

# Speed-limit Sign Recognition with GPU Computing

Pınar Muyan-Özçelik\*, Vladimir Glavtchev\*<sup>§</sup>, Jeff Ota†, John D. Owens\*

\*University of California, Davis, †BMW Group Technology Office in Palo Alto, <sup>§</sup>NVIDIA Corporation in Santa Clara

UCDAVIS  
UNIVERSITY OF CALIFORNIA



## Goal



We investigate the use of different GPU-based implementations for performing **real-time** speed limit sign recognition on a **resource-constrained** embedded system.

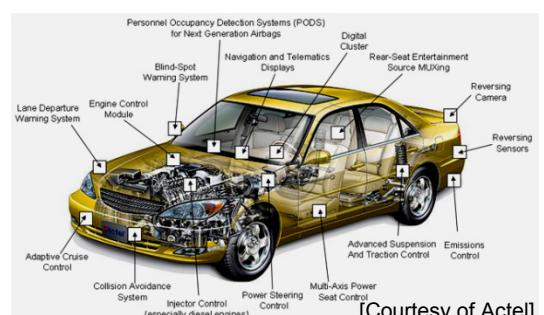
We juxtapose these alternative approaches in terms of:  
✓ success rate  
✓ run-time  
✓ how well it can be mapped to the GPU.

Also we determine:  
✓ best parameter combinations (recognition rate vs performance)  
✓ advantages of using a GPU instead of CPU  
✓ conditions which make a speed-limit sign easy/hard to recognize

## Motivation

Automotive tasks can be grouped as:  
✓ computer vision (most tasks)  
✓ signal processing  
✓ graphics  
✓ networking

GPUs are good fit for all the above except the networking tasks.



Replace above with the GPU, which:  
✓ simplifies design  
✓ is cheap  
✓ is programmable

Hence, this study serves as a **proof of concept** for the use of GPU computing in automotive tasks.

## Implementation and Results

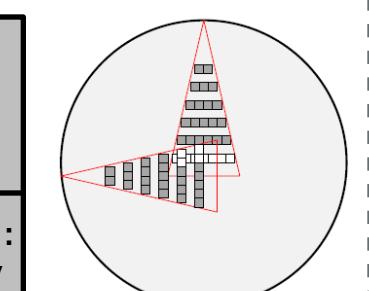
### Feature-based Approach

#### Pipeline:

New video frame

Preprocessing:

Sobel edge detect with thresholding



**Detection Step 1:**

Radial symmetry<sup>[1]</sup> voting:  
Pixels vote in the direction of their gradient. Votes accumulate as they overlap to form an overall voting image.

**Detection Step 2:**

Maximum reduction to find top speed limit sign candidate locations



[Input Image]



[Sobel result: detected edges]



[Radial symmetry voting results]



[Detected speed-limit signs]

Processor	Core Config.	Detection Runtime	Execution Rate	Speedup
Dual-Core Intel Atom 230 @ 1.6GHz	2	235 ms	4.25 fps	--
Intel Core2 6300 @ 1.86GHz	2	130 ms	7.4 fps	~1.8X
NVIDIA GeForce 9200M GS 128MB	1 : 8	65 ms	15 fps	~3.6X
NVIDIA GeForce 8600 GTS 256MB	4 : 32	23 ms	43.5 fps	~10.2X

Approach	Detection Runtime	Classification Runtime	Total Runtime <sup>◆</sup>	Success Rate <sup>◆</sup>	Mapping to GPU
Feature-based	65 ms <sup>†</sup> (15 fps)	N/A	65ms	80% detection	CUDA is used to parallelize Sobel filtering, maximum-reduction, and vertices for voting patterns. The graphics pipeline (OpenGL) is used to accelerate radial symmetry voting.
Template-based	65 ms <sup>‡</sup>	40 ms <sup>‡</sup>	125 ms <sup>‡</sup> (8 fps)	90% recognition with NO true negatives and NO false positives	Mapping the pipeline to CUDA kernels was straightforward since all stages involve data-parallel computation. CUFFT library is utilized.
SIFT-based	N/A	N/A	300 ms <sup>‡</sup> (3 fps)	75% recognition with 4 true negatives and 9 false positives	siftGPU <sup>[4]</sup> is used. Almost all stages which perform SIFT feature extraction are performed on the GPU. Matching is also done on the GPU.

<sup>†</sup>On Intel Core2 6300 @ 1.86GHz 1.7GB RAM and 128MB GeForce Quadro NVS 150M (1 SM)

<sup>‡</sup>On Intel Core2 Duo T8300 @ 2.4GHz 4GM RAM and 128MB GeForce 8400M GS (2 SM)

### Template-based Approach

(extension of Javid et al.'s approach [2])

#### Pipeline:

New video frame

Preprocessing:

CLAHE to enhance contrast

**Detection:**

FFT correlations to output candidate location and size

**Classification:**

FFT correlations to find the speed limit in the candidate location

**Temporal Integration:**

Majority voting to display speed limit in the last few frames

**Composite Filter:**

(kth-law nonlinear

MACE filter

generated offline w/ Matlab

classification composite filters w/ different in-plane rotations

detection composite filters w/ different sizes

... (also rotate around X-axis) ...

**CLAHE:**

(Contrast Limited Adaptive Histogram Equalization)

**Input Image:**

**Sobel result: detected edges**

**CLAHE:**

(Contrast Limited Adaptive Histogram Equalization)

**Before:**

**After:**

**Detected speed-limit signs:**

### SIFT-based Approach

(SIFT: Scale Invariant Feature Transform [3])

#### Pipeline:

New video frame

Preprocessing:

Extract SIFT features

**Recognition:**

Perform SIFT matching to find speed limit in the current frame

**Temporal Integration:**

Majority voting to display speed limit in the last few frames

**Extracting SIFT features:**

Detect scale-space extrema to find potential points

Select stable ones as "keypoints".

Assign pos & size

Assign one or more orientations to each keypoint

Generate descriptor (feature vector) for each keypoint to provide invariance to viewpoint and illumination change

templates represented by their SIFT features (generated offline)

60

60

60

...

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60

60