

MPC Workshop
Karlsruhe
10/7/2009

FPGA Implementation of a HOG-based Pedestrian Recognition System

Sebastian Bauer | sebastian.bauer@fh-aschaffenburg.de

Laboratory for Pattern Recognition and Computational Intelligence
Prof. Dr. Ulrich Brunsmann

Content

- Motivation
- Method
- Implementation
- Evaluation
- Summary and Outlook



Motivation

| Method | Implementation | Evaluation | Summary

■ Problem Scope

- 1,200,000 road traffic deaths per year [WHO/Worldbank 2004]
- Majority among vulnerable road users (VRU)
- **No advanced driver assistance system (ADAS) available for daylight pedestrian recognition**



Motivation

| Method | Implementation | Evaluation | Summary

■ Problem Scope

- 1,200,000 road traffic deaths per year [WHO/Worldbank 2004]
- Majority among vulnerable road users (VRU)
- **No advanced driver assistance system (ADAS) available for daylight pedestrian recognition**



■ State-of-the-Art

- HOG detection system yields best performance [Dollár 2009],[Enzweiler 2009]
- Challenge: CPU computation time: > 10s (VGA)

Motivation

| Method | Implementation | Evaluation | Summary

■ Problem Scope

- 1,200,000 road traffic deaths per year [WHO/Worldbank 2004]
- Majority among vulnerable road users (VRU)
- **No advanced driver assistance system (ADAS) available for daylight pedestrian recognition**



■ State-of-the-Art

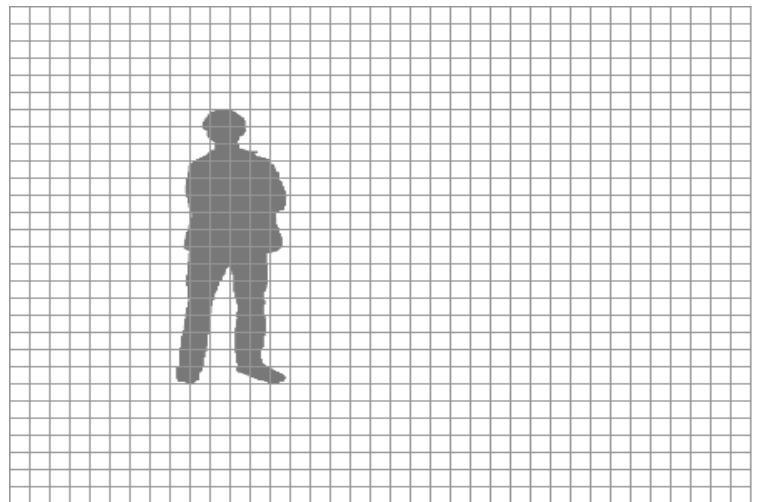
- HOG detection system yields best performance [Dollár 2009],[Enzweiler 2009]
- Challenge: CPU computation time: > 10s (VGA)
- **Goal: Real-time HOG system with FPGA**

Motivation | Method | Implementation | Evaluation | Summary

■ HOG Detection System

- **Sliding Detection Window**

- For each window:
 1. Generate **Feature Set / „Descriptor“** (HOG)
 2. **Classify** pedestrian / non-pedestrian

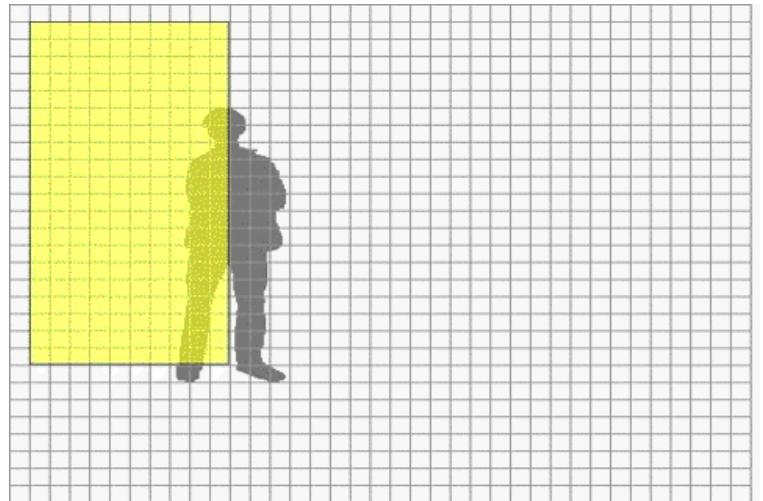


Motivation | Method | Implementation | Evaluation | Summary

■ HOG Detection System

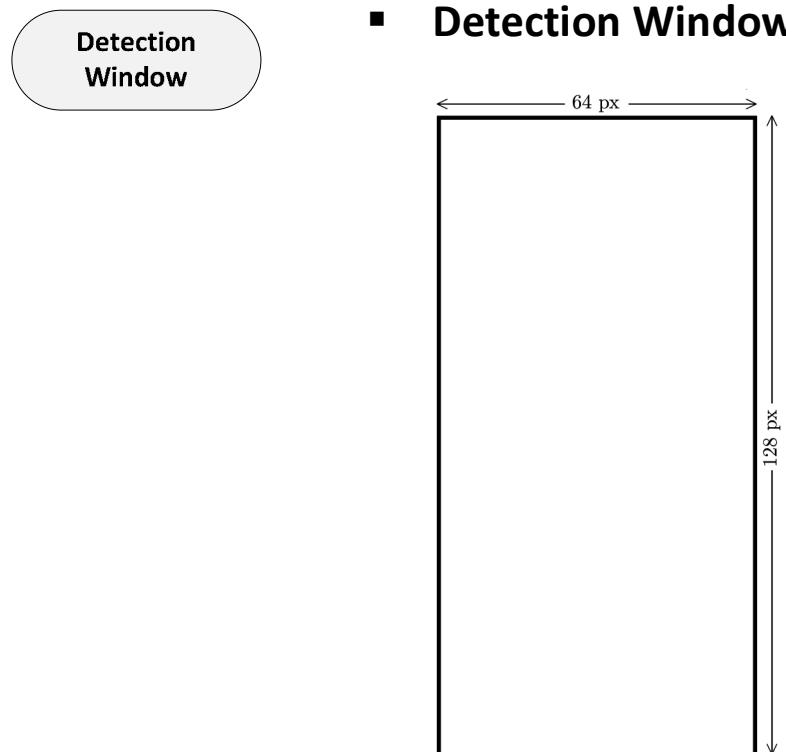
- **Sliding Detection Window**

- For each window:
 1. Generate **Feature Set / „Descriptor“** (HOG)
 2. **Classify** pedestrian / non-pedestrian



Motivation | Method | Implementation | Evaluation | Summary

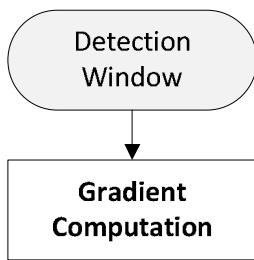
■ HOG Descriptor – Histograms of Oriented Gradients [Dalal 2005]



Method

Motivation | Implementation | Evaluation | Summary

■ HOG Descriptor – Histograms of Oriented Gradients [Dalal 2005]



■ Gradient Computation

■ Point derivatives

$$G_x = M_x * I \quad M_x = [-1 \ 0 \ 1]$$

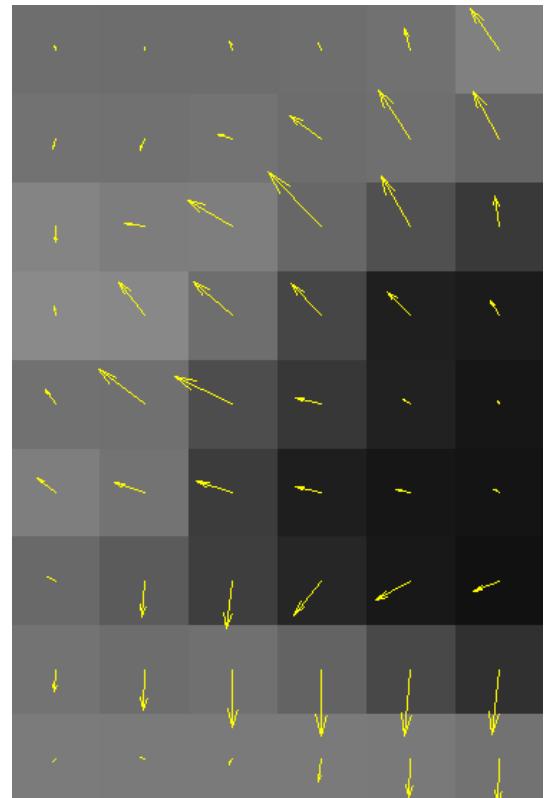
$$G_y = M_y * I \quad M_y = [-1 \ 0 \ 1]^T$$

■ Magnitude

$$|G(x, y)| = \sqrt{G_x(x, y)^2 + G_y(x, y)^2}$$

■ Orientation

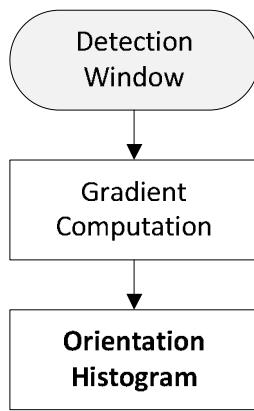
$$\tan(\phi(x, y)) = \frac{G_y(x, y)}{G_x(x, y)}$$



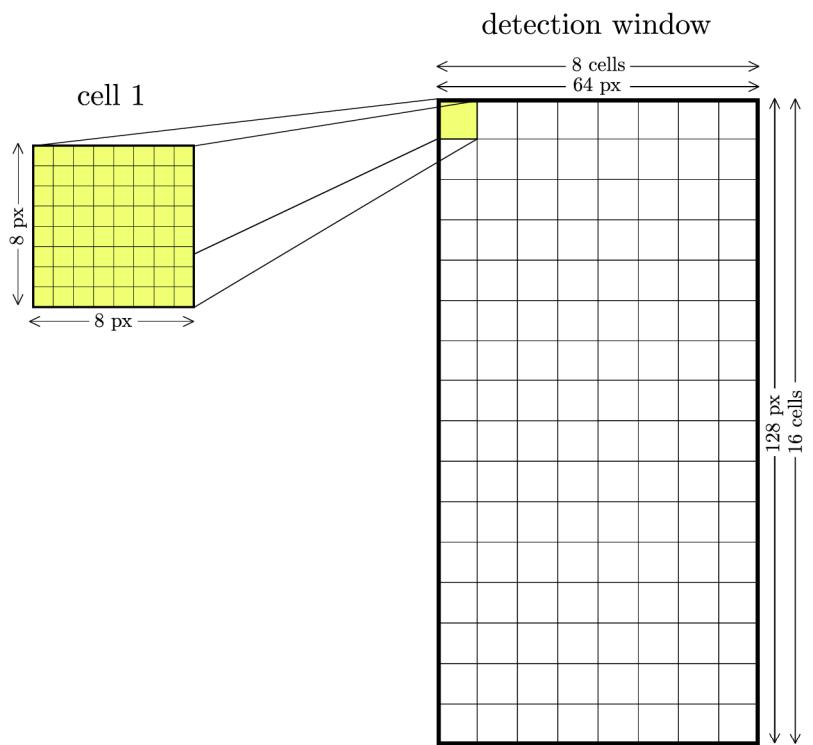
Method

Motivation | Implementation | Evaluation | Summary

■ HOG Descriptor



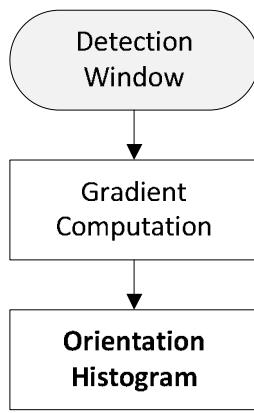
- Orientation Histogram
 - Detection Window subdivided into Cells



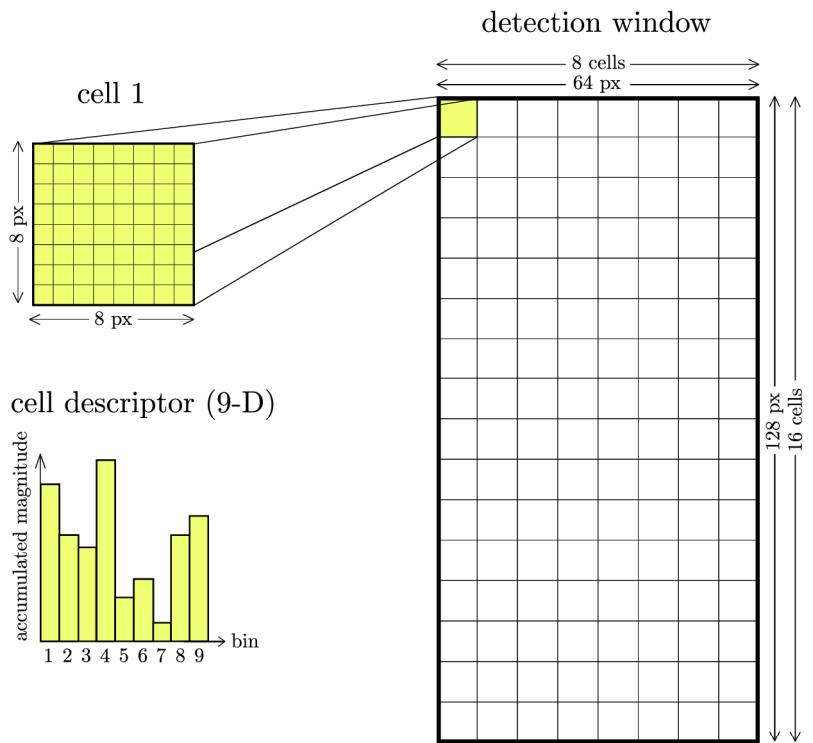
Method

Motivation | Implementation | Evaluation | Summary

■ HOG Descriptor

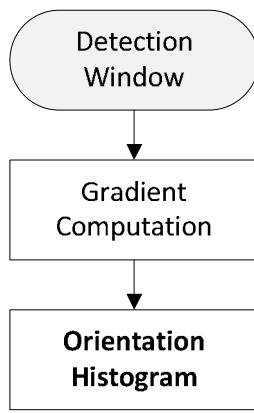


- **Orientation Histogram**
 - Detection Window subdivided into **Cells**
 - For each Cell:
Histogram Generation

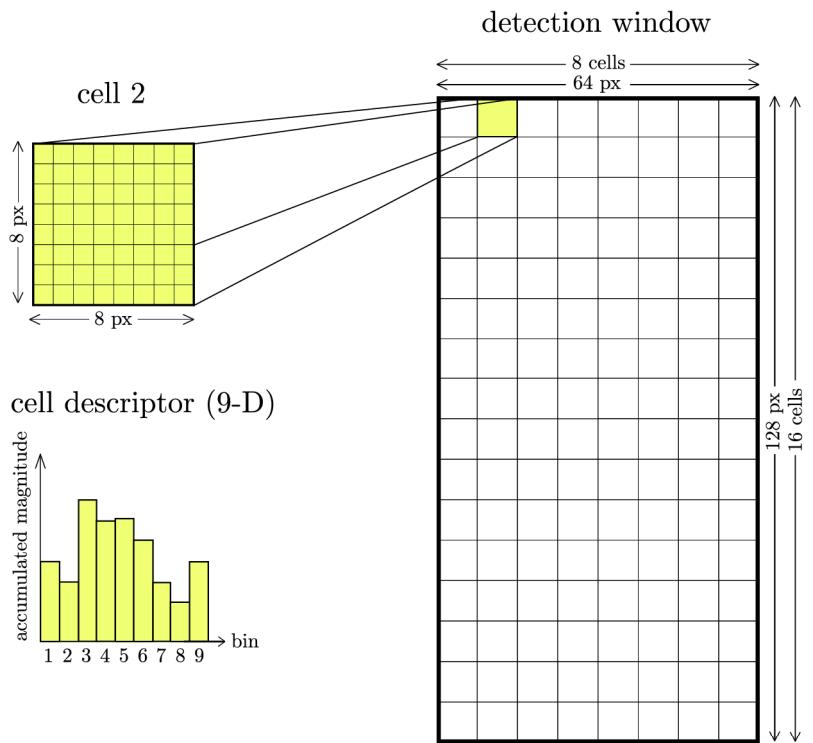


Motivation | Method | Implementation | Evaluation | Summary

■ HOG Descriptor



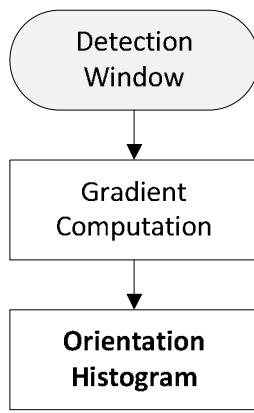
- **Orientation Histogram**
 - Detection Window subdivided into **Cells**
 - For each Cell:
Histogram Generation



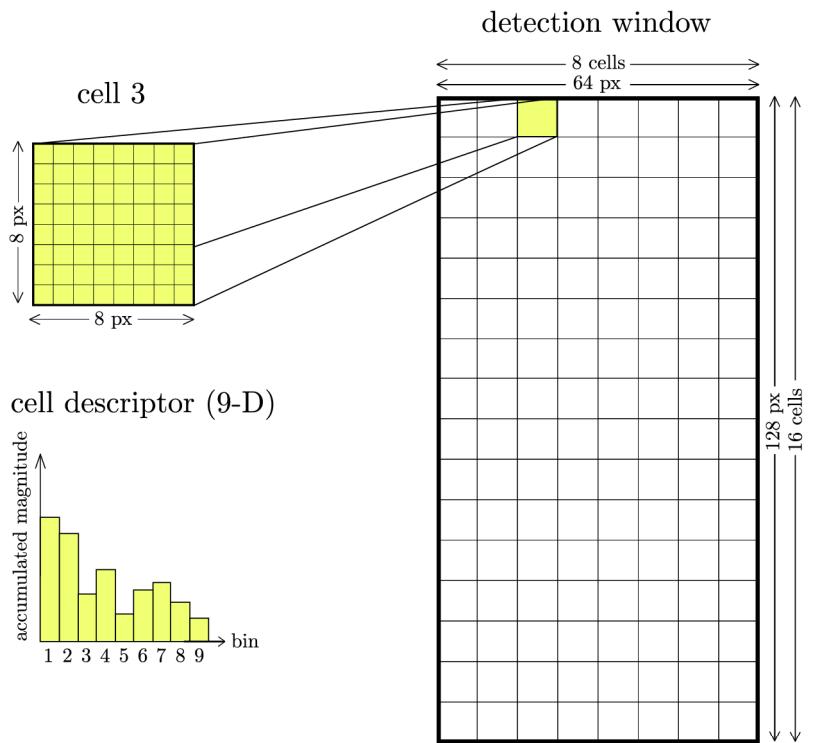
Method

Motivation | Implementation | Evaluation | Summary

■ HOG Descriptor



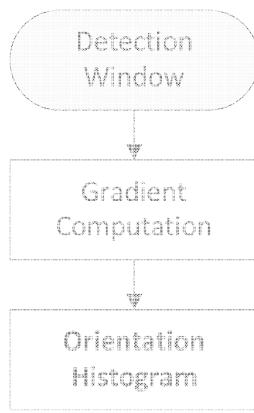
- **Orientation Histogram**
 - Detection Window subdivided into **Cells**
 - For each Cell:
Histogram Generation



Method

Motivation | Implementation | Evaluation | Summary

HOG Descriptor



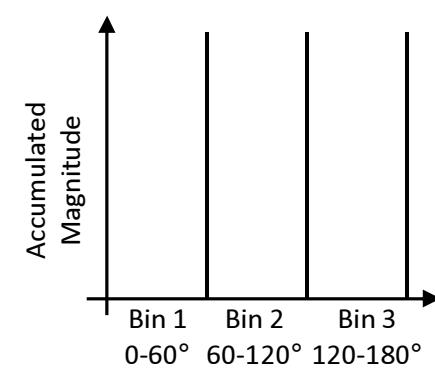
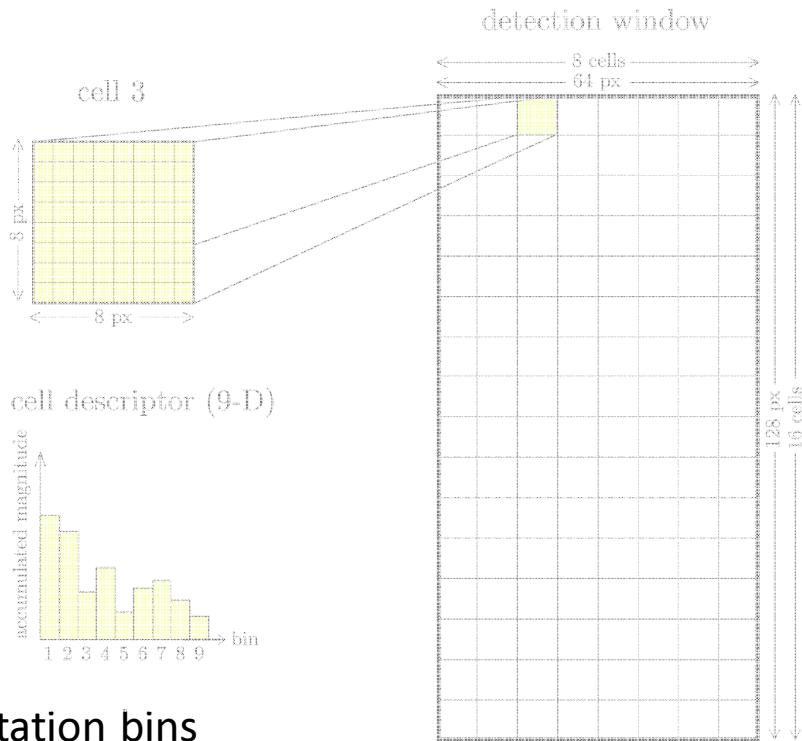
- Orientation Histogram
 - Detection Window subdivided into Cells
 - For each Cell:
 - Histogram Generation
 - Example: 4x4 px Cell, 3 orientation bins

$$|G|$$

12	42	152	232
18	91	64	31
27	48	219	102
15	43	147	160

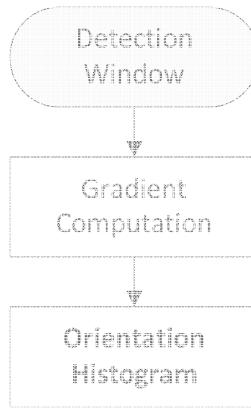
$$\phi$$

47°	21°	78°	170°
98°	94°	5°	140°
27°	77°	30°	128°
4°	11°	22°	101°



Motivation | Method | Implementation | Evaluation | Summary

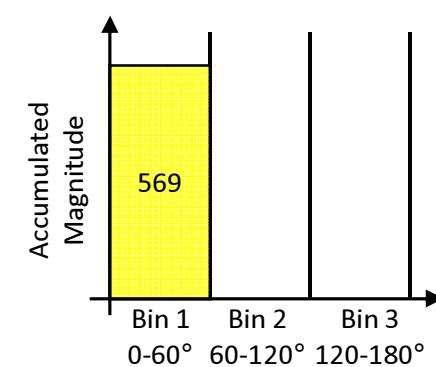
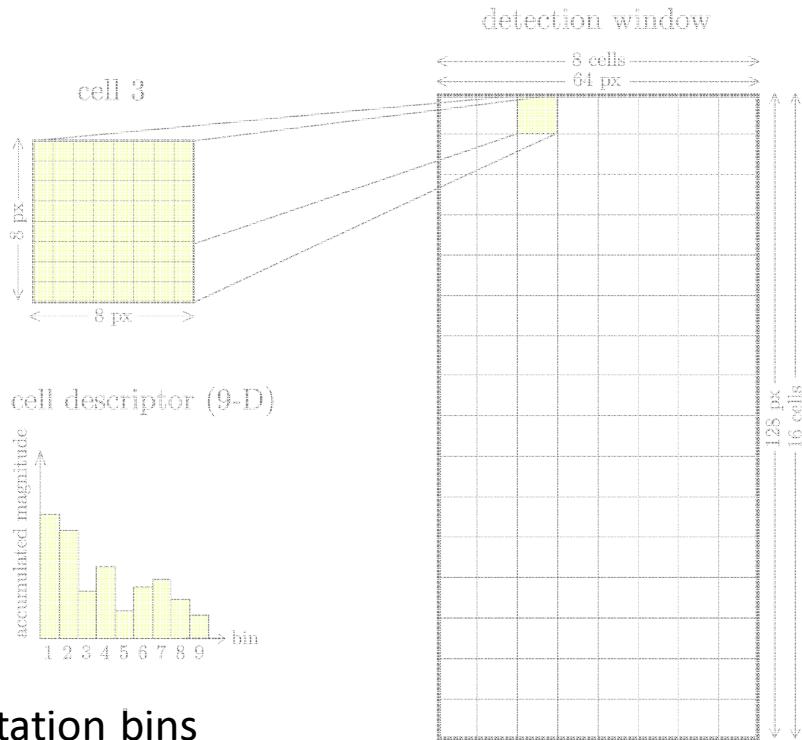
HOG Descriptor



- Orientation Histogram
- Detection Window subdivided into Cells
- For each Cell:
- Histogram Generation
- Example: 4x4 px Cell, 3 orientation bins

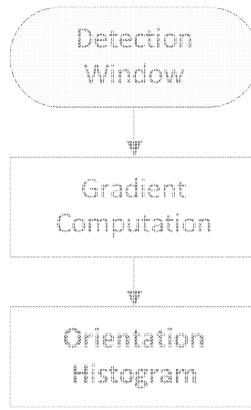
$ G $				
12	42	152	232	
18	91	64	31	
27	48	219	102	
15	43	147	160	

ϕ				
47°	21°	78°	170°	
98°	94°	5°	140°	
27°	77°	30°	128°	
4°	11°	22°	101°	

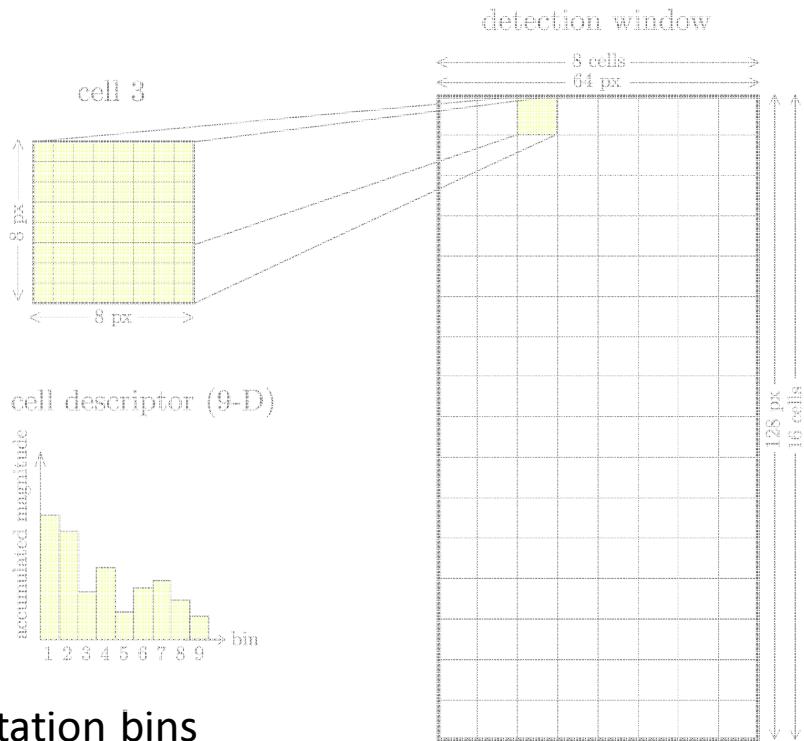


Motivation | Method | Implementation | Evaluation | Summary

HOG Descriptor

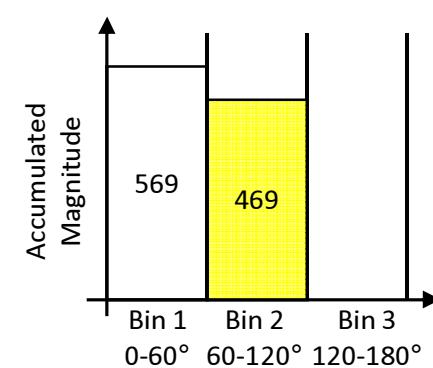


- Orientation Histogram
 - Detection Window subdivided into Cells
 - For each Cell: Histogram Generation
 - Example: 4x4 px Cell, 3 orientation bins



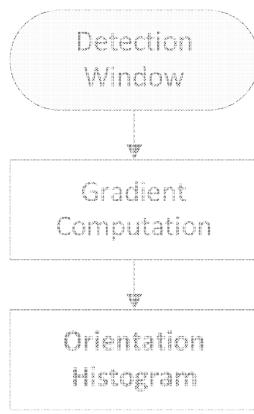
$ G $	12	42	152	232
18	91	64	31	
27	48	219	102	
15	43	147	160	

ϕ	47°	21°	78°	170°
98°	94°	5°	140°	
27°	77°	30°	128°	
4°	11°	22°	101°	

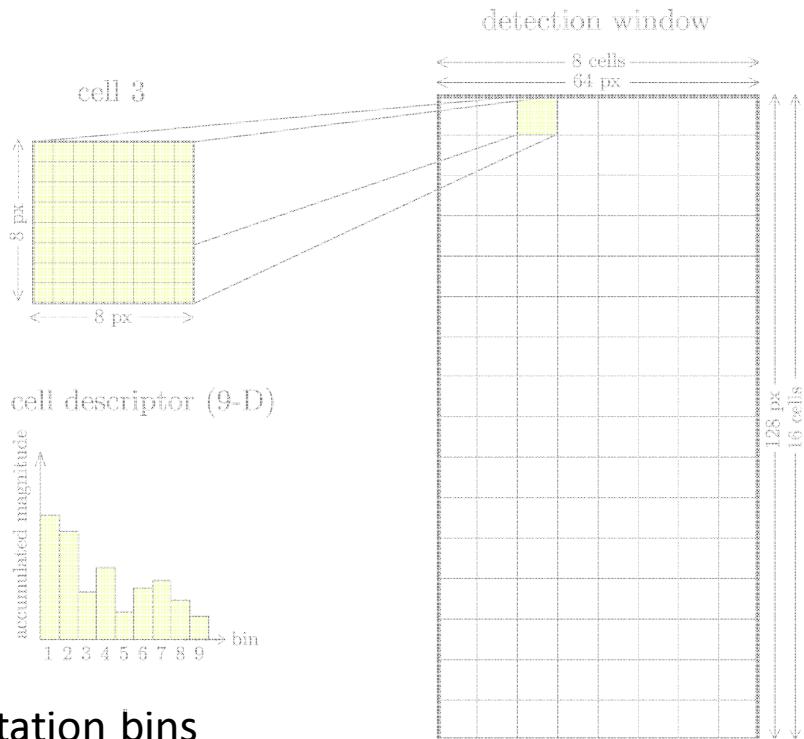


Motivation | Method | Implementation | Evaluation | Summary

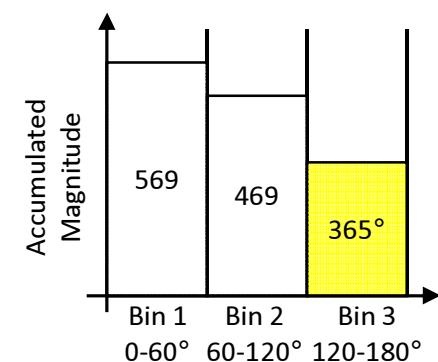
HOG Descriptor



- Orientation Histogram
- Detection Window subdivided into Cells
- For each Cell:
- Histogram Generation
- Example: 4x4 px Cell, 3 orientation bins



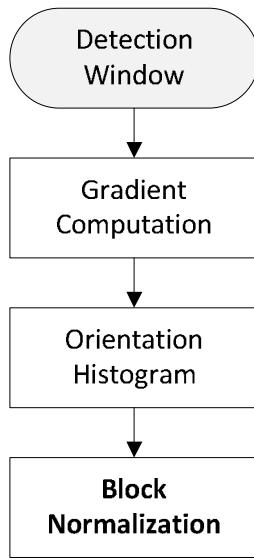
$ G $	ϕ
12	47°
42	21°
152	78°
232	170°
18	98°
91	94°
64	5°
31	140°
27	27°
48	77°
219	30°
102	128°
15	4°
43	11°
147	22°
160	101°



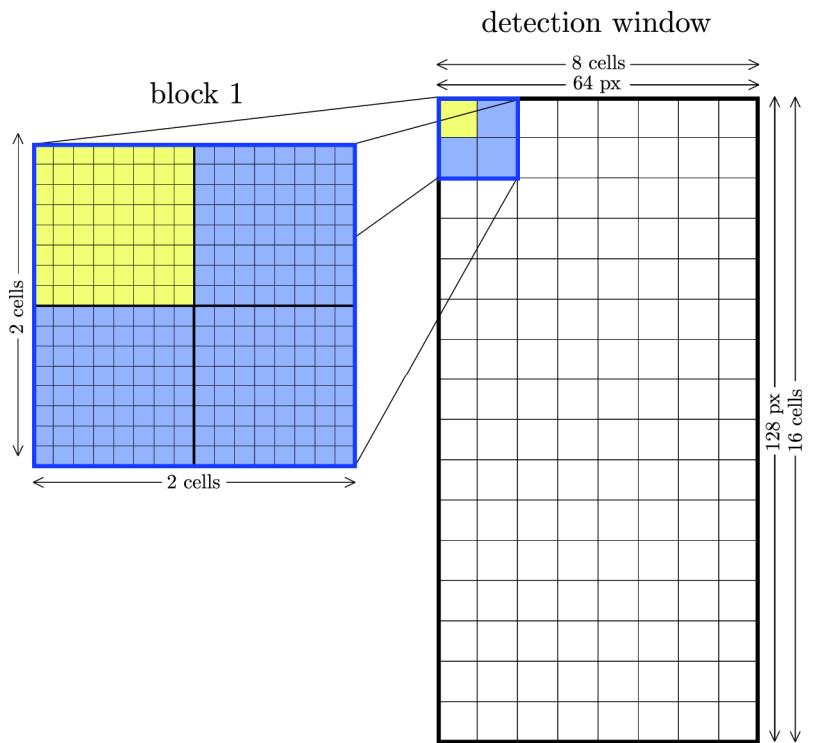
Method

Motivation | Implementation | Evaluation | Summary

■ HOG Descriptor



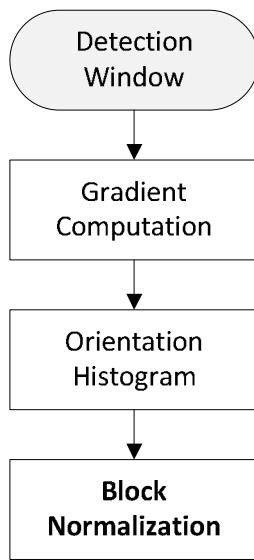
- **Block Normalization**
 - Block: 2x2 Cells



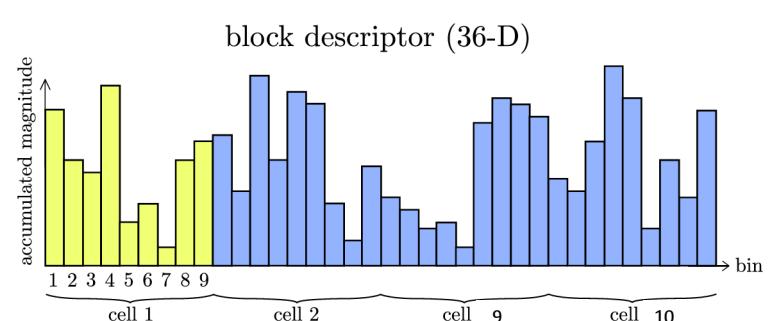
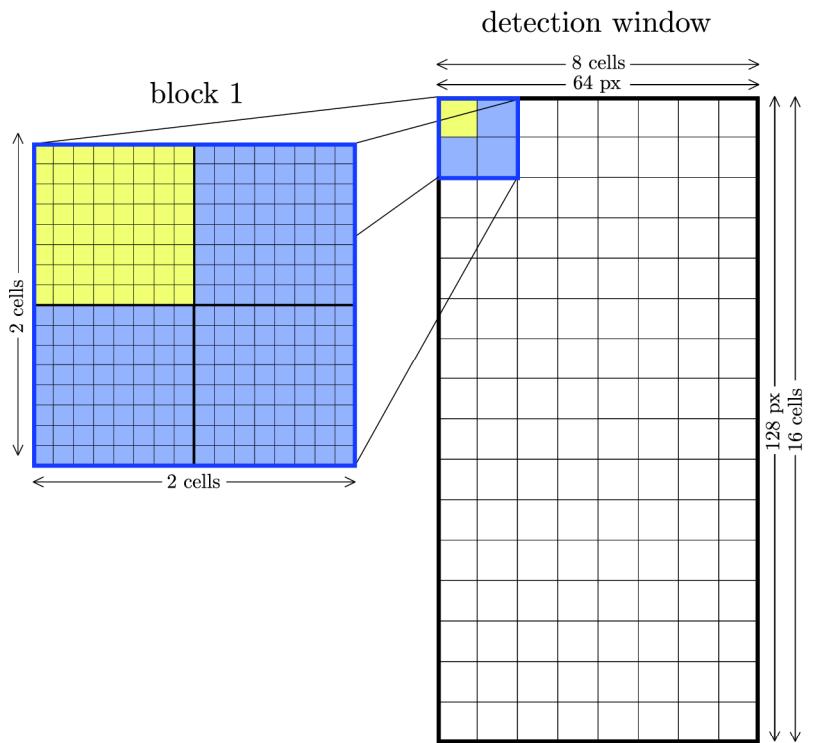
Method

Motivation | **Implementation** | Evaluation | Summary

HOG Descriptor



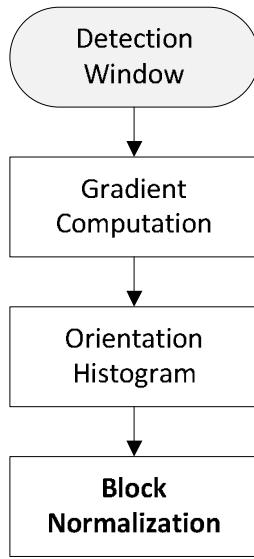
- **Block Normalization**
 - Block: 2x2 Cells
 - For each Block:
Concatenation of
Cell histograms



Method

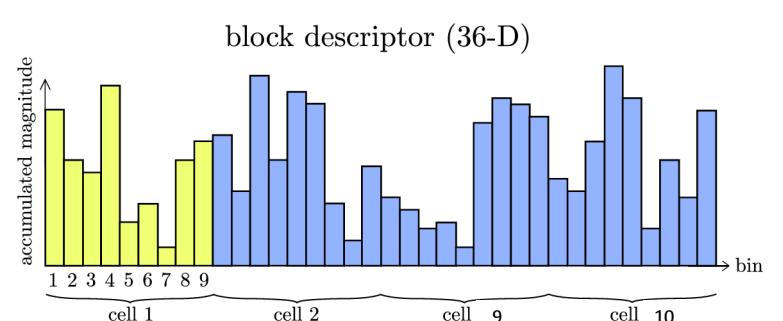
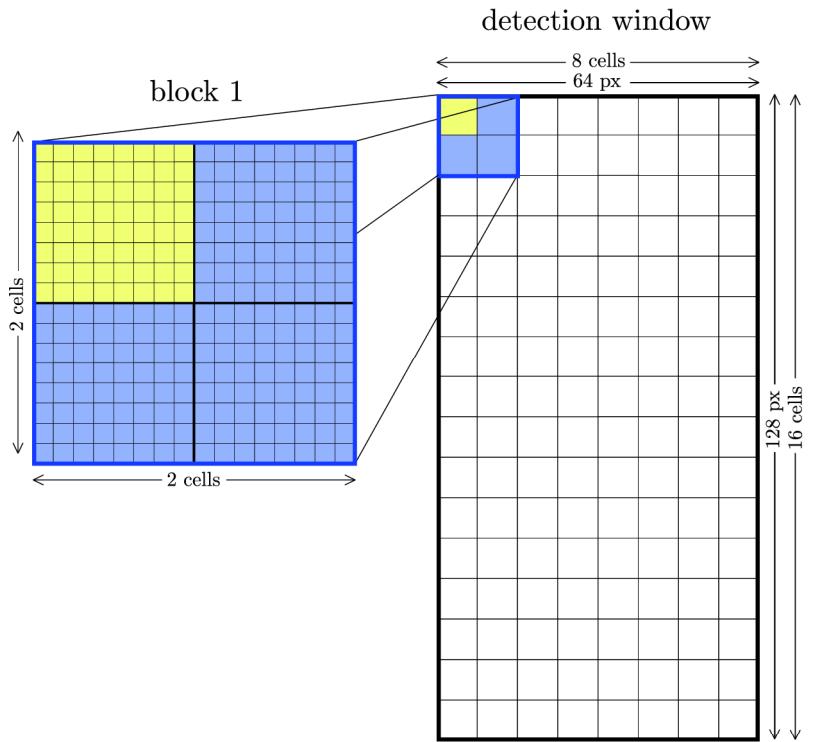
Motivation | Implementation | Evaluation | Summary

HOG Descriptor



- **Block Normalization**
 - Block: 2x2 Cells
 - For each Block:
Concatenation of
Cell histograms
- **Normalization
using L2-Norm**

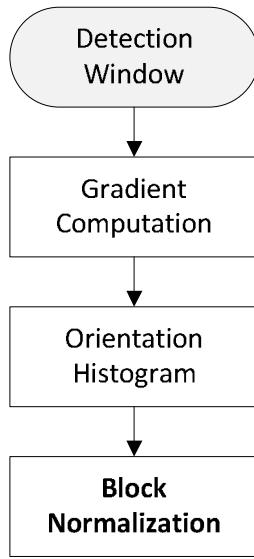
$$\mathbf{v} \rightarrow \frac{\mathbf{v}}{\sqrt{\|\mathbf{v}\|_2^2 + \epsilon^2}}$$



Method

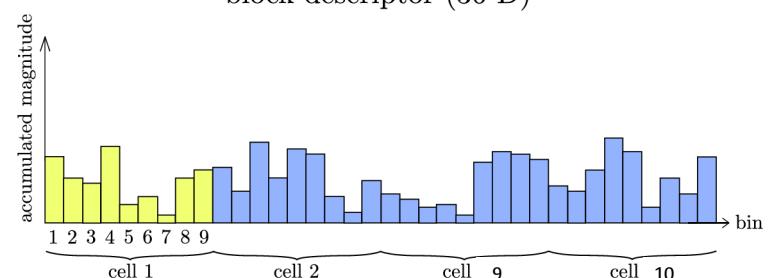
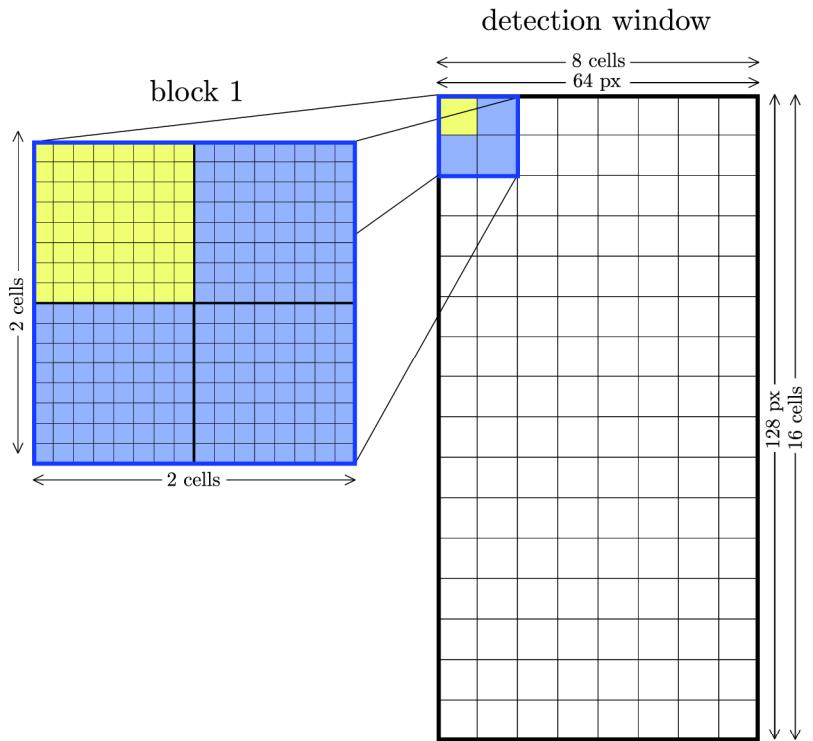
Motivation | Implementation | Evaluation | Summary

HOG Descriptor



- **Block Normalization**
 - Block: 2x2 Cells
 - For each Block:
Concatenation of
Cell histograms
- Normalization
using L2-Norm

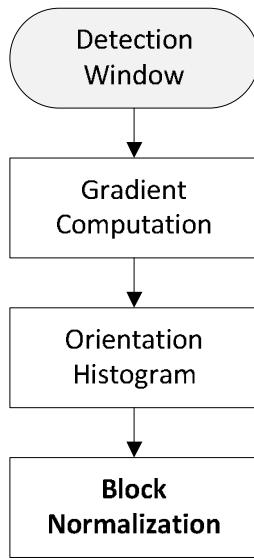
$$\mathbf{v} \rightarrow \frac{\mathbf{v}}{\sqrt{\|\mathbf{v}\|_2^2 + \epsilon^2}}$$



Method

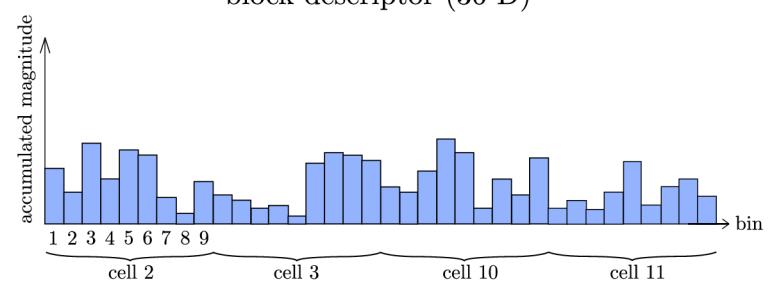
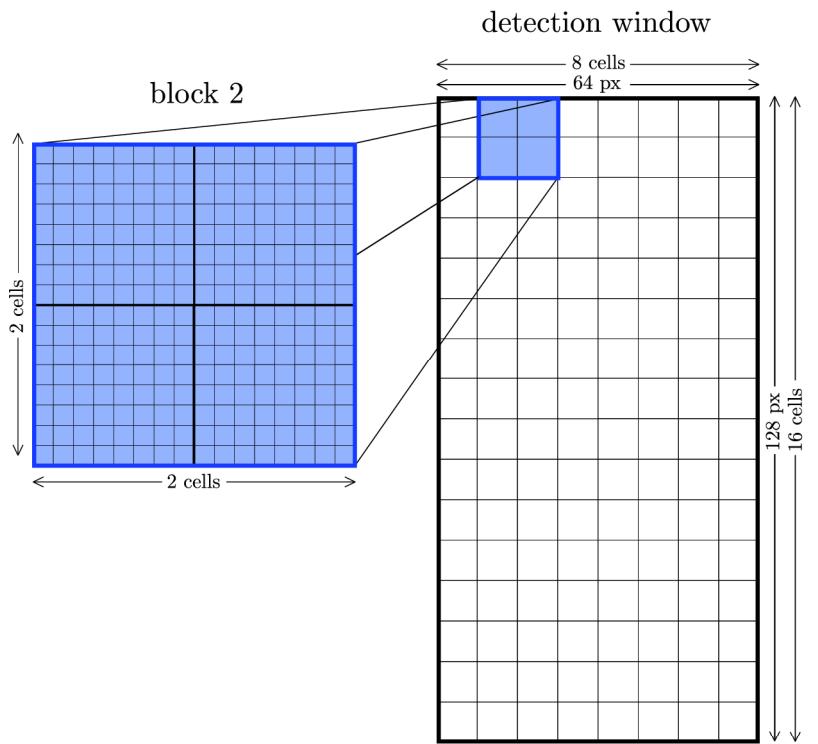
Motivation | **Implementation** | Evaluation | Summary

HOG Descriptor



- **Block Normalization**
 - Block: 2x2 Cells
 - For each Block:
Concatenation of
Cell histograms
- **Normalization using L2-Norm**

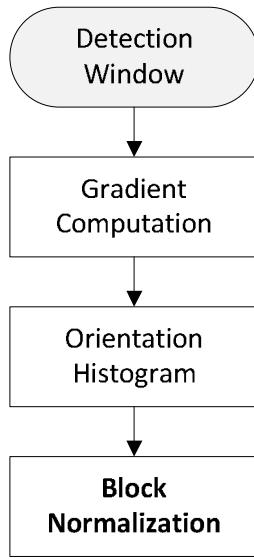
$$\mathbf{v} \rightarrow \frac{\mathbf{v}}{\sqrt{\|\mathbf{v}\|_2^2 + \epsilon^2}}$$



Method

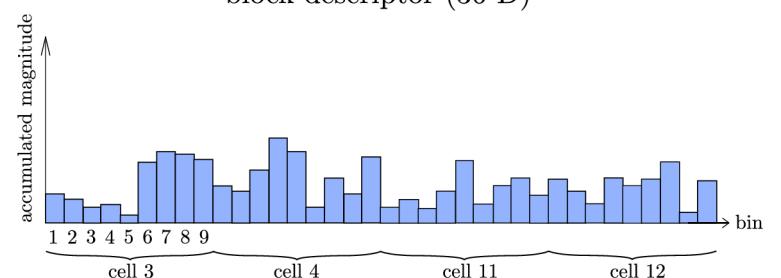
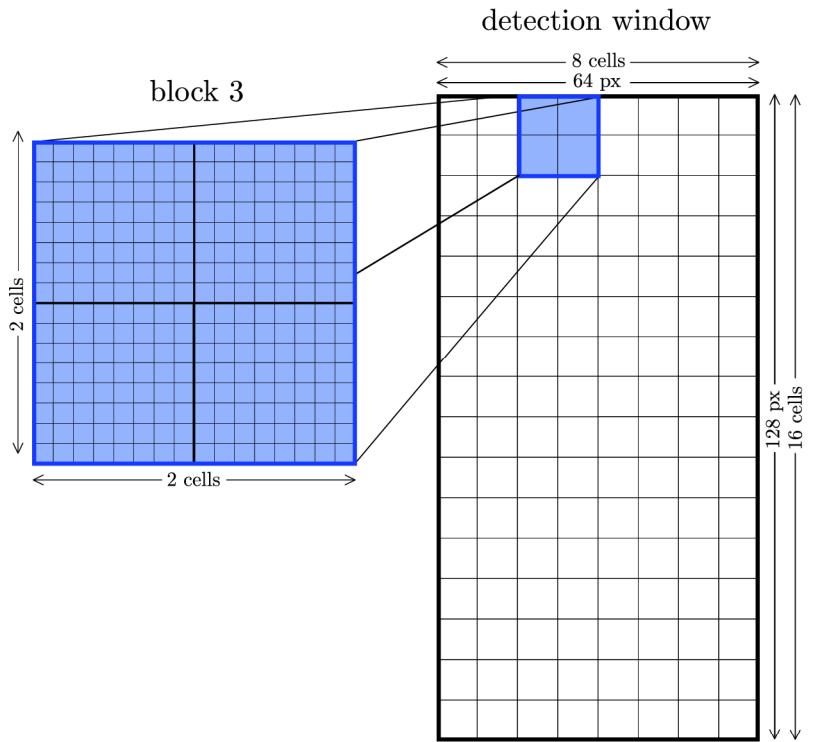
Motivation | **Implementation** | Evaluation | Summary

HOG Descriptor



- **Block Normalization**
 - Block: 2x2 Cells
 - For each Block:
Concatenation of
Cell histograms
- **Normalization using L2-Norm**

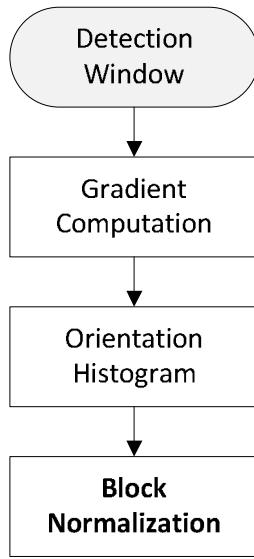
$$\mathbf{v} \rightarrow \frac{\mathbf{v}}{\sqrt{\|\mathbf{v}\|_2^2 + \epsilon^2}}$$



Method

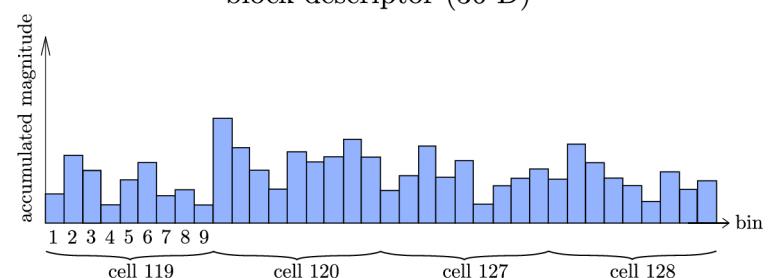
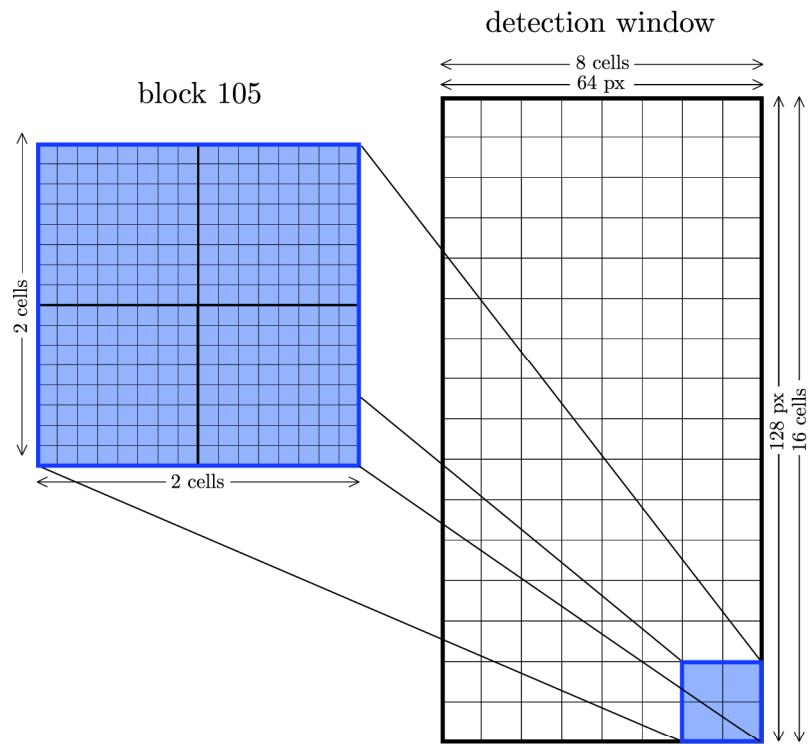
Motivation | **Implementation** | Evaluation | Summary

HOG Descriptor



- **Block Normalization**
 - Block: 2x2 Cells
 - For each Block:
Concatenation of
Cell histograms
- Normalization
using L2-Norm

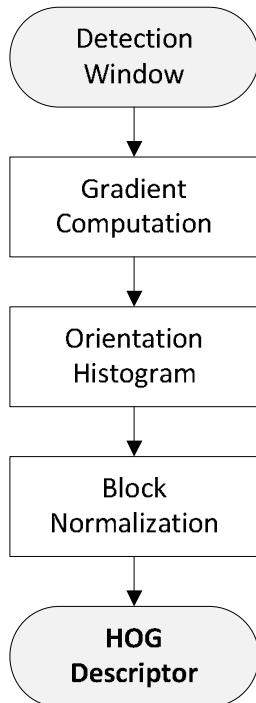
$$\mathbf{v} \rightarrow \frac{\mathbf{v}}{\sqrt{\|\mathbf{v}\|_2^2 + \epsilon^2}}$$



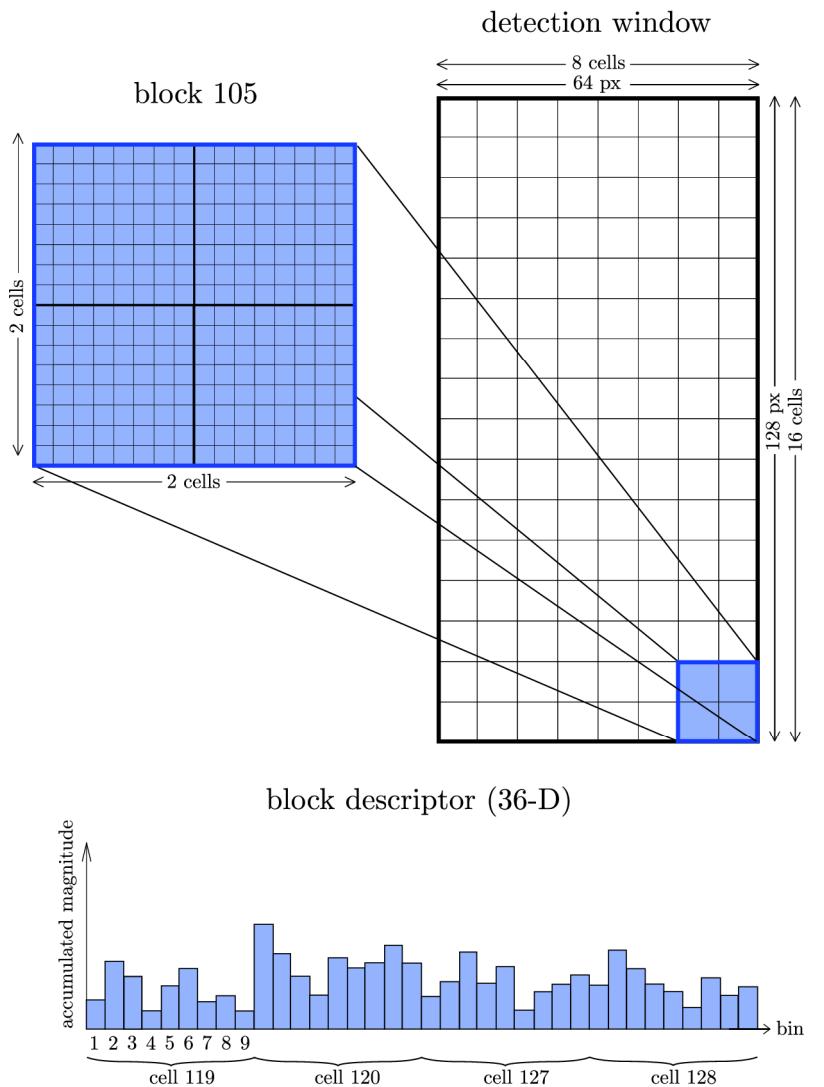
Method

Motivation | Implementation | Evaluation | Summary

■ HOG Descriptor



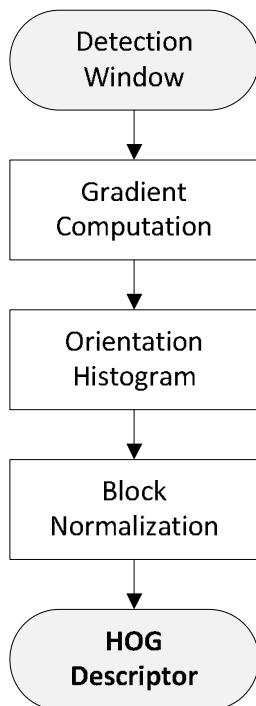
■ HOG Descriptor



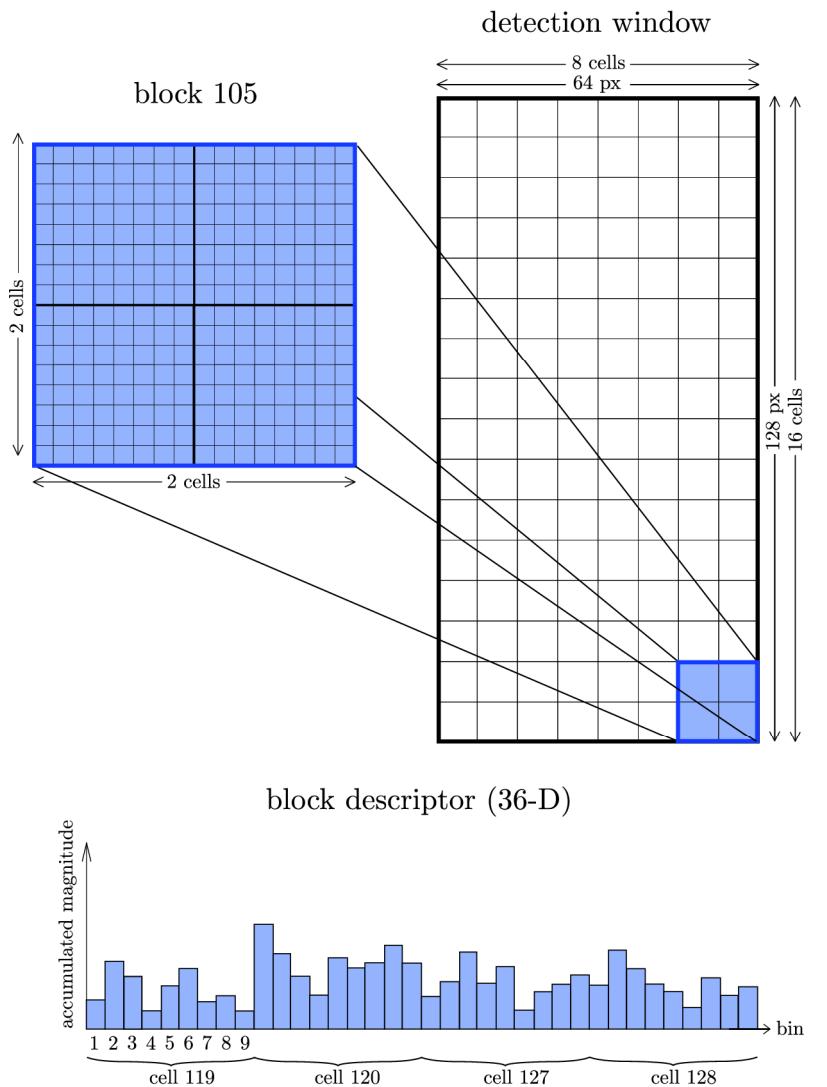
Method

Motivation | Implementation | Evaluation | Summary

HOG Descriptor



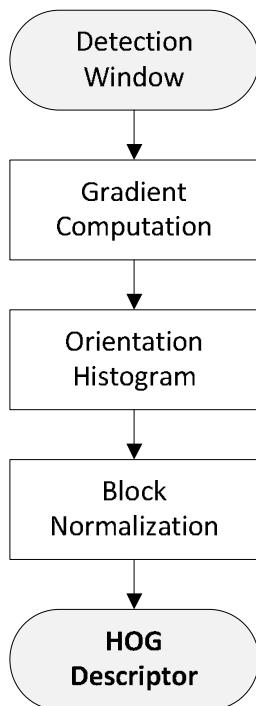
- **HOG Descriptor**
 - Concatenation of normalized Blocks
 - Feature Vector with 3780 dimensions



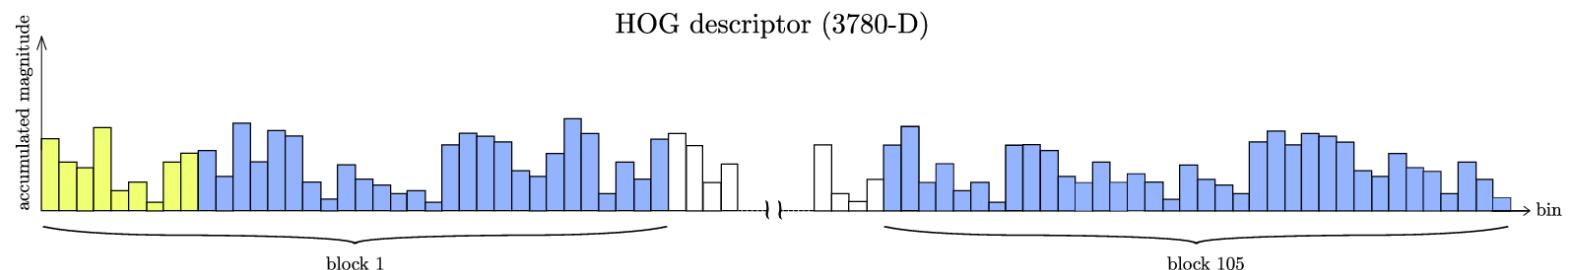
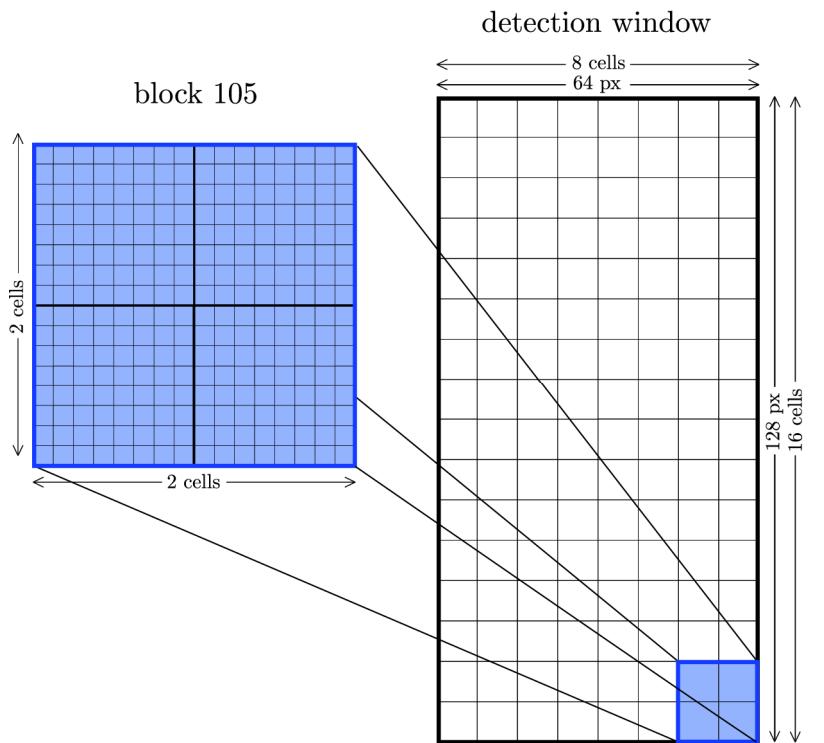
Method

Motivation | Implementation | Evaluation | Summary

HOG Descriptor

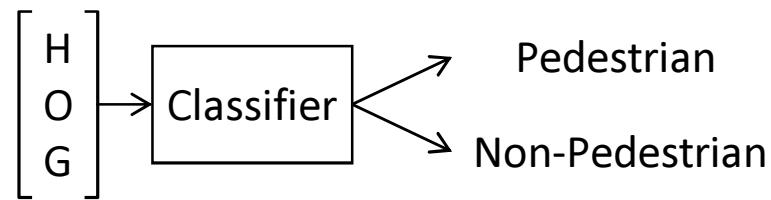


- **HOG Descriptor**
 - Concatenation of normalized Blocks
 - Feature Vector with 3780 dimensions



Motivation | Method | Implementation | Evaluation | Summary

■ Classification

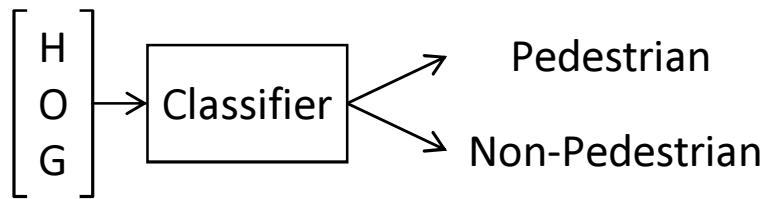


- **Machine Learning**
 - Learning from Examples (I/O pairs)
 - Established Pedestrian Datasets [INRIA]

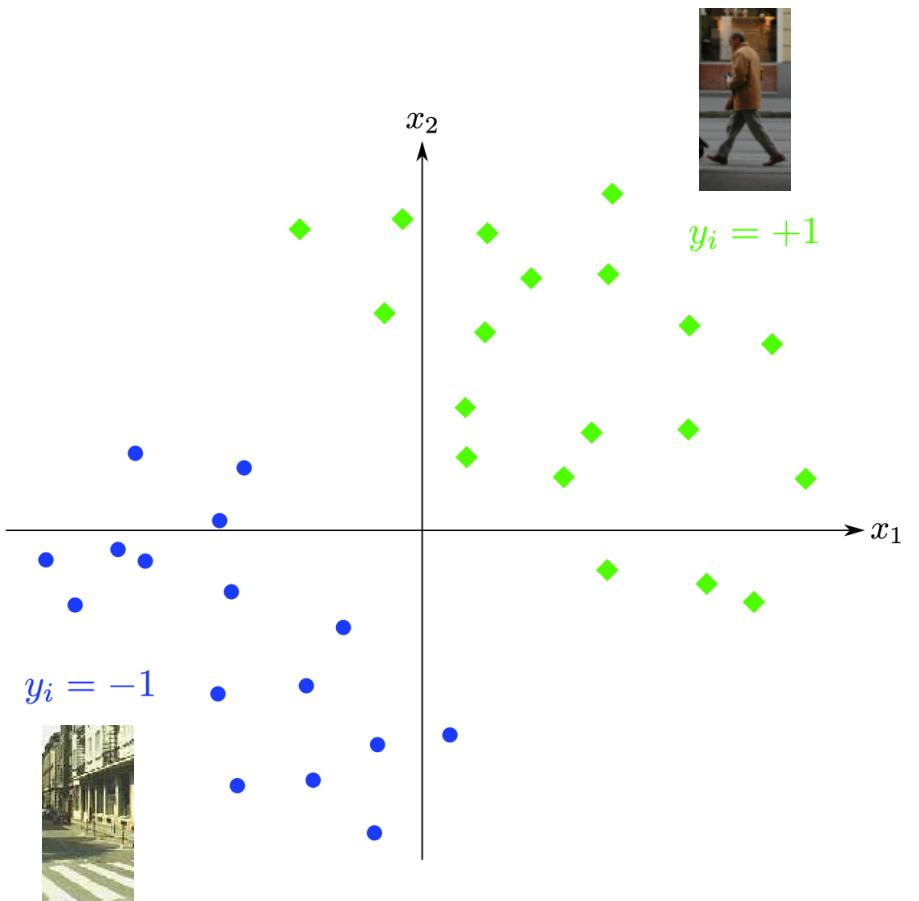
Method

Motivation | Implementation | Evaluation | Summary

Classification



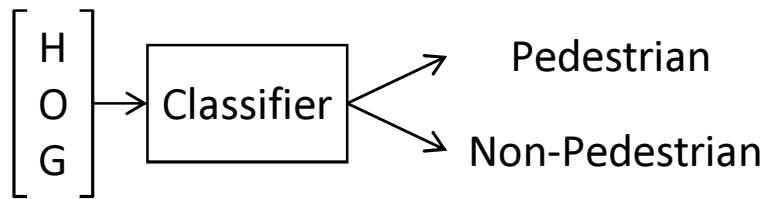
- **Machine Learning**
 - Learning from Examples (I/O pairs)
 - Established Pedestrian Datasets [INRIA]
- **Support Vector Machines (SVM)**
 - Linear Classifier
 - Decision function:
 $y = f(\mathbf{x}) = \text{sgn}(\mathbf{w} \cdot \mathbf{x})$
 - \mathbf{w} learned by training



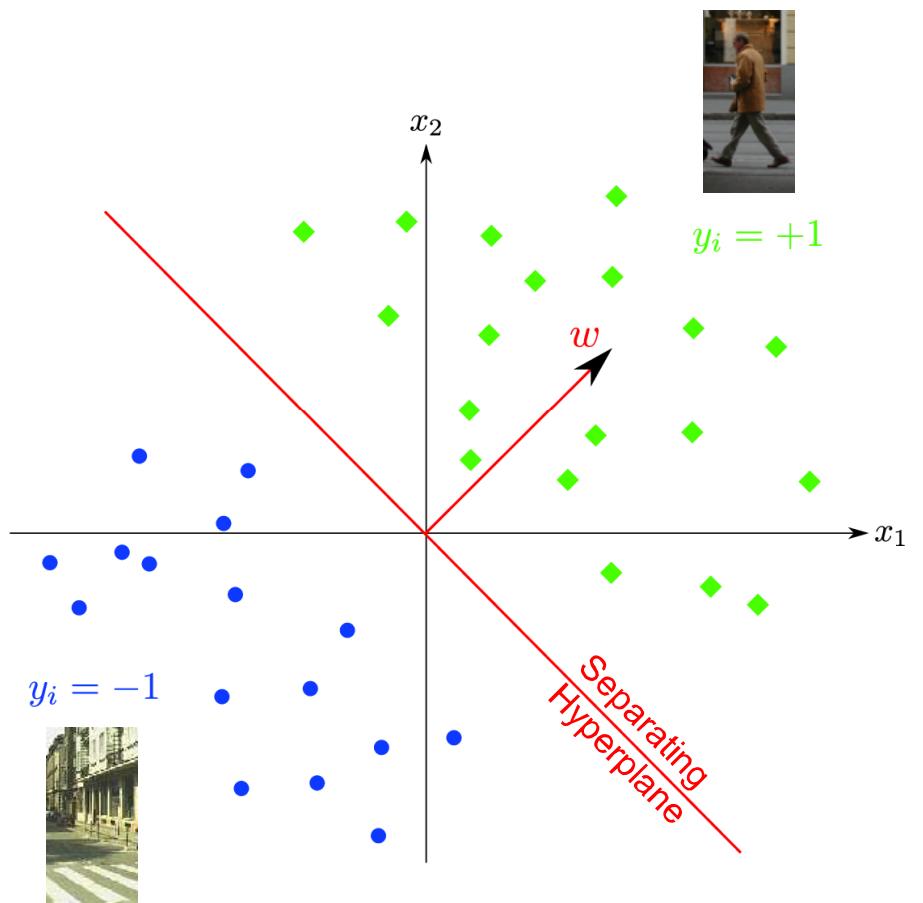
Method

Motivation | Implementation | Evaluation | Summary

Classification



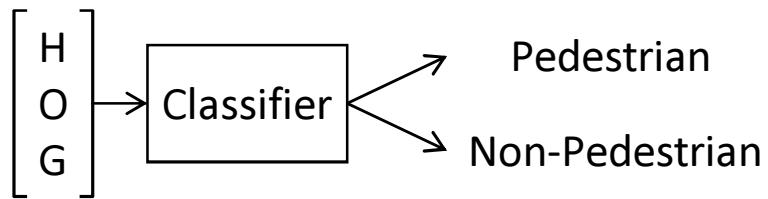
- **Machine Learning**
 - Learning from Examples (I/O pairs)
 - Established Pedestrian Datasets [INRIA]
- **Support Vector Machines (SVM)**
 - Linear Classifier
 - Decision function:
 $y = f(\mathbf{x}) = \text{sgn}(\mathbf{w} \cdot \mathbf{x})$
 - \mathbf{w} learned by training



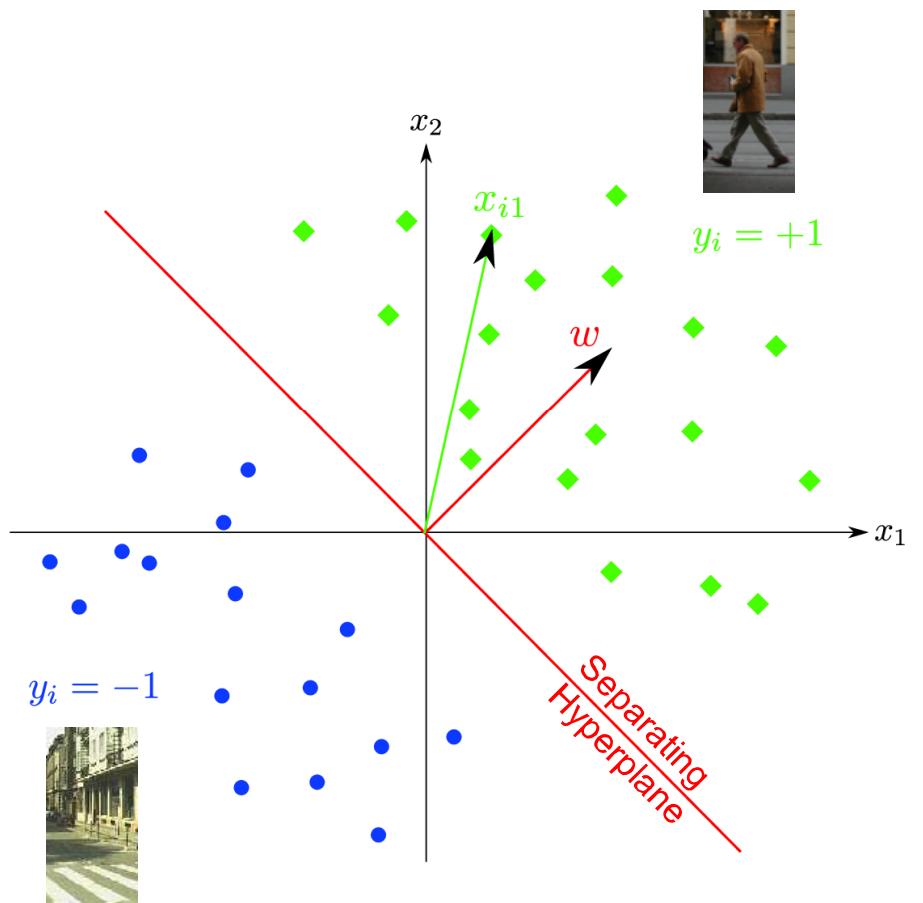
Method

Motivation | Implementation | Evaluation | Summary

Classification

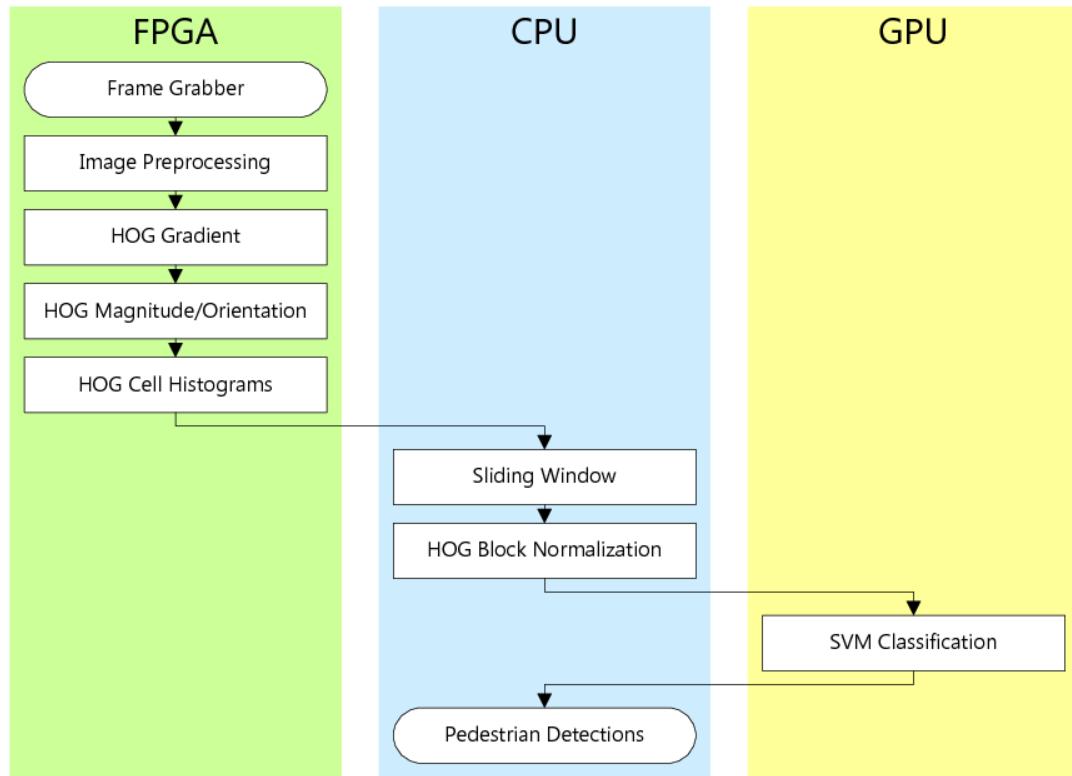


- **Machine Learning**
 - Learning from Examples (I/O pairs)
 - Established Pedestrian Datasets [INRIA]
- **Support Vector Machines (SVM)**
 - Linear Classifier
 - Decision function:
 $y = f(\mathbf{x}) = \text{sgn}(\mathbf{w} \cdot \mathbf{x})$
 - \mathbf{w} learned by training



Motivation | Method | **Implementation** | Evaluation | Summary

■ FPGA-CPU-GPU Framework

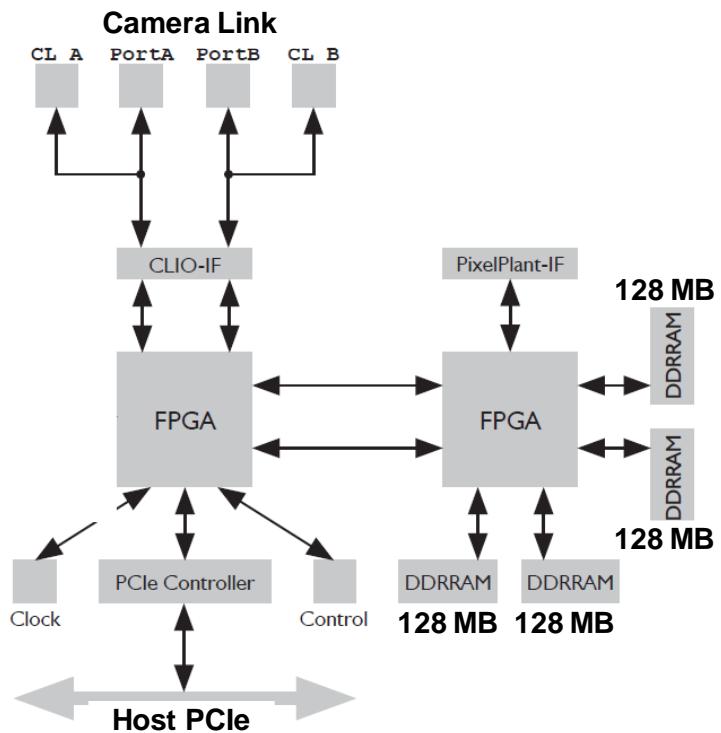
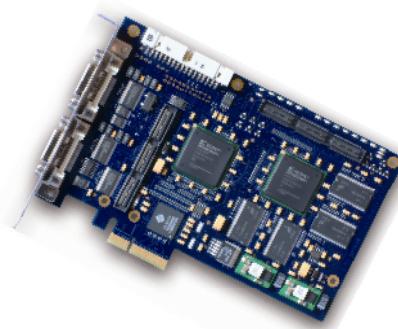


- **FPGA**
Xilinx Spartan 3 XC3S 4000
- **CPU**
Intel Core i7, 2.66 GHz
- **GPU**
NVIDIA GeForce GTX 295

Motivation | Method | Implementation | Evaluation | Summary

■ FPGA Rapid Prototyping Platform

- **microEnable IV-FULLx4** [SiliconSoftware]
 - Commercially available frame grabber board
 - FPGA for individual image pre-processing



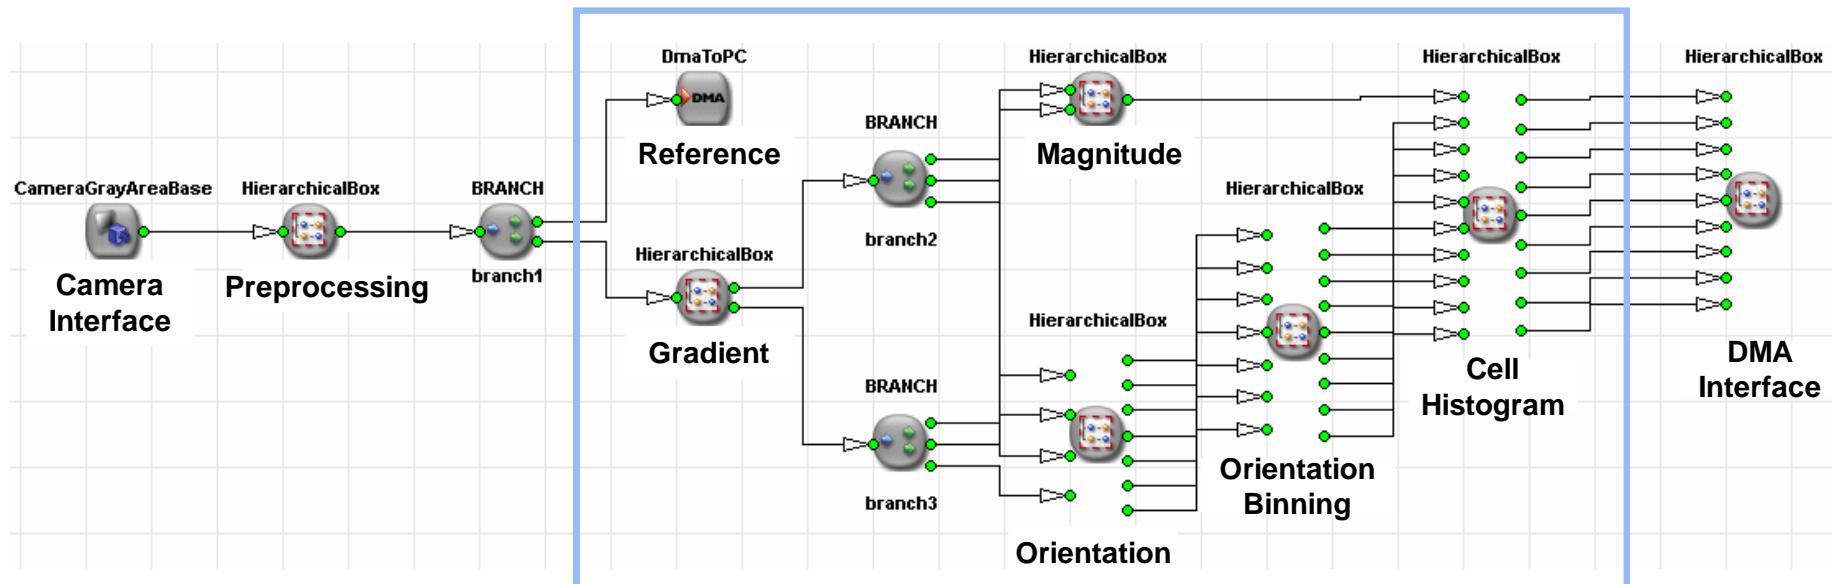
Motivation | Method | Implementation | Evaluation | Summary

■ FPGA-HOG

■ VisualApplets

- Graphic-oriented hardware development software, based on Xilinx ISE tools
- Design arranged by image processing operators and transport links

■ HOG flowchart



Motivation | Method | **Implementation** | Evaluation | Summary

■ FPGA-HOG

■ Magnitude

$$|G(x, y)| = \sqrt{G_x(x, y)^2 + G_y(x, y)^2}$$

square root extracted,
as Euclidean metric yields best performance

■ Orientation Binning

■ Example: Bin 1 (0° - 20°)

Cond. I: $G_x(x, y) > 0 \wedge G_y(x, y) > 0 \Leftrightarrow$ quadrant I

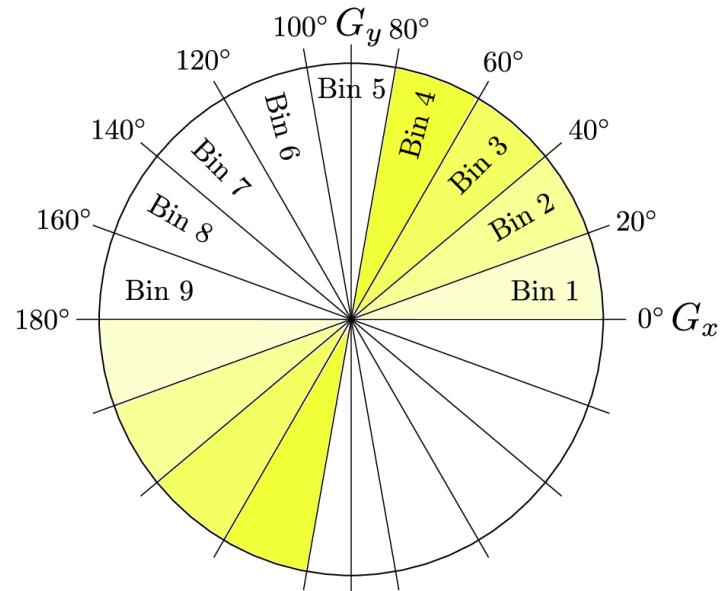
\vee $G_x(x, y) < 0 \wedge G_y(x, y) < 0 \Leftrightarrow$ quadrant III

Cond. II: $0 < \tan(\alpha) < \tan(20^\circ)$

$$0 < \left| \frac{G_y(x, y)}{G_x(x, y)} \right| < \tan(20^\circ)$$

$$0 < |G_y(x, y)| < \tan(20^\circ) \cdot |G_x(x, y)|$$

$$0 < |G_y(x, y)| < 0.364 \cdot |G_x(x, y)|$$



Motivation | Method |

Implementation

 | Evaluation | Summary

- FPGA-HOG

- Histogram Generation

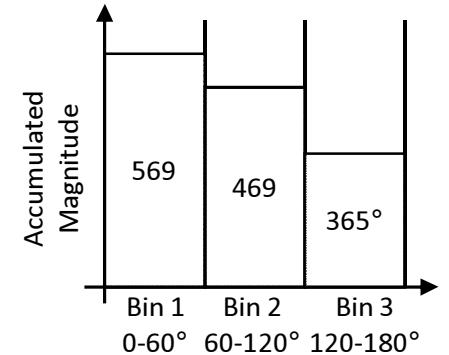
 ϕ

47°	21°	78°	170°
98°	94°	5°	140°
27°	77°	30°	128°
4°	11°	22°	101°

Example:

4x4 px Cell

3 orientation bins

 $|G|$

12	42	152	232
18	91	64	31
27	48	219	102
15	43	147	160

Implementation

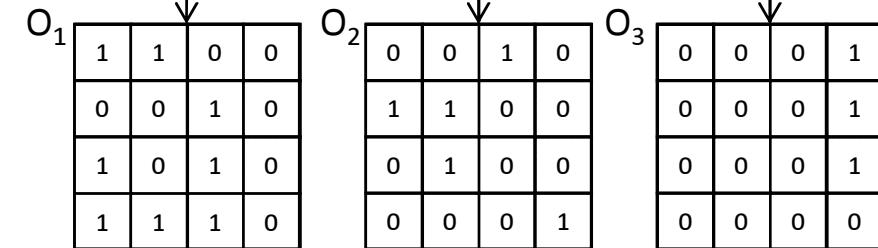
Motivation | Method | Evaluation | Summary

■ FPGA-HOG

- Histogram Generation
 - Orientation binning
 - O1: 0-60°
 - O2: 60-120°
 - O3: 120-180°

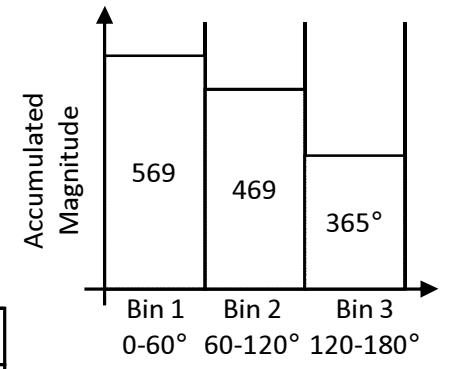
$$\phi$$

47°	21°	78°	170°
98°	94°	5°	140°
27°	77°	30°	128°
4°	11°	22°	101°



Example:

4x4 px Cell
3 orientation bins



$$|G|$$

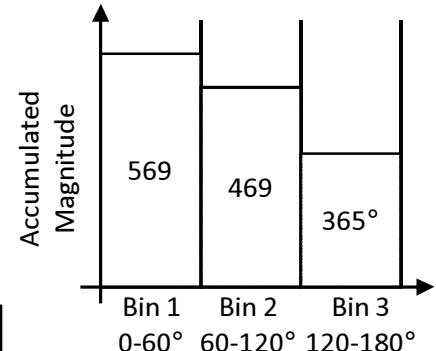
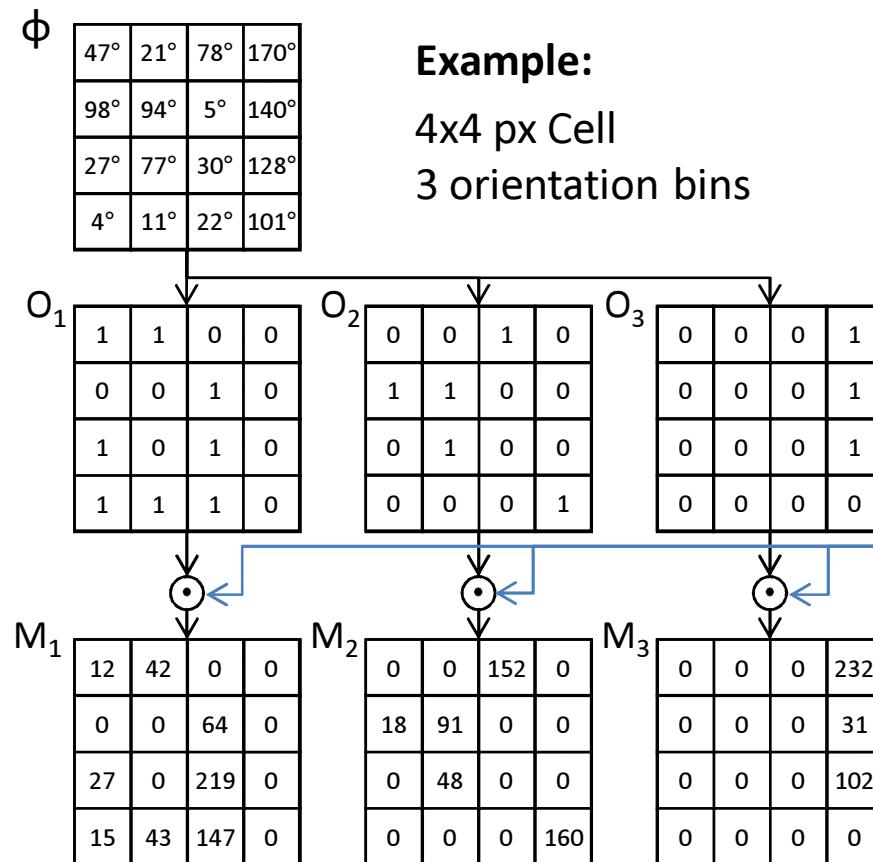
12	42	152	232
18	91	64	31
27	48	219	102
15	43	147	160

Implementation

Motivation | Method | Evaluation | Summary

FPGA-HOG

- Histogram Generation
 - Orientation binning
 - O1: 0-60°
 - O2: 60-120°
 - O3: 120-180°
 - Magnitude weighting



$|G|$

12	42	152	232
18	91	64	31
27	48	219	102
15	43	147	160

Motivation | Method | **Implementation** | Evaluation | Summary

■ FPGA-HOG

■ Histogram Generation

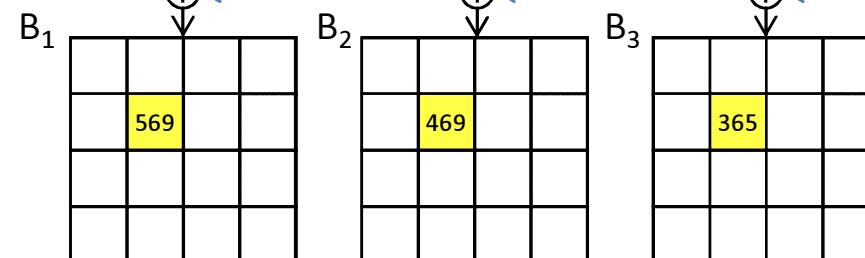
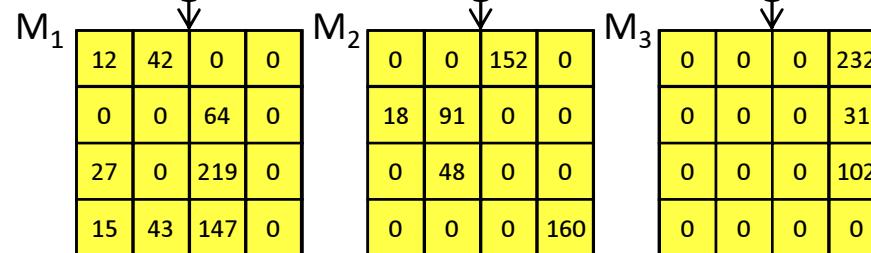
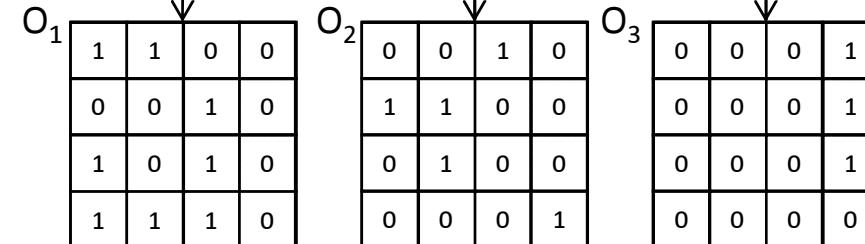
- Orientation binning
 - O1: 0-60°
 - O2: 60-120°
 - O3: 120-180°

- Magnitude weighting

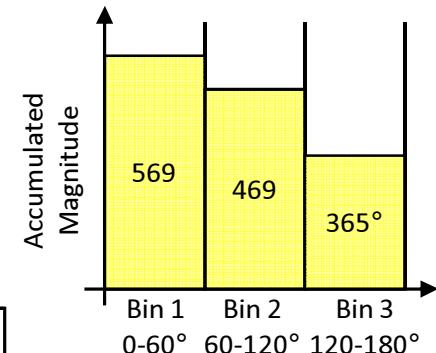
- Histogram bin accumulation

$$\phi$$

47°	21°	78°	170°
98°	94°	5°	140°
27°	77°	30°	128°
4°	11°	22°	101°



Example:
4x4 px Cell
3 orientation bins



$$|G|$$

12	42	152	232
18	91	64	31
27	48	219	102
15	43	147	160

$$K_S$$

1	1	1	1
1	1	1	1
1	1	1	1
1	1	1	1

Motivation | Method | **Implementation** | Evaluation | Summary

■ FPGA-HOG

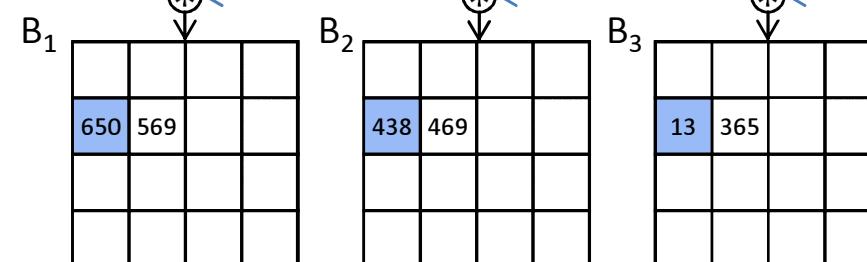
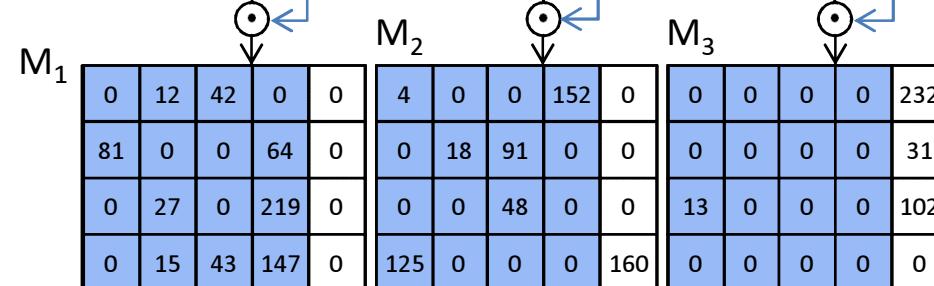
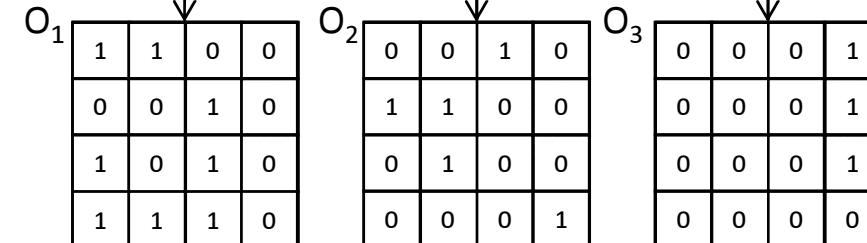
■ Histogram Generation

- Orientation binning
 - O1: 0-60°
 - O2: 60-120°
 - O3: 120-180°

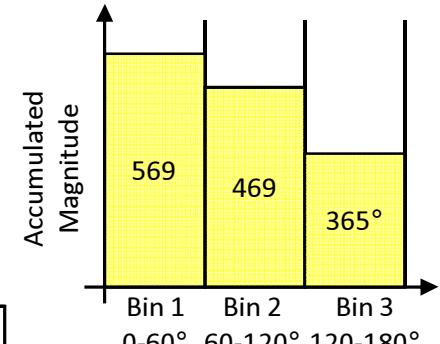
- Magnitude weighting

- Histogram bin accumulation

ϕ	47°	21°	78°	170°
	98°	94°	5°	140°
	27°	77°	30°	128°
	4°	11°	22°	101°



Example:
4x4 px Cell
3 orientation bins

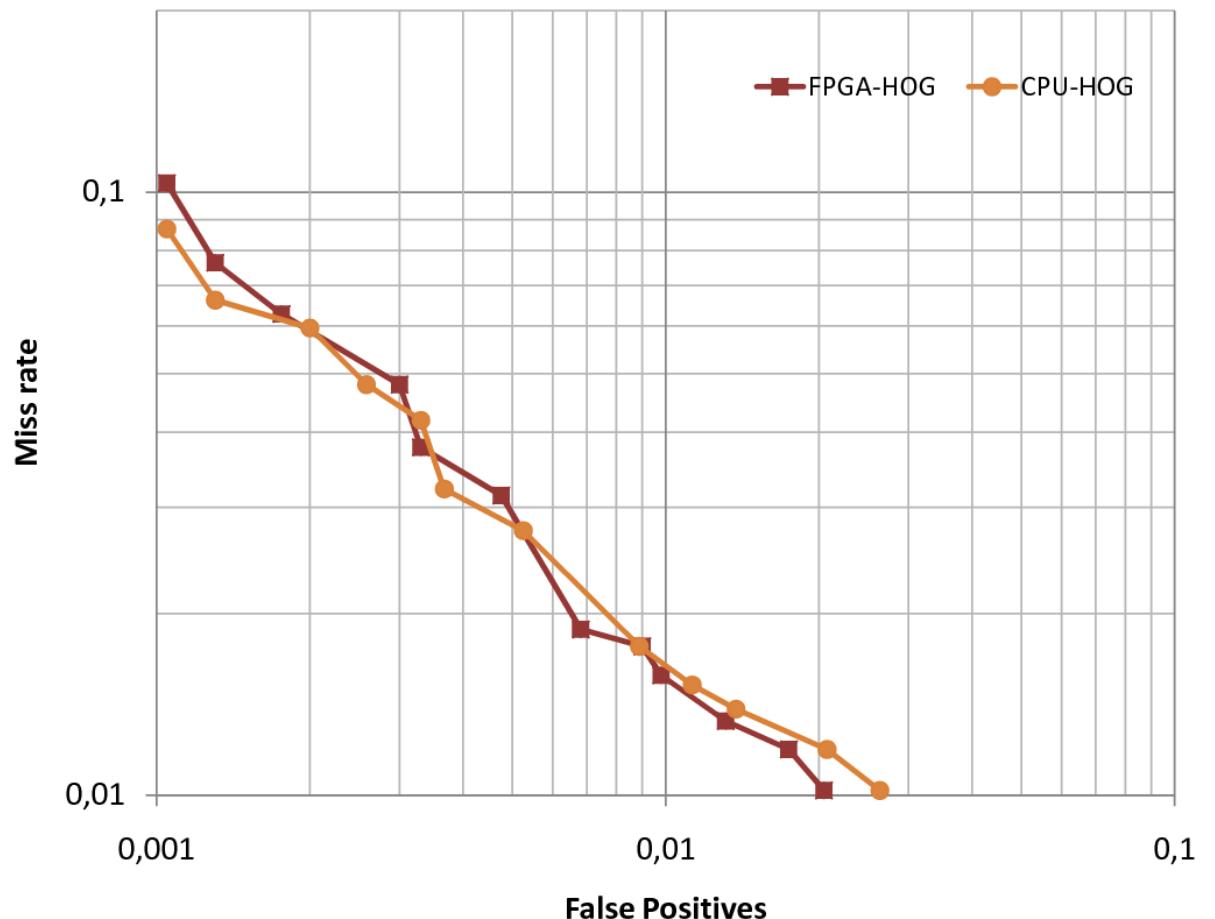


$ G $	12	42	152	232
	18	91	64	31
	27	48	219	102
	15	43	147	160

K_S	1	1	1	1
	1	1	1	1
	1	1	1	1
	1	1	1	1
	1	1	1	1

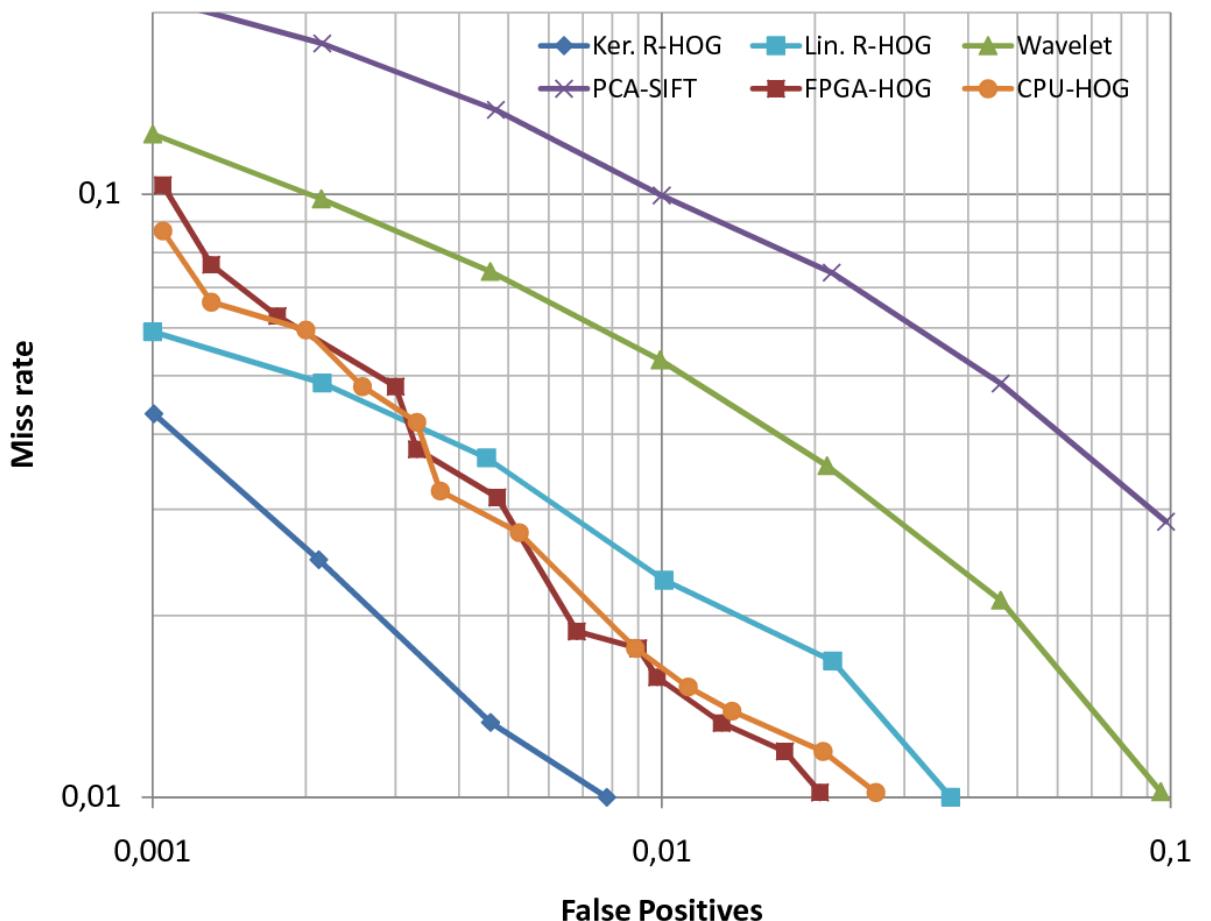
■ Classification Performance

- Evaluation on [INRIA] benchmark dataset
- Detection Error Tradeoff curves (DET)



■ Classification Performance

- Evaluation on [INRIA] benchmark dataset
- Detection Error Tradeoff curves (DET)
- Comparison to other detectors
close to reference
(Ker. R-HOG)
without retraining



Motivation | Method | Implementation | Evaluation | Summary

■ Time/Throughput

■ Latency

HOG step	Latency [μs] 800x600 px
Image Buffer	26.2
Downsampling	26.0
Gradient	52.0
Magnitude/Orientation	0.1
Histogram	207.2
TOTAL	311.6 μs

■ Time/Throughput

■ Latency

HOG step	Latency [μs] 800x600 px
Image Buffer	26.2
Downsampling	26.0
Gradient	52.0
Magnitude/Orientation	0.1
Histogram	207.2
TOTAL	311.6 μs

GPU-HOG [Wojek 2008]
13 ms (320x240 px)

■ Time/Throughput

- Latency

HOG step	Latency [μs] 800x600 px
Image Buffer	26.2
Downsampling	26.0
Gradient	52.0
Magnitude/Orientation	0.1
Histogram	207.2
TOTAL	311.6 μs

GPU-HOG [Wojek 2008]
13 ms (320x240 px)

- Throughput: full 30 fps (1600x1200 px) delivered by camera

■ Time/Throughput

▪ Latency

HOG step	Latency [μs] 800x600 px
Image Buffer	26.2
Downsampling	26.0
Gradient	52.0
Magnitude/Orientation	0.1
Histogram	207.2
TOTAL	311.6 μs

GPU-HOG [Wojek 2008]
13 ms (320x240 px)

- **Throughput:** full 30 fps (1600x1200 px) delivered by camera

■ Resources

Resource Type	Total Number	Usage Level
4 input LUTs	42,435	76%
Internal RAM	60	43%
Embedded Multipliers	18	18%

Xilinx Spartan 3 XC3S 4000
Processing resolution: 800x600

Motivation | Method | Implementation | Evaluation | **Summary**

■ What about our Goals?

- HOG descriptor in real-time on low cost FPGA
- Good classification results without retraining
- Real-time pedestrian recognition system for crossroad assistance (20 fps)



Motivation | Method | Implementation | Evaluation |

Summary

■ What about our Goals?

- HOG descriptor in real-time on low cost FPGA
- Good classification results without retraining
- Real-time pedestrian recognition system for crossroad assistance (20 fps)



■ Outlook

- Evaluation with retrained classifier
- Outsource further parts to FPGA (normalization, SVM)
- Application for recognition of other traffic participants

Thank you for your attention.



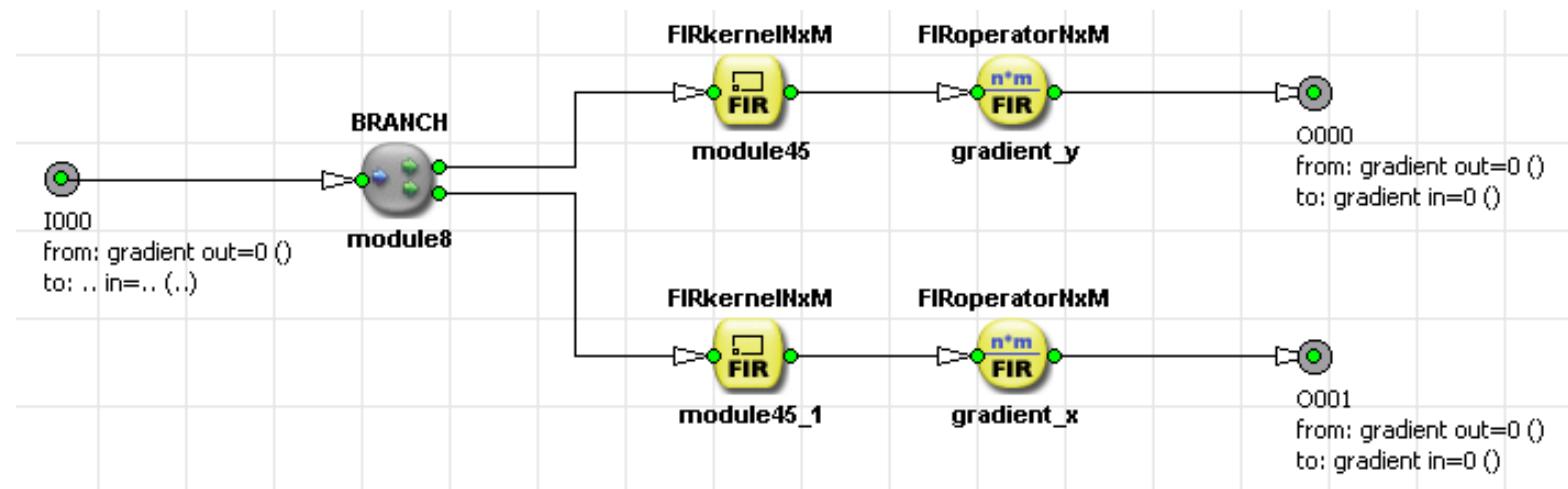
If you have any questions, feel free to ask.

References

- [Dalal::2005] *Histograms of Oriented Gradients for Human Detection.* Dalal, Triggs. CVPR 2005.
- [Cao::2008] *Real-Time Vision-based Stop Sign Detection System on FPGA.* Cao, Deng. DICTA 2008
- [Zhu::2006] *Fast Human Detection Using a Cascade of Histograms of Oriented Gradients.* Zhu, Avidan, Yeh, Cheng. CVPR 2006. MERL.

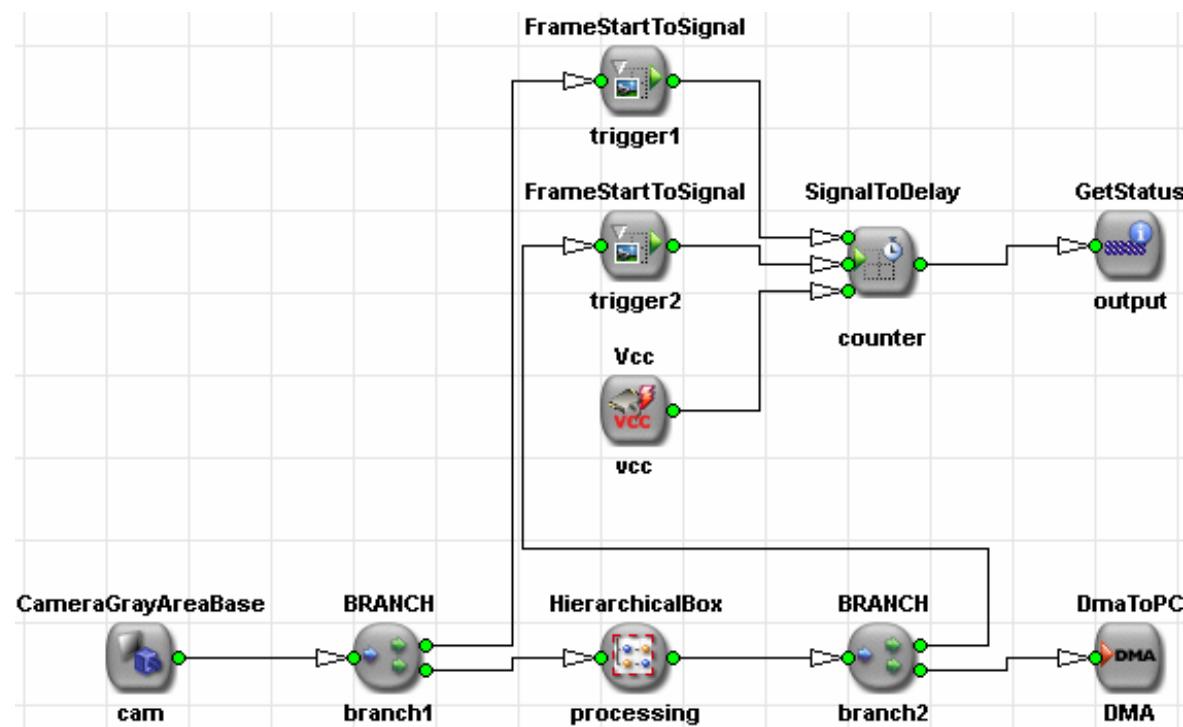
Additional Slides

■ VA: Gradient Convolution



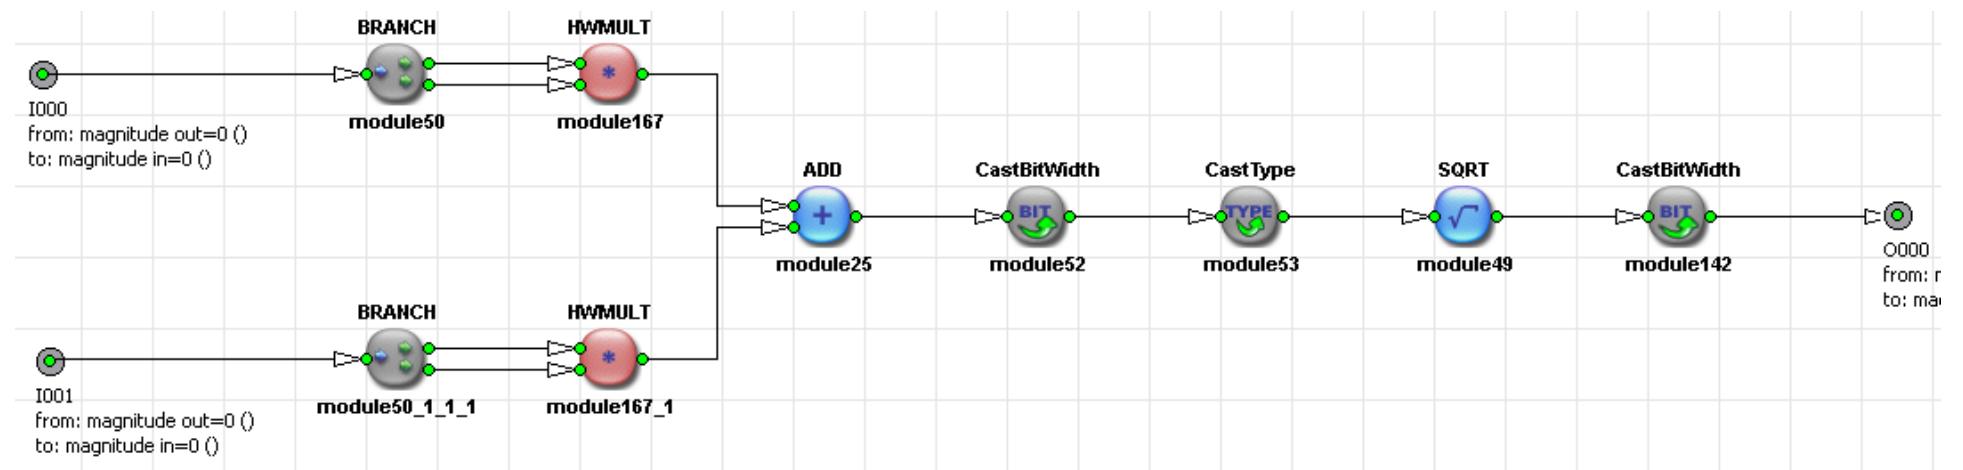
Additional Slides

■ VA: Latency Measurement



Additional Slides

- VA: Gradient Magnitude

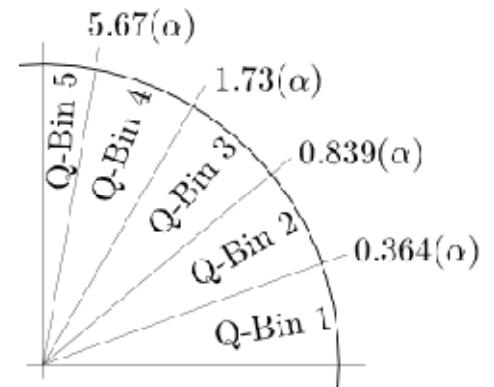
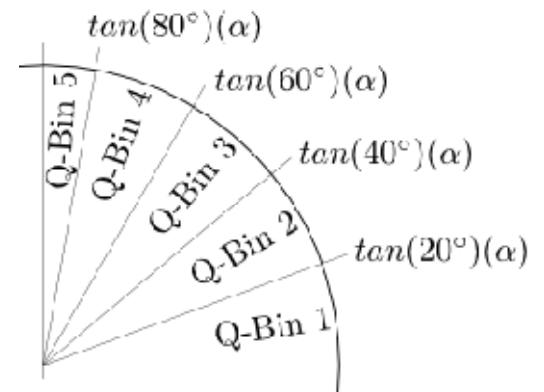


Additional Slides

- How to bypass arctan()

Cond. I: $G_x(x, y) > 0 \wedge G_y(x, y) > 0 \Leftrightarrow \text{quadrant I}$
 $\vee \quad G_x(x, y) < 0 \wedge G_y(x, y) < 0 \Leftrightarrow \text{quadrant III}$

Cond. II: $0 < \tan(\alpha) < \tan(20^\circ)$
 $0 < \left| \frac{G_y(x, y)}{G_x(x, y)} \right| < \tan(20^\circ)$
 $0 < |G_y(x, y)| < \tan(20^\circ) \cdot |G_x(x, y)|$
 $0 < |G_y(x, y)| < 0.364 \cdot |G_x(x, y)|$



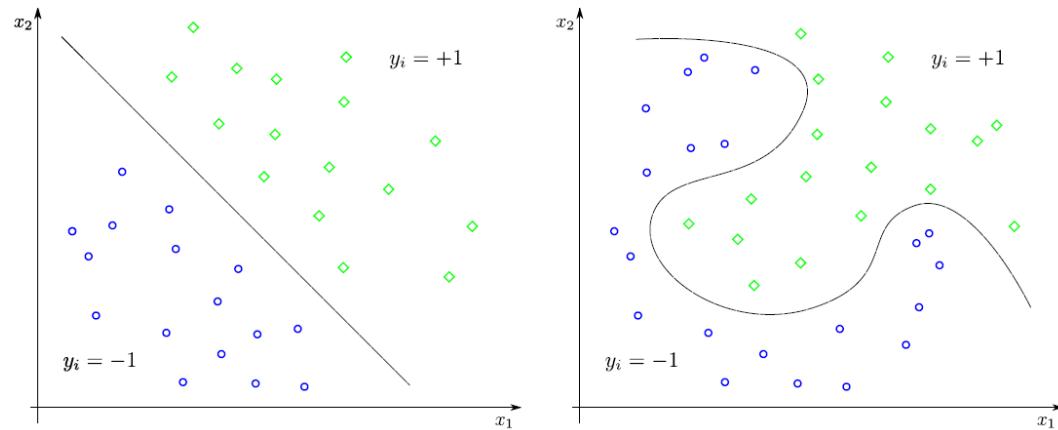
Additional Slides

■ HOG Modifications

- Grayscale images: 1.5% lower detection rate at 1E-4 FPPW (compared to color images)
- Gradient magnitude: decimal places are lost when the square root is extracted
- No Anti-Aliasing: Interpolation of votes between neighbouring histogram bin centres in both orientation and position
- No Gaussian down-weighting of pixels near the edges of the block: 1% lower detection rate at 1E-4 FPPW

Additional Slides

- Solving Nonlinear Problems with SVMs



- Kernel Trick

- Decision Function:

$$f(x) = \text{sgn}\left(\sum_{i=1}^{n_S} \alpha_i y_i K(s_i, x) + b_0\right)$$

Linear Kernels	$\langle \mathbf{x} \cdot \mathbf{y} \rangle$
Polynomial Kernels	$(\gamma \cdot \langle \mathbf{x} \cdot \mathbf{y} \rangle + \sigma)^p$
Radial Basis Function (RBF)/Gaussian Kernels	$e^{(-\gamma \ \mathbf{x} - \mathbf{y}\ ^2)}$
Sigmoid Kernels	$\tanh(\gamma \cdot \langle \mathbf{x} \cdot \mathbf{y} \rangle + \sigma)$