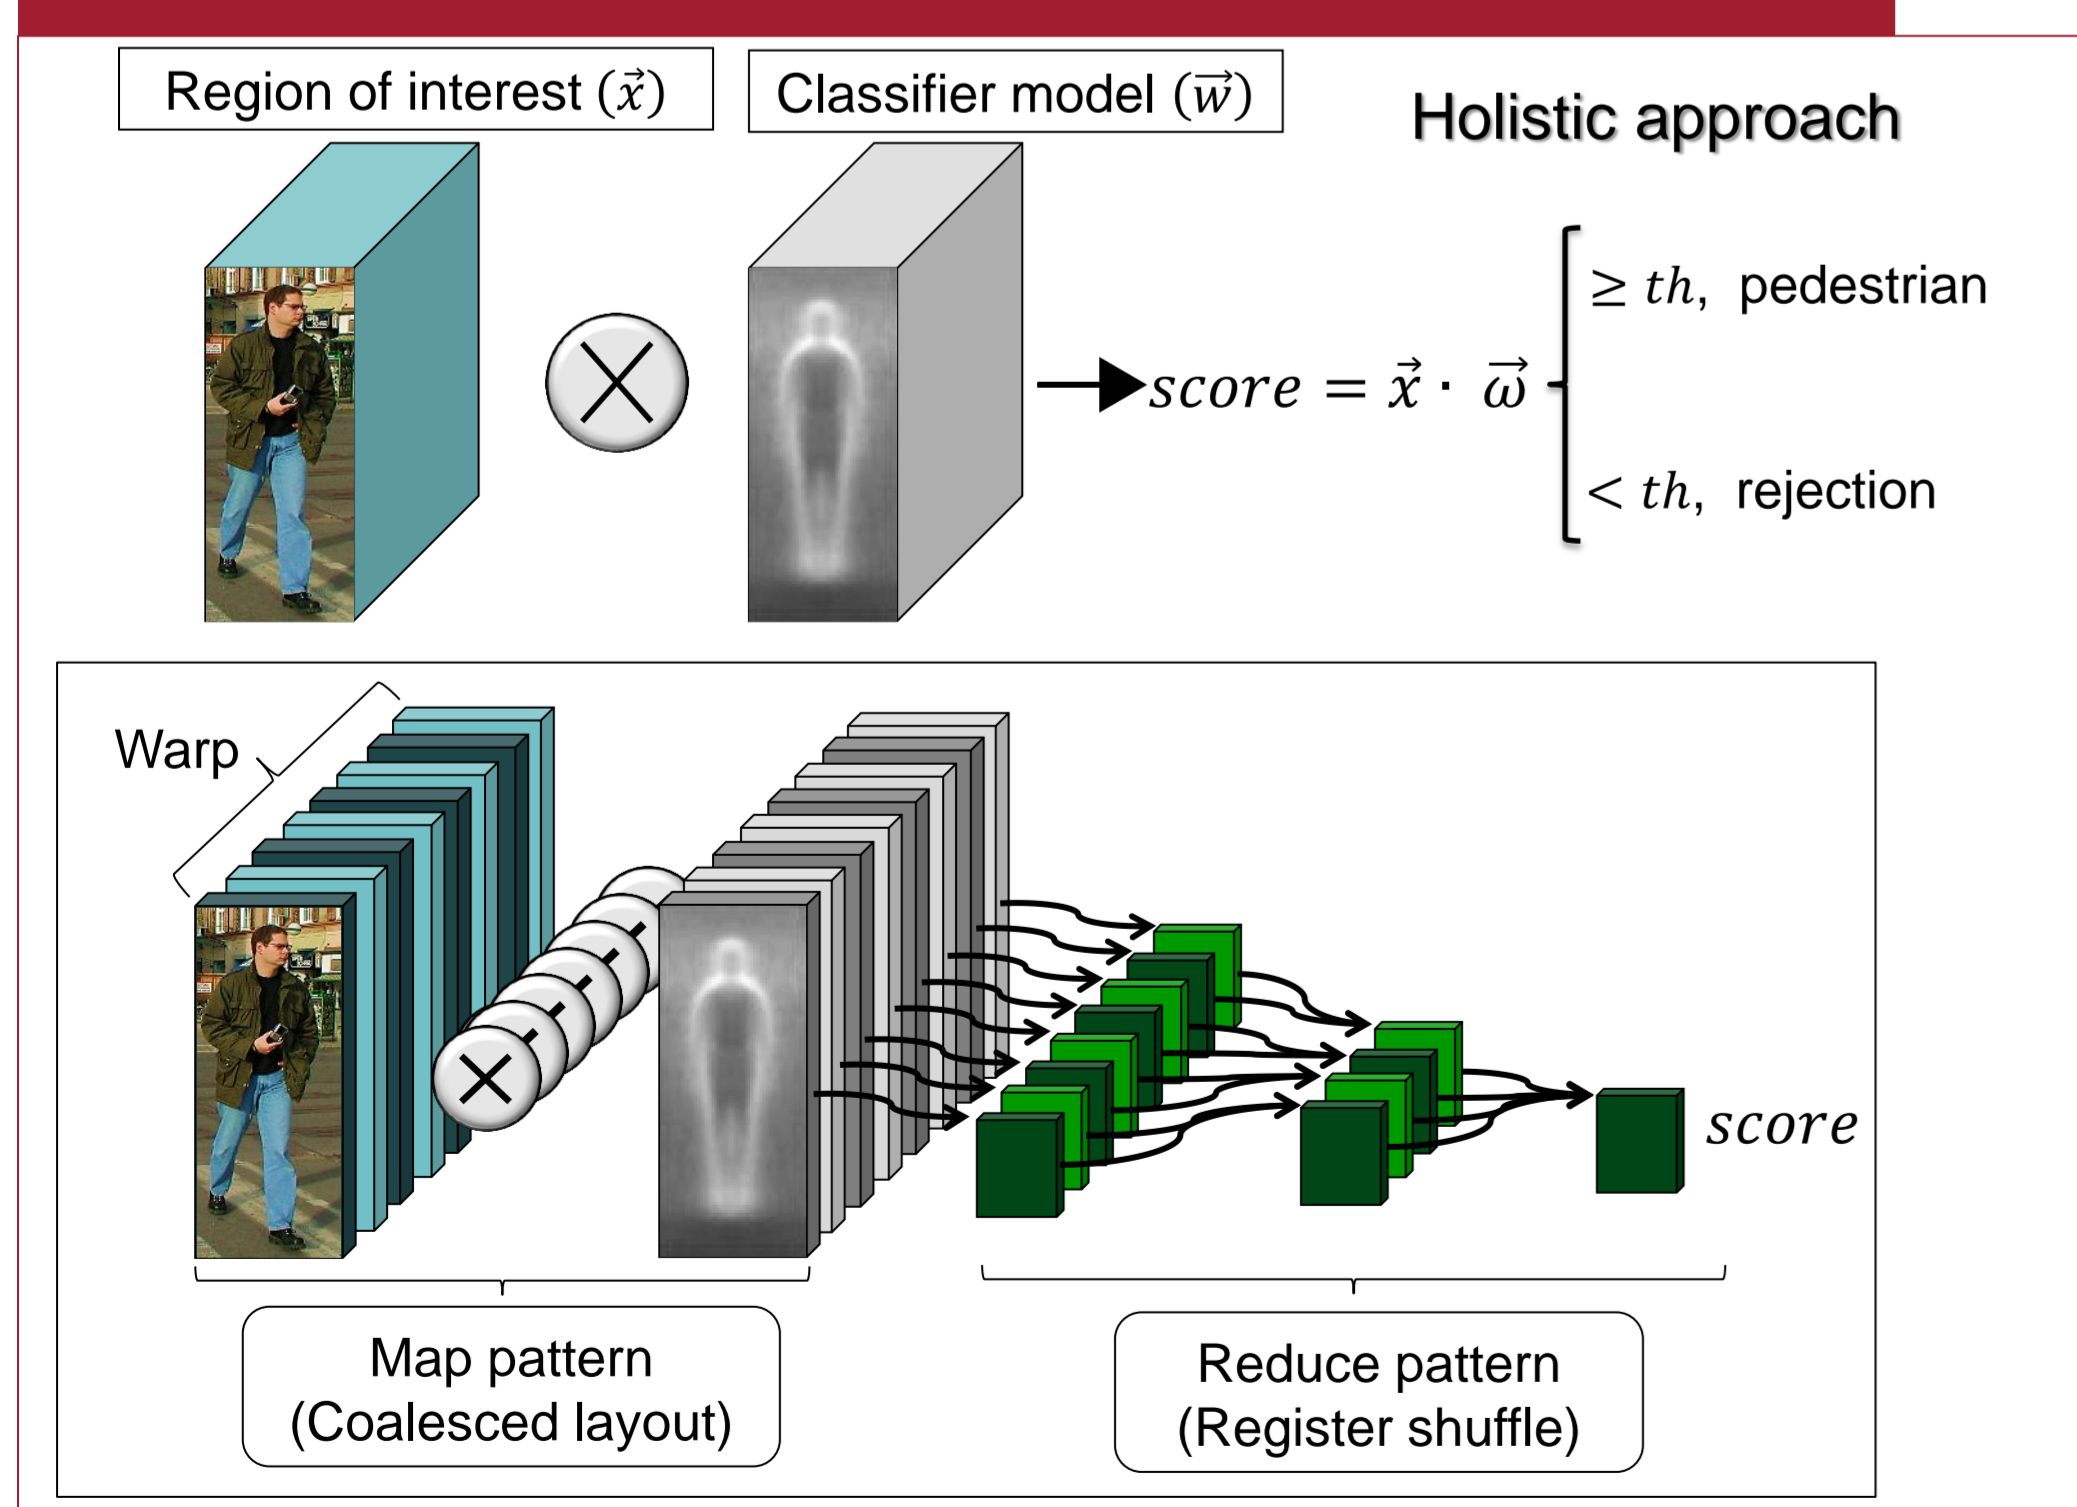
LBP + SVM	FPS	FPS/Watt
Intel i7-5930K	3,7	0,03
NVIDIA GTX 960	89	0,75
NVIDIA DrivePX	19	2,1

HOG + SVM	FPS	FPS/Watt
Intel i7-5930K	3,2	0,02
NVIDIA GTX 960	32	0,22
NVIDIA DrivePX	10	1,1

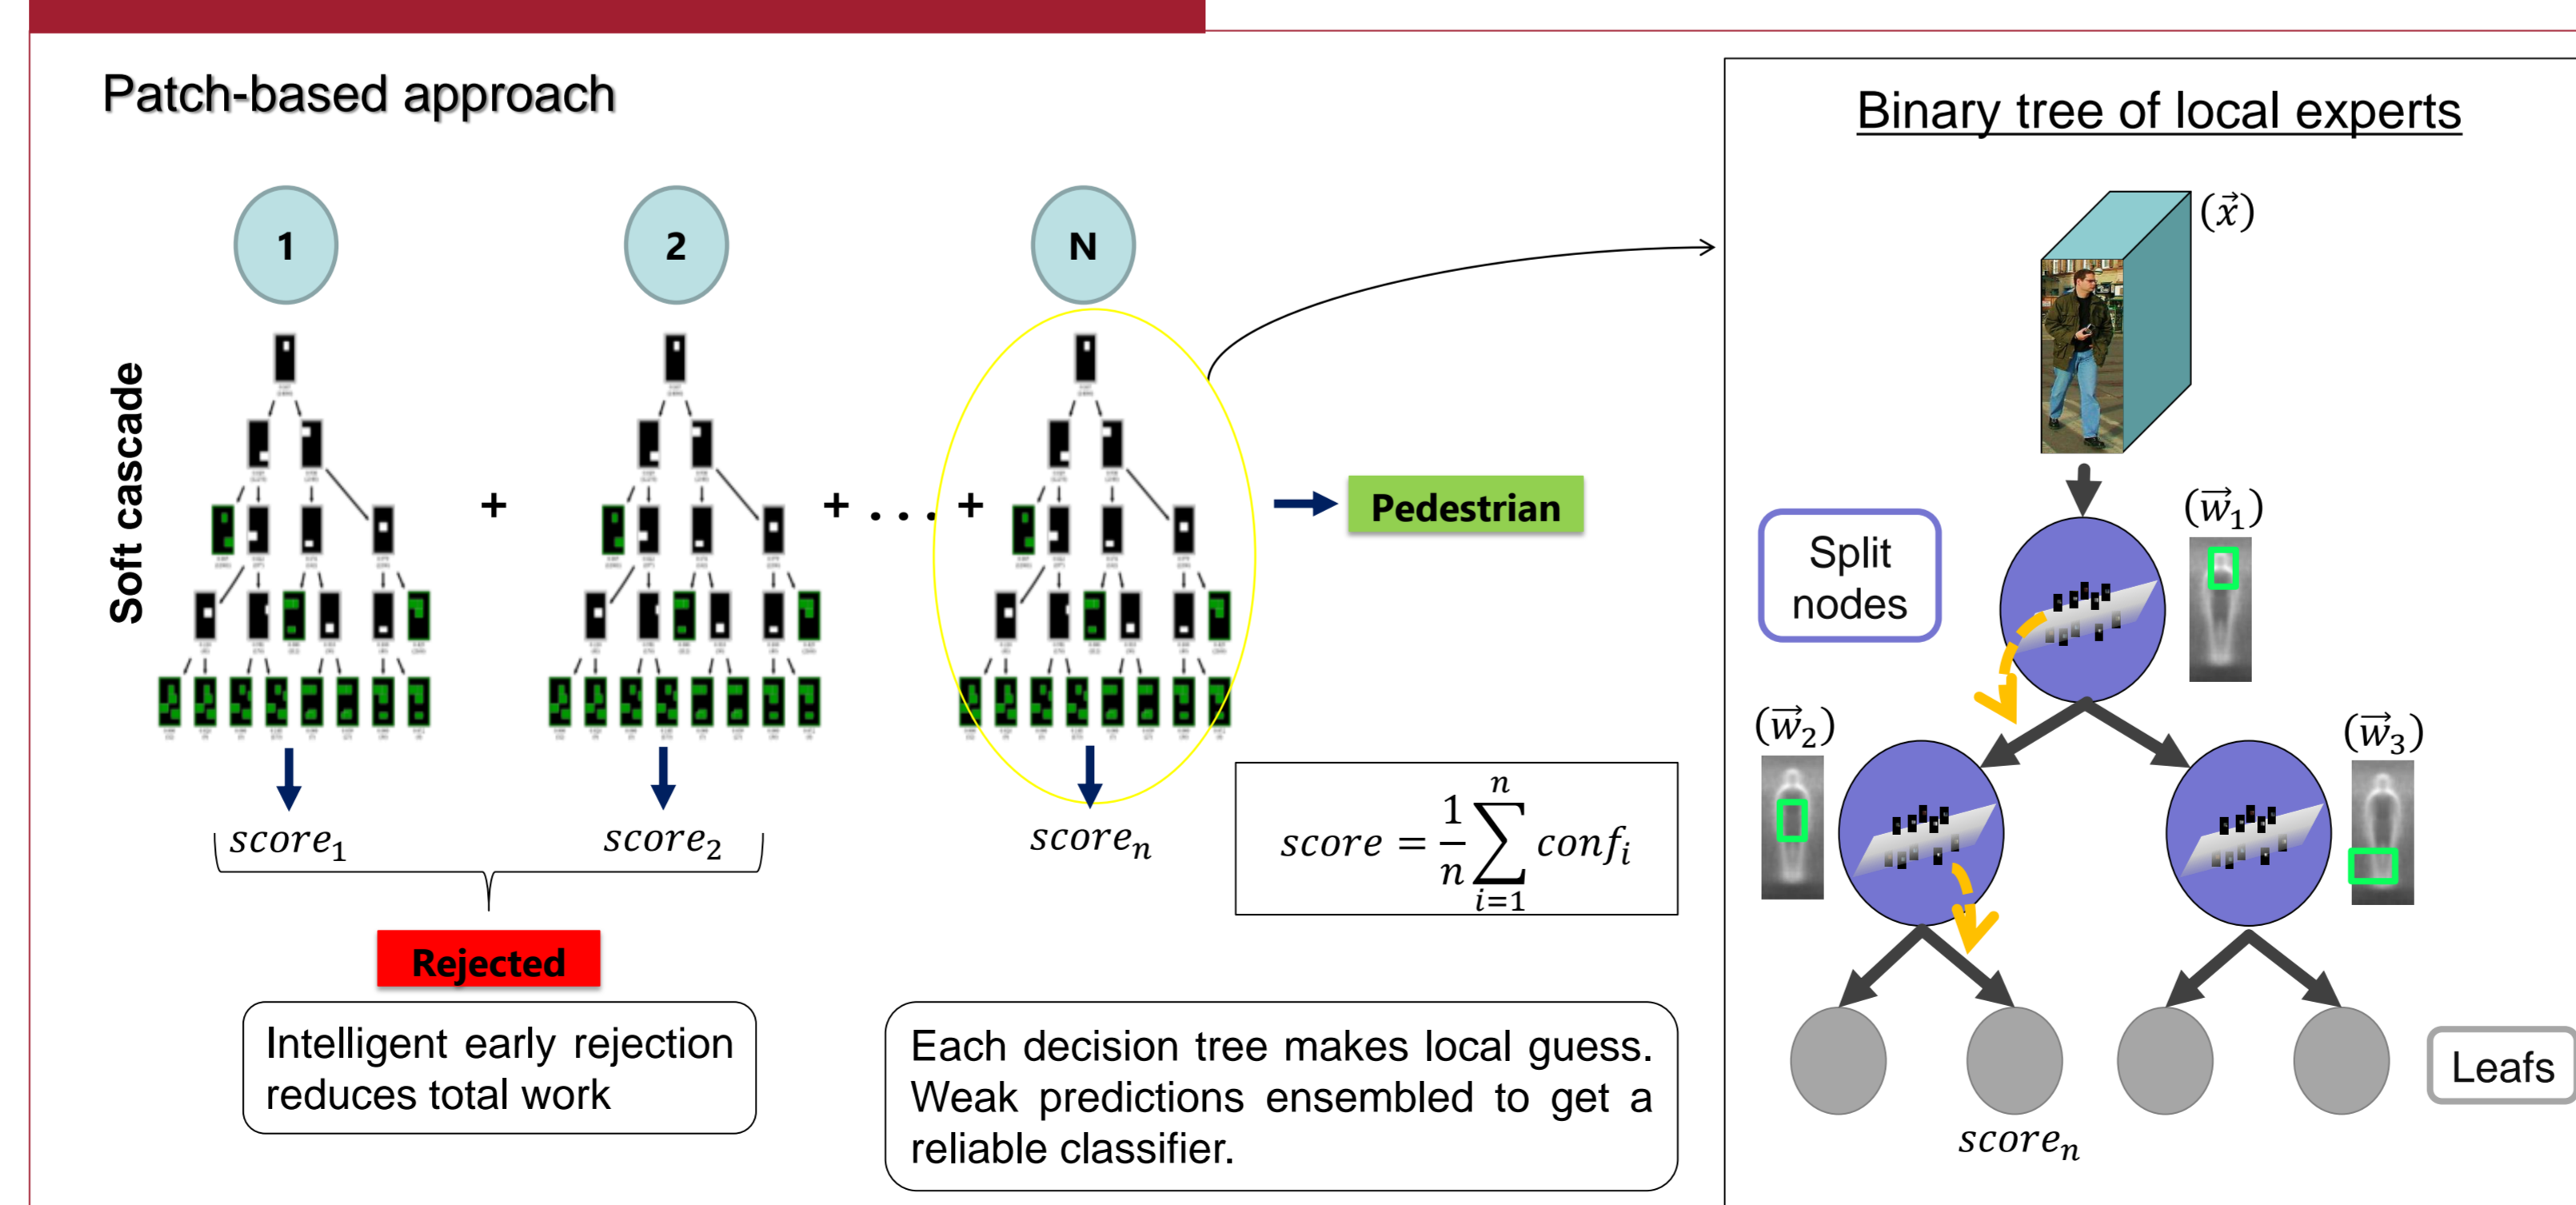
Conclusions & future work

- Our CUDA system is capable of running on real time obtaining accurate results.
- There is computational power left to integrate multiple cameras and additional ADAS systems such as lane departure warning, traffic signs recognition and 3D vision.

Classification: Support vector machine



Classification: Random forest



References:

[1] Marin, J., Vazquez, D., Lopez, A. M., Amores, J., & Leibe, B. (2013, December). Random forests of local experts for pedestrian detection. In *Computer Vision (ICCV), 2013 IEEE International Conference on* (pp.2592-2599).

Acknowledgements:

This work is supported by the Spanish MICINN projects TRA2014-57088-C2-1-R, TIN2014-53234-C2-1-R and by the Secretaria d'Universitats i Recerca del Departament d'Economia i Coneixement de la Generalitat de Catalunya (2014-SGR-1506) and DGT project SPIP2014-01352. Our research is also kindly supported by NVIDIA Corporation in the form of different GPU hardware.

