Robotics

Erwin M. Bakker| LIACS Media Lab

Universiteit Leiden

Bij ons leer je de wereld kennen

Organization and Overview

Lecturer:
Dr Erwin M. Bakker (erwin@liacs.nl)
Room 126a and LIACS Media Lab (LML)

Teaching assistant:
Mor Puigventos (email)
TBA (email)

Period: February 6th - May 22nd 2023
Time: Monday 15.15 - 17.00
Place (Rooms):
a) Gorlaeus - Lecture Hall C3
b) Sylvius - 15.31
c) Van Steenis - E0.04
d) Oortgebouw - Sitterzaal

Schedule (tentative, visit regularly):

Date	Room	Subject
6-2	a	Introduction and Overview
13-2	a	Locomotion and Inverse Kinematics
20-2	b	Robotics Sensors and Image Processing
27-2	a	SLAM + SLAM Workshop
6-3	c	Mobile Robot Challenge Introduction
13-3	a	Project Proposals I (by students)
20-3	d	Project Proposals II (by students)
27-3	d	Robotics Vision (Week 13, start 15.30)
3-4	d	Robotics Reinforcement Learning&Workshop
10-4		No Class (Eastern)
17-4	d	Project Progress I (by students)
24-4	d	Project Progress II (by students)
1-5	d	Mobile Robot Challenge I
8-5	a	Mobile Robot Challenge II
15-5	d	Project Demos I
22-5	d	Project Demos II
29-5		Whit Monday
5-6		Project Deliverables

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

Grading (6 ECTS):
- Presentations and Robotics Project (60% of grade).
- Class discussions, attendance, assignments (pass/no pass) 2 workshops (0-10) (2 x 20% = 40% of grade).
- It is necessary to be at every class and to complete every workshop and assignment.

Overview

- Sensors
- Lane Tracking
- OpenCV
- Line/Lane Tracking

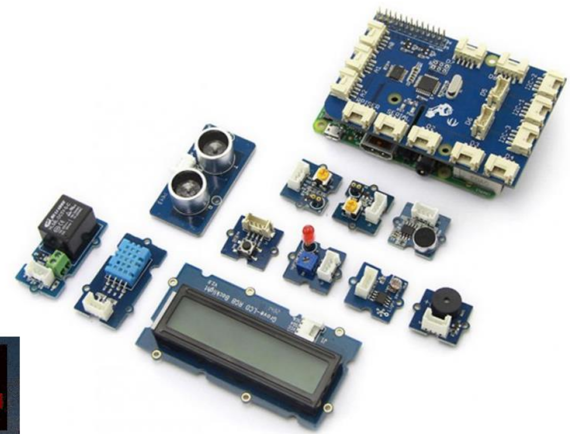

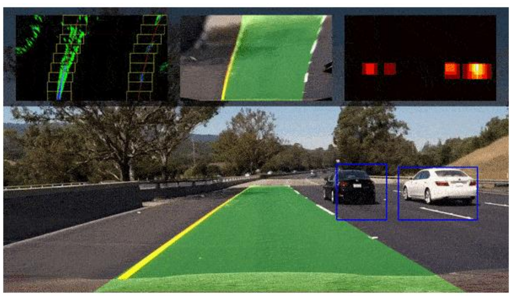

- Project Groups and Proposals
- Groups Mobile Robot Challenge

Universiteit Leiden. Bij ons leer je de wereld kennen

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	✗	✗	✗	✓

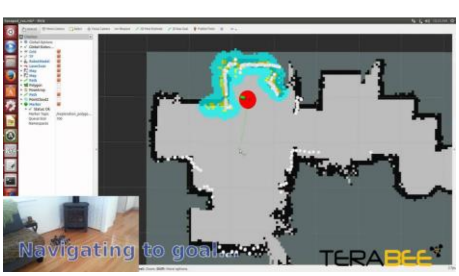

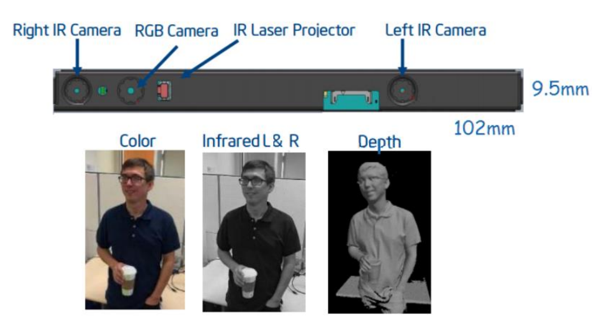

Universiteit Leiden. Bij ons leer je de wereld kennen

IROBOT CREATE 3

Sensors
1x Power button 2x User buttons
2x Front bumper zones
2x Wheel encoders
4x IR cliff sensors
7x IR obstacle sensors
1x Downward optical flow sensor for odometry
1x 3D gyroscope
1x 3D accelerometer
1x Battery level monitor

Technical Details
- 14.29 cm x 40.64 cm x 47.31 cm, 4.68 kg

Connectivity
- Wi-Fi or Ethernet over USB Host: Connect and control the robot using ROS 2.
- Bluetooth® Low Energy: Connect to and control the robot using iRobot Education's Python Web Playground.

Battery
- 1800 mAh Li-Ion rechargeable battery

Actuators
- 2x Drive motors
- 6x RGB LED ring
- 1x Speaker

Charging and Expansion
•1x Docking port
•1x USB-C™ port, 3A at regulated 5V
•1x Payload power, 2A at unregulated (nominal 14.4V)

Structured Light

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

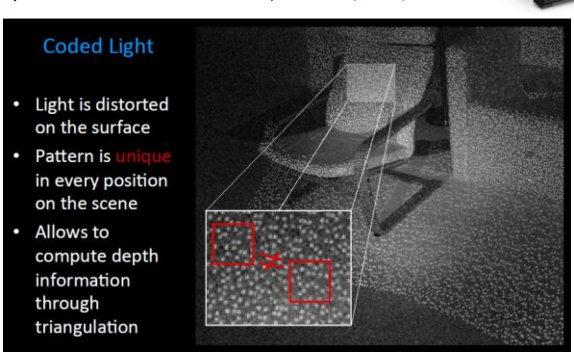

From: Anyline presentation by Peter Sperl

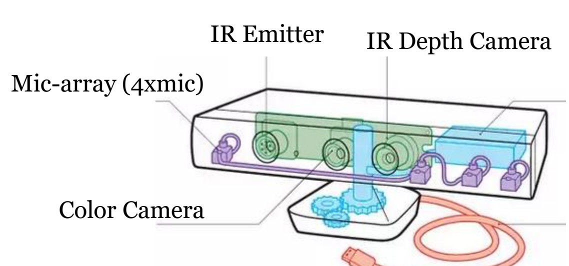

LIDAR Explanation

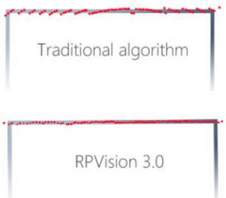

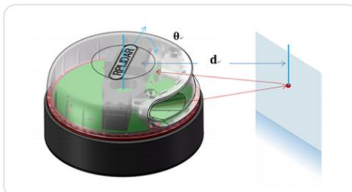

http://www.slamtec.com/en/lidar/A3

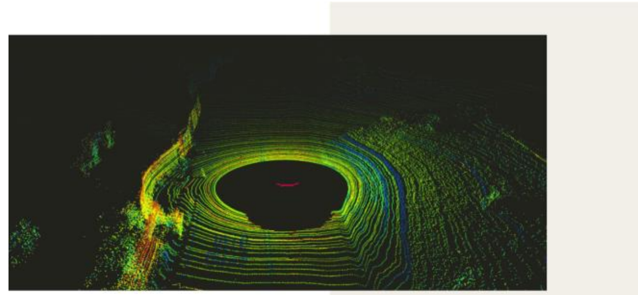

https://news.voyage.auto/an-introduction-to-lidar-the-key-self-driving-car-sensor-a7e405590cff

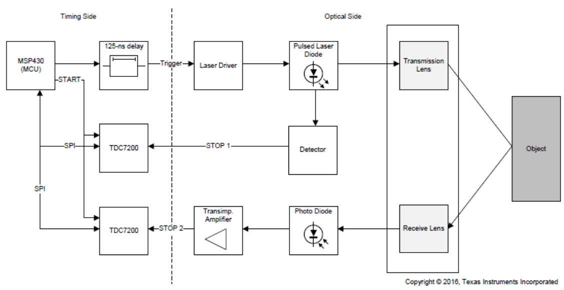

Texas Instruments LIDAR Pulsed Reference Design

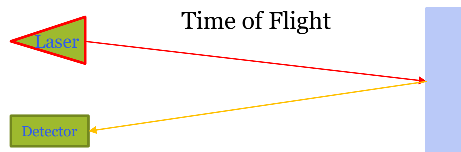

- Speed of light ~ 3 x 10^8 m/s
- In 1 picosecond (= 10^{-12} sec) light travels ~ 3 x 10^{-4} m = 0.3 mm
- During 33 picoseconds light travels ~ 1cm

VLP-16, 16 line Lidar

Robotics Self Driving Cars

Google Firefly, 2015

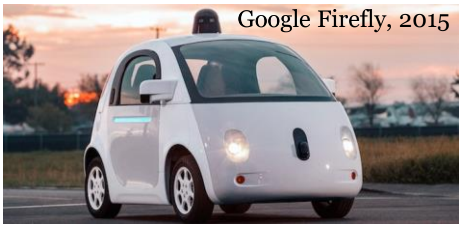

Google, 2009

Waymo OEM, 2021 – Waymo Driver

Ended in Dec. 2020 -> Aurora Innovation: Aurora Driver

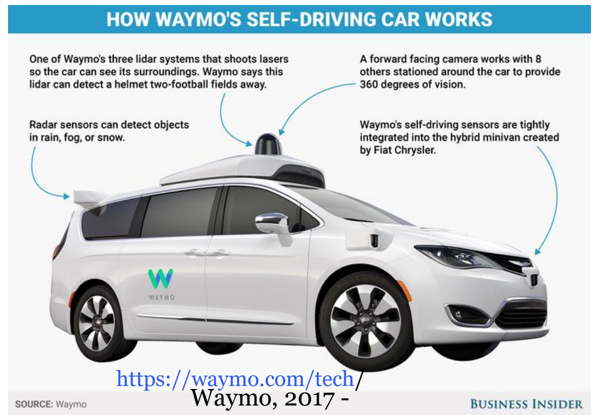

https://waymo.com/tech/
Waymo, 2017 -

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**

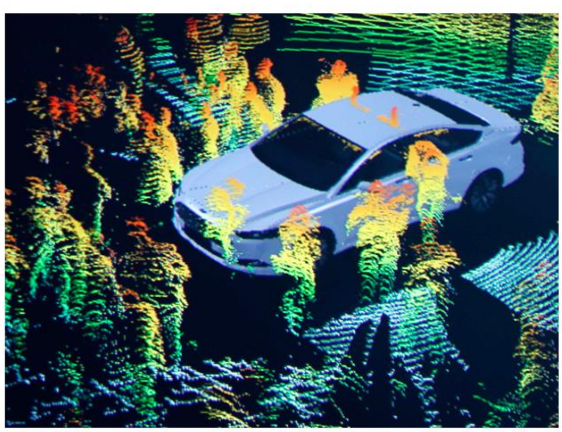

Photo: Jeff Kowalsky/Corbis

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

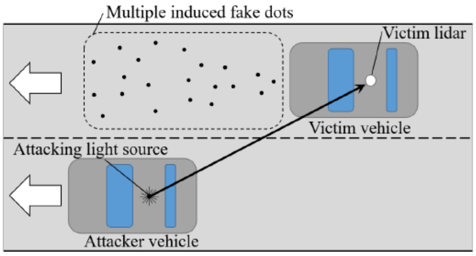

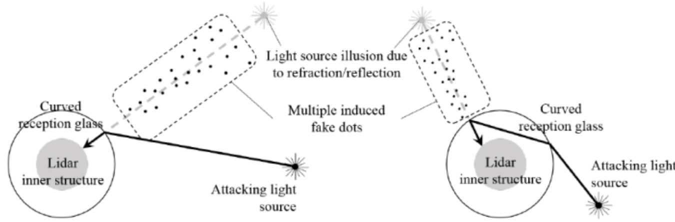

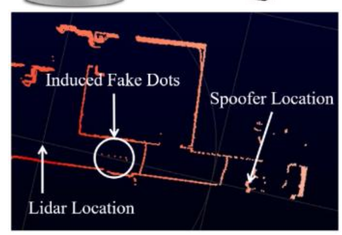

Adversarial Sensor Attack on LiDAR-based Perception in Autonomous Driving

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

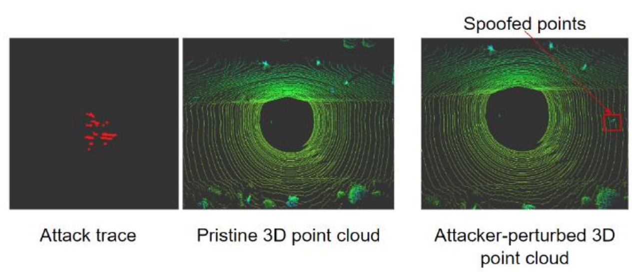

Figure 5: Generating the attacker-perturbed 3D point cloud by synthesizing the pristine 3D point cloud with the attack trace to spoof a front-near obstacle 5 meters away from the victim AV.

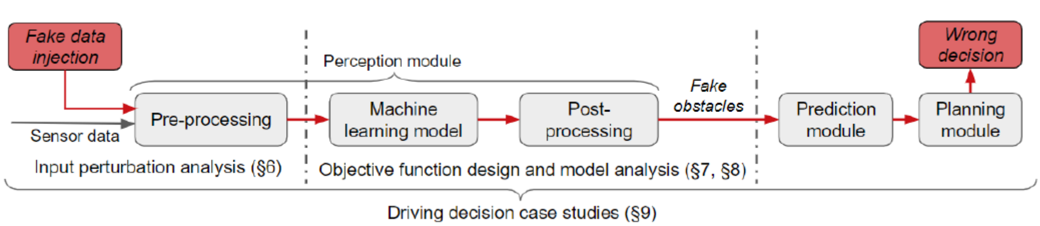

Universiteit Leiden. Bij ons leer je de wereld kennen

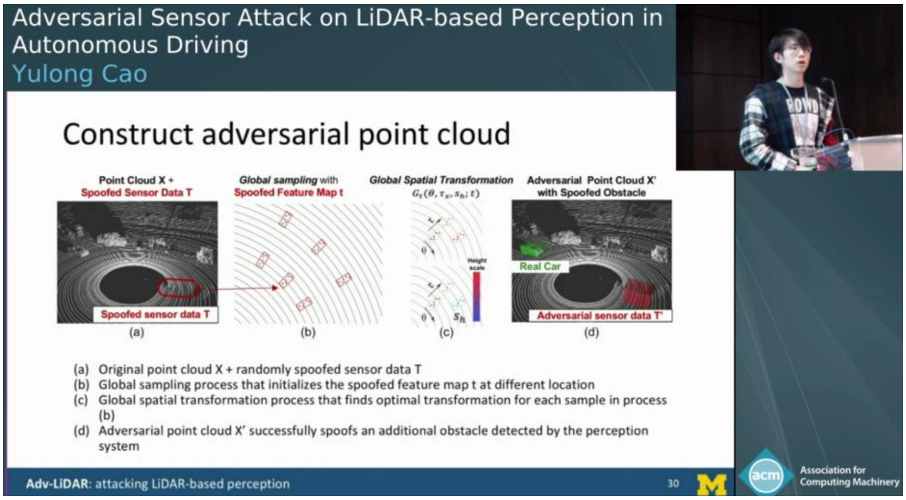

Many more studies since, also attacks against multimodal systems, e.g.,
Y. Chao et al., Invisible for both Camera and LiDAR: Security of Multi-Sensor Fusion based Perception in Autonomous Driving Under Physical-World Attacks, June 2021

RaspberryPi Sensors Kit

GrovePi+ Board for Raspberry Pi

De ATMEGA328 microcontroller communicates with the Raspberry Pi.

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

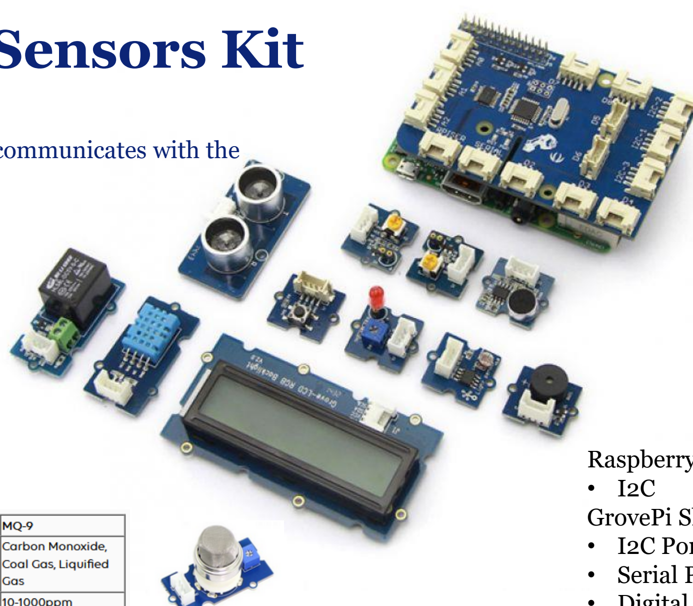

Symbol	MQ-2	MQ-3	MQ-5	MQ-9
Detect Gas	Combustible Gas, Smoke	Alcohol Vapor	LPG, Natural Gas, Town Gas	Carbon Monoxide, Coal Gas, Liquified Gas
Detect Concentration	300-10000ppm	0.04-4mg/L Alcohol	300-10000ppm	10-1000ppm CO;100-10000PPm Gas

Raspberry Pi -> GrovePi
- I2C

GrovePi Shield:
- I2C Ports
- Serial Ports
- Digital Ports
- Analog Ports

Lane Tracking

Autonomous Driving &
Driver Assist Systems

- Steering Angle Prediction
- Lane Detection
- Lane Tracking

Noisy and partial observations =>

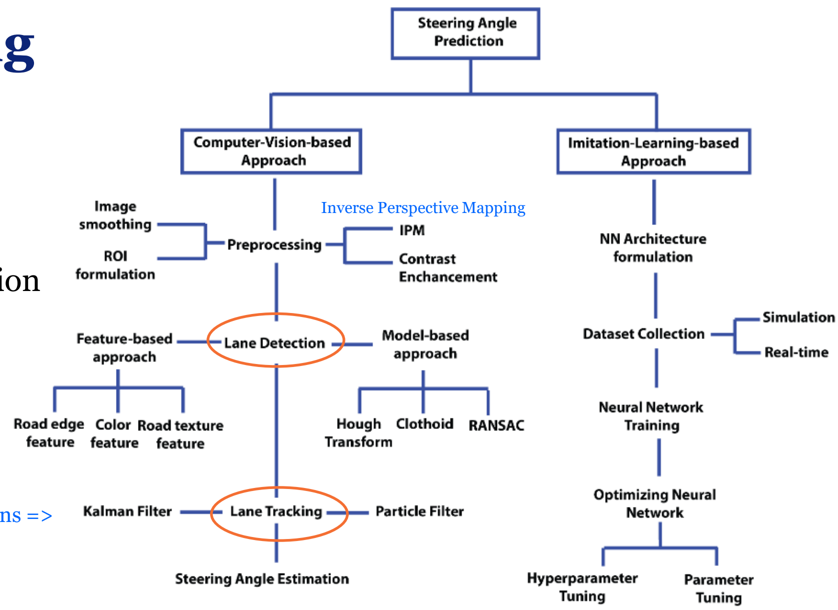

H. Saleem et al., Steering Angle Prediction Techniques for Autonomous Ground Vehicles: A Review. IEEE Access, June 2021

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 (1362 citations)
- 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 (879 citations)
- 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

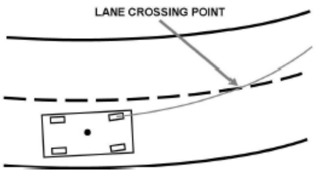

(a) lane departure warning

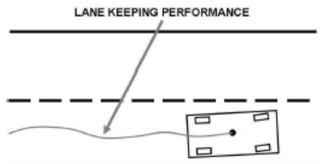

(b) driver attention monitoring

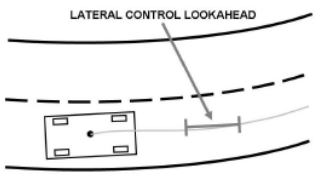

(c) vehicle control


(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

Universiteit Leiden. Bij ons leer je de wereld kennen

Lane Tracking by Day and Night

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

Universiteit Leiden. Bij ons leer je de wereld kennen

General System

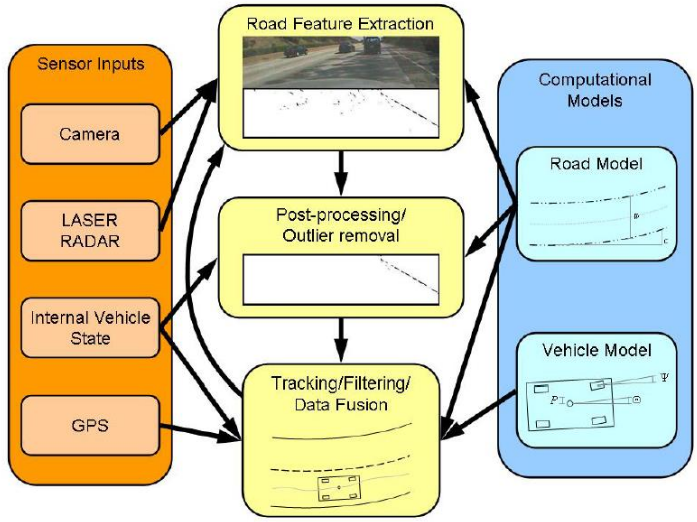

1. camera and vision sensors
2. internal vehicle state sensors
   e.g., speed, steering angle, etc.
3. line detectors, etc.
4. LASER RADAR sensors
5. GPS sensors

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

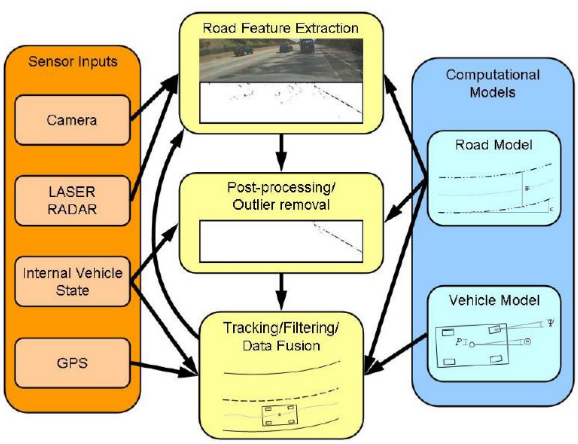

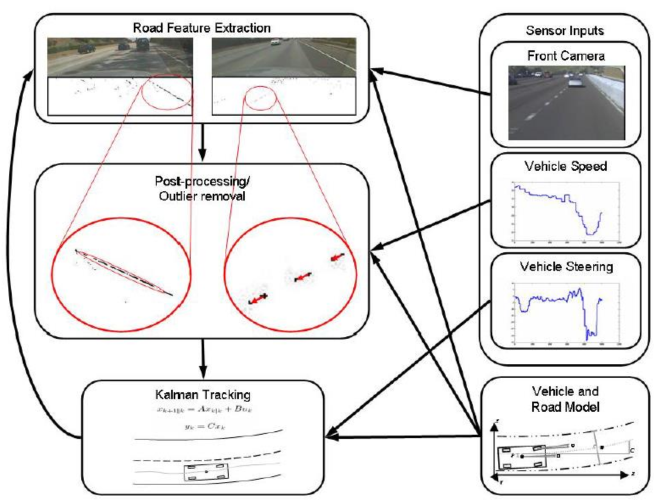

A generalized flow chart for lane position detection systems combining multiple modalities an iterative detection/tracking loop and road and vehicle models.

Vehicle and Road Models

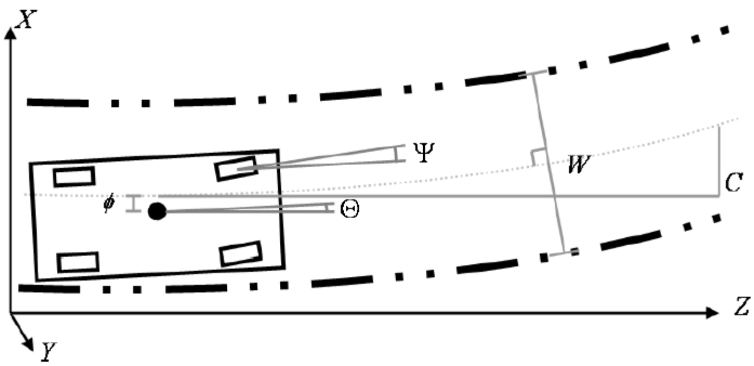

- A constant curvature road model
- Linearized vehicle dynamics
- Kalman Filter Applied

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.

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	

Steerable Filters

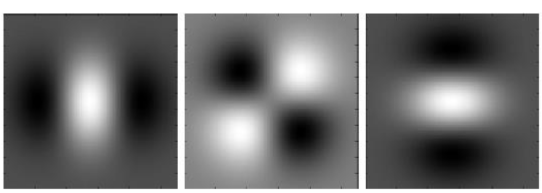

Fig. 7. A basis set for steerable filters based on the second derivatives of a two-dimensional Gaussian.

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

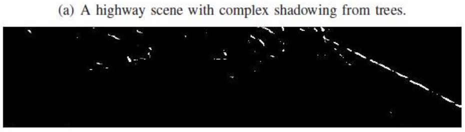

(b) Detection results for lines tuned to the lane angle.

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.

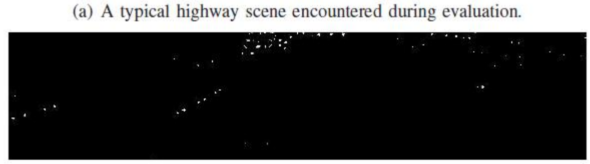

(b) Results of filtering for circular reflectors.

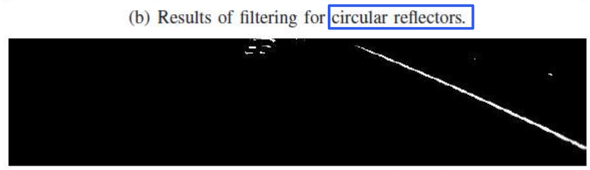

(c) Results from filter for a line tuned to the lane angle.

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

Inverse Perspective Warping (IPM) and Template Matching

IPM: 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 or matching easier.

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

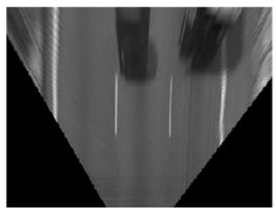

Inverse Perspective Warping

(a) Detected lanes with curvature overlaid onto image

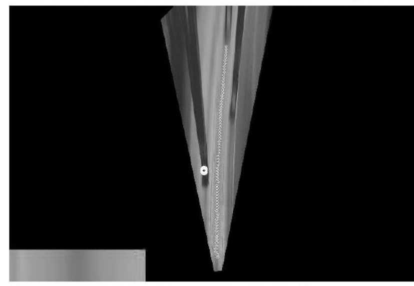

(b) Inverse perspective warping showing curvature detection (small white dots) and template (lower left corner)

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.

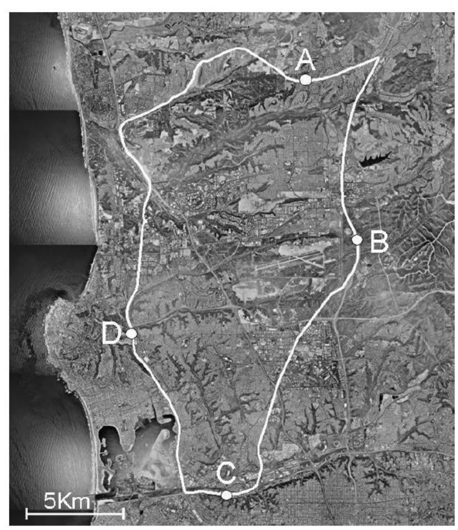

Fig. 14. The 65Km route in San Diego used in the evaluation. The route is overlayed on aerial photography. Points A, B, C, and D are sections of road used in the evaluation (photography courtesy USGS)

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.

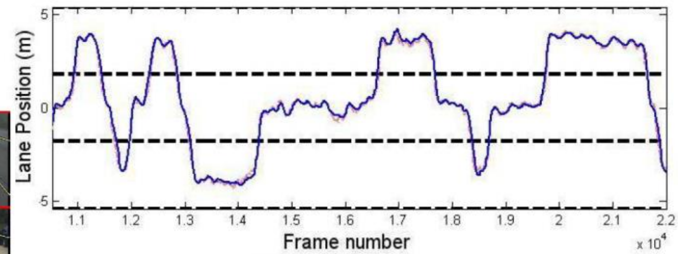

ig. 12. Detected lateral position in meters (solid blue) superimposed on ground truth (dashed pink) plotted vs. frame number with dashed lines marking ıe position of lane boundaries for an 11,000 frame (slightly over 6 minute) sequence.

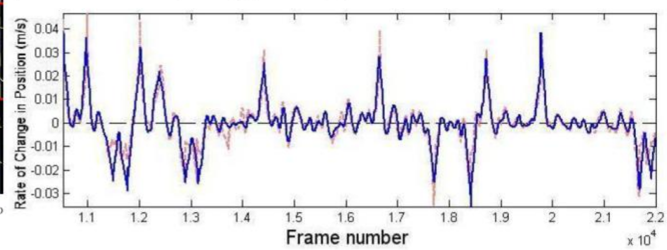

Fig. 13. Detected departure rate in m/s^2 (solid blue) superimposed on ground truth (dashed pink) plotted vs. frame number with dashed line marking the abscissa for the same sequence shown in figure 12.

Detected lanes = Solid blue
Ground truth = Pink dashed line

Challenges: Occlusions and Highlights

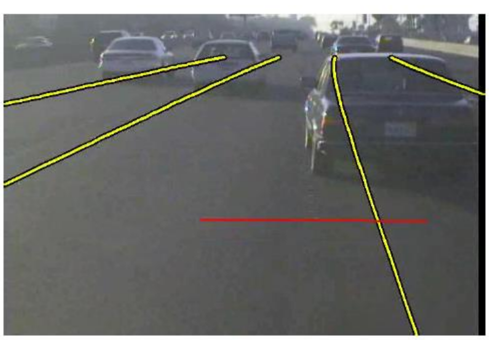

Fig. 16. Error due to occlusion of the road by a vehicle on the dusk dataset on road segment C. The red horizontal line shows the proximity of the occluding vehicle detected by the in-vehicle LASER RADAR sensors.

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.

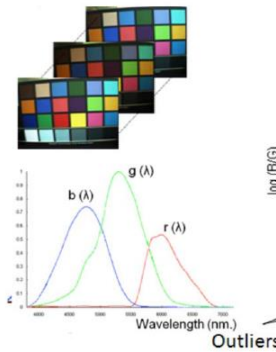

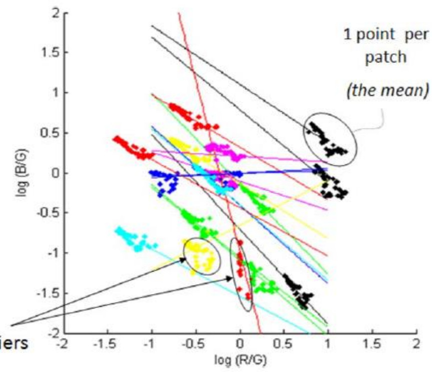

Road Detection Based on Illuminant Invariance

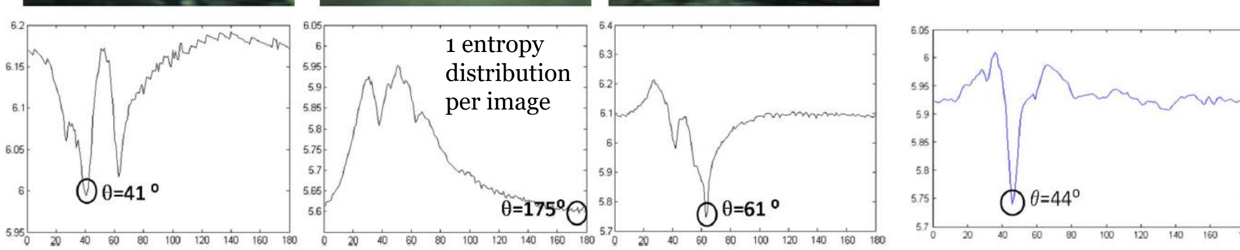

Entropy based camera calibration: can be used for onboard self calibration and image content.

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

Road Detection Based on Illuminant Invariance

Idea: Projection in log(R/G) log(B/G) space along angle theta (Θ)
=>
Gives axis where each level corresponds to a chromaticity independent of illumination

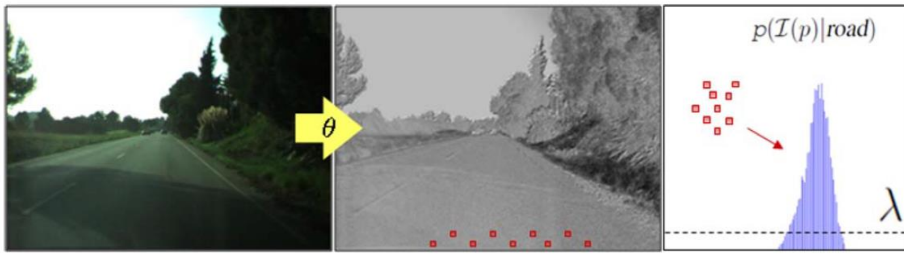

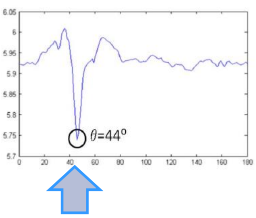

Entropy based camera calibration resulted in Θ is the invariant angle that allows us to compute I a grayscale image invariant to lighting variations.

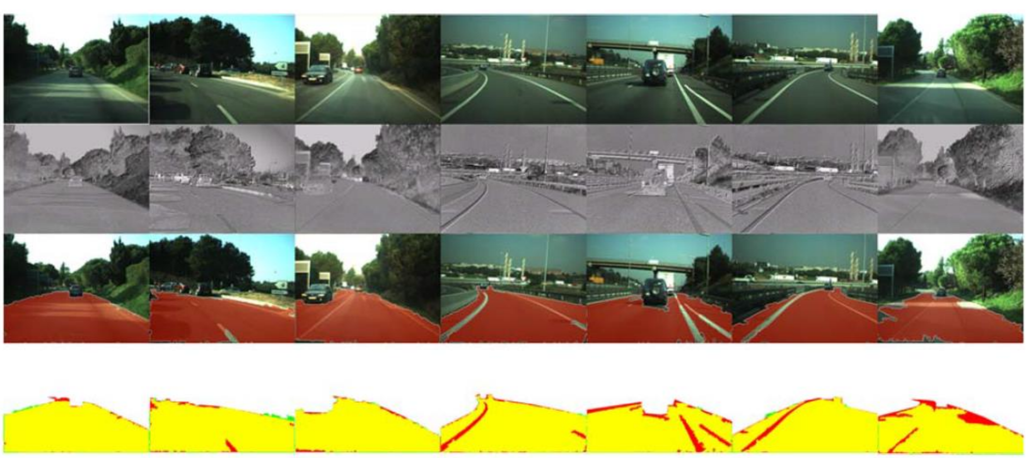

Fig. 11. Road-detection examples. (Top row) Original image. (Second row) Illuminant-invariant image. (Third row) Detected road. (Bottom row) Comparison against hand-segmented result. (Yellow) Correctly classified pixels. (Green) Falsely detected road pixels. (Red) False background pixels.

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.

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

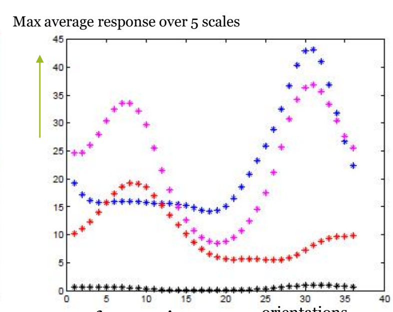

Gabor complex responses for 4 points.

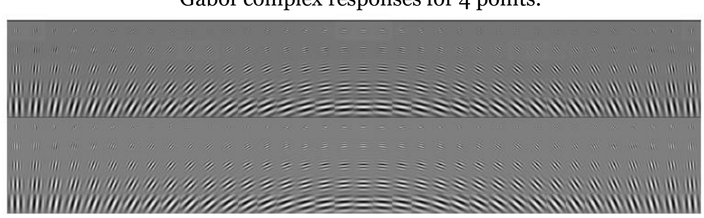

Gabor kernels: real part (row 1-5), imaginary part (row 6-10)

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

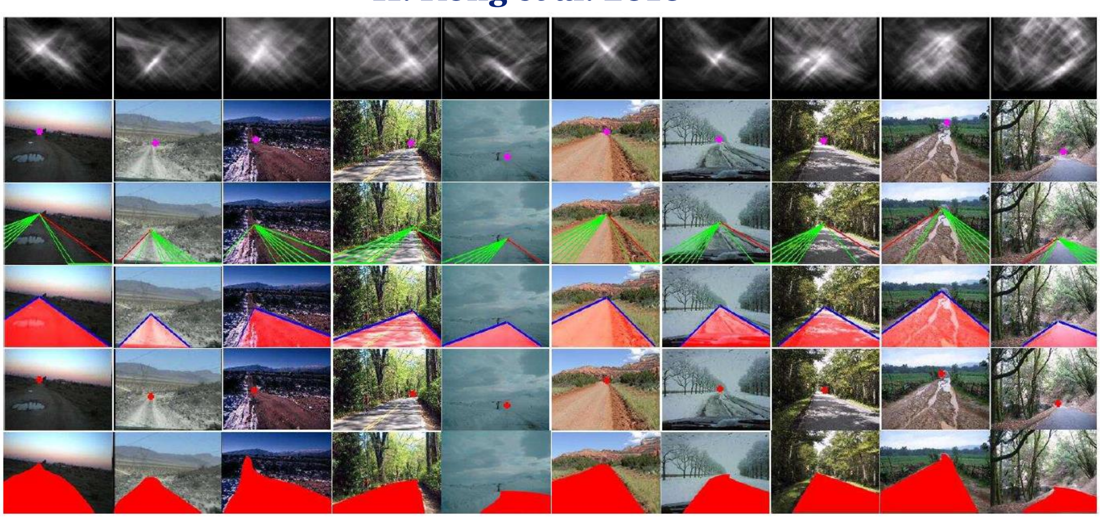

Last row: ground truth..

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

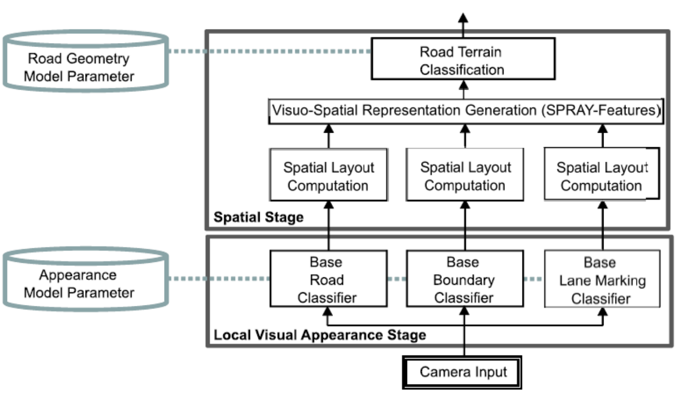

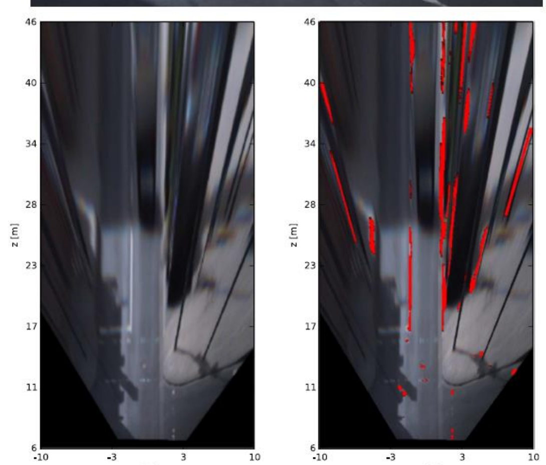

Ground Truth

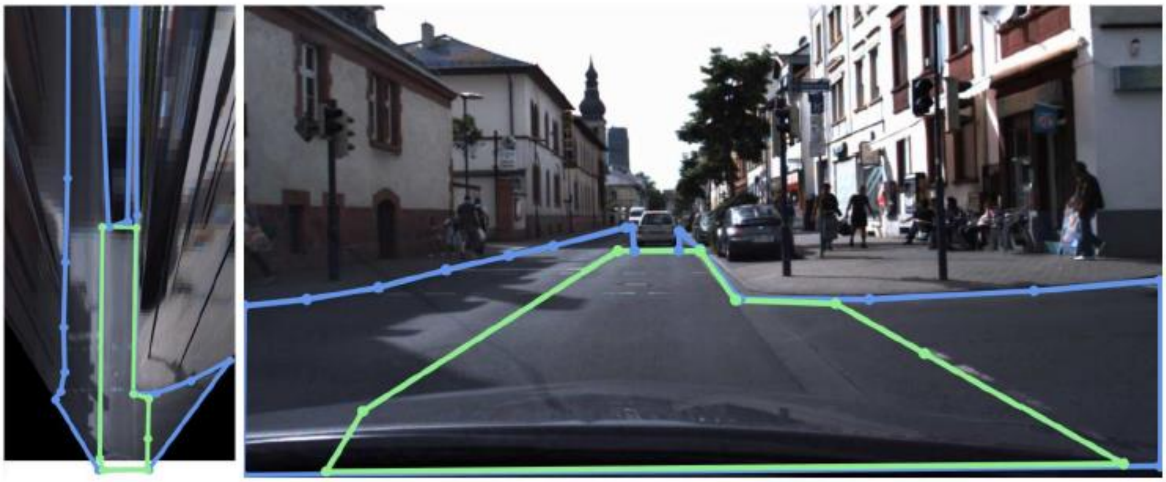

Slow Feature Analysis (SFA)

Slow Feature Analysis (SFA)

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

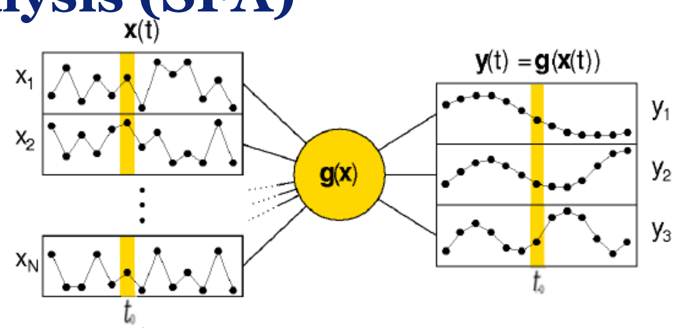

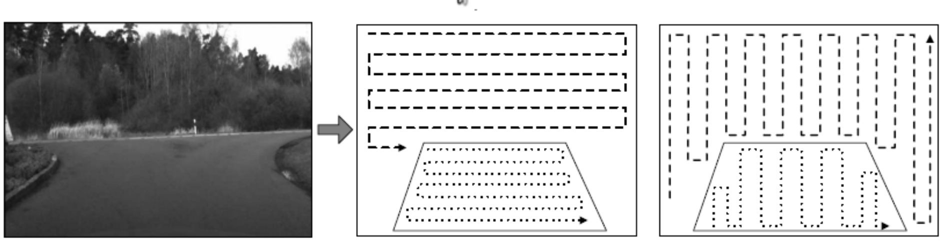

Spatial patch sequence extraction for SFA training: on the left the horizontal path and on the right the vertical path is illustrated. The paths are partitioned into road (dotted) and non-road (dashed) sections.

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.

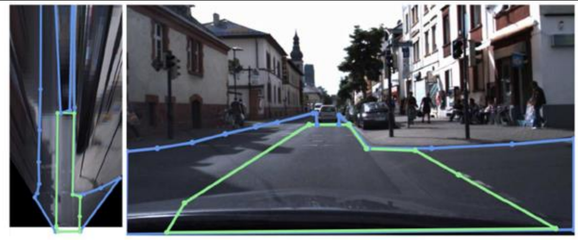

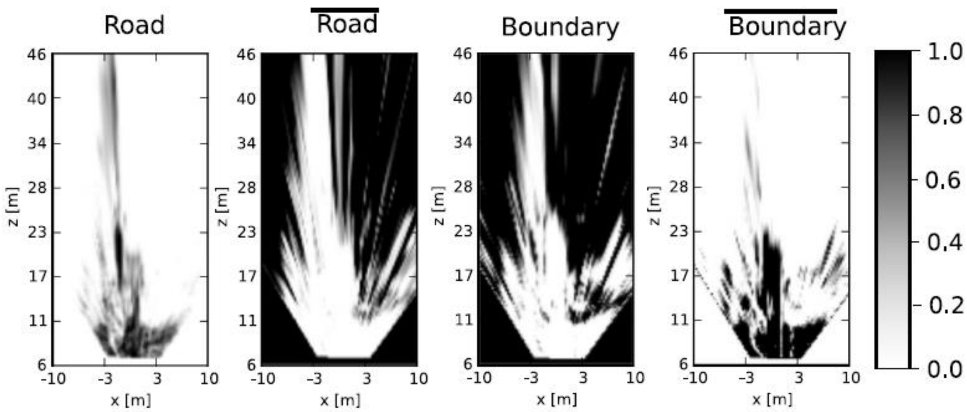

Feature vectors:

- Slow Feature Analysis (SFA) features
- Color features (RGB)
- Walsh Hadamard texture features

Ground Truth

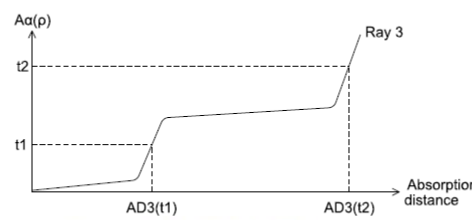

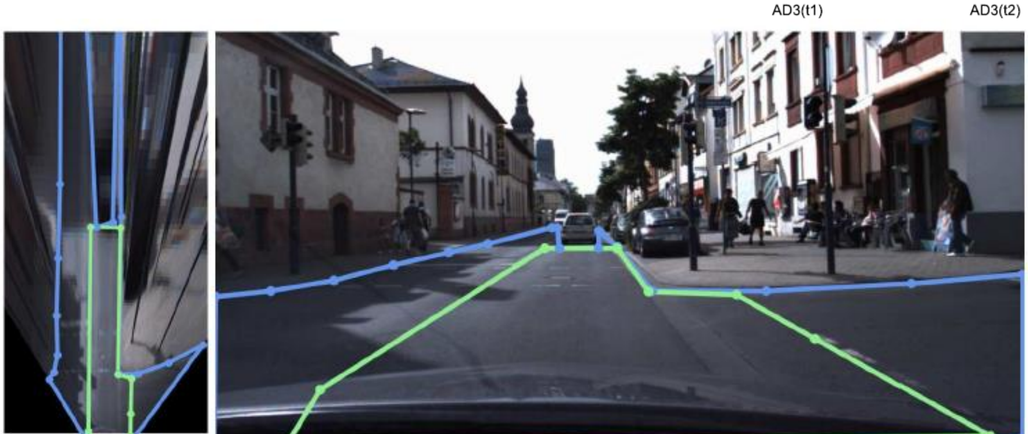

Results

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)

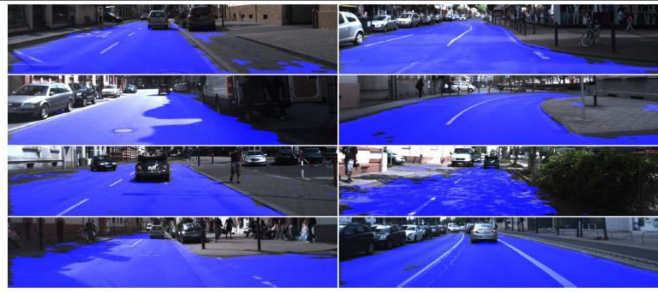

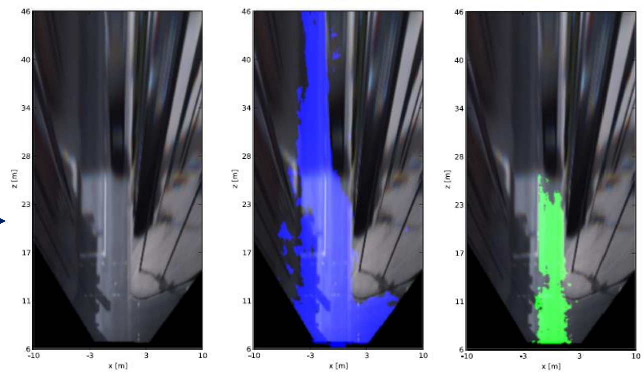

Results

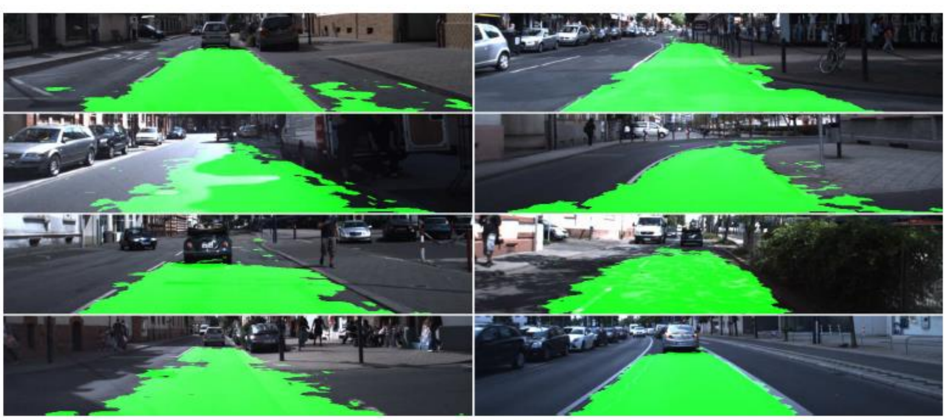

- Ego Lanes

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)

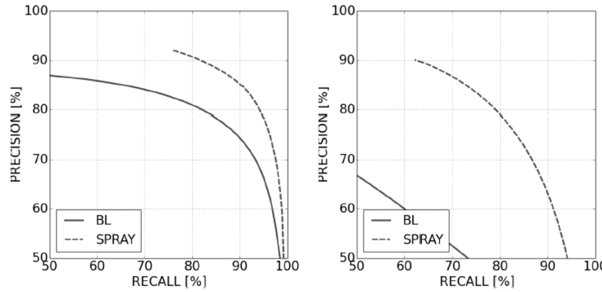

(a) Perspective evaluation (b) Metric evaluation

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

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

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

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

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

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

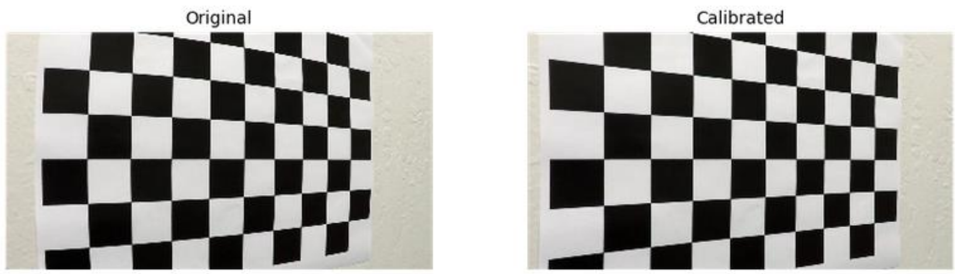

Original vs. calibrated images

... 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

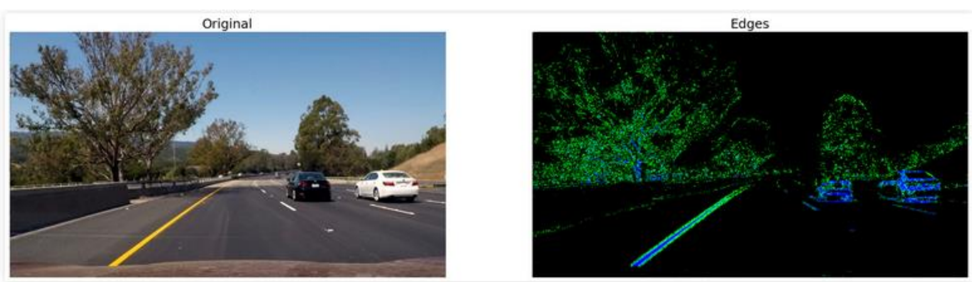

Original vs. highlighted edges

Gradient Absolute Values, Ranges within certain magnitudes, Gradient Directions

- Sobel Operator (using a convolutional Kernel)

Color Ranges

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

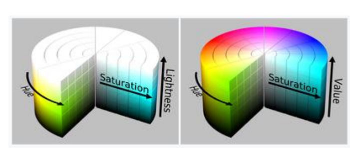

Fig. 2a. HSL cylinder. Fig. 2b. HSV cylinder.

Universiteit Leiden. Bij ons leer je de wereld kennen

Image Processing using Convolutional Kerne

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

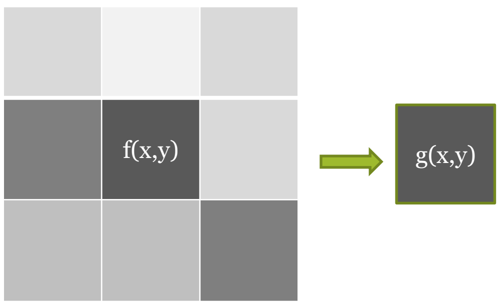

Operation	Kernel ω	Image result g(x,y)
Identity	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	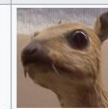
Edge detection	$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{bmatrix}$	
	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	
	$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$	
Sharpen	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 &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

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

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.
- ...

Mobile Robot Challenge

Lecturer:
Dr Erwin M. Bakker (erwin@liacs.nl)
Room 126a and LIACS Media Lab (LML)

Teaching assistant:
Mor Puigventos (email)
TBA (email)

Period: February 6th - May 22nd 2023
Time: Monday 15.15 - 17.00
Place (Rooms):
a) Gorlaeus - Lecture Hall C3
b) Sylvius - 15.31
c) Van Steenis - E0.04
d) Oortgebouw - Sitterzaal

Schedule (tentative, visit regularly):

Date	Room	Subject
6-2	a	Introduction and Overview
13-2	a	Locomotion and Inverse Kinematics
20-2	b	Robotics Sensors and Image Processing
27-2	a	SLAM + SLAM Workshop
6-3	c	Mobile Robot Challenge Introduction
13-3	a	Project Proposals I (by students)
20-3	d	Project Proposals II (by students)
27-3	d	Robotics Vision (Week 13, start 15.30)
3-4	d	Robotics Reinforcement Learning&Workshop
10-4		No Class (Eastern)
17-4	d	Project Progress I (by students)
24-4	d	Project Progress II (by students)
1-5	d	Mobile Robot Challenge I
8-5	a	Mobile Robot Challenge II
15-5	d	Project Demos I
22-5	d	Project Demos II
29-5		Whit Monday
5-6		Project Deliverables

3x

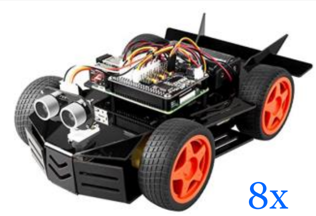

8x

6x

2x

Group Size 4 students

Please form your group in the coming week.

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

Robotics Project Proposals Presentations

Monday 13-3 and 20-3 2023

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 15^{th} and May 22^{rd} 2023.

Note: Groups of 1-5 members are allowed. (Please form your project group in the coming week.)

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