

# Histogram of Oriented Gradients for Human Detection

Navneet Dalal and Bill Triggs  
(presented by Lupeng and Yuduo)

# Introduction

Appearance / Clothing

Background

Illumination

Scales

Poses

...



# Introduction

**Objective:** Feature sets for robust visual object recognition used in detecting human in images

**Method:** Linear SVM as classifier based human detection grids of Histograms of Oriented Gradient (HOG) descriptors

# Descriptor Processing Chain

Person / non-person Input Image

Normalize gamma & colour

Compute gradients

Weighted vote into spatial & orientation cells

Contrast normalize over overlapping spatial blocks

Collect HOG's over detection window

Linear SVM

Person / non-person classification



(d)

# Descriptor Processing Chain

Person / non-person Input Image

Normalize gamma & colour

Compute gradients

Weighted vote into spatial & orientation cells

Contrast normalize over overlapping spatial blocks

Collect HOG's over detection window

Linear SVM

Person / non-person classification

**Tested with**

- RGB
- LAB (three axis color system)
- Grayscale

**Gamma Compression**

- Square root
- Logarithm

# Descriptor Processing Chain

Person / non-person Input Image

Normalize gamma & colour

Compute gradients


Weighted vote into spatial & orientation cells

Contrast normalize over overlapping spatial blocks

Collect HOG's over detection window

Linear SVM

Person / non-person classification



(a)

# Descriptor Processing Chain

Person / non-person Input Image

Normalize gamma & colour

Compute gradients


Weighted vote into spatial & orientation cells

Contrast normalize over overlapping spatial blocks

Collect HOG's over detection window

Linear SVM

Person / non-person classification



(e)

(f)

(g)

# Descriptor Processing Chain

Person / non-person Input Image

Normalize gamma & colour

Compute gradients


Weighted vote into spatial & orientation cells

Contrast normalize over overlapping spatial blocks

Collect HOG's over detection window

Linear SVM

Person / non-person classification





# Data Sets

- **MIT Pedestrian Database:**

509 training and 200 test images

Front and back views with **limited** poses



- **INRIA:**


1805 images

**No particular bias** on poses



Figure 2. Some sample images from our new human detection database. The subjects are always upright, but with some partial occlusions and a wide range of variations in pose, appearance, clothing, illumination and background.

# Methodology



- selected 2478 images as **positive** training set
- randomly sampled 12080 person-free - **negative**
- **retraining** using 12080 + “hard examples”





# Results



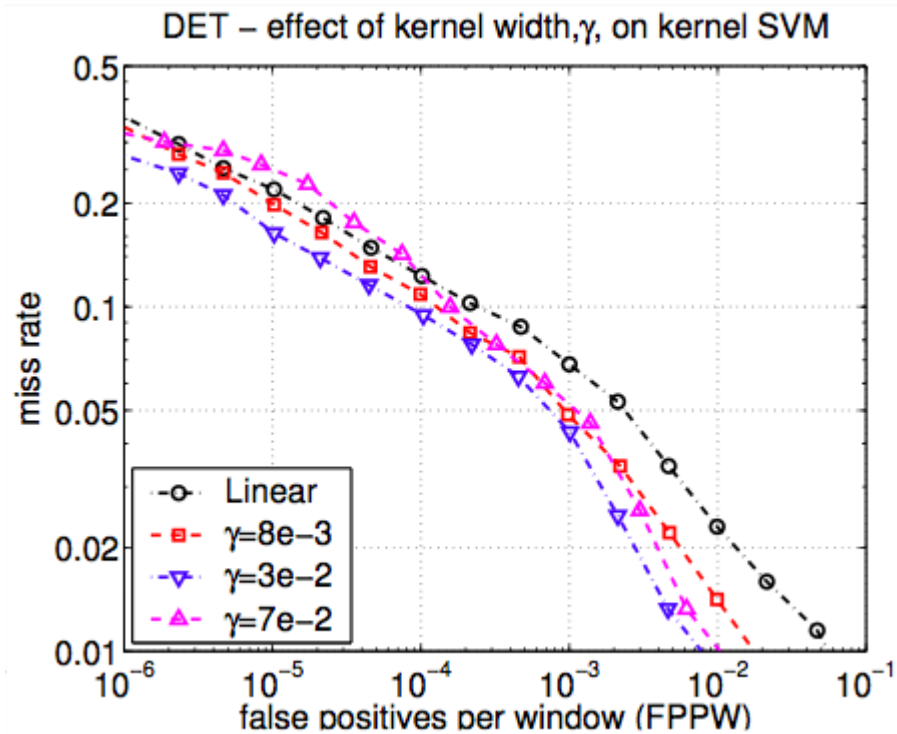

# Implementation



# Implementation



# Implementation



# Performance Analysis

1. fine-scale gradients
2. fine orientation binning
3. relatively coarse spatial binning
4. high-quality local contrast normalization in overlapping descriptor blocks

# Conclusion

- **Locally Normalized HOG Orientation Features**
  - good results for person detection
  - significantly reduce false positive rate
- **Influence of Descriptor Parameters**
  - gradients
  - orientation / spatial binning
  - local contrast normalization in descriptor blocks
- **More Challenging Pedestrian Database**
  - publicly available



# Future Work

- Optimize speed for efficiency
- Incorporate motion information
- Include parts-based model for performance

**Q & A**  
**Thank you!**