Robotics

Erwin M. Bakker| LIACS Media Lab

Universiteit
Leiden

Bij ons leer je de wereld kennen

Organization and Overview

Period: February 7th – May 23rd 2022
Time: Monday 16.15 – 18.00
Place: Room 407 - 409
Lecturer: Erwin M. Bakker (erwin@liacs.nl)
Assistant: Hainan Yu (h.yu@liacs.leidenuniv.nl)

NB Register on Brightspace

Schedule:

7-2	Introduction and Overview
14-2	Locomotion and Inverse Kinematics
21-2	**Robotics Sensors and Image Processing**
28-2	SLAM + SLAM Workshop
7-3	**Mobile Robot Challenge Introduction**
14-3	**Project Proposals I (presentation by students)**
21-3	**Project Proposals II (presentation by students)**
28-3	Robotics Vision
4-4	Robotics Reinforcement Learning
11-4	Robotics Reinforcement Learning Workshop II
18-4	No Class (Eastern)
25-4	Project Progress I (presentations by students)
2-5	Project Progress II (presentations by students)
9-5	**Mobile Robot Challenge**
16-5	**Project Demos I**
23-5	**Project Demos II**

Website: http://liacs.leidenuniv.nl/~bakkerem2/robotics/

Grading (6 ECTS):
- Presentations and Robotics Project (60% of grade).
- Class discussions, attendance, workshops and assignments (40% of grade).
- It is necessary to be at every class and to complete every workshop and assignment.

Universiteit Leiden. Bij ons leer je de wereld kennen

Overview

- Sensors
- Lane Tracking
- OpenCV
- Line Tracking

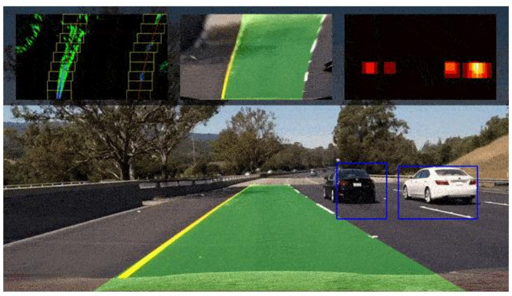

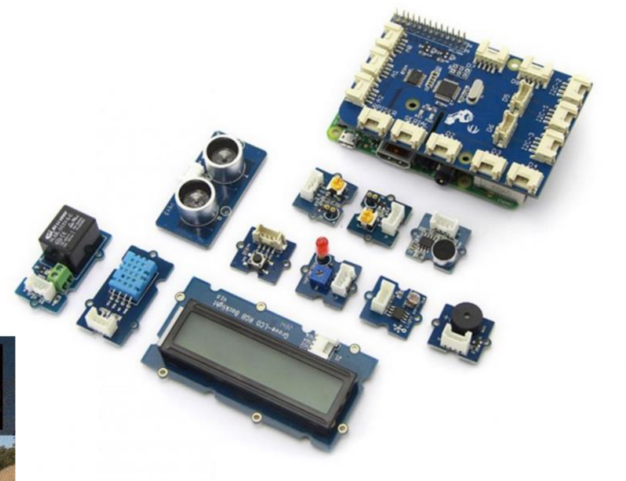

- Project Proposals
- Robotics Discussion Session

ROBOTICS SENSORS

- Bumper switches
- Acceleration, Orientation, Magnetic
- **IR/Visible Light**
- **Pressure**, Force
- **Ultrasonic, Lidar**, Radar
- **Camera's**, stereo camera's
- **Structured Light Camera's**

	Ultrasound	Infrared Triangulation	Laser	TeraRanger Time-of-Flight
High reading frequency	✗	✗	✓	✓
Long range	✗	✗	✓	✓
Minimal weight	✓	✓	✗	✓
Small form factor	✓	✓	✗	✓
Eye safety	✓	✓	Class 1 lasers only	✓
Use with multiple sensors	✗	✗	✗	✓

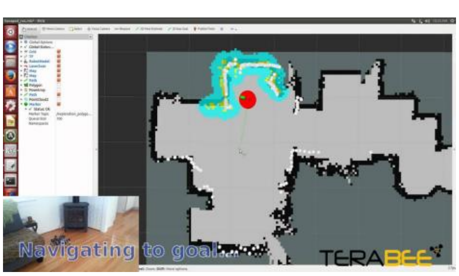

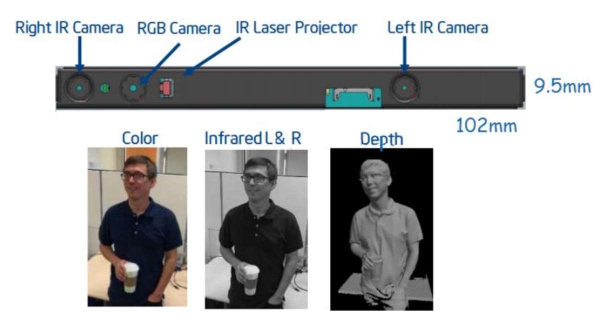

Universiteit Leiden. Bij ons leer je de wereld kennen

Structured Light

1. Primesense based Occipital
2. Asus X-tion Pro Live
3. Microsoft Kinect v1, v2
4. Intel RealSense F200, R200 (blue)

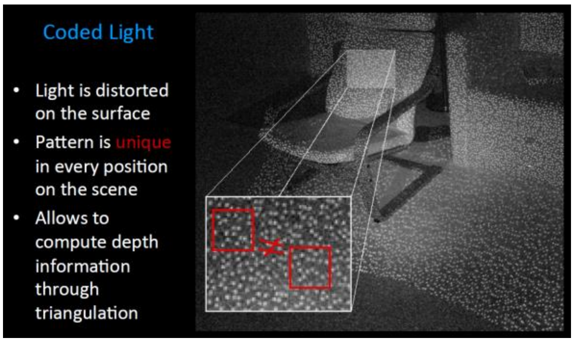

From: Anyline presentation by Peter Sperl

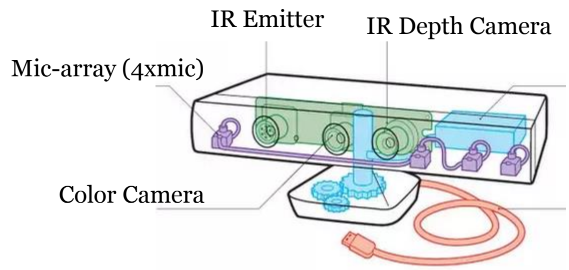

LIDAR Explanation

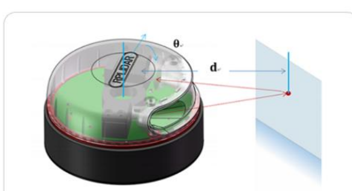

http://www.slamtec.com/en/lidar/A3

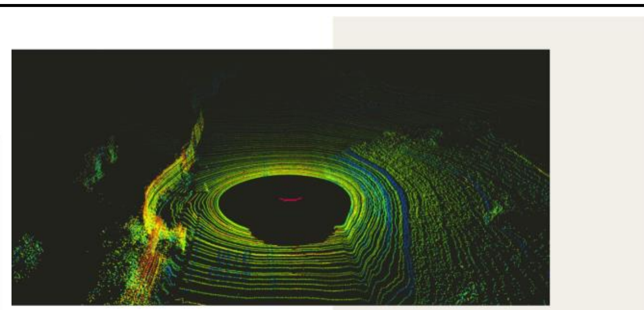

https://news.voyage.auto/an-introduction-to-lidar-the-key-self-driving-car-sensor-a7e405590cff

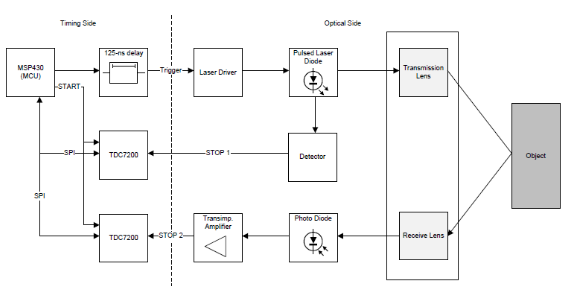

Texas Instruments LIDAR Pulsed Reference Design

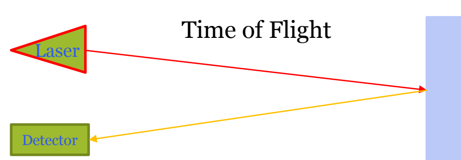

- Speed of light ~ 3 x 10^{8} m/s
- In 1 picosecond (= 10^{-12} sec) light travels ~ 3 x 10^{-4} m = 0.3 mm
- During 33 picoseconds light travels ~ 1cm

Robotics Self Driving Cars

Google Firefly, 2015

Google, 2009

Waymo OEM, 2021

Ended in Dec. 2020 -> Aurora Inovation

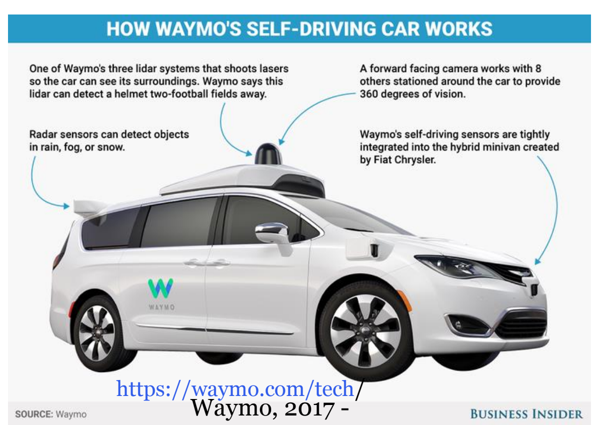

https://waymo.com/tech/

Waymo, 2017 -

CES2019

BMW Self Driving Car

InnovizOne Solid-state Lidar (goal: sub $1000 sensor)

- Angular resolution 0.1° x 0.1°
- FOV 120° x 25°
- 25 FPS
- Range 250m

Perception Capabilities

- Object detection and classification
- Lane detection
- Object Tracking
- SLAM

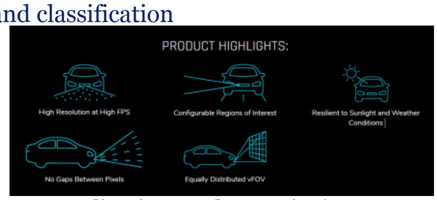

SLAM (Simultaneous Localization And Mapping)

INDY Autonomous Challenge (CES2021)

https://www.indyautonomouschallenge.com

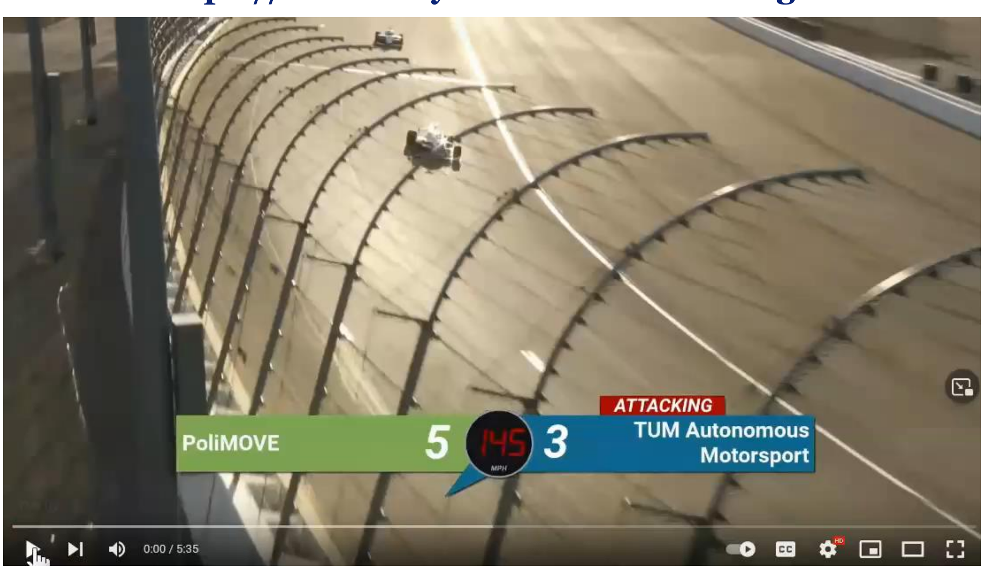

Researcher Hacks Self-driving Car Sensors

$60 lidar spoofing device generates fake cars, pedestrians and walls (2015)

By **Mark Harris**

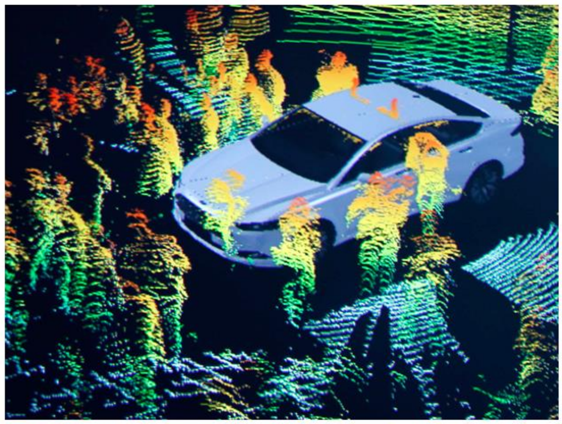

Photo: Jeff Kowalski/Corbis

https://spectrum.ieee.org/cars-that-think/transportation/self-driving/researcher-hacks-selfdriving-car-sensors

Illusion and Dazzle: Adversarial Optical Channel Exploits against Lidars for Automotive Applications (2017)

Hocheol Shin, Dohyun Kim, Yujin Kwon, and Yongdae Kim

Korea Advanced Institute of Science and Technology, Dajeon, Republic of Korea,
{h.c.shin, dohyunjk, dbwls8724, yongdaek}@kaist.ac.kr

Abstract. With the advancement in computing, sensing, and vehicle electronics, autonomous vehicles are being realized. For autonomous driving, environment perception sensors such as radars, lidars, and vision sensors play core roles as the eyes of a vehicle; therefore, their reliability cannot be compromised. In this work, we present a spoofing by relaying attack, which can not only induce illusions in the lidar output but can also cause the illusions to appear closer than the location of a spoofing device. In a recent work, the former attack is shown to be effective, but the latter one was never shown. Additionally, we present a novel saturation attack against lidars, which can completely incapacitate a lidar from sensing a certain direction. The effectiveness of both the approaches is experimentally verified against Velodyne's VLP-16.

Keywords: attack, autonomous car, sensor, lidar, saturating, spoofing

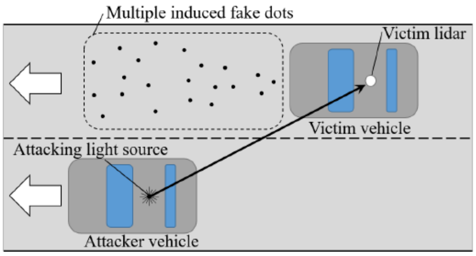

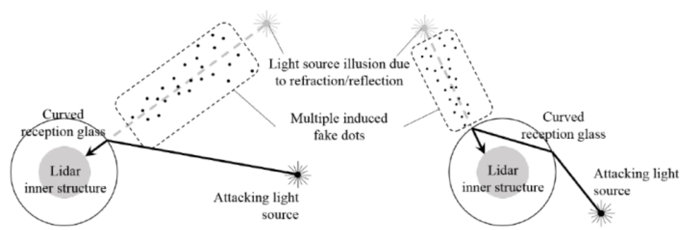

Adversarial Sensor Attack on LiDAR-based Perception in Autonomous Driving

by Yulong Cao et al.
CCS ’19, November 11–15, 2019, London, UK.

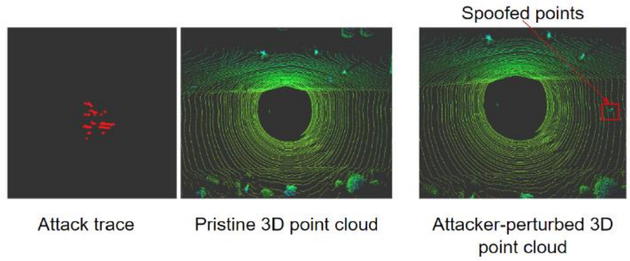

Figure 5: Generating the attacker-perturbed 3D point cloud by synthesizing the pristine 3D point cloud with the attack trace to spoof a front-near obstacle 5 meters away from the victim AV.

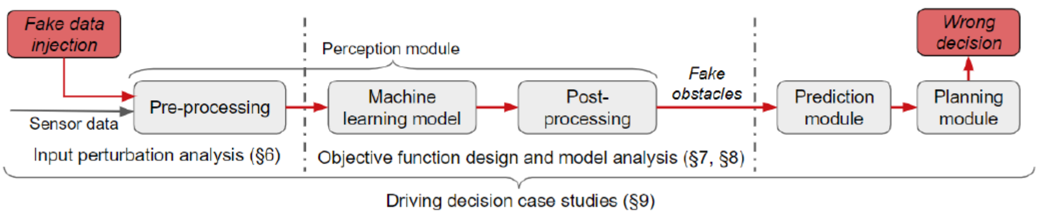

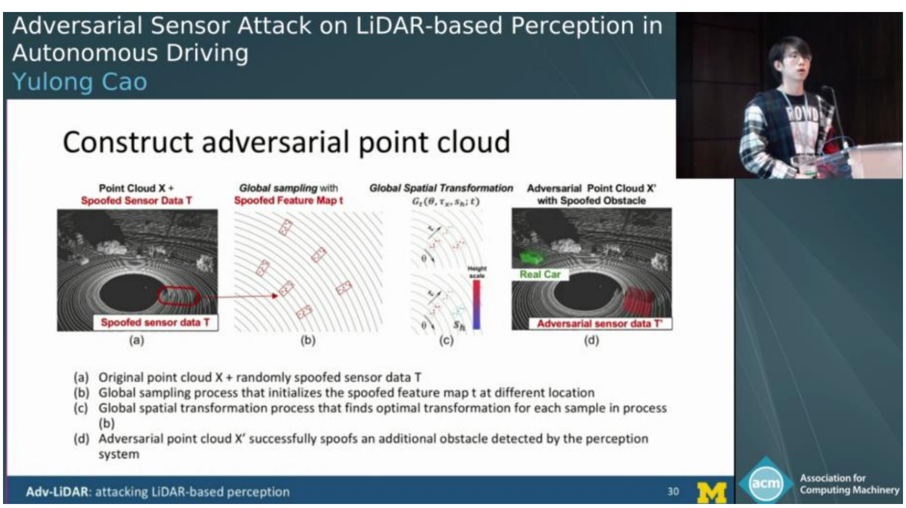

Many more studies since, also attacks against multimodal systems, e.g.,
Y. Chao et al., Invisible for both Camera and LiDAR: Security of Multi-Sensor Fusion based Perception in Autonomous Driving Under Physical-World Attacks, June 2021

Universiteit Leiden. Bij ons leer je de wereld kennen

The perfect anti-collision solution for any environment

Technology Comparison
distance sensors for robotics

	Ultrasound	Infrared Triangulation	Laser	TeraRanger Time-of-Flight
High reading frequency	✗	✗	✓	✓
Long range	✗	✗	✓	✓
Minimal weight	✓	✓	✗	✓
Small form factor	✓	✓	✗	✓
Eye safety	✓	✓	Class 1 lasers only	✓
Use with multiple sensors	✗	✗	✗	✓

RaspberryPi Sensors Kit

GrovePi+ Board for Raspberry Pi

De ATMEGA328 microcontroller communicates with the Raspberry Pi.

- Sound Sensor
- Temperature & Humidity
- Light Sensor
- Button
- UItrasonic Ranger
- Rotary Angle Sensor
- Gas Sensor

Symbol	MQ-2	MQ-3	MQ-5	MQ-9
Detect Gas	Combustible Gas, Smoke	Alcohol Vapor	LPG, Natural Gas, Town Gas	Carbon Monoxide, Coal Gas, Liquified Gas
Detect Concentration	300-10000ppm	0.04-4mg/L Alcohol	300-10000ppm	10-1000ppm CO;100-10000PPm Gas

Universiteit Leiden. Bij ons leer je de wereld kennen

Lane Tracking

- Joel C. McCall and Mohan M. Trivedi, **Video Based Lane Estimation and Tracking for Driver Assistance: Survey, System, and Evaluation.** IEEE Transactions on Intelligent Transportation Systems, 2006
- A. Bar Hillel, R. Lerner, D. Levi, G. Raz, **Recent progress in road and lane detection: a survey. Machine Vision and Applications** (2014) 25:727–745
- J. Fritsch, T. Kühnl, F. Kummert, **Monocular Road Terrain Detection by Combining Visual and Spatial Information.** IEEE Transactions on Intelligent Transportation Systems, 2014.
- J. Sattar, J. Mo, SafeDrive: **A Robust Lane Tracking System for Autonomous and Assisted Driving Under Limited Visibility.** January 31, 2017

(https://arxiv.org/abs/1701.08449)

Some example project for detecting road features using OpenCV:

https://navoshta.com/detecting-road-features/ by Alex Staravoitau

Lane Tracking

Joel C. McCall and Mohan M. Trivedi, **Video Based Lane Estimation and Tracking for Driver Assistance: Survey, System, and Evaluation.** IEEE Transactions on Intelligent Transportation Systems, 2006

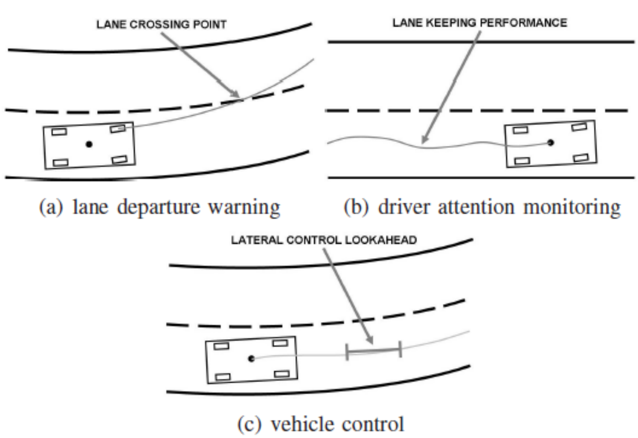

(a) lane departure warning (b) driver attention monitoring

(c) vehicle control

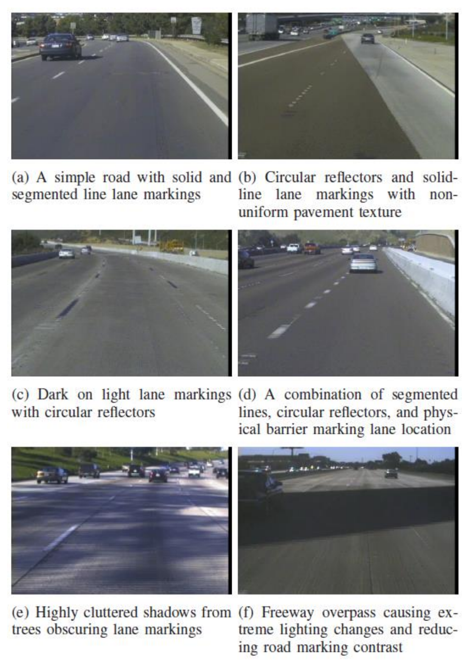

(a) A simple road with solid and segmented line lane markings

(b) Circular reflectors and solid-line lane markings with non-uniform pavement texture

(c) Dark on light lane markings with circular reflectors

(d) A combination of segmented lines, circular reflectors, and physical barrier marking lane location

(e) Highly cluttered shadows from trees obscuring lane markings

(f) Freeway overpass causing extreme lighting changes and reducing road marking contrast

Lane Tracking by Day and Night

Fig. 3. Images of the same stretch of road shown in the daytime and nighttime

General System

1. camera and vision sensors
2. internal vehicle state sensors
   e.g., speed, steering angle, etc.
3. line detectors, etc.
4. LASER RADAR sensors
5. GPS sensors

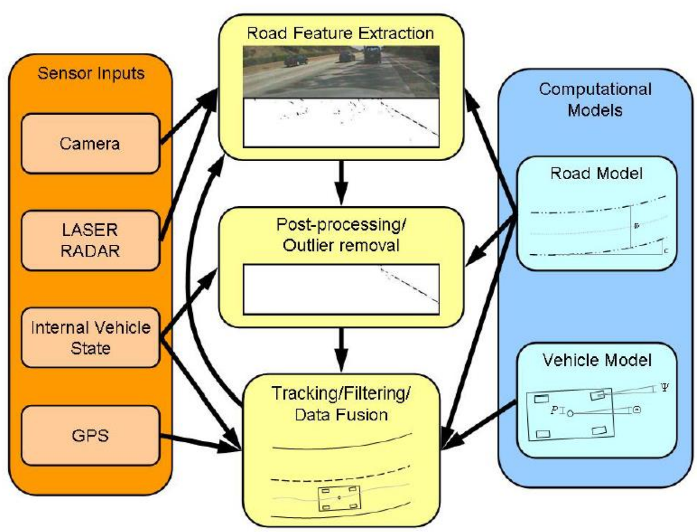

A generalized flow chart for lane position detection systems combining multiple modalities an iterative detection/tracking loop and road and vehicle models.

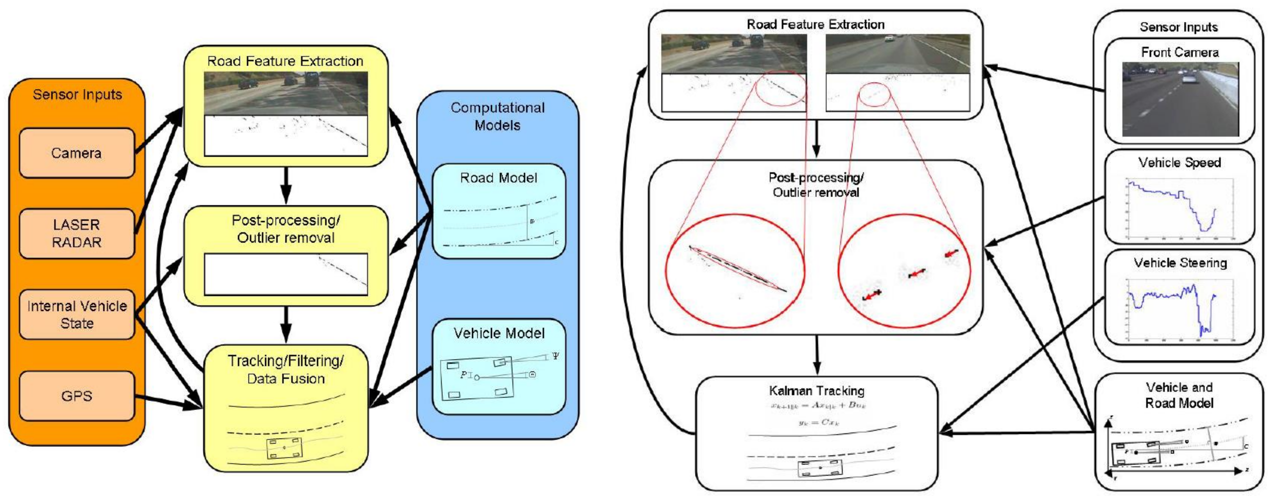

A generalized flow chart for lane position detection systems combining multiple modalities an iterative detection/tracking loop and road and vehicle models.

Vehicle and Road Models

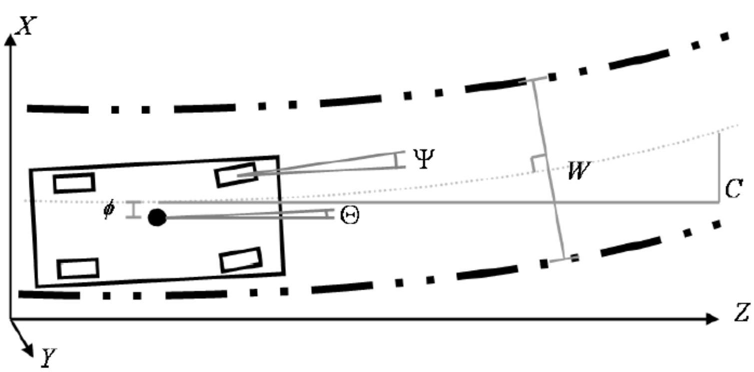

Fig. 6. Vehicle and road models used in the system. We are using a constant curvature road model and linearized vehicle dynamics for use in a Kalman filter.

Kalman Filter: Linear Quadratic Estimation to cope with noisy data.

System	Use[a]	Road Model	Feature Extraction	Postprocessing	Tracking	Evaluation	Comments
VaMoRs (1992) [16]	A	Clothoid Model with vertical curvature	Edge Elements	eliminates points which are not collinear	Linear vehicle dynamics model	Single frame images	Limited processing power. Simple edge detection fails in difficult situations.
YARF (1995) [33]	A	Circular road segments on flat plane	Hue based segmentation and edge detection	Averaging and linear median squares estimation	Operation on single frame	Positive detection rates for feature extraction, single frame images	Multiple detectors. Limited to yellow and white stripes.
ALVINN (1996) [19], [36]	A	Flat road model for generating training data	Image intensity	Neural Network	None	Road tests, various error measure associated with neural networks	Neural network makes it difficult to decouple control from detection, requires lots of training
RALPH (1996) [25]	A B	Constant curvature on flat plane	scan line matched to template	Template matching to slowly evolving near template and fast evolving far template	No inter-frame tracking described	Single frame images	template methods can fail near construction zone or areas where the road has changed. Shows limited quantitative results
GOLD (1998) [20]	C	Constant lane width on flat plane	Adaptive thresholding of pixel differences	Morphological widening	Operation on single frame	Single frame images	Assumes line markings on dark road. some robustness to lighting and occlusion
LOIS (1998) [34]	B A	Parabolic approximation on flat plane	Edge magnitudes and orientations	Maximum a posteriori estimation evaluated by Metropolis algorithm	Kalman filtering	Error histogram from one drive. Standard deviation of error 13cm	Robust to shadowing in presence of strong lane markings. Otherwise untested.
LANA (1999) [24]	B A	Parabolic approximation on flat plane	DCT coefficients for diagonally dominant edges	Maximum a posteriori estimation	Operation on single frame	Single frame images, comparison to LOIS shown	Only using diagonal DCT coefficients limits detection based on orientation of vehicle
Taylor et al. (1999) [12]	A	Constant curvature on flat plane	Template matching	Hough transform	Kalman Filter input into various control schemes	Performance of controllers shown	Focussed on controller performance. Limited real-world testing.
Ma et al. (2000) [13]	A B C	Circular road model on flat plane	Likelihood based on gradient image	Fusion on radar and optical images	Operation on single frame	Single frame images	Designed for elevated or bordered rural roads.

Universiteit Leiden. Bij ons leer je de wereld kennen

System	Use[a]	Road Model	Feature Extraction	Postprocessing	Tracking	Evaluation	Comments
Southall et al. (2001) [30]	C	Curvature and rate of change of curvature	Threshold both pixel values and cross-correlation to dark-bright-dark function	Factored sampling for particle filter	Particle Filtering via CONDENSATION	Estimates shown for an image sequence, no ground truth or quantitative results	Very limited results and testing. Unclear whether feature extraction will work in difficult situations.
Kwon and Lee (2002) [4], [31]	B	Piecewise linear	multiple "feature transformation modules"	combined with data fusion and constraint satisfaction, heuristic departure warning function	nonlinear filtering	analysis of departure warning system given	Good architecture for sensor fusion. Testing limited to false alarm rate of departure warning.
DARVIN (2002) [5]	A B	DGPS based maps of roads	Image gradient	match to DGPS data	nonlinear filtering	selected frames from experimentation	Directed towards urban driving. Heavy reliance on GPS data.
Lee et al. (2003) [37], [38]	B	Straight road on flat plane	Edge distribution function	Hough transform to extract lanes	Not discussed	Detection rate of lane departure warning	Robust to lighting. Will not work for circular reflectors.
Apostoloff et al. (2003) [29]	C	Not discussed	lane markers, road edge, color, width	Cue scheduling to determine which cues are used	Particle Filtering via Distillation	Success rate, mean absolute error for position, yaw, and road width.	Possibly fail in conditions of strong cues that contradict each other (i.e. fig. 2b)
Kang et al. (2003) [28]	D	Straight road on flat plane	Edge direction and magnitude	Connected-component analysis, Dynamic programming	Single frame operation	Qualitative comparison to hough transform based techniques, Single images shown	Focusses on showing visual comparison to hough transform based technique.
Nedevschi et al. (2004) [22]	D	3D model based on clothoids and roll angle	edge detection	outlier removal based on 3D location found with stereo camera system, roll angle detected	Kalman filtering	single images from road scenes with clearly marked lane boundaries	Simple edge detection not robust to shadows, occlusions
This paper (2004)	C B	Parabolic approximation on flat plane	Steerable filters, adaptive road template	Statistical and motion based outlier removal	Kalman Filtering	Extensive error evaluation described in section V-B	

Universiteit Leiden. Bij ons leer je de wereld kennen

Steerable Filters

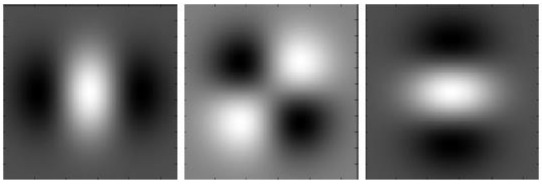

Fig. 7. A basis set for steerable filters based on the second derivatives of a two-dimensional Gaussian.

(a) A highway scene with complex shadowing from trees.

(b) Detection results for lines tuned to the lane angle.

Fig. 9. Filter results when lane markings are shadowed with complex shadows and non-uniform road materials.

(a) A typical highway scene encountered during evaluation.

(b) Results of filtering for circular reflectors.

(c) Results from filter for a line tuned to the lane angle.

Fig. 8. Application of Steerable filter road marking recognition for circular reflectors and solid lines on a highway

Inverse Perspective Warping and Template Matching

A perspective transformation, can be used to obtain an overview of the road ahead of the vehicle. This can make the problem of lane boundaries extraction easier.

- Curvature detection done by using an intensity template of past images in order to detect the curvature of the road ahead.

Inverse Perspective Warping

(a) Detected lanes with curvature overlaid onto image

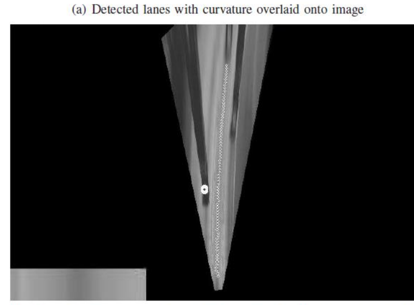

(b) Inverse perspective warping showing curvature detection (small white dots) and template (lower left corner)

Fig. 10. Curvature detection in the VioLET lane tracking Document1 - Word

A. System Test-bed Configuration and Test Conditions

Fig. 11. The LISA-Q intelligent vehicle test bed. Inset are close up views of the front camera (left inset) used for detection and tracking and side camera (right inset) used for generating ground truth.

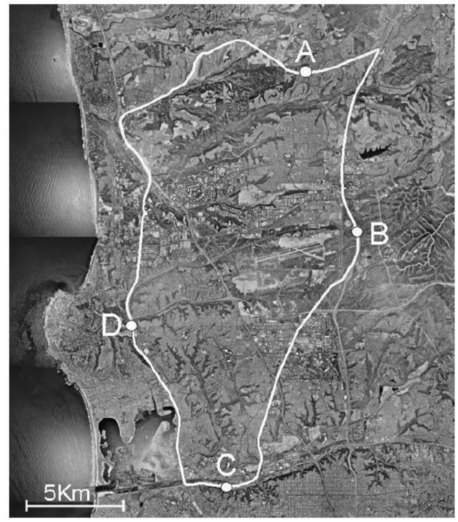

Fig. 14. The 65Km route in San Diego used in the evaluation. The route is overlayed on aerial photography. Points A, B, C, and D are sections of road used in the evaluation (photography courtesy USGS)

Results

Fig. 15. Scenes from dawn (row 1), daytime (row 2), dusk (row 3), and nighttime (row 4) data runs for each of the four sections of road (A, B, C, and D in figure 14). These scenes show the environmental variability caused by road markings and surfaces, weather, and lighting.

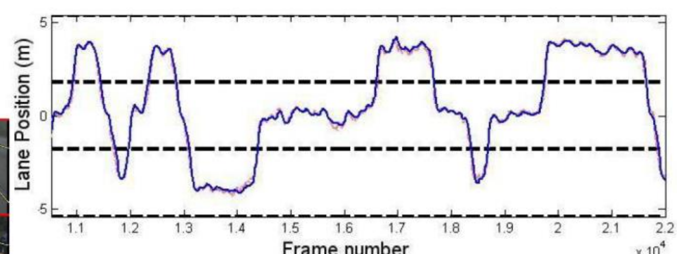

ig. 12. Detected lateral position in meters (solid blue) superimposed on ground truth (dashed pink) plotted vs. frame number with dashed lines marking e position of lane boundaries for an 11,000 frame (slightly over 6 minute) sequence.

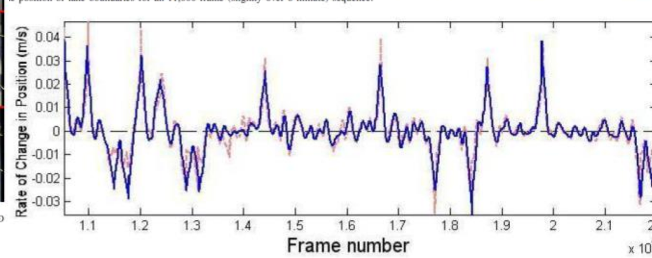

Fig. 13. Detected departure rate in m/s^2 (solid blue) superimposed on ground truth (dashed pink) plotted vs. frame number with dashed line marking the abscissa for the same sequence shown in figure 12.

Ground truth = pink dashed line

Challenges: Occlusions and Highlights

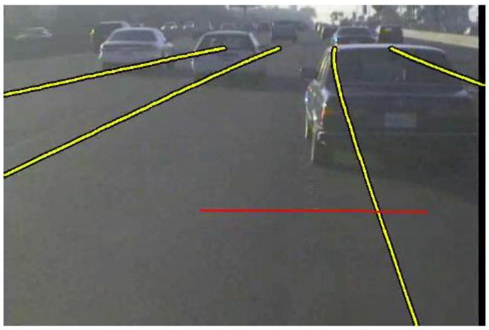

Fig. 16. Error due to occlusion of the road by a vehicle on the dusk dataset on road segment C. The red horizontal line shows the proximity of the occluding vehicle detected by the in-vehicle LASER RADAR sensors.

Fig. 17. Scenes from the special case scenarios of complex shadowing (top row) and tunnels (bottom row). These scenes highlight the extreme variability that can occur within short sections of road.

Road Detection

1. H. Kong, et al. General road detection from a single image, IEEE Transactions on Image Processing, Vol. 19, Issue 8, Aug. 2010.
2. J.M. Alvarez, et al., Road Detection Based on Illuminant Invariance, IEEE Transactions on Intelligent Transportation Systems, 2010.

Road Detection Based on Illuminant Invariance

Goal: Identify all pixels that belong to the road.

Challenges:

- Images from a mobile platform, outdoor: shadows, sunlight, etc.
- Changing background: trees, buildings and many (moving) objects: cars, bikes, pedestrians, dogs, signs, etc.
- Road types in different shapes and forms.

Sensors:

- Monocular => Features: Color, Texture, etc.

J.M. Alvarez, et al. 2010

Road Detection Based on Illuminant Invariance

Challenges: Shadows, sunlight, cloudy, etc.

Camera calibration using Macbeth color checker, can not be used for onboard self calibration and image content.

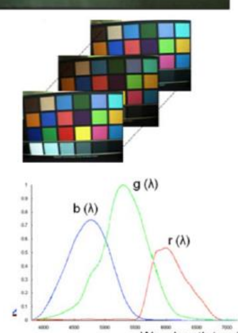

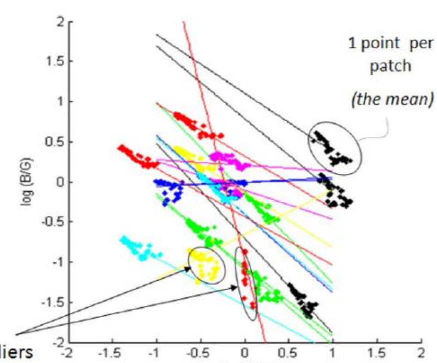

Road Detection Based on Illuminant Invariance

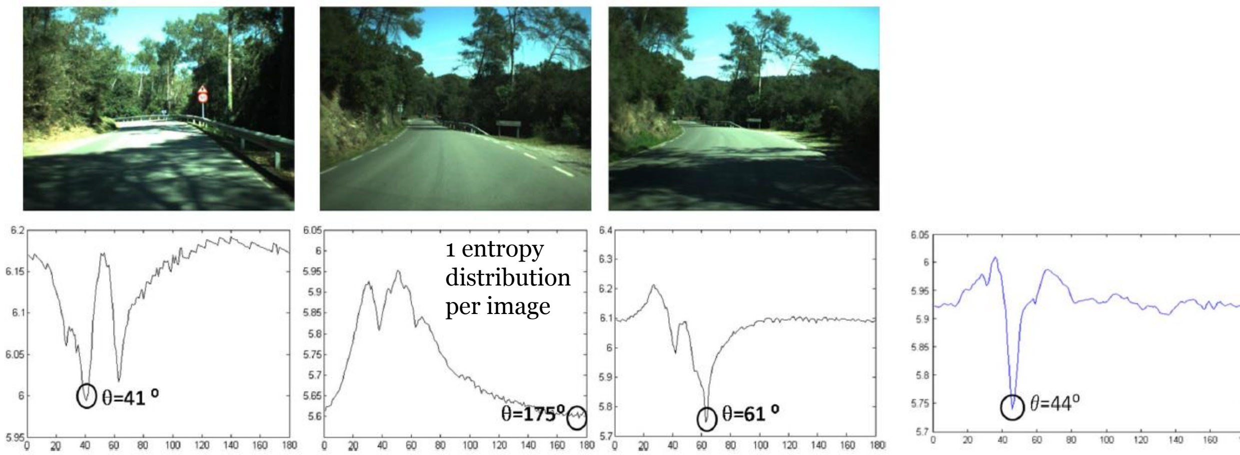

Entropy based camera calibration: can be used for onboard self calibration and image content.

Θ is the absolute minimum of the average distribution of entropy values for Image$_{RGB}$ (minimization procedure)

Road Detection Based on Illuminant Invariance

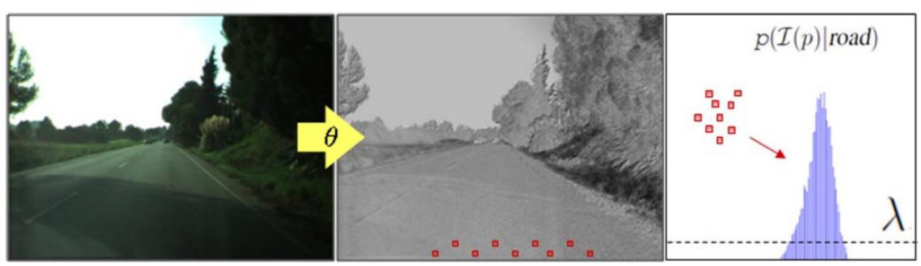

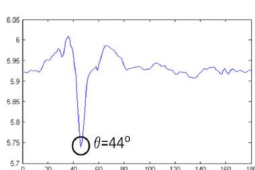

Fig. 7. (Left) Color image I_{RGB}. (Middle) Corresponding $\mathcal{I}$ computed using the invariant angle θ. (Right) Nonparametric road model built using the normalized histogram formed with the surrounding region of several seeds (nine in this case) placed at the bottom part of $\mathcal{I}$.

Entropy based camera calibration resulted in Θ is the invariant angle

that allows us to compute I a grayscale image invariant to lighting variations.

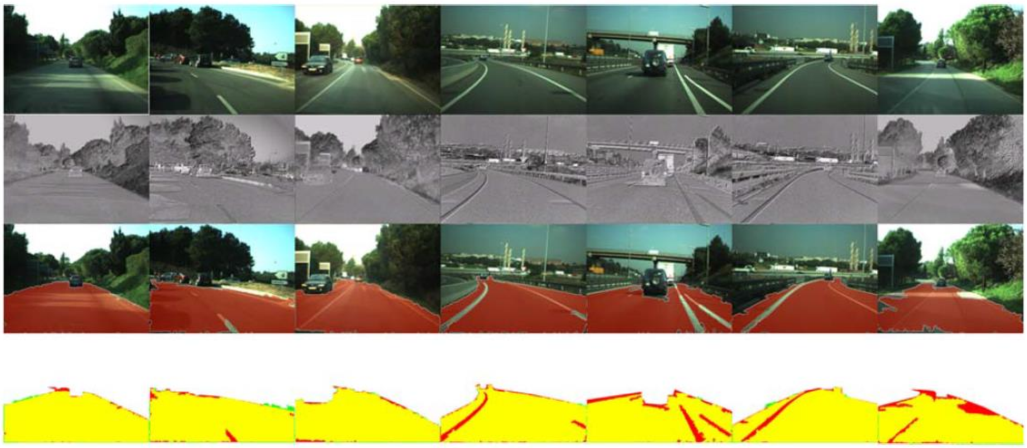

Fig. 11. Road-detection examples. (Top row) Original image. (Second row) Illuminant-invariant image. (Third row) Detected road. (Bottom row) Comparison against hand-segmented result. (Yellow) Correctly classified pixels. (Green) Falsely detected road pixels. (Red) False background pixels.

Fig. 12. Road-detection results. Example results include the following: (First column) Images from a sunny day with annoying shadows and traffic; (second column) images from a wet road without puddles with dim shadows and traffic; (third column) different scenarios; (fourth column) different road shapes; (last column) presence of other vehicles.

Universiteit Leiden. Bij ons leer je de wereld kennen

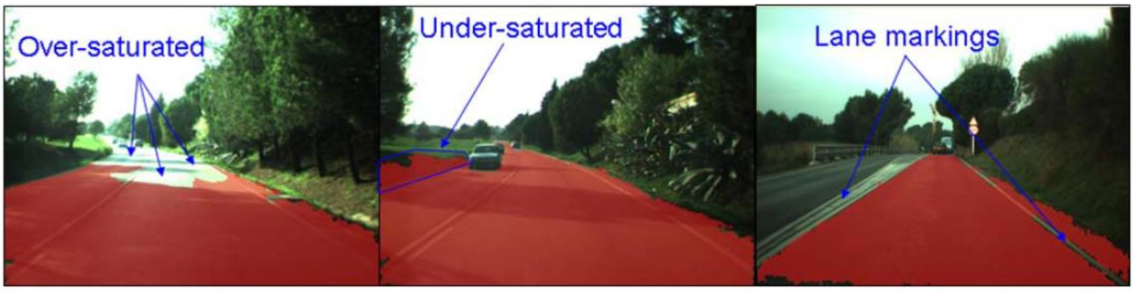

Fig. 13. Examples where our road-detection method currently fails.

General road detection from a single image

H. Kong et al. 2010

Road in different shapes and forms under different lighting with changes in colors, and textures.

General road detection from a single image

H. Kong et al. 2010

Road in different shapes and forms under different lighting with changes in colors, and textures.

Detection

1. Vanishing point estimation of main road
2. Segmentation based on this vanishing point

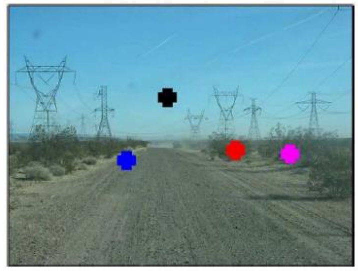

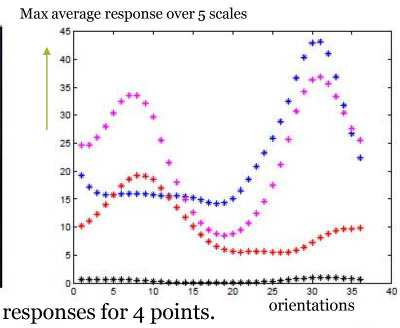

Gabor complex responses for 4 points.

Voting scheme:

- High-confidence voters using texture orientations resulting from Gabor filters (5 scales 36 orientations)
- Vanishing-point-constrained edge detection for road boundaries

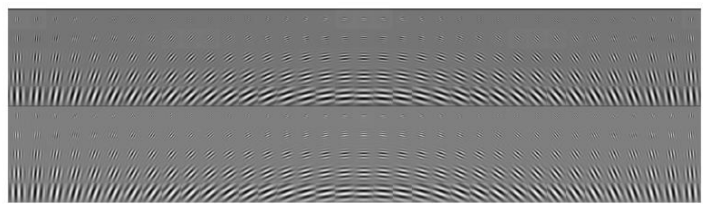

Gabor kernels: real part (row 1-5), imaginary part (row 6-10)

General road detection from a single image

H. Kong et al. 2010

Road in different shapes and forms under different lighting with changes in colors, and textures.

Detection

1. Vanishing point estimation of main road
2. Segmentation based on this vanishing point

Line segments

vanishing point with rays

orientation consistent rays

compute color difference of neighboring regions

Voting scheme:

- High-confidence voters using texture orientations resulting from Gabor filters (5 scales 36 orientations)
- Vanishing-point-constrained edge detection for road boundaries

dominant border (red) => vanishing point adaptation

General road detection from a single image

H. Kong et al. 2010

Result

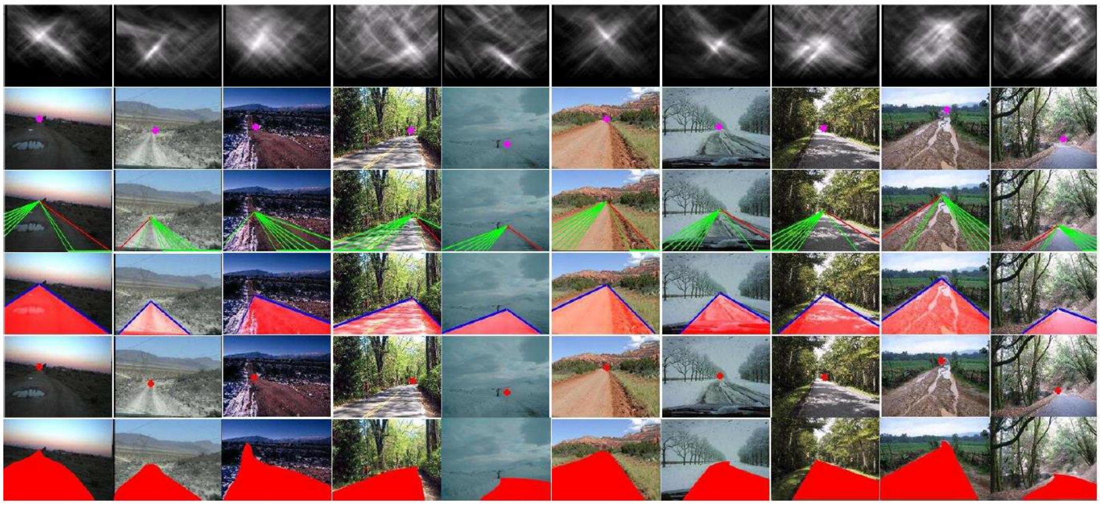

Last row: ground truth..

Lane Tracking

J. Fritsch, T. Kühnl, F. Kummert, **Monocular Road Terrain Detection by Combining Visual and Spatial Information.** IEEE Transactions on Intelligent Transportation Systems, 2014.

Overview System

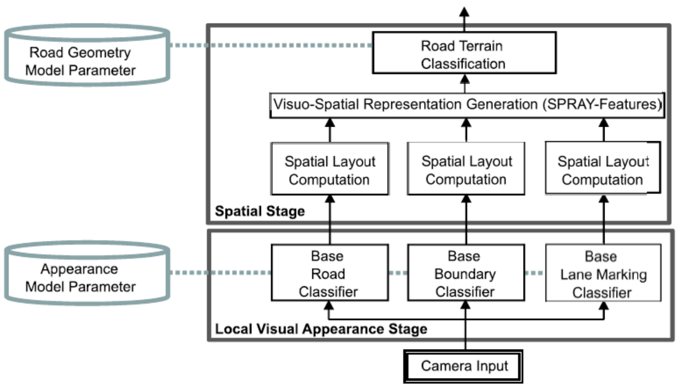

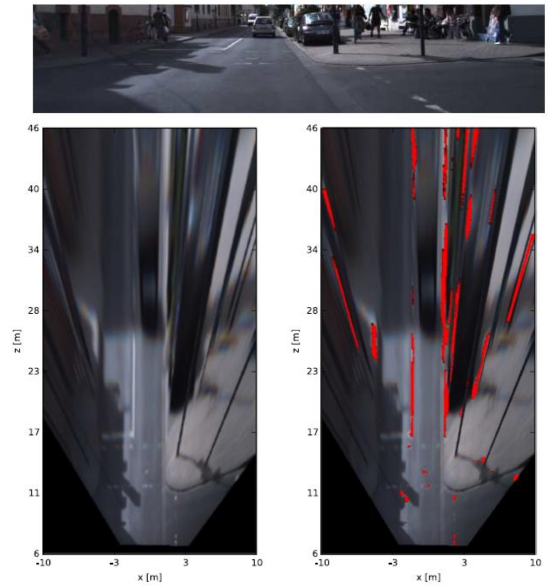

Ground Truth

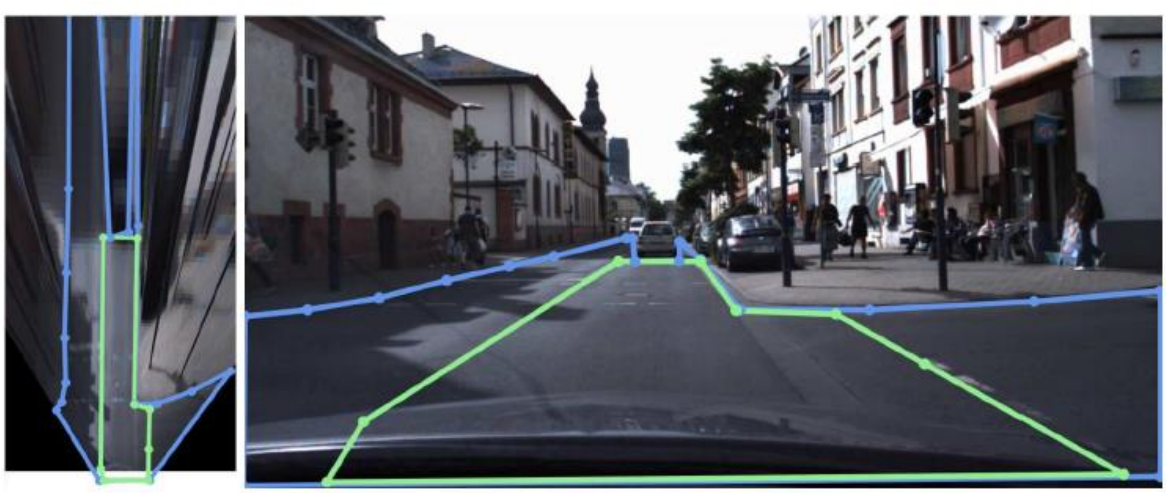

Slow Feature Analysis (SFA)

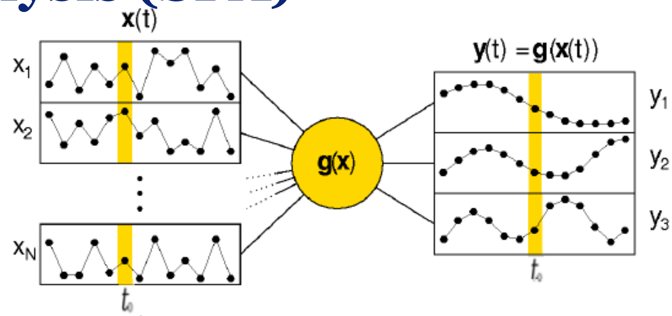

Slow Feature Analysis (SFA)

Generating the slowest varying output functions $y_j(t)$ from a multidimensional input signal x(t)

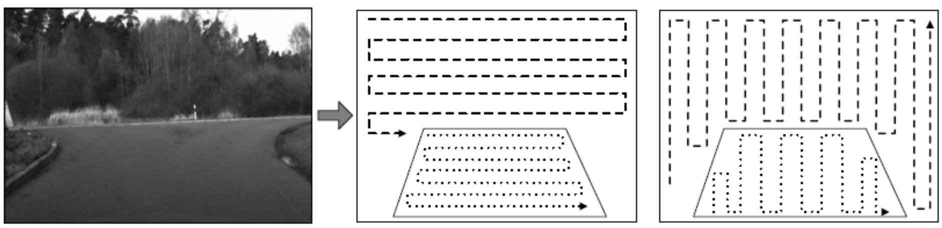

Spatial patch sequence extraction for SFA training: on the left the horizontal path and on the right the vertical path is illustrated. The paths are partitioned into road (dotted) and non-road (dashed) sections.

Result of the base road and base boundary classification for the given ground truth scene.

From left to right the positive and negative part of the confidence values is depicted for each base classifier.

Dark points denote high confidence of the classification.

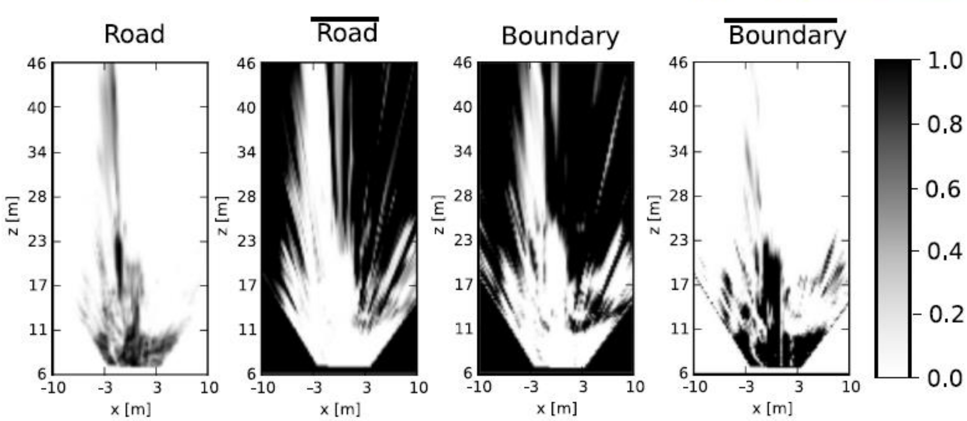

Feature vectors:

- SFA features
- Color features (RGB)
- Walsh Hadamard texture features

Ground Truth

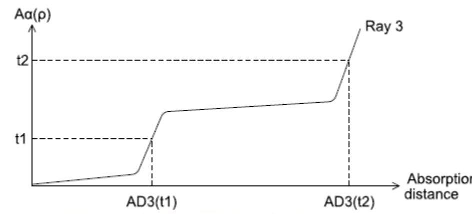

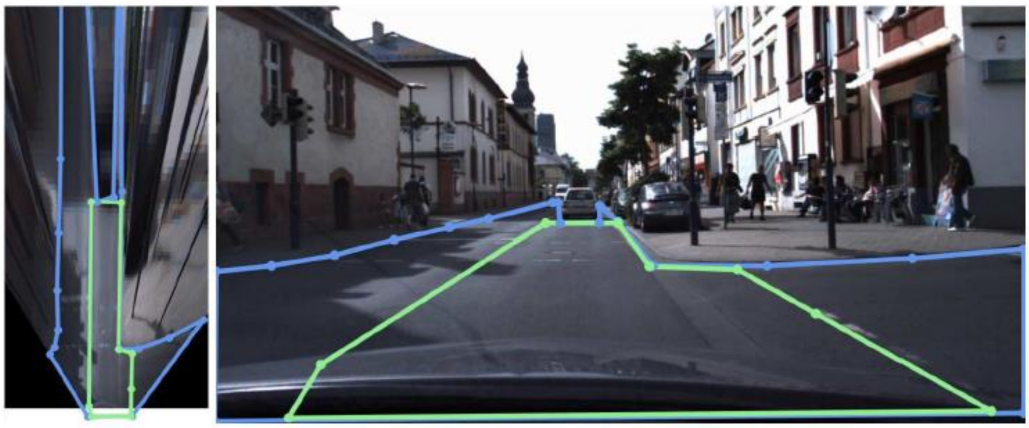

Results

Result of the road terrain classification for the ground truth (above):

- The classification result for road area (middle blue).
- The classification result for the ego-lane (right green)

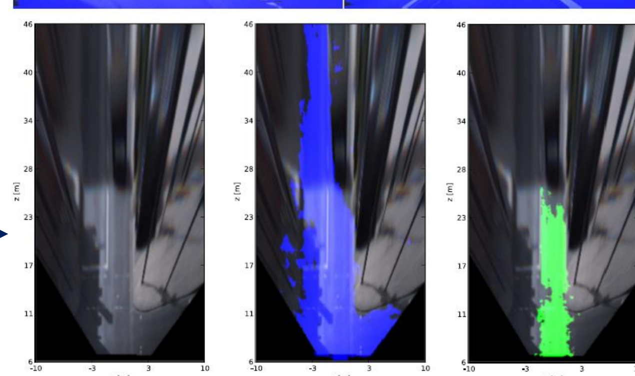

Results

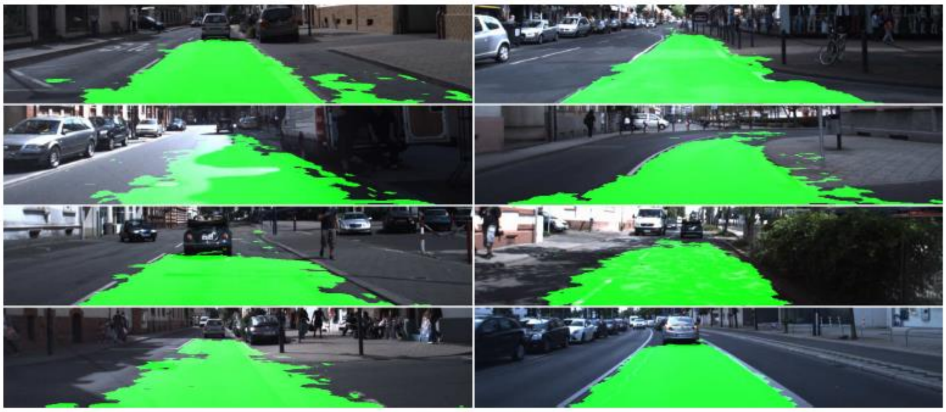

- Ego Lanes

Results

RESULTS OF PIXEL-BASED EVALUATION.

	AP	$F_{\max}$	Prec.	Recall	Acc	FPR	Q_{test}
perspective road area							
BL	89.1	85.6	79.4	92.8	78.9	50.4	74.8
SPRAY	95.6	94.5	94.0	95.0	92.5	12.8	89.5
metric road area							
BL	70.0	66.3	56.4	80.5	68.1	39.7	49.6
SPRAY	89.8	87.0	87.1	86.9	89.9	8.2	77.0
perspective ego-lane							
BL	80.1	81.7	76.4	87.7	90.2	9.0	69.1
SPRAY	85.2	87.6	84.7	90.6	93.6	5.4	77.9
metric ego-lane							
BL	61.7	60.3	56.6	64.6	92.5	4.8	43.2
SPRAY	78.9	79.5	79.6	79.4	96.4	2.0	66.0

(BL = Baseline)

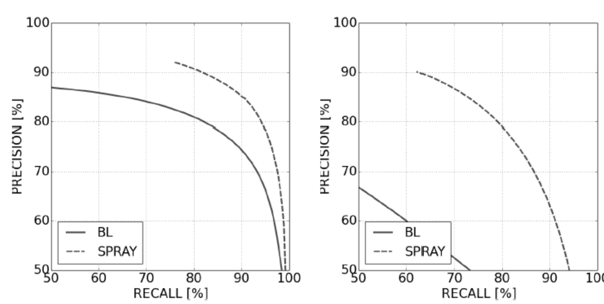

(a) Perspective evaluation (b) Metric evaluation

$$\text{Precision} = \frac{tp}{tp + fp}$$

$$\text{Accuracy} = \frac{tp + tn}{tp + tn + fp + fn}$$

$$\text{Recall} = \frac{tp}{tp + fn}$$

$$F = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

Harmonic mean of precision and recall.

Image Processing using OpenCV

Core module: the basic building blocks of this library for manipulating the images on a pixel level.

Imgproc module: the image processing (manipulation) functions inside OpenCV.

High Level GUI and Media (highgui module)

Image Input and Output (imgcodecs module)

Video Input and Output (videoio module)

Camera calibration and 3D reconstruction (calib3d module)

2D Features framework (feature2d module): feature points detectors, descriptors and matching framework found inside OpenCV.

Video analysis (video module) algorithms usable on your video streams like motion extraction, feature tracking and foreground extractions.

Object Detection (objdetect module) face detectors, etc.

Deep Neural Networks (dnn module)

Machine Learning (ml module) machine learning classes for statistical classification, regression and clustering of data.

Graph API (gapi module)

Computational photography (photo module) for advanced photo processing.

Images stitching (stitching module) create photo panoramas and more with OpenCV stitching pipeline.

GPU-Accelerated Computer Vision (cuda module); OpenCV iOS:

Lane Tracking

Some example project for detecting road features using OpenCV:

https://navoshta.com/detecting-road-features/ by Alex Staravoitau

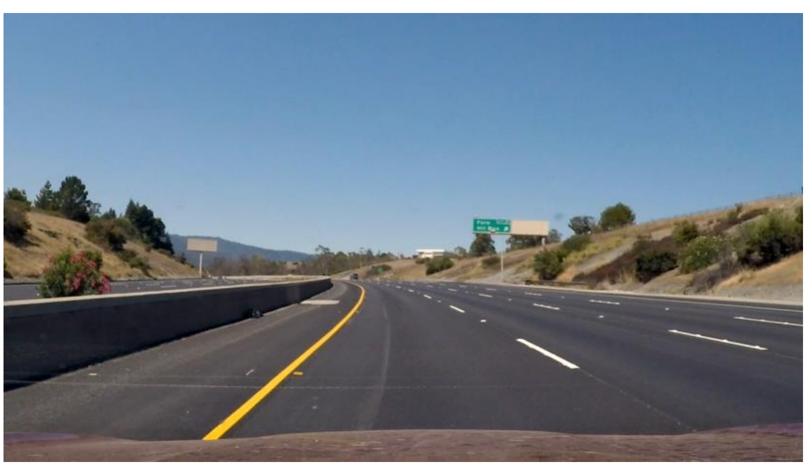

Universiteit Leiden. Bij ons leer je de wereld kennen

Overview Processing Pipeline

Camera calibration

- Each camera gives image distortions, these can be rectified using information from a camera calibration. OpenCV has functionality to calibrate and correct camera images. Calibration is done using chessboard images.

Edge detection

- OpenCV has many different edge detectors using gradient and color information. These edges can be used for the detection of structures such as lines etc.

Perspective transformation

- A perspective transformation, can be used to obtain an overview of the road ahead of the vehicle. This can make the problem of lane boundaries extraction easier.

Fitting boundary lines

- The resulting frame pixels are determined that may belong to lane boundaries.
- These are then used to approximate lines, road properties and vehicle position.
- Furthermore a rough estimate on road curvature and vehicle position within the lane is determined using known road dimensions.

Processing Pipeline: Camera Calibration

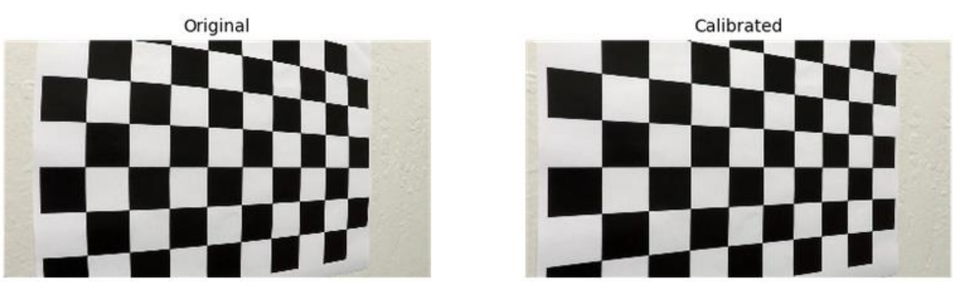

Original vs. calibrated images

... cv2.findChessboardCorners(image, (9, 6), None) // Inner corners 9x6

... cv2.calibrateCamera(pattern_points, image_points, (image.shape[1], image.shape[0]), None, None)

corrected_image = cv2.undistort(image, self.camera_matrix, self.dist_coefficients, None, self.camera_matrix)

Processing Pipeline: Edge Detection

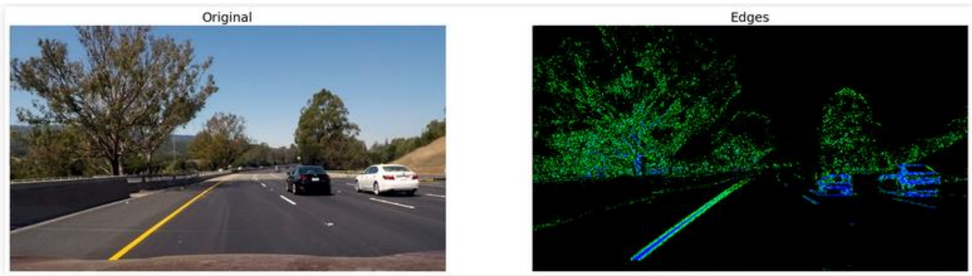

Original vs. highlighted edges

Gradient Absolute Values, Ranges within certain magnitudes, Gradient Directions

- Sobel Operator (using a convolutional Kernel)

Color Ranges

- HLS Color Space: Hue, Lightness, and Saturation (for road detection, etc.)

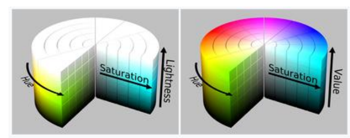

Fig. 2a. HSL cylinder. Fig. 2b. HSV cylinder.

Universiteit Leiden. Bij ons leer je de wereld kennen

Image Processing using Convolutional Kernel

$$g(x,y) = \omega * f(x,y) = \sum_{dx=-a}^{a} \sum_{dy=-b}^{b} \omega(dx,dy) f(x+dx, y+dy)$$

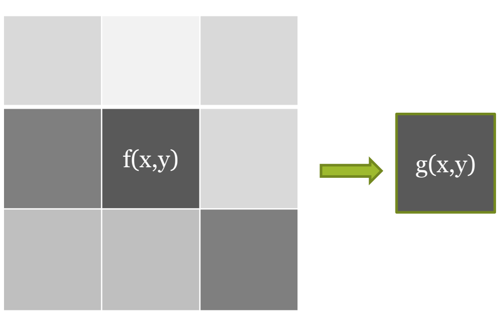

Operation	Kernel ω	Image result g(x,y)
Identity	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	
Edge detection	$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{bmatrix}$	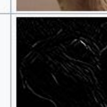
	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	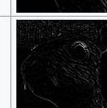
	$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$	
Sharpen	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	

https://en.wikipedia.org/wiki/Kernel_(image_processing)

Universiteit Leiden. Bij ons leer je de wereld kennen

Processing Pipeline: Perspective Transformation

Original vs. bird's eye view

.... transform_matrix = cv2.getPerspectiveTransform(source, destination)

.... image = cv2.warpPerspective(image, transform_matrix, (w, h))

Processing Pipeline: Perspective Transformation

Boundary detection pipeline

Left: The *original* image after the camera calibration and perspective transform.

Right: After edge detection with edges highlighted in **green** and **blue**.

Scanning windows boundaries for areas with pixel that may belong to lines are highlighted in **yellow**,

A second order polynomial approximation of the collected points in **red**.

Lane and Vehicle Tracking

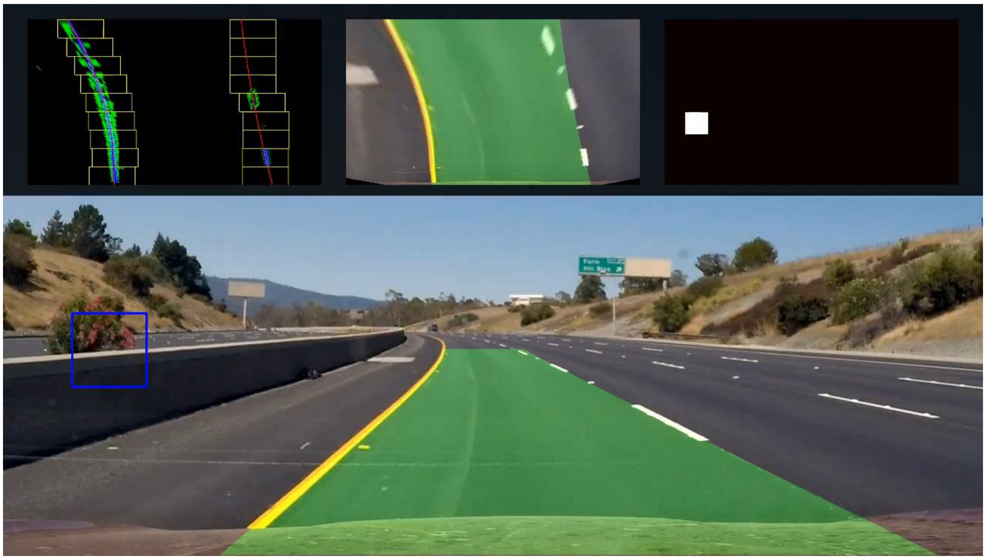

Some remarks

Alex Staravoitau:

“This clearly is a very naive way of detecting and tracking road features, and wouldn’t be used in real world application as-is, since it is likely to fail in too many scenarios: “

- Going up or down the hill.
- Changing weather conditions.
- Worn out lane markings.
- Obstruction by other vehicles or vehicles obstructing each other.
- Vehicles and vehicle positions different from those classifier was trained on.
- …

Organization and Overview

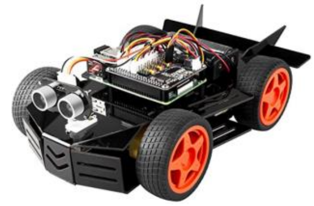

Period: February 7th – May 23rd 2022
Time: Monday 16.15 – 18.00
Place: Room 407 - 409
Lecturer: Erwin M. Bakker (erwin@liacs.nl)
Assistant: Hainan Yu (h.yu@liacs.leidenuniv.nl)

NB Register on Brightspace

Schedule:

7-2	Introduction and Overview
14-2	Locomotion and Inverse Kinematics
21-2	Robotics Sensors and Image Processing
28-2	SLAM + SLAM Workshop
7-3	**Mobile Robot Challenge Introduction**
14-3	**Project Proposals I (presentation by students)**
21-3	**Project Proposals II (presentation by students)**
28-3	Robotics Vision
4-4	Robotics Reinforcement Learning
11-4	Robotics Reinforcement Learning Workshop II
18-4	No Class (Eastern)
25-4	Project Progress I (presentations by students)
2-5	Project Progress II (presentations by students)
9-5	**Mobile Robot Challenge**
16-5	**Project Demos I**
23-5	**Project Demos II**

Website: http://liacs.leidenuniv.nl/~bakkerem2/robotics/

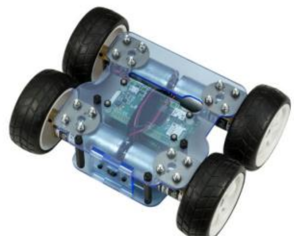

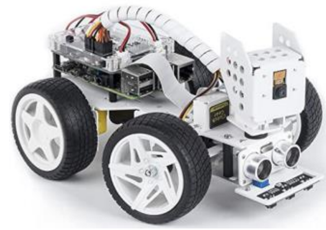

Grading (6 ECTS):
- Presentations and Robotics Project (60% of grade).
- Class discussions, attendance, workshops and assignments (40% of grade).
- It is necessary to be at every class and to complete every workshop and assignment.

Universiteit Leiden. Bij ons leer je de wereld kennen

Robotics Project Proposals Presentations

Monday 14-3 2022

Present your Robotics Project Proposal during a **5 minute (max)** talk. Clearly state the title of your project, the team members, your goals, how you will pursue them, what are the challenges and what at least can and should be delivered on the demo day on May 16th and May 23rd 2022.

Note: Groups of 1-5 members are allowed.

The presentation should contain slides for:
1. Title and group members.
2. Goal of the project.
3. How will you pursue these goals.
4. What are the challenges.
5. What at least can and should be delivered on the demo days on May 16th and May 23rd 2022.

The LIACS Media Lab can support your project with some materials for your project. Please clearly state any materials that you would need for your proposal. Note that these materials are limited so project goals may need to be adjusted accordingly.

Each presentation will be followed by a short class discussion.

Universiteit Leiden. Bij ons leer je de wereld kennen

Previous Projects

1. Evolutionary Locomotion
2. Nao plays Tic-Tac-Toe
3. Slam Robot Project.
4. Dolphin Drone: Drone Recognition and Maneuvering with Hoops.
5. Delivery Drone.
6. Programming a NAO to play a tune using a xylophone.
7. Floor mapping with Swarm Robotics
8. Tootballing Yetiborg
9. Cat Flap Opening Based on Audio/Video/RFID
10. DrawBot
11. Traffic coordination (simulation).
12. Plane filling curves (simulation).

1. AimBot
2. Artificial Muscles
3. Ball Tracking Car
4. BorrelBot
5. Fetch Bot
6. Floor Mapping Robot
7. Gesture Control Pachenko
8. Hexapod
9. Nao Pose
10. Position Estimation
11. Race Car Training
12. Face Touch

Traffic coordination (simulation).

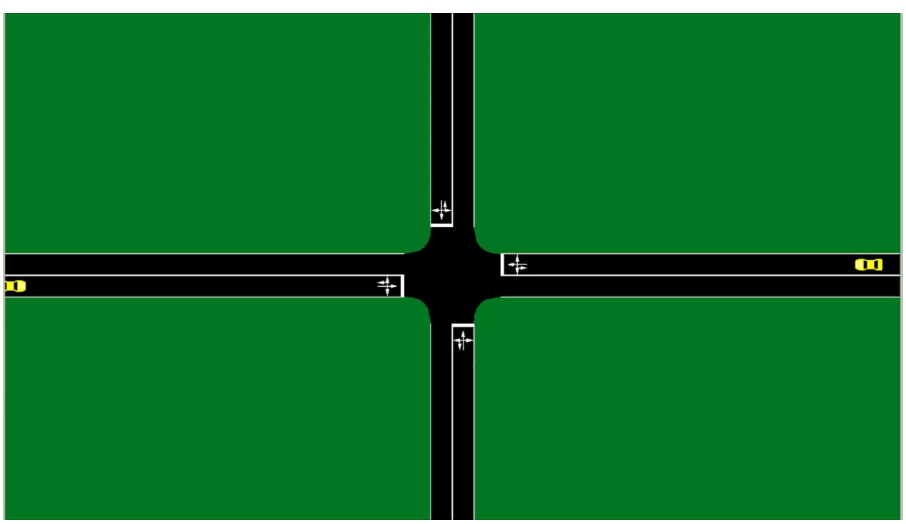

Roombas with Brains: Neural Networks for Coverage Path Planning

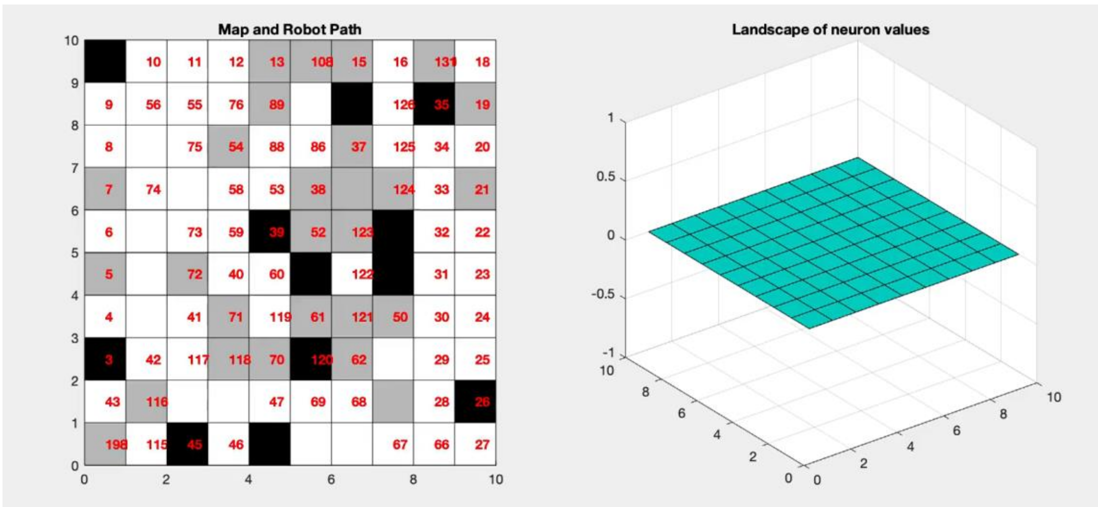

Gesture Controlled Pachenko

Gesture Controlled Pachinko Game

Aaron Dunlea, Nathan van der Putten, Malte Wilhelm

liacs Leiden Institute of Advanced Computer Science

Universiteit Leiden The Netherlands

Traffic coordination (simulation).

DrawBot
AimBot
FetchBot

Cellular Automaton: https://youtu.be/maC1eo8_-lI
Dragon Curve: https://youtu.be/9ilDP_pvDEk

https://vimeo.com/425043774

S.P.I.N. - Spider Python INator

Marcel Huijben (s1780107)
Martijn Swenne (s1923889)
Sebastiaan Alvarez Rodriguez (s1810979)
Robin Voetter (s1835130)

References

1. Joel C. McCall and Mohan M. Trivedi, Video Based Lane Estimation and Tracking for Driver Assistance: Survey, System, and Evaluation. IEEE Transactions on Intelligent Transportation Systems, 2006
2. A. Bar Hillel, R. Lerner, D. Levi, G. Raz, Recent progress in road and lane detection: a survey. Machine Vision and Applications (2014) 25:727–745
3. J. Fritsch, T. Kühnl, F. Kummert, Monocular Road Terrain Detection by Combining Visual and Spatial Information. IEEE Transactions on Intelligent Transportation Systems, 2014.
4. J. Sattar, J. Mo, SafeDrive: A Robust Lane Tracking System for Autonomous and Assisted Driving Under Limited Visibility. January 31, 2017 (https://arxiv.org/abs/1701.08449)
5. https://navoshta.com/detecting-road-features/ by Alex Staravoitau
6. OpenCV.org

Robotics

Universiteit Leiden

Robotics Discussion Session

Wednesday 24-2 at 15.15
Robotics Kaltura Room

During this session we discuss some practical aspects of robotics in an informal and interactive setting.

Especially for people who did not work with microcontrollers, servo's, sensors etc. before.

Team 1
Joost Mollen
Marianne Bossema
Martijn Wester
Jordy van Miltenburg

Team 2
Yannick Ligthart
Aaron Kannangara
Renzo Baasdam
Orson Peters
Mihai Ghidoveanu

Team 3
Rachel Losacco
Ioannis Politopoulos
Stella Tsilia
Alex Tripsas
Micha Heilman

Team 4
Abhishek Sira Chandrashekar
Ankita Parashar
Hassan Ibrahim
Thomas Gmelig Meyling

Team 5
Jesse Jonathan ('Nathan') van der Putten
Aaron Dunlea
Malte Wilhelm
Sophie van der Bliek

Team 6
Liyang He
Yuxin Xiong
Hansha Leng
Jiakun Sun
Jincheng Li

Team 7
Robin Voetter
Sebastiaan Alvarez-Rodriguez
Martijn Swenne
Marcel Huijben

Team 8
Rick Boks
Rens Dofferhoff
Levi Vos

Team 9
Elgar van der Zande
Luc de Jonckheere

SPRAY Features: Spatial Rays

1) 2) 3)

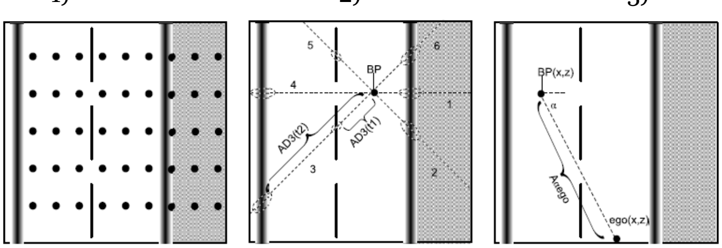

1. Distribution of base points in metric space
2. The SPRAY feature generation procedure for one base point (BP)
3. The ego SPRAY feature.

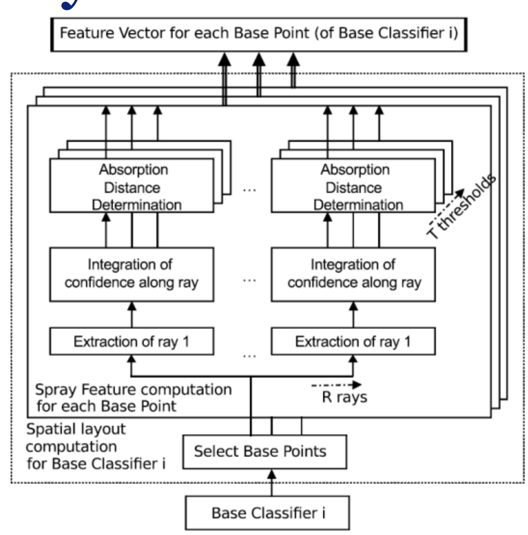

Universiteit Leiden. Bij ons leer je de wereld kennen

J. Sattar, J. Mo, SafeDrive: A Robust Lane Tracking System for Autonomous and Assisted Driving Under Limited Visibility. January 31, 2017 (https://arxiv.org/abs/1701.08449)

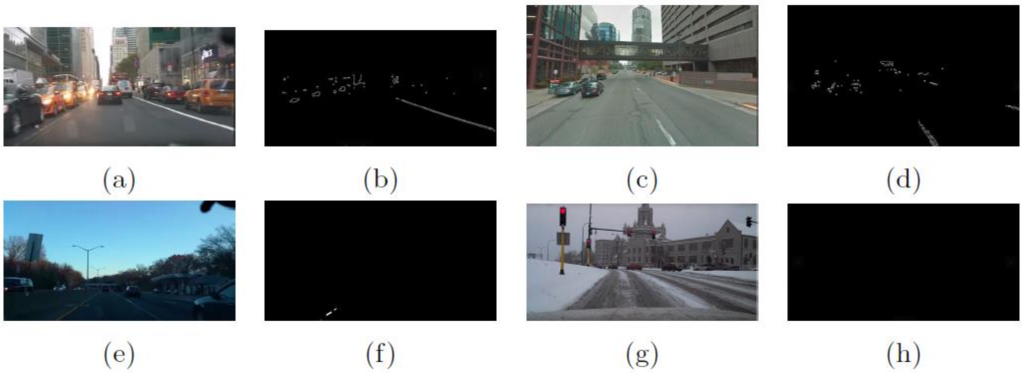

Visual lane tracking on several urban scenes from YouTubeTM videos. Snapshot (1a) (output in (1b)): lane markers not distinct in the center, though side markers are detectable. Snapshot (1c) (output in (1d)): lane markers mostly washed out. Snapshot (1e) (output in (1f)): evening drive, low-light conditions make the lane markers almost undetectable. Snapshot (1g) (output in (1h)): snow-covered roads, no lane markers visible.

System Overview

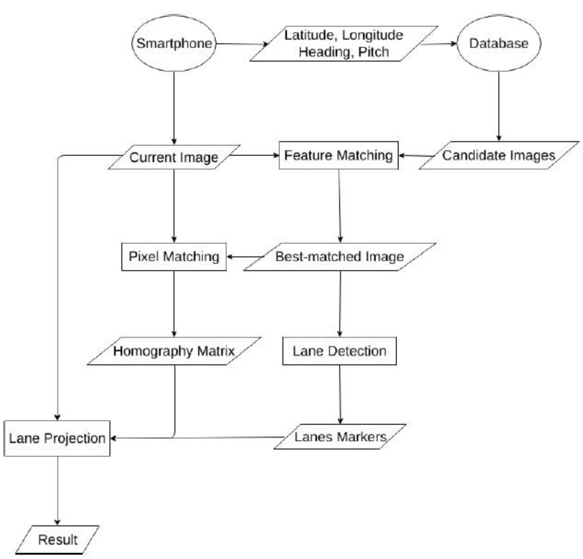

The process of extracting pixels with common" visual content.

The feature-based matching (in red lines) are used to choose the point features, and for each feature point, a square sub-window is extracted from the candidate image, centered on that feature point.

Stitching together all these windows results in an image with most uncommon" visual elements removed.

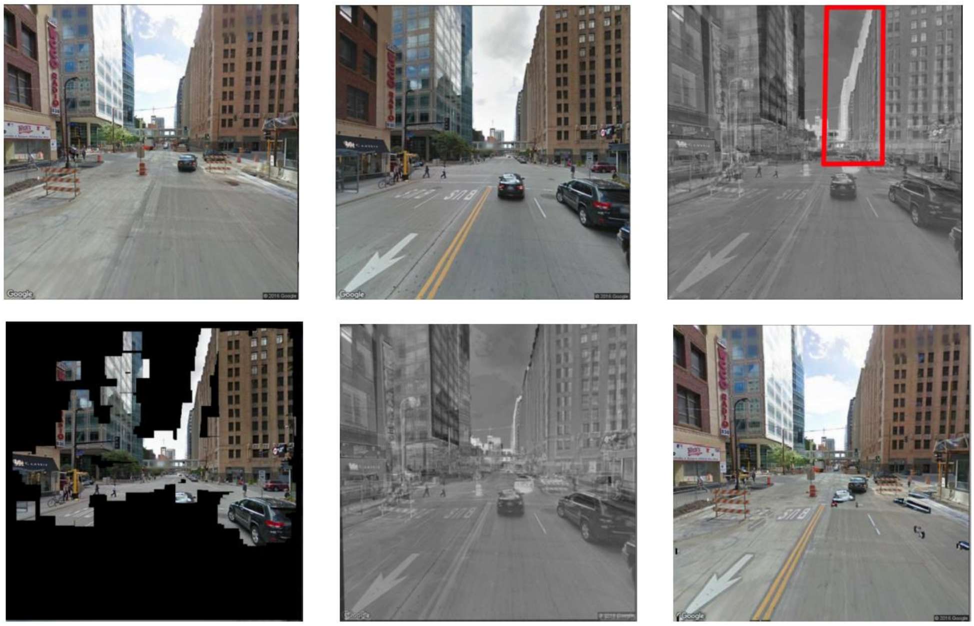

Road Detection

1. H. Kong, et al. General road detection from a single image, IEEE Transactions on Image Processing, Vol. 19, Issue 8, Aug. 2010.
2. J.M. Alvarez, et al., Road Detection Based on Illuminant Invariance, IEEE Transactions on Intelligent Transportation Systems, 2010.

Universiteit Leiden. Bij ons leer je de wereld kennen

Road Detection Based on Illuminant Invariance

Goal: Identify all pixels that belong to the road.

Challenges:

- Images from a mobile platform, outdoor: shadows, sunlight, etc.
- Changing background: trees, buildings and many (moving) objects: cars, bikes, pedestrians, dogs, signs, etc.
- Road types in different shapes and forms.

Sensors:

- Monocular => Features: Color, Texture, etc.

J.M. Alvarez, et al. 2010

Universiteit Leiden. Bij ons leer je de wereld kennen

Road Detection Based on Illuminant Invariance

Challenges: Shadows, sunlight, cloudy, etc.

Camera calibration using Macbeth color checker, can not be used for onboard self calibration and image content.

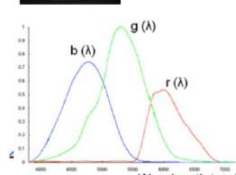

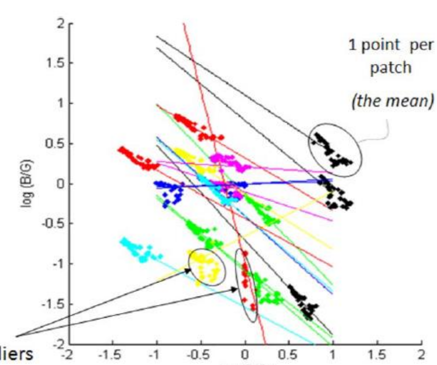

Universiteit Leiden. Bij ons leer je de wereld kennen

Road Detection Based on Illuminant Invariance

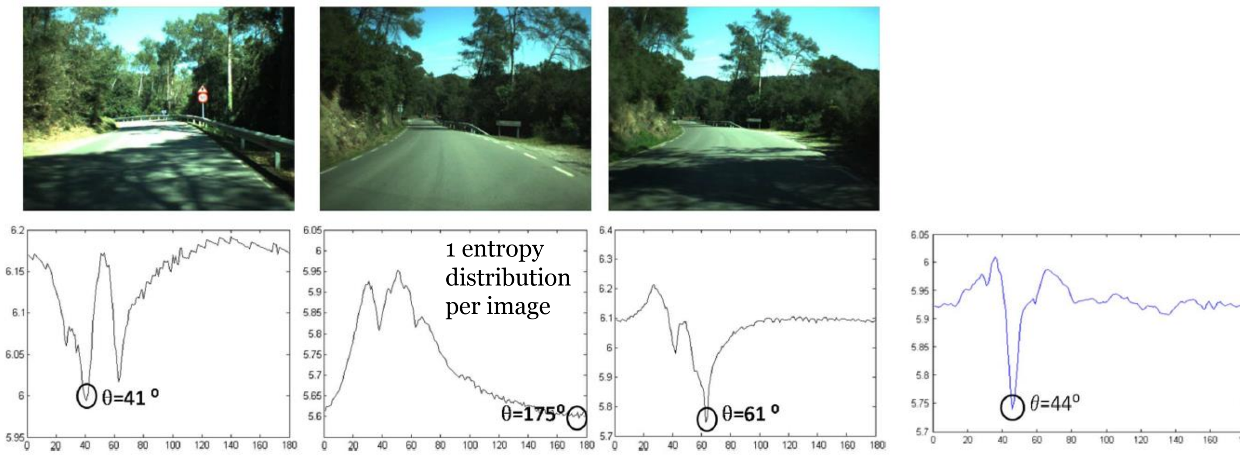

Entropy based camera calibration: can be used for onboard self calibration and image content.

Θ is the absolute minimum of the average distribution of entropy values for $\text{Image}_{\text{RGB}}$ (minimization procedure)

Universiteit Leiden. Bij ons leer je de wereld kennen

Road Detection Based on Illuminant Invariance

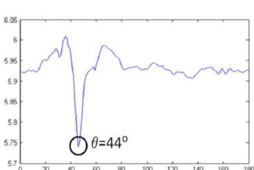

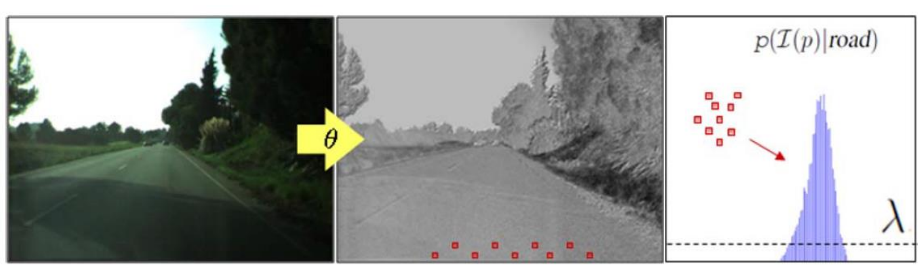

Fig. 7. (Left) Color image I_{RGB}. (Middle) Corresponding $\mathcal{I}$ computed using the invariant angle θ. (Right) Nonparametric road model built using the normalized histogram formed with the surrounding region of several seeds (nine in this case) placed at the bottom part of $\mathcal{I}$.

Entropy based camera calibration: can be used for onboard self calibration and image content.

Θ is the absolute minimum of the average distribution of entropy values for $\text{Image}_{\text{RGB}}$ (minimization procedure)

Universiteit Leiden. Bij ons leer je de wereld kennen

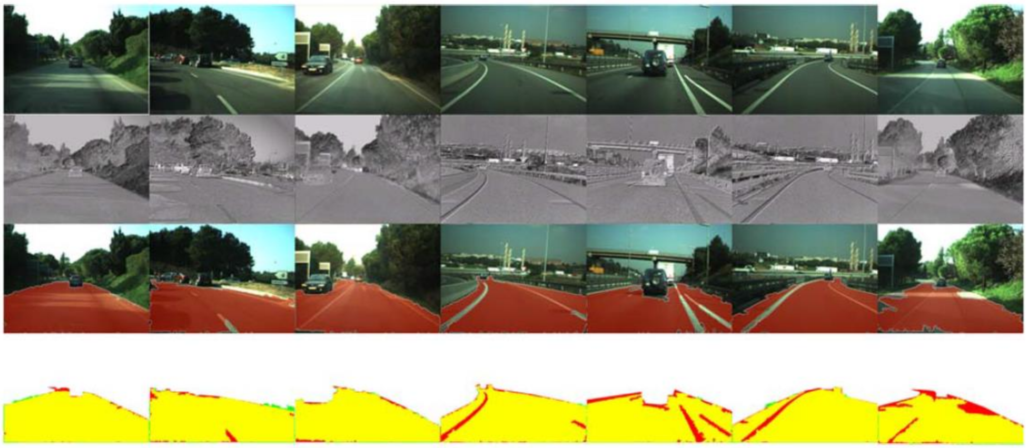

Fig. 11. Road-detection examples. (Top row) Original image. (Second row) Illuminant-invariant image. (Third row) Detected road. (Bottom row) Comparison against hand-segmented result. (Yellow) Correctly classified pixels. (Green) Falsely detected road pixels. (Red) False background pixels.

Fig. 12. Road-detection results. Example results include the following: (First column) Images from a sunny day with annoying shadows and traffic; (second column) images from a wet road without puddles with dim shadows and traffic; (third column) different scenarios; (fourth column) different road shapes; (last column) presence of other vehicles.

Universiteit Leiden. Bij ons leer je de wereld kennen

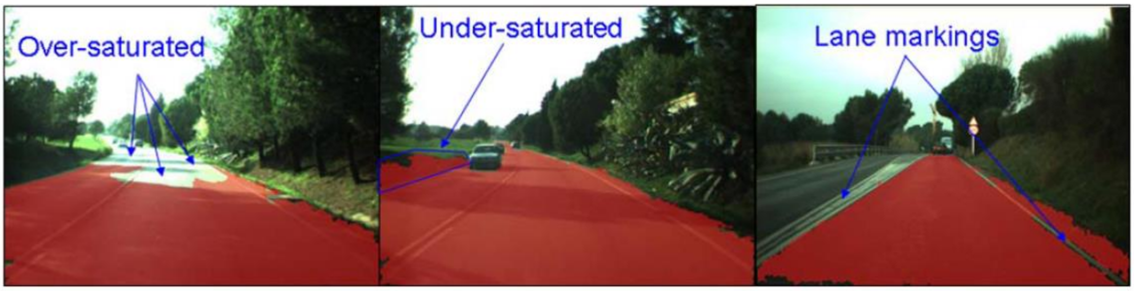

Fig. 13. Examples where our road-detection method currently fails.

General road detection from a single image

H. Kong et al. 2010

Road in different shapes and forms under different lighting with changes in colors, and textures.

General road detection from a single image
H. Kong et al. 2010

Road in different shapes and forms under different lighting with changes in colors, and textures.

Detection

1. Vanishing point estimation of main road
2. Segmentation based on this vanishing point

Voting scheme:

- High-confidence voters using texture orientations resulting from Gabor filters (5 scales 36 orientations)
- Vanishing-point-constrained edge detection for road boundaries

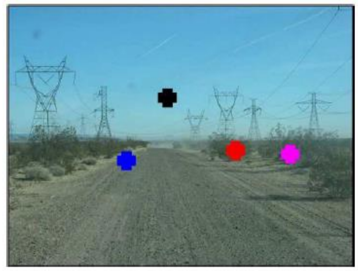

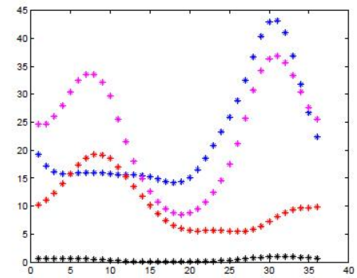

Gabor complex responses for 4 points.

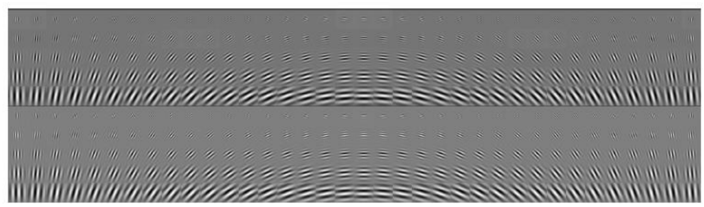

Gabor kernels: real part (row 1-5), imaginary part (row 6-10)

General road detection from a single image
H. Kong et al. 2010

Road in different shapes and forms under different lighting with changes in colors, and textures.

Detection

1. Vanishing point estimation of main road
2. Segmentation based on this vanishing point

Voting scheme:

- High-confidence voters using texture orientations resulting from Gabor filters (5 scales 36 orientations)
- Vanishing-point-constrained edge detection for road boundaries

Line segments

vanishing point with rays

orientation consistent rays

compute color difference of neighboring regions

dominant border (red) => vanishing point adaptation

General road detection from a single image

H. Kong et al. 2010

Result

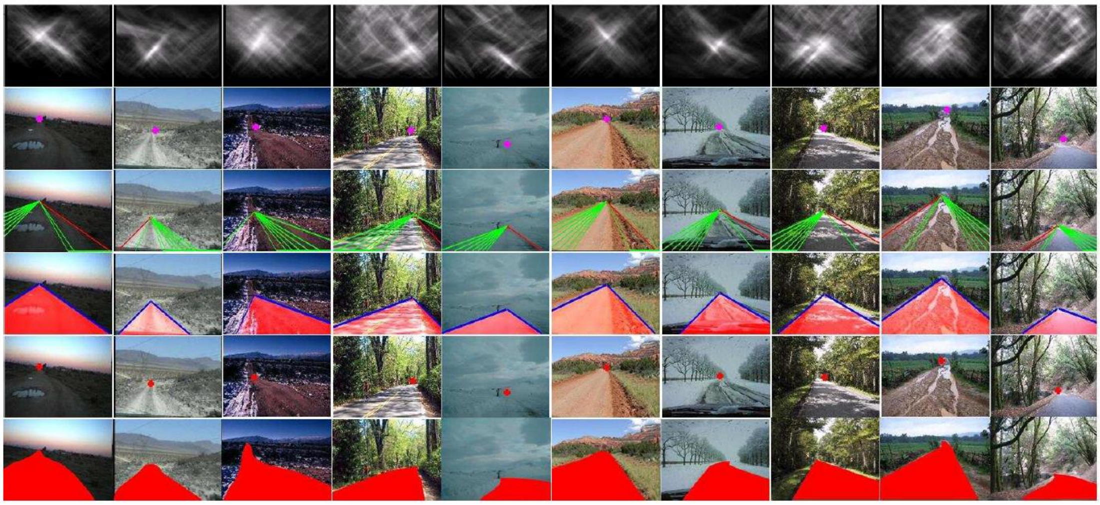

Last row: ground truth..