team MIT DARPA Grand Challenge

Development of a Self-Driving Car for the DARPA Urban Challenge

Seth Teller
CS and AI Laboratory
EECS Department
MIT

Joint work with: Matt Antone, **David Barrett**, Mitch Berger, Bryt Bradley, Ryan Buckley, Stefan Campbell, Alexander Epstein, Gaston Fiore, Luke Fletcher, Emilio Frazzoli, **Jonathan How**, Albert Huang, Troy Jones, Sertac Karaman, Olivier Koch, Yoshi Kuwata, Victoria Landgraf, **John Leonard**, Keoni Maheloni, David Moore, Katy Moyer, Edwin Olson, Andrew Patrikalakis, Steve Peters, Stephen Proulx, Nicholas Roy, Chris Sanders, Justin Teo, Robert Truax, Matthew Walter, Jonathan Williams

CSAIL

"Desert" Grand Challenges

- Military interest in autonomous land vehicles
  - Congressional mandate (H.R. 4205/P.L. 106-398, Oct. 2000): "one third of operational ground combat vehicles to be **unmanned** by 2015"
- DGC 1: March 2004
  - 142 miles in 10 hours, $1M prize
  - 106 entering teams; no finishers
  - Dense, pre-mapped GPS corridor
  - No moving obstacles (static world)
  - One vehicle at a time
    - Whenever two robots came close, one was manually paused
- DGC 2: October 2005
  - 132 miles in 10 hours, $2M prize
  - Dense, pre-mapped graded roadway
  - One vehicle at a time, as in DGC 1
  - 195 teams, 5 finishers

Urban Challenge (2007)

- Novel elements:
  - **Urban road network**
  - **Moving traffic**
    - Human and robotic!
  - **No course inspection**
  - 60 miles in 6 hours
    - Scored by speed, safety
    - $3.5M prize pool
  - 89 entering teams
- Program goals
  - Safe (collision-free, polite) driving at up to 30mph
  - Capable (turns, stops, intersections, merging, parking, …)
  - Robust (blocked roads, erratic drivers, sparse waypoints, GPS degradation and outages, …)

Source: DARPA Urban Challenge Participants Briefing, May 2006

Why tackle *this* problem?

- **Fatalities and injuries** from driving accidents
  - Tens of thousands of fatalities per year in U.S.
  - Hundreds of thousands of injuries annually
- **Productivity lost** to commuting, travel
  - Billions of person-hours per year "spent" driving
- **Energy inefficiency** of braking and idling
  - Could do much better with cooperating vehicles
- **Mandate** from U.S. Congress
  - 1/3 of military ground vehicles unmanned by 2015
- **Sheer appeal** of designing a robotic vehicle that exhibits human-level driving capability!

Program Scope

- In scope:
  - Following
  - Intersections
  - Passing, Merging
  - Parking, U-turns
  - Emergency stops
  - Timely left turns across traffic
  - Potholes, construction sites
  - Blockages, replanning
- Out of scope:
  - Pedestrians
  - High speed (> 30 mph)
  - Traffic signals, signage
  - Difficult off-road terrain
  - Highly inclement weather

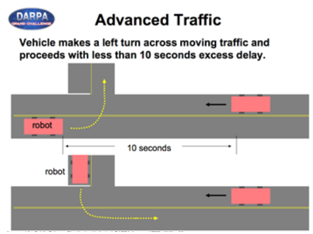

DARPA-Provided Inputs

- USB stick with two data files:
- RNDF: Road Network Description File
  - Provided 48 hours ahead
  - Topology of road network
  - "Sparse" GPS waypoints
  - Geometry of intersection, parking zone boundaries
- MDF: Mission Description File
  - Provided 5 minutes ahead
  - List of RNDF waypoints to be traversed by autonomous car

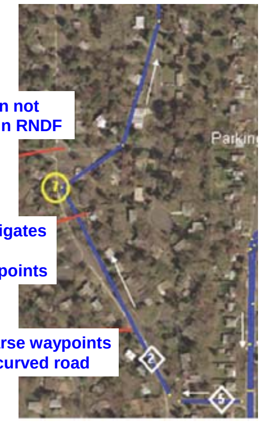

Source: DARPA Participants Briefing, May 2006

Fundamental Questions

- Autonomous driving includes four key problems:
  - Where is the road?
    - Identify drivable road surface at fine grain
  - Where are the static obstacles?
    - Hard: curbs, potholes, signposts, buildings
    - Soft: lane markings, shoulders, vegetation
  - Where are the other vehicles?
    - Where might they move in the near future?
  - How should the vehicle behave?
    - Codify (non-algorithmic) rules of safe, legal, "human-like" driving
- Solve all of the above, with available (uncertain) sensor data, in real time (without killing anyone).

Related Work

- Partial Autonomous Driving Systems
  - Limited domain (highway lane; traffic-free road)
  - Require human to: stage control handoff, monitor operation, and take over in emergency situations
  - Munich's VaMoRs (1985-2004), VAMP (1993-2004); CMU's NAVLAB (1985); Penn (Southall & Taylor 2001)
- Assistive Driving Technologies
  - Limited duty cycle (cruising, emergencies, staged parking) and actuation (e.g. none, or brakes only)
  - Require human handoff and resumption of control
  - Automakers' ABS, cruise control, self-parking systems
  - Lane departure warnings (Mobileye, Iteris, ANU)

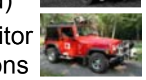

Assessment and Strategy

- Human-level urban driving not achievable with existing algorithms / systems as of 2006
  - Key issues: uncertainty; sensing/CPU resources; safety
  - Example: if vehicle is unsure where the road is, and/or where it is with respect to the road, identifying a safe, appropriate traffic behavior (at speed!) is very difficult
- Strategy
  - Technical footprint for success covers many disciplines → interdisciplinary approach integrating EECS & MechE
  - Spiral design approach → figure out how to solve the problem while designing the system at the same time

Compressed Timeline

- Bring up rapid prototype vehicle (Ford Escape) summer/fall ‘06
  - Gain experience with sensors, DBW dynamics, coding, configuration
- Bring up competition vehicle (LandRover LR3) spring ‘07
  - Develop mature algorithms, tune for qualifying rounds and final event
- Compete in Challenge, fall 2007

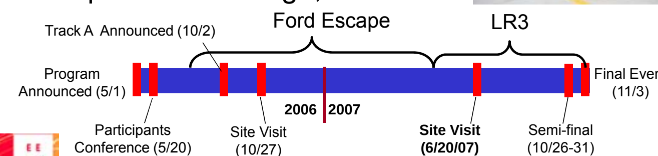

Our Approach

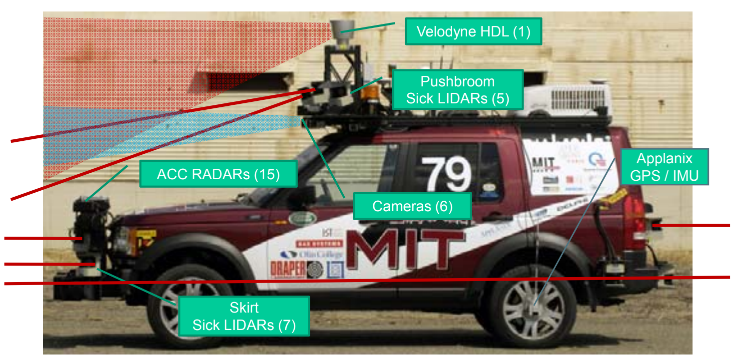

Design Strategy

- Sensor-rich, CPU- and I/O-intensive architecture
  - Many sensors to interpret surroundings “live”
  - Intensive use of live and logged data visualization
  - Many resources, to avoid premature optimization
- Redundancies:
  - Sensor type and spatial coverage
  - Closed-loop multi-level planning and control
  - Computation failover at process level
  - Firmware-mediated actuator control
- Failsafe behaviors
  - If no progress, relax perceived constraints

Planar Laser Range Scanners

False colored by height

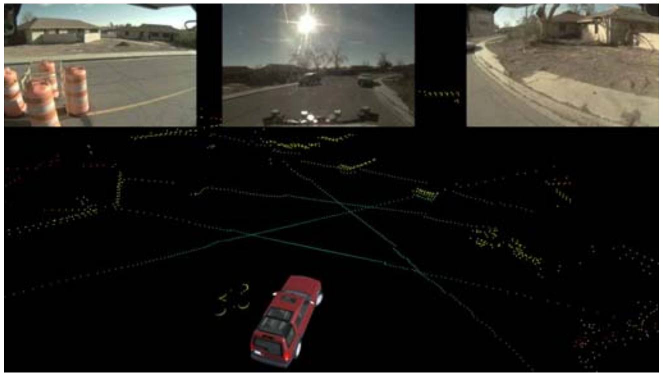

Velodyne Lidar

- 64 lasers, 360° HFOV
- All spinning at 15 Hz
- Vertical FOV spans -24° to +2° w.r.t level
- Redundant (albeit relatively noisy) lidar

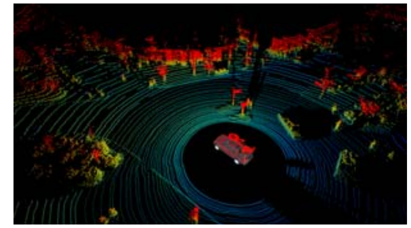

Sample Velodyne Data

False colored by height

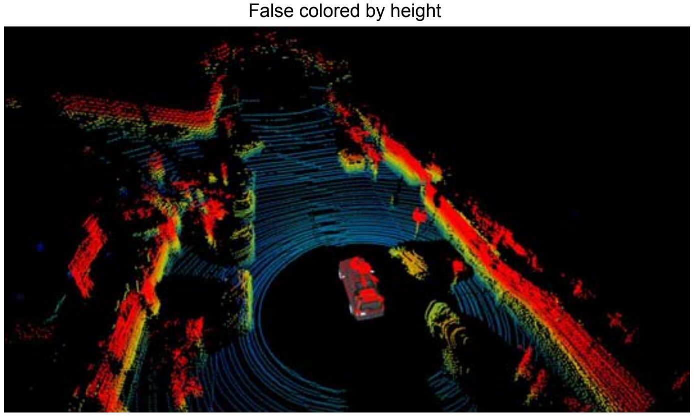

Detection of Static Obstacles

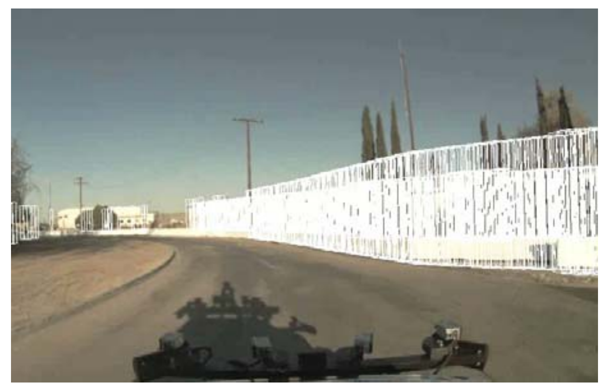

Automotive Radars

- 15 Delphi automotive radars
- Doppler range, bearing, closing speed of 20 objects @ 10Hz
- Narrow beam-width (~18°)
- Good far-field car detectors

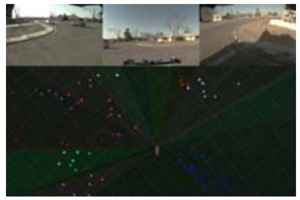

Sample Radar Data

Raw range, bearing, range rate data, false-colored by radar ID

Vehicle Detection and Tracking

Video Cameras

- 5 Firewire Cameras
  - Point Grey Firefly MV
- 720x480 8bpp Bayer pattern @ 22.8 fps
- ~40 MB/s (2.5 GB/min)
  Lots of data!
- Purpose: Detection of painted lane markings

Forward left Forward center Forward right

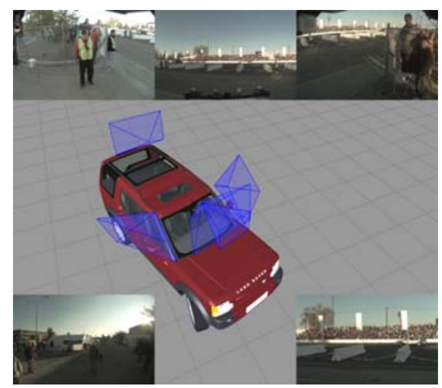

Rear view Narrow forward view

Lane Estimation

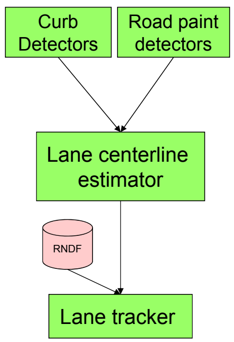

Closed-Loop Lane Tracking

Playback speed: 2x

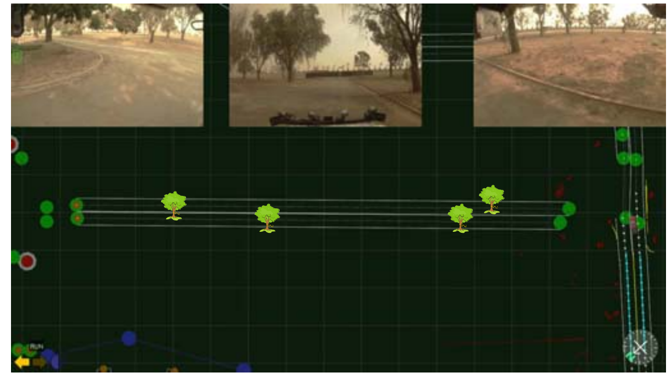

Driving on a one-way, two-lane road
RNDF-interpolated estimate goes through trees and bushes!

System Architecture

- Perception
  - Vehicle surroundings
  - Vehicle location w.r.t. surroundings and RNDF
- Planning & Control
  - Codified driving rules
  - How to reach the goal

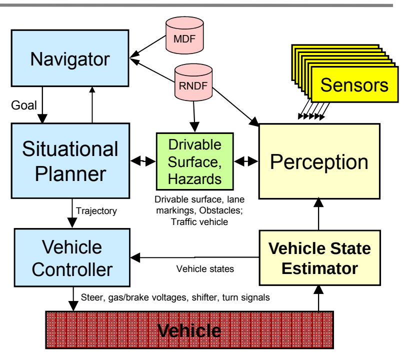

- AEVIT Vehicle Conversion (EMC) control unit
  - Continuous actions (steering, gas/brake)
  - Discrete actions (turn signals, gear shift)

Trajectory Planning

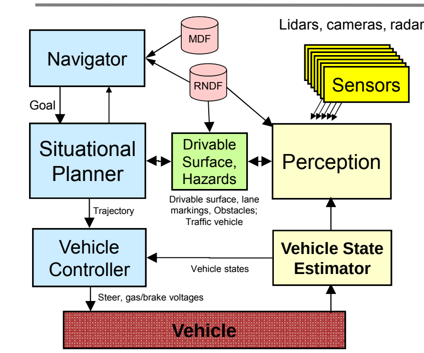

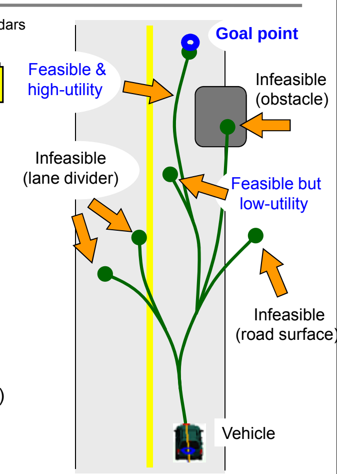

Electrically-controlled drive-by-wire (DBW)
- Continuous elements (steering, gas/brake)
- Discrete actions (turn signals, gear shift)
- Actuation via CAN bus and servomotors

Navigator

- Mission planner
- Sets high-level goals
  - "Carrot" for the motion planner
- Replan around blocked roads
- Knob on constraints in drivability map
  - Perception algorithms are not perfect
  - If car is stuck and isn't making progress, start invoking failsafe levels (by violating "soft" constraints such as lane markings)

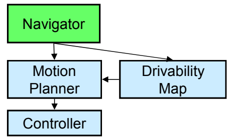

Competition Vehicle: LR3

- Linux blade cluster with two fast interconnection networks
  - 10 blades each with 2.33GHz quad-core processor → 40 cores
  - Approximately 80 driving-related processes steady-state
- Many sensors
  - Applanix IMU/GPS
  - Hi-res odometry
  - 12 SICK Lidars
  - Velodyne (~64 Lidars)
  - 15 automotive radars
  - 5 video cameras
- Roof-mounted AC
- Total power consump-tion was
- Internal gas generator

LR3 Mobile Machine Room

Autonomous Driving Test Site

- South Weymouth Naval Air Station
  - About 40 min. from MIT off 3S
  - Usually $10K/day; free to our team when no paying customer
- Large tarmac area
  - Can create arbitrary (flat) road networks
  - Environmentally sensitive:
    - Obstacles: traffic cones
    - Lane markings: only flour
    - Traffic: team members' cars

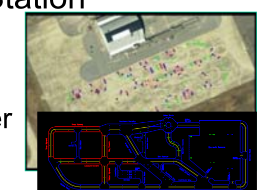

NQE and Competitors (10/07)

Fine-Grained Prior Map Data

- Every other team, to our knowledge, *manually* “densified” the provided RNDF (map) data during the 48-hour pre-competition period
  - Used high-resolution geo-referenced aerial imagery
  - Added precise position and / or curvature samples at dense intervals (every few meters) along roads
- Rationales we heard for this approach:
  - “All the other teams are densifying” (not correct)
  - Such “data infrastructure” will be widely available soon via commercial efforts (e.g. NavTeq, Google)
- DARPA implicitly blessed this strategy
  - 48-hour RNDF distribution; no surprise road segments

NQE Area B

- Our very first NQE run.

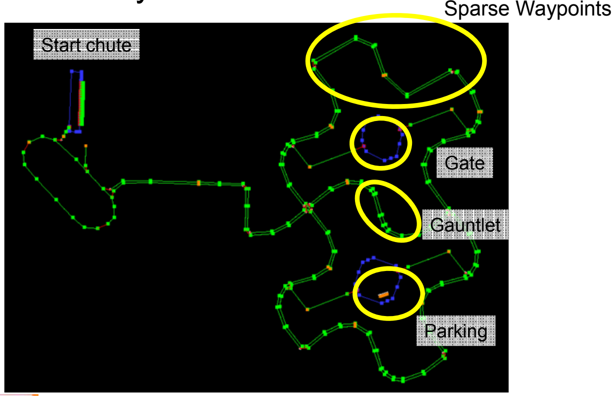

NQE Area B

Playback speed: 3.3x

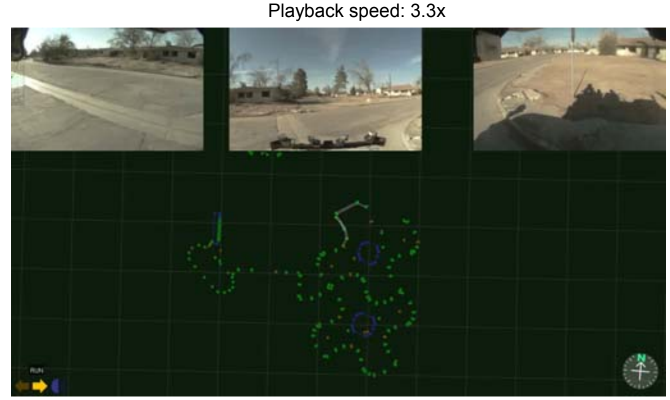

MIT was one of only two teams to complete Area B on the first attempt.

Area B Parking Test

- Parking: RNDF target position was blocked

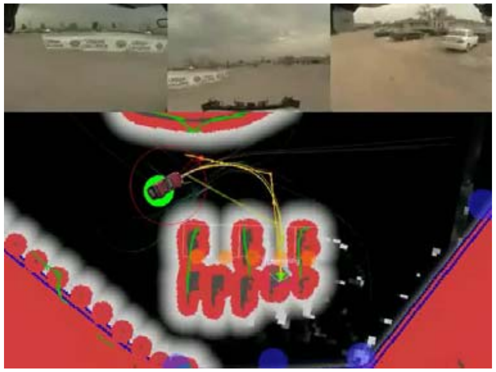

NQE Area A

- Advanced traffic capabilities
  - Merging into traffic
  - Left turn across oncoming traffic
  - Excessive delay (> 10 sec.) prohibited
- ~10 traffic vehicles moving at 10mph.

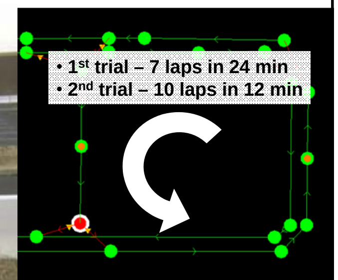

NQE Area C

- Objectives
  - Intersection precedence (turn-taking)
  - Blocked check points (replanning, 3-point turns)

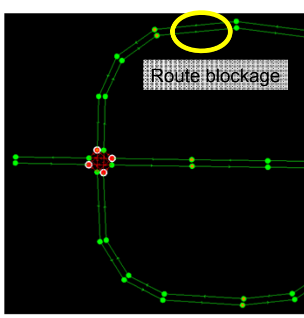

Urban Challenge Event (11/07)

- DARPA selected only 11 of 35 teams, due to safety concerns
- 50 human-driven traffic vehicles
- 3 MDF missions totaling ~60 miles

Change in DARPA's Emphasis

- Focus very different from specification, NQE
  - Allowed human pre-inspection of road network
    - Opportunity to validate manually densified RNDFs
  - Simplified competition setup
    - No DARPA-generated road blockages
    - No passing or merging at speed
    - No dense human traffic or planned human challenges
    - No other vehicles (parked or moving) in parking lots
  - No GPS degradation or outages
    - RF sources turned off (plasma display, chase video)
- **Other robot cars**
  - Highly unpredictable

Starting Line

Degraded Performance

- Phantom curbs on the dirt segment

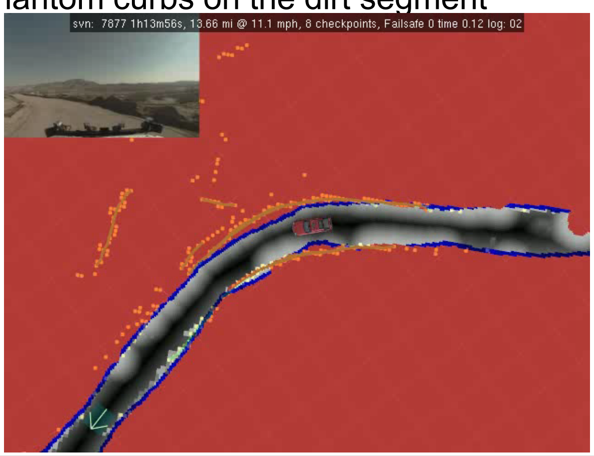

Accident with CarOLO

- Accident had several contributory causes:
  - CarOLO drove into us; it was damaged, removed from competition
  - Our moving-object detector didn't integrate over long time scales
  - Our evasive maneuver was poor (we should have stayed right)

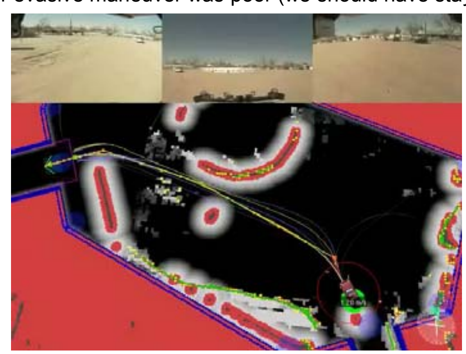

The first bot-on-bot car accident in history!

Accident with Cornell

- Cornell vehicle:
  - Stopped, then reversed toward the intersection
  - Started moving as we passed
- MIT vehicle:
  - Tried passing to Cornell's left
  - Returned to right lane too quickly

➔ DARPA: ruled it a "no fault" incident, then allowed both teams to continue

The second bot-on-bot car accident in history!

"Traffic Jam" Rule

- If any car fails to move within 10 seconds, "traffic jam" is declared

Legend:

PX = position occupied
PC = position clear
PU = position unknown
IX = intersection occupied
IC = intersection clear
IU = intersection unknown

"High-Speed" Section

- MDF speed limit: 30mph
  - Braking distance = 36m (with 2.5m/s^2 deceleration)
  - Standoff distance = 10m
  - Requires reliable detection range: 50m

➔ Capped at 25mph by our software

Competition Results

- 6 teams finished; other 5 removed from competition by DARPA officials due to:
  - At-fault collisions
  - Near misses
  - Excessive delay

CMU **Stanford** **Virginia Tech**

- Many race-day firsts for us:
  - More than 20 miles in one day
  - Steep dirt (unpaved) segment
  - Mile-long, wide lanes at 25mph
  - Interaction with other robots
- We drove safely
  - No processes died
  - Our chase driver: "your vehicle was always safe, in my opinion"

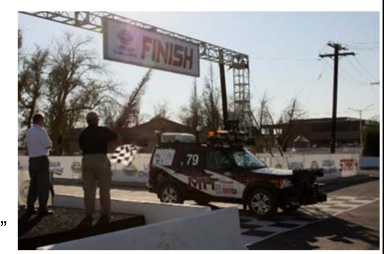

Failure Modes

- Perception limitations
  - Hallucinated curbs (at detection size threshold)
  - Vulnerability to shadows, sun blinding
  - Sensitivity of lidar interpretation to vehicle pitch
  - Inability to track slow-moving objects (< 3mph)
- Control / planning limitations
  - Occasionally failed to achieve target orientation
  - Caused over-correction, unsafe maneuvers
- Failsafe strategy
  - Unclear whether to observe or relax constraints
  - Example: U-turn at roadblock, or drive around?

Teamwork!

Achievements

- Respectable rookie showing
  - First time in DGC for MIT team
- Fourth place overall
  - One of only 6 teams (of 89 initially entering) to complete UCE course
- Completed all NQE missions without manual annotation of provided RNDF

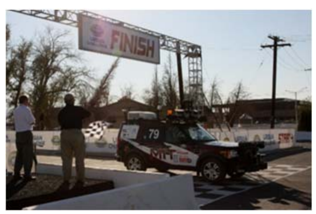

Lessons Learned

- About competing effectively in the real world
- About DARPA's expectations
- About the autonomous driving task
- About differing long-term approaches

More info: http://dgc.mit.edu