Master in Computer Vision Barcelona

UAB UOC UPC upf

M3 - Machine Learning for Computer Vision

Traffic Sign Detection and Recognition

Adrià Ciurana

Guim Perarnau

Pau Riba

Index

Correctly crop dataset
Bootstrap

Dataset generation

Extract features
Normalization
Dimensionality reduction

Data pre-processing

Sliding window
Detection
Recognition

Sign detection and recognition

Get metrics (F1-Score, AUC)
Visualize data

Evaluation

Introduction

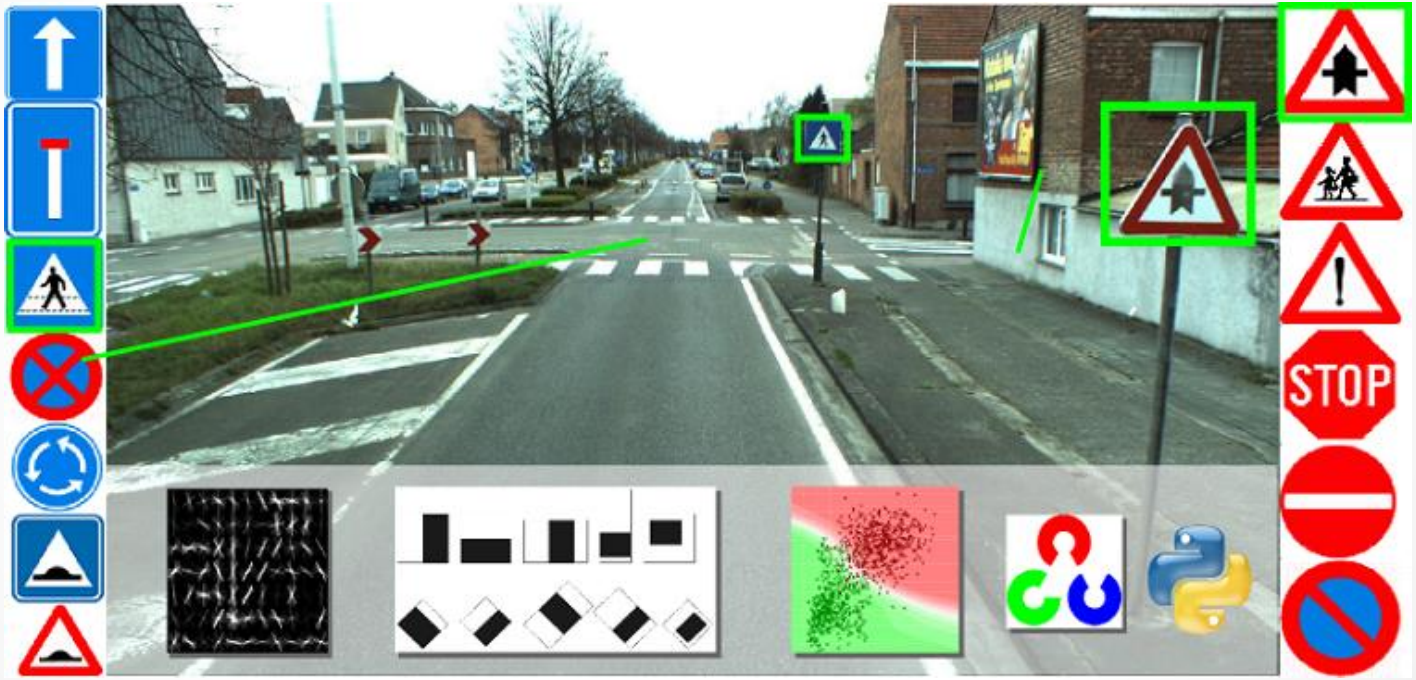

Motivation

Module 1 project segmentation

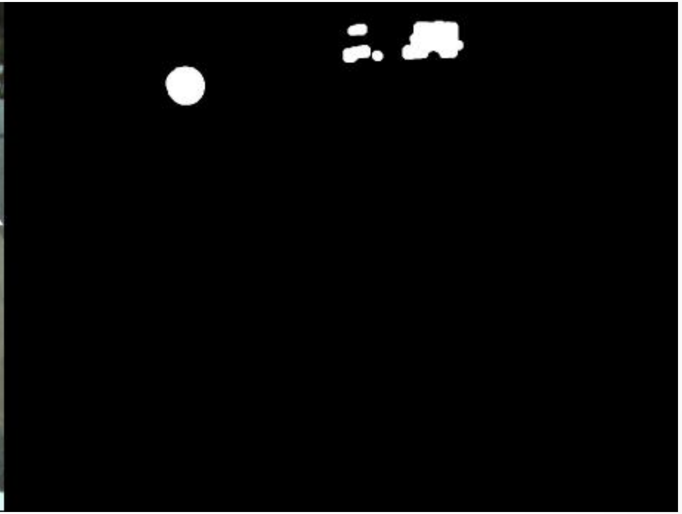

Per window results (669 images):

Precision	Accuracy	Recall	F1-Score	Time / frame
47.88%	38.25%	65.55%	55.34%	0.73 s

Pipeline

Bootstrap
Sliding Window Framework
Segmentation
Detection
Recognition
Evaluation

Image → Initial Detector

Round sign?

Triangular sign?

Square sign?

New background dataset

= False Positive

Pipeline

Bootstrap
Sliding Window Framework
Segmentation
Detection
Recognition
Evaluation

Pipeline

Bootstrap
Sliding Window Framework
Segmentation
Detection
Recognition
Evaluation

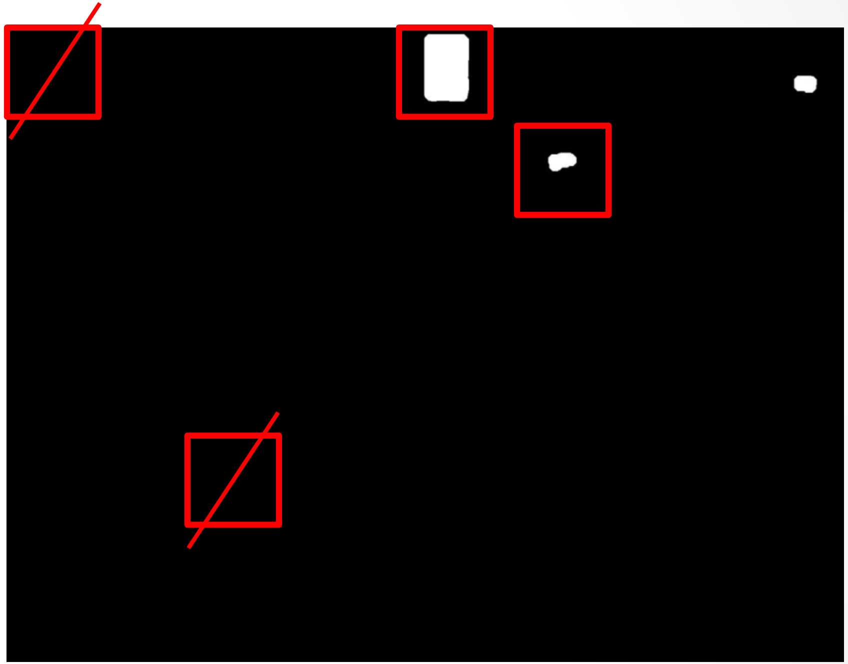

Pipeline

Bootstrap
Sliding Window Framework
Segmentation
Detection
Recognition
Evaluation

Pipeline

Bootstrap
Sliding Window Framework
Segmentation
Detection
Recognition
Evaluation

Pipeline

Bootstrap
Sliding Window Framework
Segmentation
Detection
Recognition
Evaluation

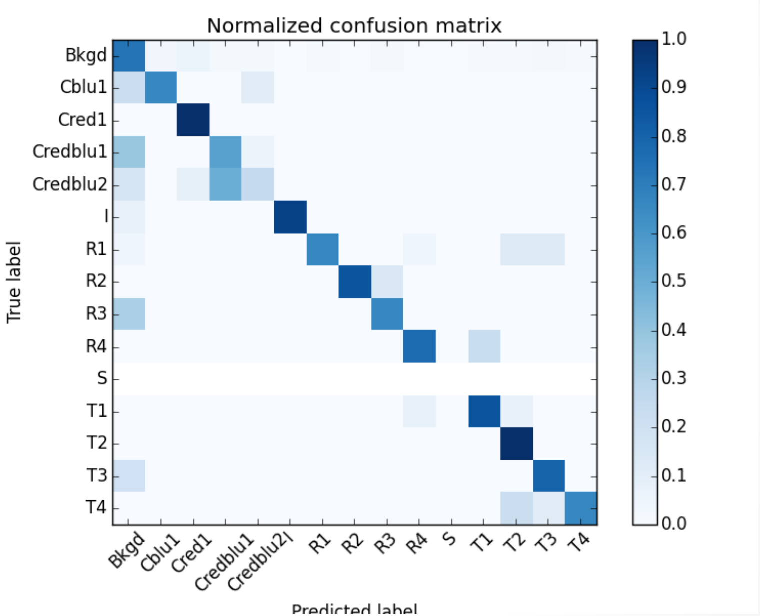

Dataset

1. http://btsd.ethz.ch/shareddata/

Dataset used: reduced ***BelgiumTS Dataset***[1] (62 classes)

Problems found:

- Traffic signs in (supposedly) only background images:   

- Traffic signs not labeled but correctly detected: 

Assumption:

- Do Not Care Object : types of signs that we will ignore (**No** penalization, **No** gain).

Crop training dataset

BelgiumTS Dataset already cropped images:

Problem:

1. Cropped images need to have a canonical size.
2. All signs must have the same height (vertical padding).

Crop training dataset

Solution: make our own 32x32 crops with 4 vertical padding pixels.

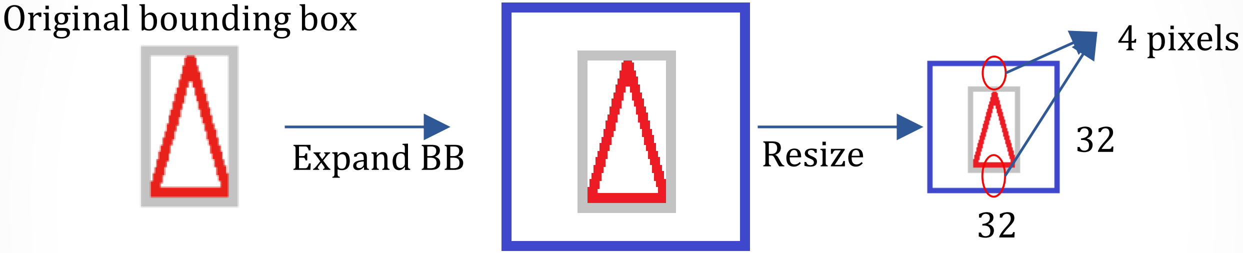

Results: 

Special case: sign is at image boundary → add boundary padding

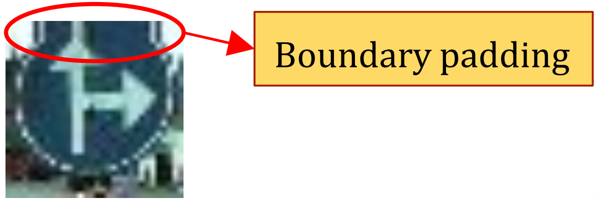

Bootstrap

= False Positive

Background Images

Initial Detector

Square sign?

Round sign?

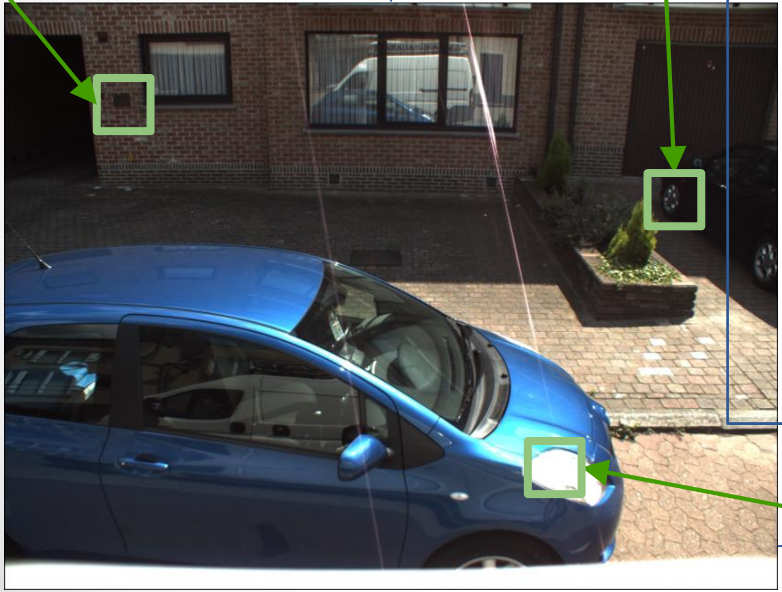

Triangular sign?

Train a new model adding False Positives

Result

New background dataset

Hard negatives

	Total
Initial	9863
Hard Negatives	11647
Total	21510

Segmentation

Original Image

Segmentation using YCbCr color space

Possible sign

Morphology

Advantages

- Speed up SW
- Reduces False Positives

Segmentation

However...

...we miss some signs!

Sliding window

Input image

Gaussian pyramid

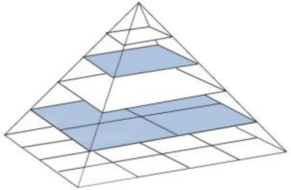

For each level of GP:

Segmentation

Integral Image

Sliding window of the image and the integral image

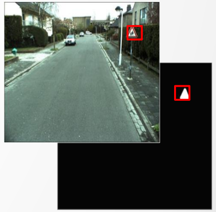

Possible sign region

Data augmentation

Idea: Generate more positive samples for each class.

- Flip samples:
  - Add more positive samples:

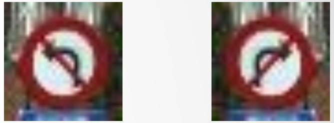

  - Flip not desired in some cases:

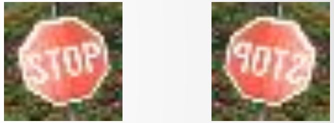

- Blur samples:
  - Smooth sudden changes. Gives the shape.

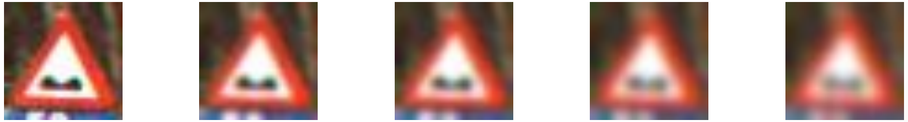

Original (3,3) (5,5) (7,7) (9,9)

Dataset division for detection

First idea: Background vs Signs

Problem: Very different kinds of signs. Separation is not easy.

Solution: Divide signs according to its shape:

Up-triangle

Down-triangle

Horizontal rectangle

No-flip

Vertical rectangle

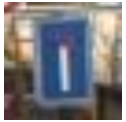

Parking

No-flip

Round

Stop

No-flip

Diamond

Detection

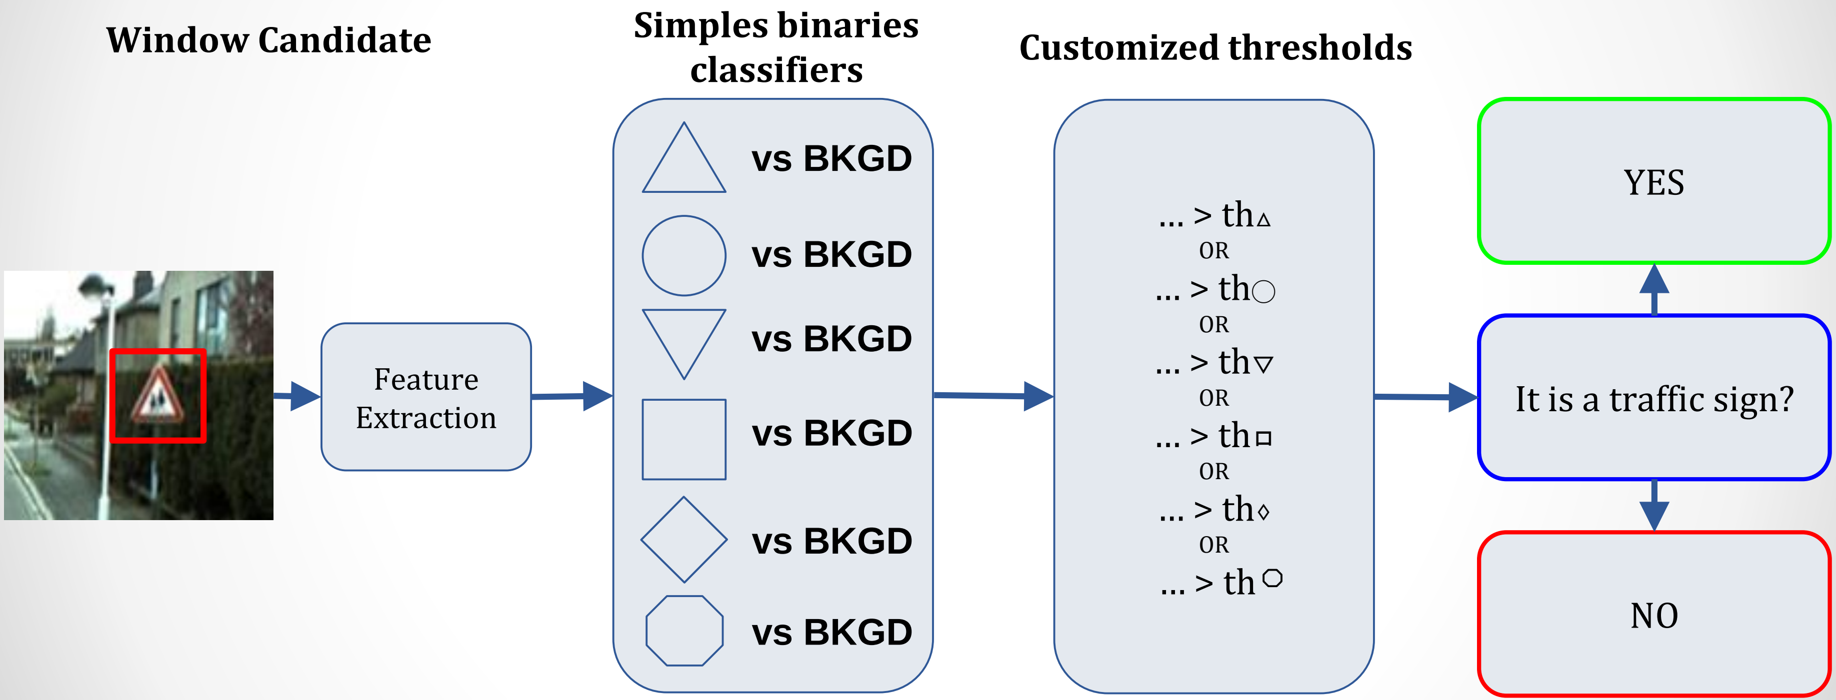

Window Candidate
Feature Extraction
Simples binaries classifiers
vs BKGD
vs BKGD
vs BKGD
vs BKGD
vs BKGD
vs BKGD
Customized thresholds
... > th△
OR
... > th○
OR
... > th▽
OR
... > th□
OR
... > th◊
OR
... > th⬡
It is a traffic sign?
YES
NO

Non maximum suppression

Multiple detection:

- **Red**: Ground truth
- **Green**: Detections

Combine detections: Overlap > threshold → Keep the best score.

- Pascal Vallotton (Pascal) $$\text{overlap} = \frac{\text{area}(\text{intersection}(A,B))}{\text{area}(\text{union}(A,B))}$$

- Pedro Felzenwalb (Pedro) $$\text{overlap} = \frac{\text{area}(\text{intersection}(A,B))}{\text{area}(A)}$$

score(A)<score(B)

- Technische Universität Darmstadt (TUD) $$\text{overlap} = \frac{\text{area}(\text{intersection}(A,B))}{\min(\text{area}(A),\text{area}(B))}$$

Recognition

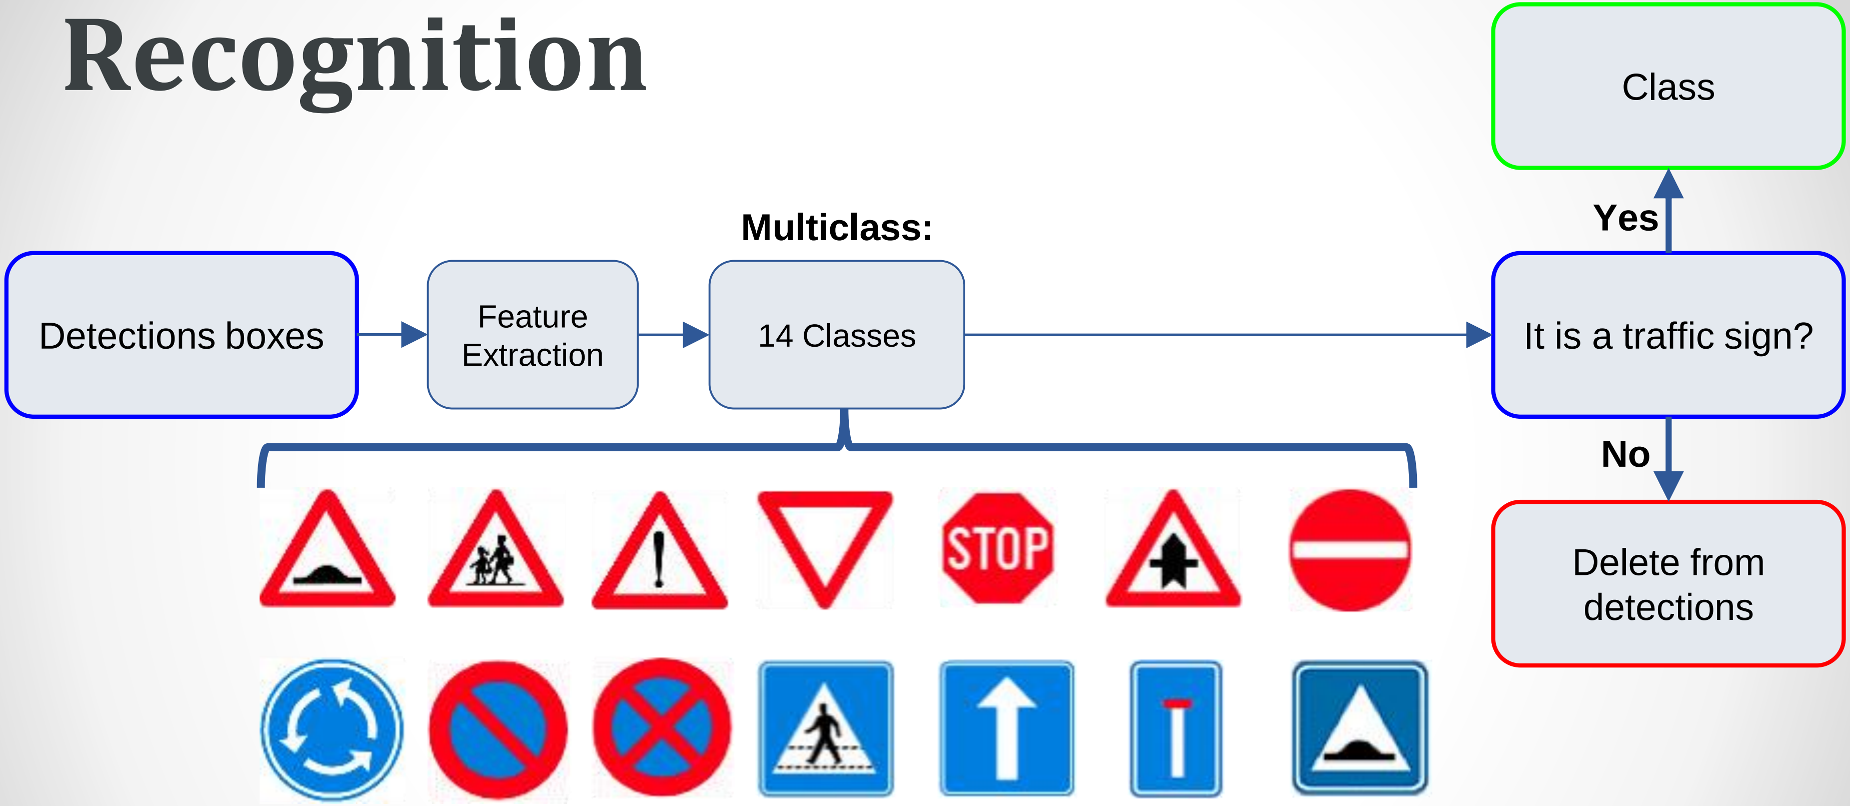

+ **Background (refinement step)**

Evaluation

- Train and test: cropped images
- Signs are centered
- Same scale

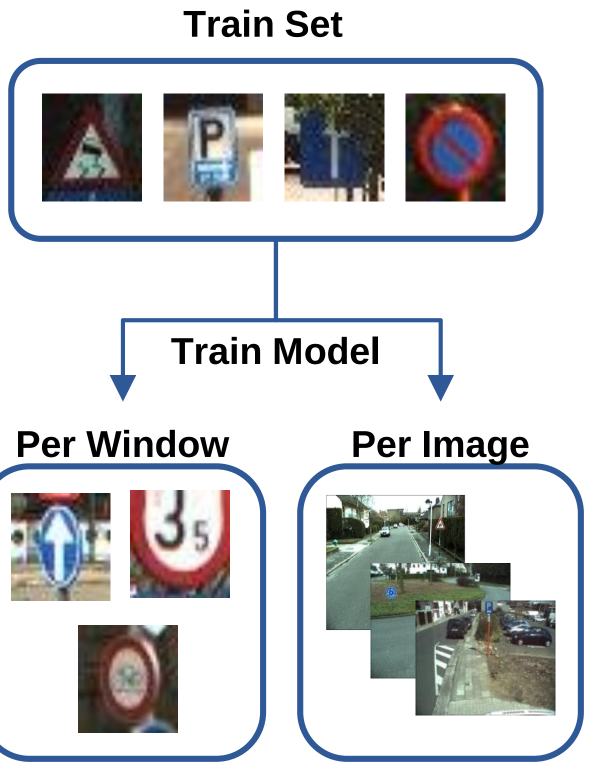

- Test: Sliding window
- Translation
- Different scale
- Multiple detections

Evaluation - Detection

Per window results:

Classifier	Solver	Feature Descriptor	Dimension Reduction	Data Normalization	F1-Score
SVM	Faster! Linear	HOG (4x4 pxc)	No	Yes	98.63%
		HOG (8x8 pxc)	No	Yes	97.95%
		HOG (8x8 pxc)	Yes (PCA)	Yes	97.30%
		HOG+ColorHist	Yes (PCA)	Yes	97.26%
	Slower RBF	HOG (8x8 pxc)	No	Yes	97.66%
		HOG+LBP	No	Yes	97.31%
		HOG Color Multichannel	No	Yes	96.98%
LDA	SVD	LBP	No	Yes	96.26%

Color is not important

Evaluation - Detection

Per image results:

Blurring the images is key

Classifier	Solver	Feature descriptor	Dimension reduction	Segmentation	Blur images	F1-Score
SVM	Linear	HOG	Yes (LDA)	Yes	Yes	55.17%
			Yes (LDA)	Yes	No	44.89%
			Yes (LDA)	No	No	24.86%
			No	No	No	21.49%
Cascade boosted classifiers	-	Haar + Adaboost	No	No	No	27.61%

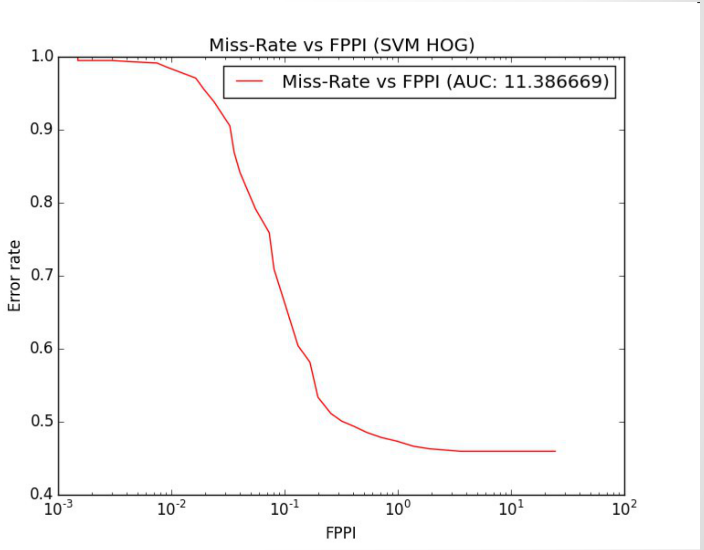

LDA and segmentation improve results and speed

Evaluation - Recognition

Model	Solver	Feature descriptor	Dimension reduction	F1-Score (per window)	F1-Score (per image)
ECOC + SVM (one vs rest)	Linear	HOG	Yes (LDA)	82.50%	64.68%
Neural Network	-	HOG	No	-	75.68%

	F1-Score	Precision	Recall
Mean	56.22%	52.52%	76.15%
Weighted	75.68%	81.05%	73.02%

SVM →

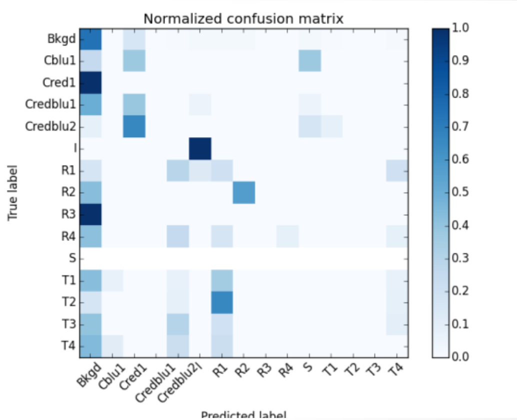

NN →

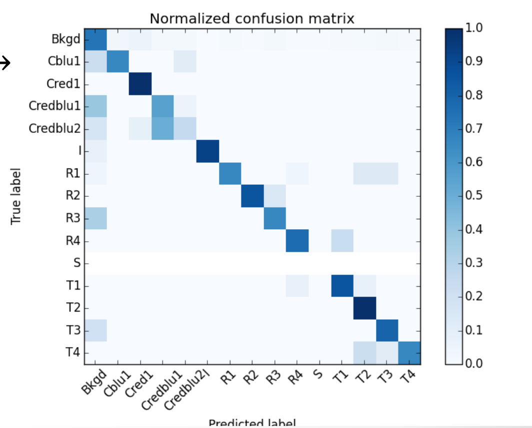

Evaluation - Whole Pipeline

Video

= Ground truth

= Estimated sign

= Do not care object

Note: this video shows the final output of the recognition given the detection, not the detection by itself.

Conclusions

- ✓ Color segmentation and parallelization saved us time.
- ✓ LDA improves performance (both speed and results).
- Tricks learned:
  - ✓ Correctly cropping the dataset
  - ✓ Bootstrap
  - ✓ Data augmentation
- ✗ Low results.

	M1	M3
F1-Score	55.34%	55.17%

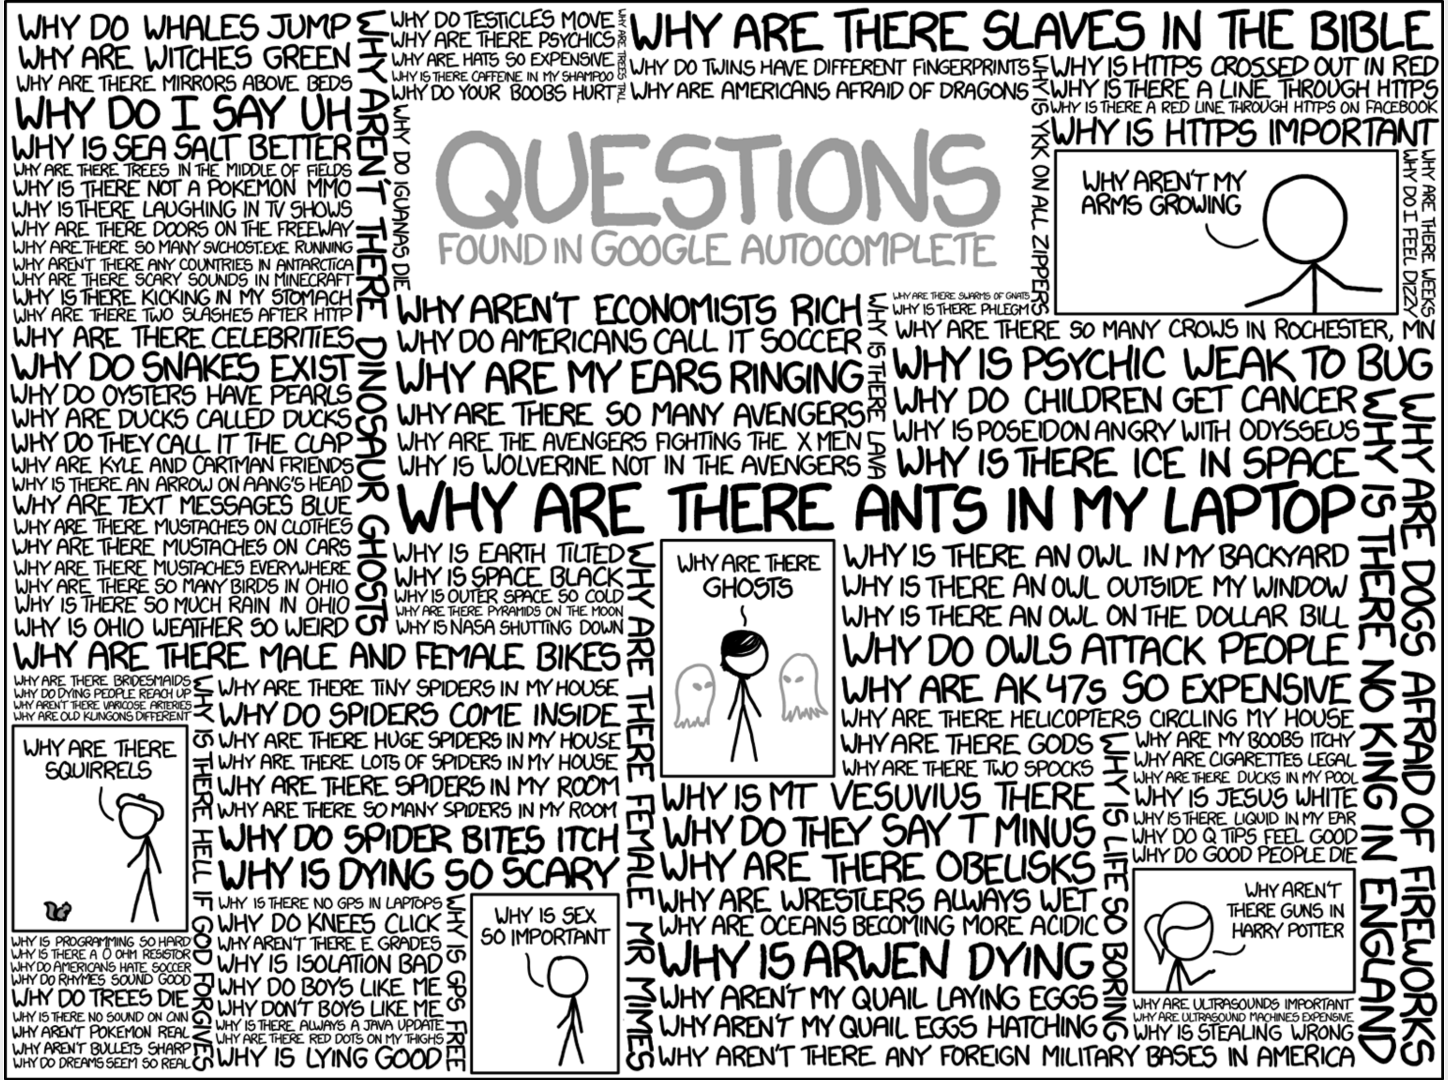
QUESTIONS
FOUND IN GOOGLE AUTOCOMPLETE
WHY DO WHALES JUMP
WHY ARE WITCHES GREEN
WHY ARE THERE MIRRORS ABOVE BEDS
WHY DO I SAY UH
WHY IS SEA SALT BETTER
WHY ARE THERE TREES IN THE MIDDLE OF FIELDS
WHY IS THERE NOT A POKEMON MMO
WHY IS THERE LAUGHING IN TV SHOWS
WHY ARE THERE DOORS ON THE FREEWAY
WHY ARE THERE SO MANY SVCHOST.EXE RUNNING
WHY AREN'T THERE ANY COUNTRIES IN ANTARCTICA
WHY ARE THERE SCARY SOUNDS IN MINECRAFT
WHY IS THERE KICKING IN MY STOMACH
WHY ARE THERE TWO SLASHES AFTER HTTP
WHY ARE THERE CELEBRITIES
WHY DO SNAKES EXIST
WHY DO OYSTERS HAVE PEARLS
WHY ARE DUCKS CALLED DUCKS
WHY DO THEY CALL IT THE CLAP
WHY ARE KYLE AND CARTMAN FRIENDS
WHY IS THERE AN ARROW ON AANG'S HEAD
WHY ARE TEXT MESSAGES BLUE
WHY ARE THERE MUSTACHES ON CLOTHES
WHY ARE THERE MUSTACHES ON CARS
WHY ARE THERE MUSTACHES EVERYWHERE
WHY ARE THERE SO MANY BIRDS IN OHIO
WHY IS THERE SO MUCH RAIN IN OHIO
WHY IS OHIO WEATHER SO WEIRD
WHY AREN'T THERE DINOSAUR GHOSTS
WHY DO TESTICLES MOVE
WHY ARE THERE PSYCHICS
WHY ARE HATS SO EXPENSIVE
WHY IS THERE CAFFEINE IN MY SHAMPOO
WHY DO YOUR BOOBS HURT
WHY ARE TREES TALL
WHY DO IGUANAS DIE
WHY ARE THERE SLAVES IN THE BIBLE
WHY DO TWINS HAVE DIFFERENT FINGERPRINTS
WHY ARE AMERICANS AFRAID OF DRAGONS
WHY IS YKK ON ALL ZIPPERS
WHY IS HTTPS CROSSED OUT IN RED
WHY IS THERE A LINE THROUGH HTTPS
WHY IS THERE A RED LINE THROUGH HTTPS ON FACEBOOK
WHY IS HTTPS IMPORTANT
WHY AREN'T MY ARMS GROWING
WHY ARE THERE WEEKS
WHY DO I FEEL DIZZY
WHY AREN'T ECONOMISTS RICH
WHY DO AMERICANS CALL IT SOCCER
WHY ARE MY EARS RINGING
WHY ARE THERE SO MANY AVENGERS
WHY ARE THE AVENGERS FIGHTING THE X MEN
WHY IS WOLVERINE NOT IN THE AVENGERS
WHY IS THERE LAVA
WHY ARE THERE SWARMS OF GNATS
WHY IS THERE PHLEGM
WHY ARE THERE SO MANY CROWS IN ROCHESTER, MN
WHY IS PSYCHIC WEAK TO BUG
WHY DO CHILDREN GET CANCER
WHY IS POSEIDON ANGRY WITH ODYSSEUS
WHY IS THERE ICE IN SPACE
WHY ARE THERE ANTS IN MY LAPTOP
WHY IS THERE NO KING IN ENGLAND
WHY ARE DOGS AFRAID OF FIREWORKS
WHY IS EARTH TILTED
WHY IS SPACE BLACK
WHY IS OUTER SPACE SO COLD
WHY ARE THERE PYRAMIDS ON THE MOON
WHY IS NASA SHUTTING DOWN
WHY ARE THERE GHOSTS
WHY IS THERE AN OWL IN MY BACKYARD
WHY IS THERE AN OWL OUTSIDE MY WINDOW
WHY IS THERE AN OWL ON THE DOLLAR BILL
WHY DO OWLS ATTACK PEOPLE
WHY ARE AK 47s SO EXPENSIVE
WHY ARE THERE HELICOPTERS CIRCLING MY HOUSE
WHY ARE THERE GODS
WHY ARE THERE TWO SPOCKS
WHY IS LIFE SO BORING
WHY ARE MY BOOBS ITCHY
WHY ARE CIGARETTES LEGAL
WHY ARE THERE DUCKS IN MY POOL
WHY IS JESUS WHITE
WHY IS THERE LIQUID IN MY EAR
WHY DO Q TIPS FEEL GOOD
WHY DO GOOD PEOPLE DIE
WHY AREN'T THERE GUNS IN HARRY POTTER
WHY ARE ULTRASOUNDS IMPORTANT
WHY ARE ULTRASOUND MACHINES EXPENSIVE
WHY IS STEALING WRONG
WHY ARE THERE MALE AND FEMALE BIKES
WHY ARE THERE FEMALE MR MIMES
WHY ARE THERE BRIDESMAIDS
WHY DO DYING PEOPLE REACH UP
WHY AREN'T THERE VARICOSE ARTERIES
WHY ARE OLD KLINGONS DIFFERENT
WHY ARE THERE SQUIRRELS
WHY IS THERE HELL IF GOD FORGIVES
WHY ARE THERE TINY SPIDERS IN MY HOUSE
WHY DO SPIDERS COME INSIDE
WHY ARE THERE HUGE SPIDERS IN MY HOUSE
WHY ARE THERE LOTS OF SPIDERS IN MY HOUSE
WHY ARE THERE SPIDERS IN MY ROOM
WHY ARE THERE SO MANY SPIDERS IN MY ROOM
WHY DO SPIDER BITES ITCH
WHY IS DYING SO SCARY
WHY IS PROGRAMMING SO HARD
WHY IS THERE A 0 OHM RESISTOR
WHY DO AMERICANS HATE SOCCER
WHY DO RHYMES SOUND GOOD
WHY DO TREES DIE
WHY IS THERE NO SOUND ON CNN
WHY AREN'T POKEMON REAL
WHY AREN'T BULLETS SHARP
WHY DO DREAMS SEEM SO REAL
WHY IS THERE NO GPS IN LAPTOPS
WHY DO KNEES CLICK
WHY AREN'T THERE E GRADES
WHY IS ISOLATION BAD
WHY DO BOYS LIKE ME
WHY DON'T BOYS LIKE ME
WHY IS THERE ALWAYS A JAVA UPDATE
WHY ARE THERE RED DOTS ON MY THIGHS
WHY IS LYING GOOD
WHY IS GPS FREE
WHY IS SEX SO IMPORTANT
WHY IS MT VESUVIUS THERE
WHY DO THEY SAY T MINUS
WHY ARE THERE OBELISKS
WHY ARE WRESTLERS ALWAYS WET
WHY ARE OCEANS BECOMING MORE ACIDIC
WHY IS ARWEN DYING
WHY AREN'T MY QUAIL LAYING EGGS
WHY AREN'T MY QUAIL EGGS HATCHING
WHY AREN'T THERE ANY FOREIGN MILITARY BASES IN AMERICA

Master in
Computer Vision
Barcelona

UAB UOC UPC upf.