ORT (other random topics)

CSE P 576
Larry Zitnick (larryz@microsoft.com)

Autonomous vehicles

- Navlab (1990's)
- Stanley (Offroad, 2004)
- Boss (Urban, 2007)

Navlab (1985-2001)

Navlab 1

Navlab 2

Navlab (1985-2001)

Navlab 5

Navlab 6

Navlab (1985-2001)

Navlab 10

Navlab (1992)

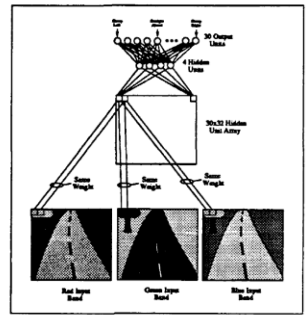

Neural Network Perception for Mobile Robot Guidance, Dean **A.** Pomerleau**, 1992**

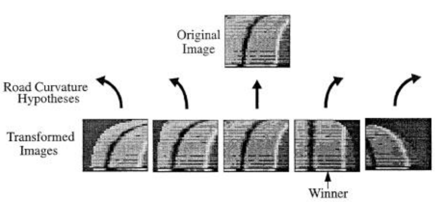

Figure 2: Technique for determining road curvature by "straightening" image features.

RALPH: Rapidly Adapting Lateral Position Handler, *Dean Pomerleau, 1995*

No Hands Across America

- **2797/2849 miles (98.2%)**
- The researchers handled the throttle and brake.
- When did it fail?

NO HANDS ACROSS AMERICA

Stanley

Following slides courtesy of Sebastian Thrun

2004: Barstow, CA, to Primm, NV

- 150 mile off-road robot race across the Mojave desert
- Natural and manmade hazards
- No driver, no remote control
- No dynamic passing
- Fastest vehicle wins the race (and 2 million dollar prize)

Grand Challenge 2005: 195 Teams

Stanford Racing Team

Final Result: Five Robots finished!

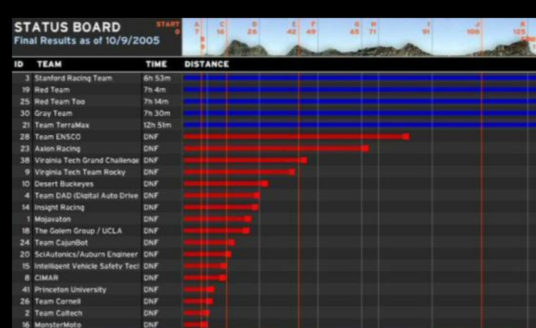

Stanford Racing Team

Manual Offroad Driving

Stanford Racing Team

Software Architecture

Stanford Racing Team

Stanley Software Architecture

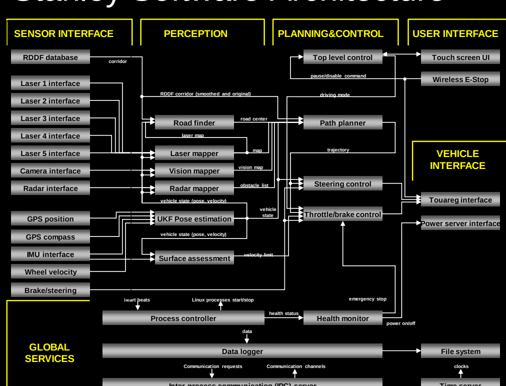

Planning and Steering Control

Stanford Racing Team

Low-Level Steering Control

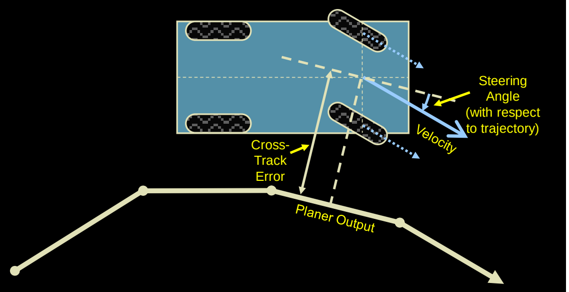

Discuss Kalman Filter

To the whiteboard…

Parameterizing Search Space

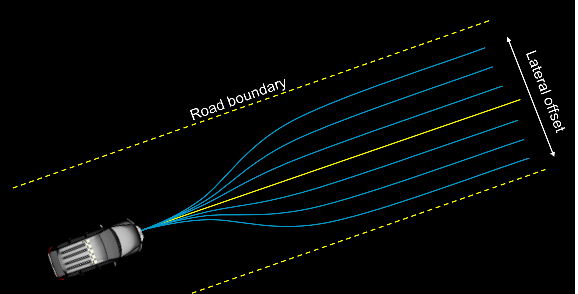

Planning = Rolling out Trajectories

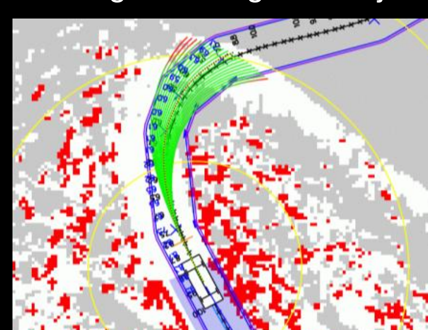

Lateral Offset Profiles

Swerves

- step changes in desired lateral offset
- avoidance of frontal obstacles

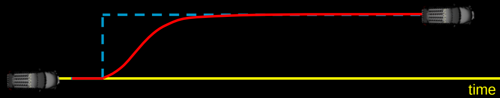

Nudges

- ramp changes in desired lateral offset
- Road centering

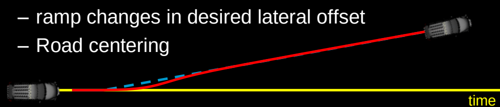

Smooth Driving at 25mph

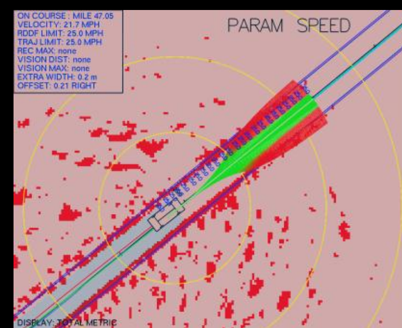

Laser Terrain Mapping

UKF Position Estimation

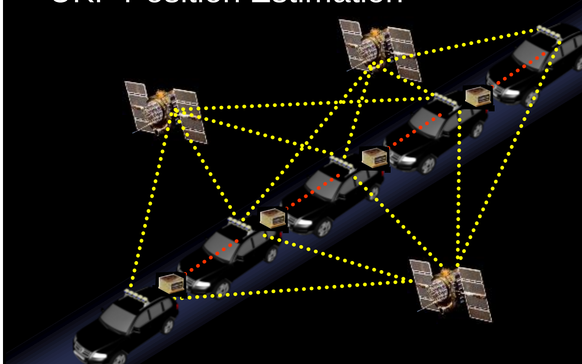

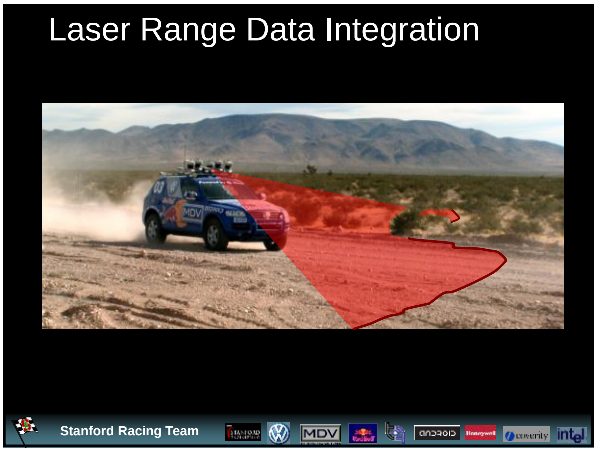
Laser Range Data Integration
Stanford Racing Team

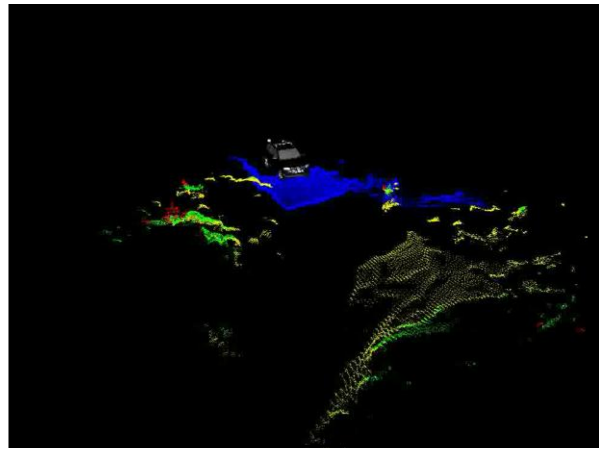

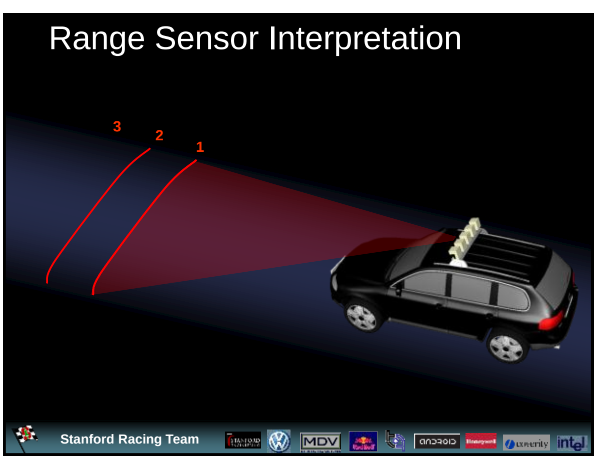
Range Sensor Interpretation
3
2
1
Stanford Racing Team

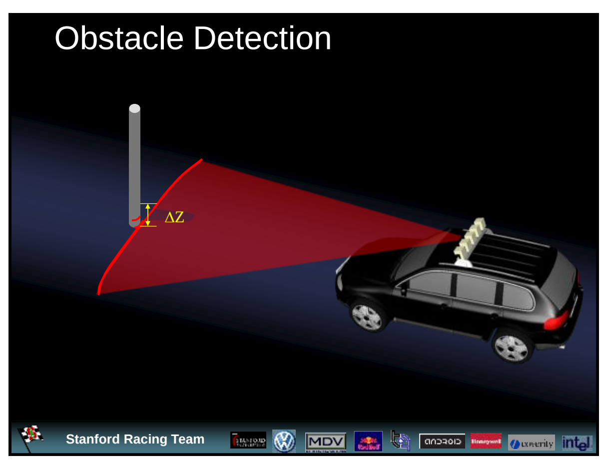
Obstacle Detection
ΔZ
Stanford Racing Team

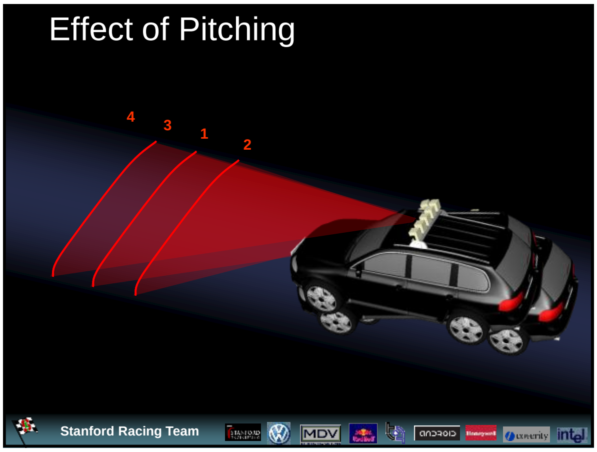
Effect of Pitching
4
3
1
2
Stanford Racing Team

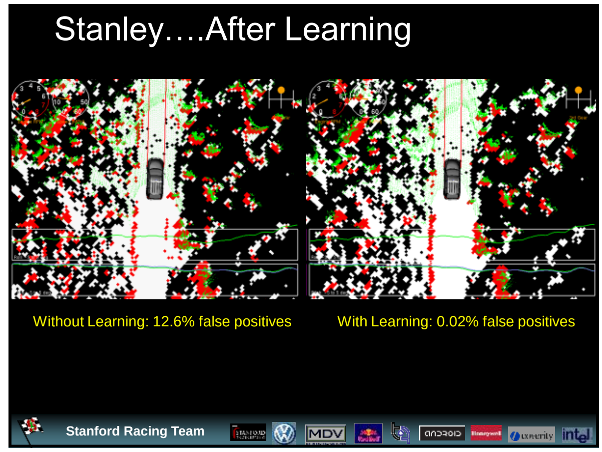
Stanley….After Learning
Without Learning: 12.6% false positives
With Learning: 0.02% false positives
Stanford Racing Team

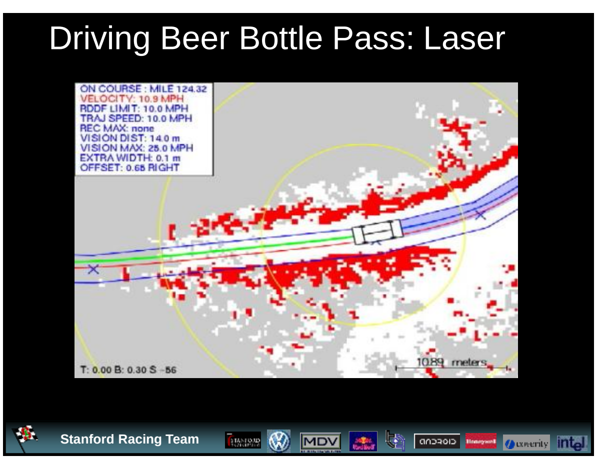
Driving Beer Bottle Pass: Laser
ON COURSE : MILE 124.32
VELOCITY: 10.9 MPH
RDDF LIMIT: 10.0 MPH
TRAJ SPEED: 10.0 MPH
REC MAX: none
VISION DIST: 14.0 m
VISION MAX: 25.0 MPH
EXTRA WIDTH: 0.1 m
OFFSET: 0.65 RIGHT
T: 0.00 B: 0.30 S -56
10.89 meters
Stanford Racing Team

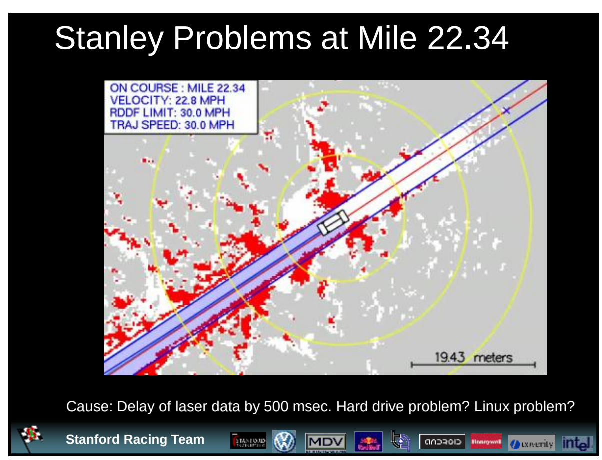
Stanley Problems at Mile 22.34
ON COURSE : MILE 22.34
VELOCITY: 22.8 MPH
RDDF LIMIT: 30.0 MPH
TRAJ SPEED: 30.0 MPH
19.43 meters
Cause: Delay of laser data by 500 msec. Hard drive problem? Linux problem?
Stanford Racing Team

Computer Vision Terrain Mapping

Limits of lasers

Lasers see 22m = 25mph

They needed to go 35mph to finish the race.

Stanford Racing Team

What Defines A Road?

Stanford Racing Team

Idea: Continual Terrain Adaptation

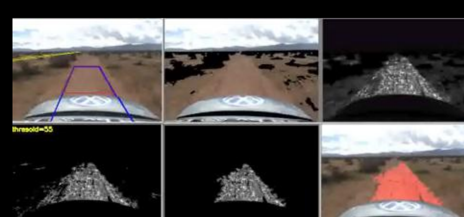

Fast adaptation: Mean & covariance of Gaussian, exponential forgetting
Slow learning: memory of k past Gaussians

Stanford Racing Team

Adaptive Vision In Action (NQE)

Stanford Racing Team

Adaptive Vision in Mojave Desert

Stanford Racing Team

Driving Beer Bottle Pass: Vision

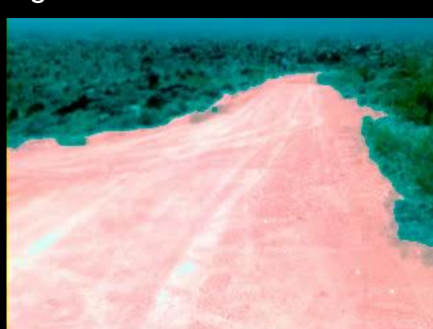

Stanford Racing Team

Speed Control

Stanford Racing Team

Speed Controllers

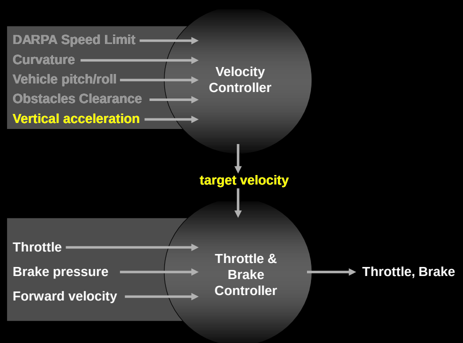

Stanford Racing Team

Controlling speed

If you hit a bump, slow down (that first pothole really hurts…)

If you haven't hit a bump in awhile linearly increase speed.

Slow down on hills.

Stanford Racing Team

How Fast Do Humans Drive

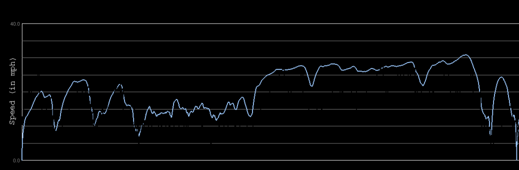

Stanford Racing Team

Learning To Drive Like a Person

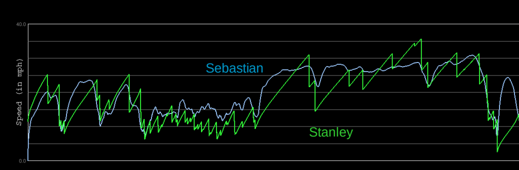

Stanford Racing Team

More info

For full detail read the paper:

Stanley: The Robot that Won the DARPA Grand Challenge, Sebastian et al., Journal of Field Robotics, 2006

Stanford Racing Team

Boss

http://www.tartanracing.org/blog/index.html#22

DARPA Urban Challenge

- 36 teams invited to National Qualification Event.
- 11 teams invited to Urban Challenge Final Event

Suddenly, the vehicle did a U-turn and headed directly at Tether's vehicle. "Five of us in the vehicle were all yelling 'pause!'" Tether recalled, referring to the pause command that DARPA could send to a vehicle.

http://www.tartanracing.org/blog/index.html#22

Boss

2007 Chevrolet Tahoe

CompactPCI chassis with 10 2.16-GHz Core2Duo processors, each with 2 GB of memory

Sensor	Characteristics
Applanix POS-LV 220/420 GPS/IMU (APLX)	• Submeter accuracy with Omnistar VBS corrections • Tightly coupled inertial/GPS bridges GPS outages
SICK LMS 291-S05/S14 LIDAR (LMS)	• 180/90 deg × 0.9 deg FOV with 1/0.5-deg angular resolution • 80-m maximum range
Velodyne HDL-64 LIDAR (HDL)	• 360 × 26-deg FOV with 0.1-deg angular resolution • 70-m maximum range
Continental ISF 172 LIDAR (ISF)	• 12 × 3.2 deg FOV • 150-m maximum range
IBEO Alasca XT LIDAR (XT)	• 240 × 3.2 deg FOV • 300-m maximum range
Continental ARS 300 Radar (ARS)	• 60/17 deg × 3.2 deg FOV • 60-m/200-m maximum range
Point Grey Firefly (PGF)	• High-dynamic-range camera • 45-deg FOV

Motion planning

- Structured driving (road following)
- Unstructured driving (maneuvering in parking lots)

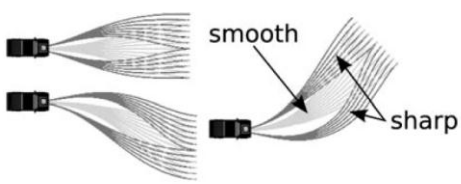

Figure 4. Smooth and sharp trajectories. The trajectory sets are generated to the same endpoints but differ in their initial commanded curvature.

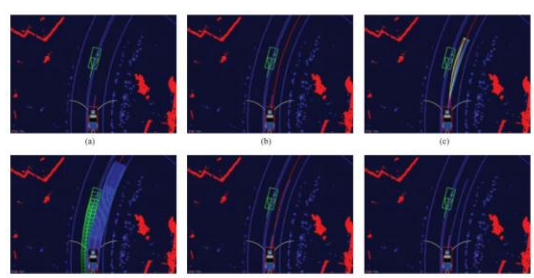

Figure 5. A single timeframe following a road lane from the DARPA Urban Challenge. Shown is the centerline path extracted from the lane (b), the trajectories generated to track this path (c), and the evaluation of one of these trajectories against both static and dynamic obstacles (d and e).

Where am I?

- GPS + inertial + wheel encoder = 0.1m, but if you go under a tree you loose the signal. 30 minutes to reacquire.
- Lane markers are found using SICK lasers.

Particle Filters

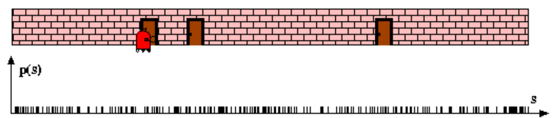

Particle filter slides courtesy of Sebastian Thrun

Sensor Information: Importance Sampling

$$Bel(x) \leftarrow \alpha\, p(z \mid x)\, Bel^{-}(x)$$

$$w \leftarrow \frac{\alpha\, p(z \mid x)\, Bel^{-}(x)}{Bel^{-}(x)} = \alpha\, p(z \mid x)$$

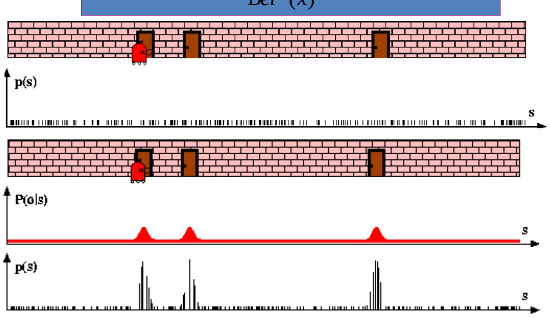

Robot Motion

$$Bel^{-}(x) \leftarrow \int p(x \mid u, x')\, Bel(x')\, \mathrm{d}x'$$

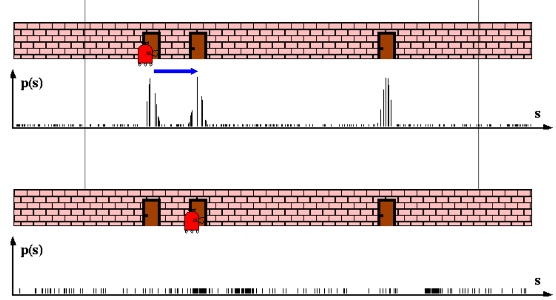

Sensor Information: Importance Sampling

$$Bel(x) \leftarrow \alpha\, p(z \mid x)\, Bel^{-}(x)$$

$$w \leftarrow \frac{\alpha\, p(z \mid x)\, Bel^{-}(x)}{Bel^{-}(x)} = \alpha\, p(z \mid x)$$

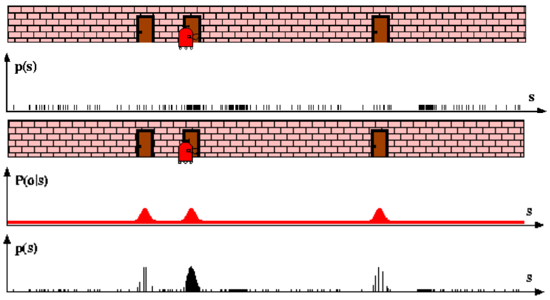

Robot Motion

$$Bel^{-}(x) \leftarrow \int p(x \mid u, x')\, Bel(x')\, \mathrm{d}x'$$

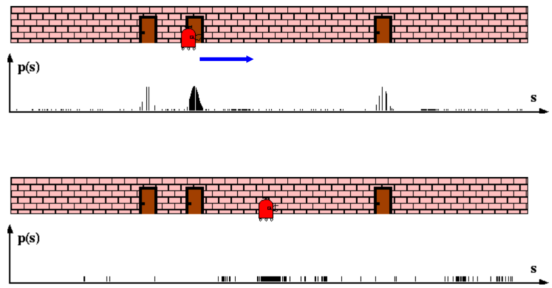

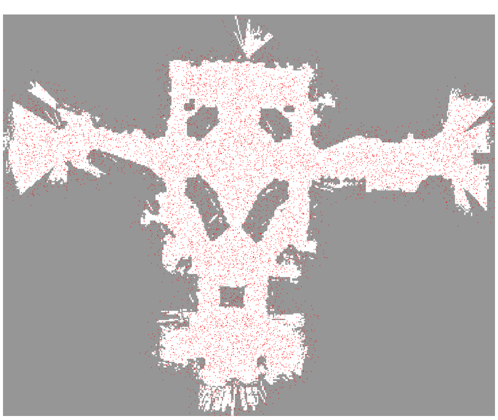

59

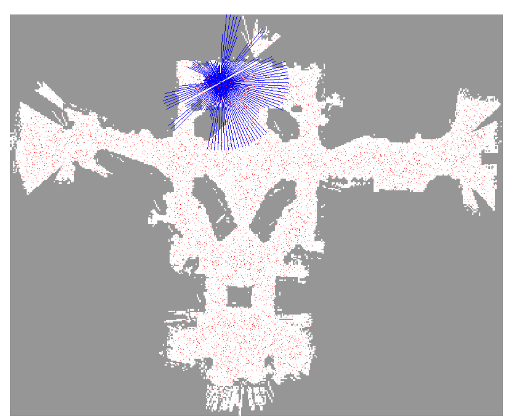

60

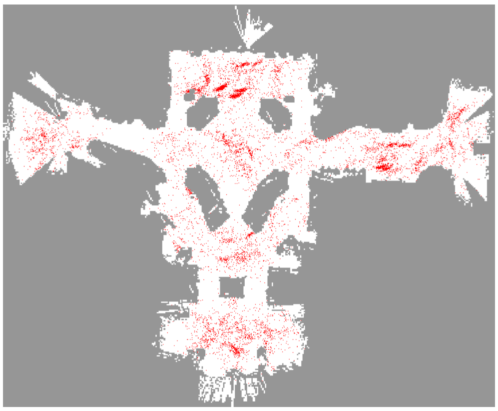

61

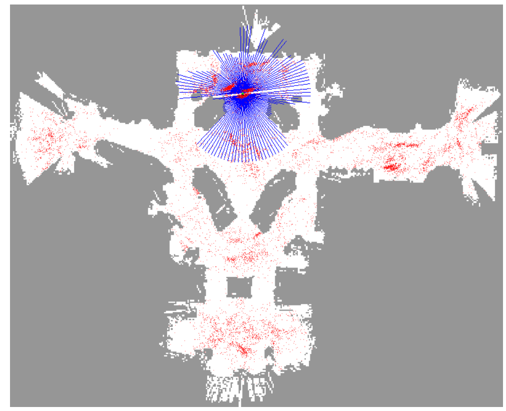

62

63

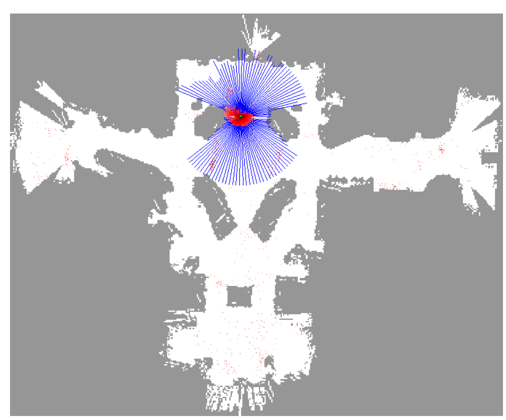

64

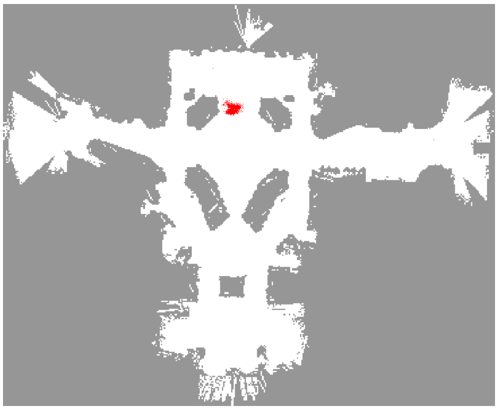

Road estimation

Figure 13. Example showing the road shape estimate (parallel curves) for an off-road scene. Obstacles and berms are illustrated by pixels.

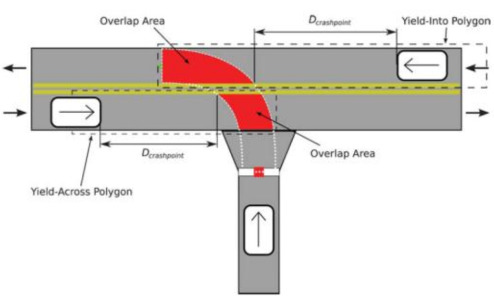

Figure 16. Typical tee intersection with yield lanes.

$$T_{\text{required}} = T_{\text{action}} + T_{\text{delay}} + T_{\text{spacing}}$$

Failures

Figure 28. Data replay shows how the incorrectly extrapolated path of a vehicle (shaded rectangles) and the wall (pixels to the right of Boss) create a space that Boss believes is too narrow to drive through (indicated by the arrow).

Failures

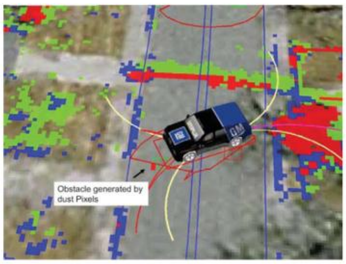

Figure 29. The false obstacles generated by dust and the bush behind Boss prevented Boss from initially completing its U-turn.

Failures

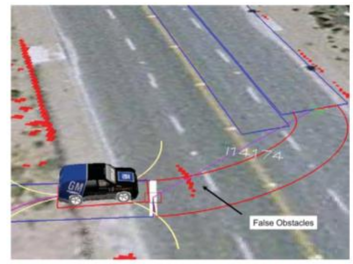

Figure 30. False obstacles that caused Boss to stutter when leaving a dirt road.

Failures

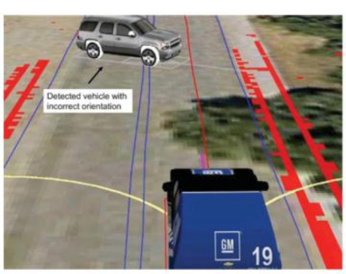

Figure 31. This incorrect estimate of an oncoming vehicle's orientation caused Boss to swerve and almost to become irrevocably stuck.

More info

Autonomous Driving in Urban Environments: Boss and the Urban Challenge, Urmson et al., Journal of Field Robotics, 2008

A Fast Approximation of the Bilateral Filter using a Signal Processing Approach

Sylvain Paris and Frédo Durand

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology

CSAIL

Link with Linear Filtering
Introducing a Convolution

space: 1D Gaussian
× range: 1D Gaussian

combination: 2D Gaussian

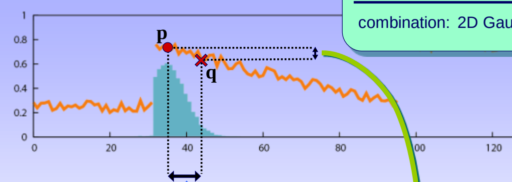

$$\begin{pmatrix} W^{bf}_{\mathbf{p}} I^{bf}_{\mathbf{p}} \\ W^{bf}_{\mathbf{p}} \end{pmatrix} = \sum_{\mathbf{q}\in\mathcal{S}} \underbrace{G_{\sigma_s}(\|\mathbf{p}-\mathbf{q}\|)}_{\text{space}} \underbrace{G_{\sigma_r}(|I_{\mathbf{p}}-I_{\mathbf{q}}|)}_{\text{range}} \begin{pmatrix} W_{\mathbf{q}} I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

Link with Linear Filtering
Introducing a Convolution

space: 1D Gaussian
× range: 1D Gaussian

combination: 2D Gaussian

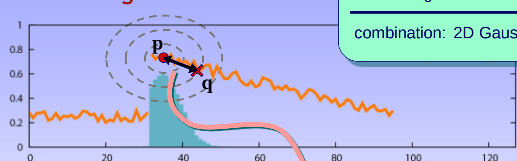

$$\begin{pmatrix} W^{bf}_{\mathbf{p}} I^{bf}_{\mathbf{p}} \\ W^{bf}_{\mathbf{p}} \end{pmatrix} = \sum_{\mathbf{q}\in\mathcal{S}} \underbrace{G_{\sigma_s}(\|\mathbf{p}-\mathbf{q}\|)\, G_{\sigma_r}(|I_{\mathbf{p}}-I_{\mathbf{q}}|)}_{\text{space x range}} \begin{pmatrix} W_{\mathbf{q}} I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

Corresponds to a 3D Gaussian on a 2D image.

Link with Linear Filtering
Introducing a Convolution

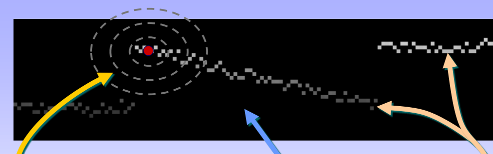

sum all values

black = zero

$$\begin{pmatrix} W^{bf}_{\mathbf{p}} I^{bf}_{\mathbf{p}} \\ W^{bf}_{\mathbf{p}} \end{pmatrix} = \boxed{\sum_{(\mathbf{q},\zeta)\in\mathcal{S}\times\mathcal{R}}} \ \text{space-range Gaussian} \begin{pmatrix} W_{\mathbf{q}} I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

sum all values multiplied by kernel ⇨ convolution

Link with Linear Filtering
Introducing a Convolution

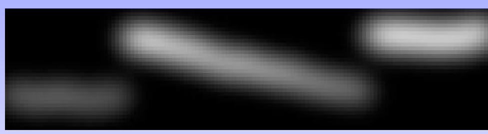

result of the convolution

$$\begin{pmatrix} W^{\mathrm{bf}}_{\mathbf{p}} \; I^{\mathrm{bf}}_{\mathbf{p}} \\ W^{\mathrm{bf}}_{\mathbf{p}} \end{pmatrix} = \sum_{(\mathbf{q},\zeta)\in\mathcal{S}\times\mathcal{R}} \text{space-range Gaussian} \begin{pmatrix} W_{\mathbf{q}} \; I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

Link with Linear Filtering
Introducing a Convolution

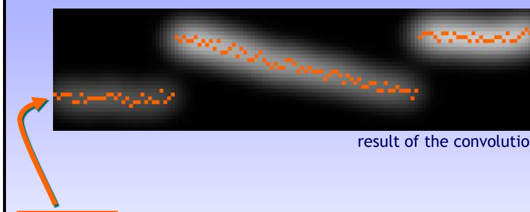

result of the convolution

$$\begin{pmatrix} \boxed{W^{\mathrm{bf}}_{\mathbf{p}} \; I^{\mathrm{bf}}_{\mathbf{p}}} \\ W^{\mathrm{bf}}_{\mathbf{p}} \end{pmatrix} = \sum_{(\mathbf{q},\zeta)\in\mathcal{S}\times\mathcal{R}} \text{space-range Gaussian} \begin{pmatrix} W_{\mathbf{q}} \; I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

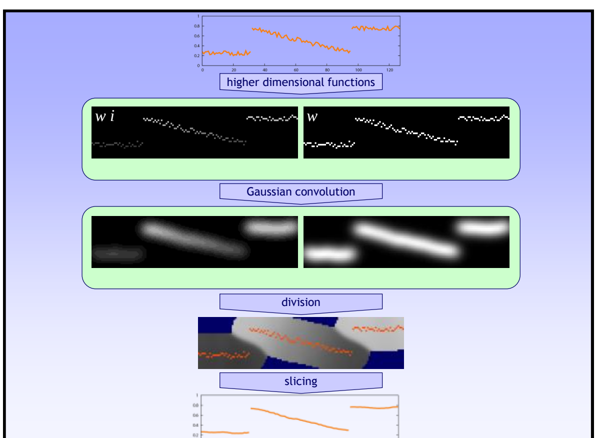

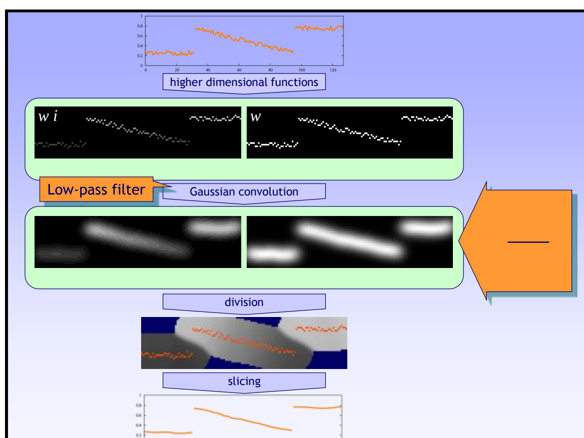

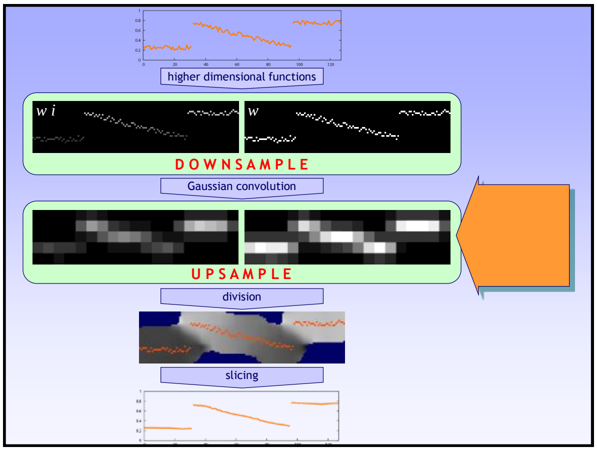

Accuracy versus Running Time

- Finer sampling increases accuracy.
- More precise than previous work.

Digital photograph
1200 × 1600

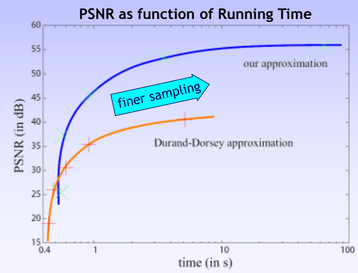

Straightforward implementation is over 10 minutes.

Visual Results

- Comparison with previous work [Durand 02]
  - running time = 1s for both techniques

1200 × 1600

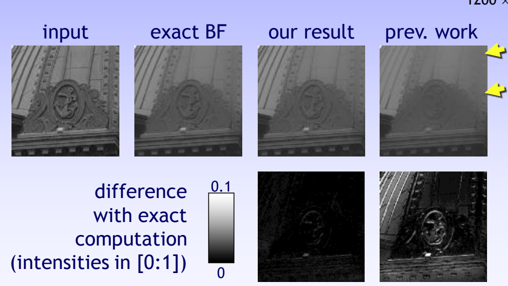

More advanced approaches:

http://graphics.stanford.edu/papers/gkdtrees/

http://graphics.stanford.edu/papers/permutohedral/