Comp 790-058 : Multi-Agent Simulation for Crowds & Autonomous Driving

Sahil Narang & Andrew Best

August 22, 2017

University of North Carolina at Chapel Hill

Multi-Agent Simulation

- Multiple robots in shared environments

Kiva Systems

Multi-Agent Simulation

- Multi-agent simulation in entertainment

Multi-Agent Simulation

- Multi-agent simulation as biological entities

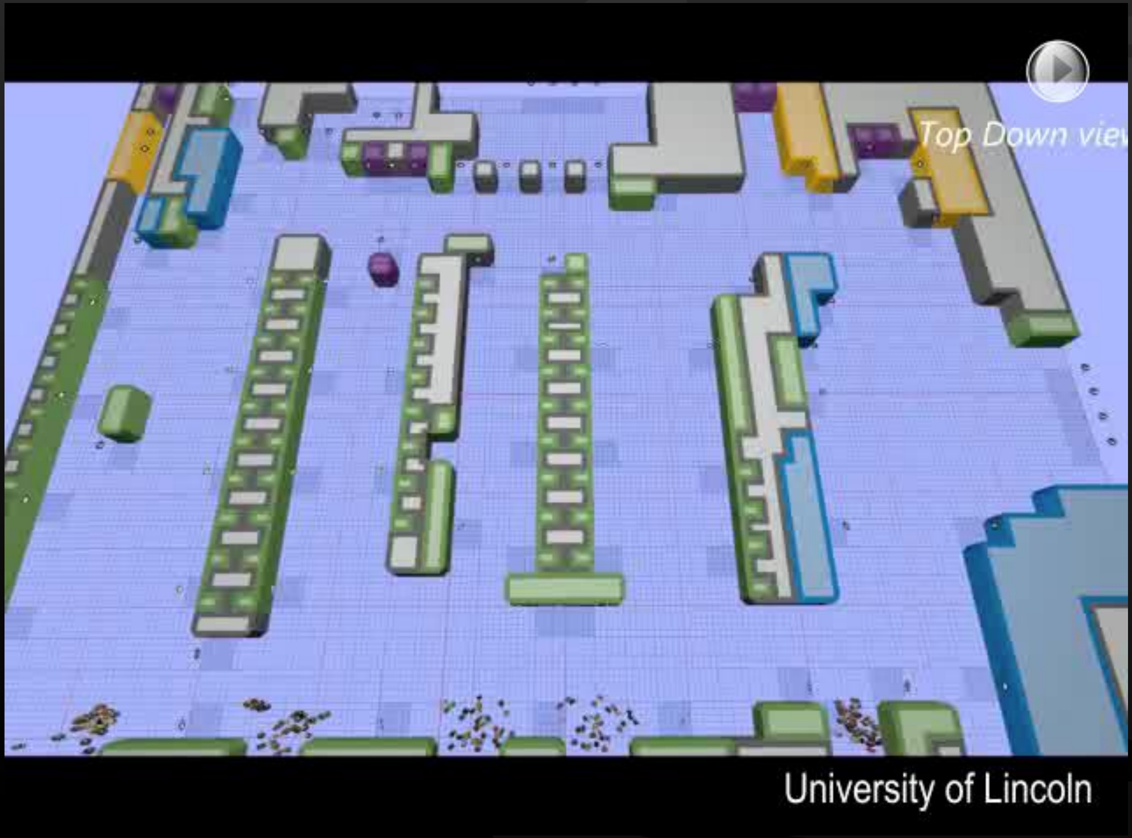

Structure

- Introduction
- Course details
- Background
- Multi-agent simulation
  - Crowd simulation
  - Pedestrian tracking
- Autonomous Driving

Multi-Agent Simulation, Crowds and Autonomous Driving

- COMP 790-058 (Fall 2017)
  - Tue 11-1:30 in SN 115
- Instructor: Dinesh Manocha (dm@cs.unc.edu)
- Co-instructors:
  - Aniket Bera
  - Andrew Best
  - Sahil Narang
- Website
  - http://gamma.cs.unc.edu/courses/planning-f17/

What is this course about?

- Underlying geometric concepts of motion planning
  - Configuration space
- Character motion in virtual environments
- Multi-agent and Crowd simulation
- Autonomous driving navigation and coordination
- Local and global collision avoidance
- Pedestrian tracking and path prediction

Do I have the right background?

- Undergraduate algorithms course
- Exposure to geometric concepts
- Basic physics and dynamics
- Willingness to read about new concepts and applications!

Course Load & Grading

- 3-4 assignments (30%)
  - Geometric concepts (problems)
  - Multi-agent simulation: programming assignments
  - Autonomous driving: problems and programming
- Class participation and a lecture (20%)
  - Lecture topic (consult the instructor)
- Course Project (45%)

Course Project

- Any topic related to multi-agent simulation, crowds, and autonomous driving
- Must have some novelty to it!
- Can work by yourself or in small groups (2-3)
- Can combine with course projects in other courses
- Start thinking now of possible course project

Course Project Schedule

- Project topic proposal (October 03)
- Monthly updates
- Mid semester project update (early November)
- Final project presentation (During the finals week)
- Scope for extra credit + publications!

Course Schedule (Tentative)

- August 22, 2017: Course Introduction and Overview (Andrew and Sahil)
- August 29, 2017: Graph Searches and Global Navigation (Dinesh)
- Sep. 05, 2017: Local Navigation Methods (Dinesh)
- Sep. 12, 2017: High-DOF Motion Planning & Configuration Spaces (Dinesh)
- Sep. 19, 2017: Overview of Autonomous Driving (Andrew and Sahil)
- Sep. 26, 2017: Autonomous Driving: Dynamics and Navigation (Andrew and Sahil)
- Oct. 03, 2017: Project Proposals

Course Schedule (Tentative)

- Oct. 10, 2017: Pedestrian Tracking and vision methods (Aniket)
- Oct. 17, 2017: Path Prediction and Anomaly Detection (Aniket)
- Oct. 24, 2017: Autonomous Driving Perception (Andrew and Sahil)
- Oct. 31: Student lectures
- Nov. 07: Student Lectures
- Nov. 14: Project Update
- Nov. 21: Student Lectures
- Nov. 28: Student Lectures
- Dec. 05: Course Wrapup

Structure

- Introduction
- Course details
- Background
- Multi-agent simulation
  - Crowd simulation
  - Pedestrian tracking
- Autonomous Driving

Robotic Paradigm : Primitives

- Sense
  - Takes raw data from sensors and produces information
- Plan
  - Takes information and produces tasks
- Act
  - Functional components which carry out the task

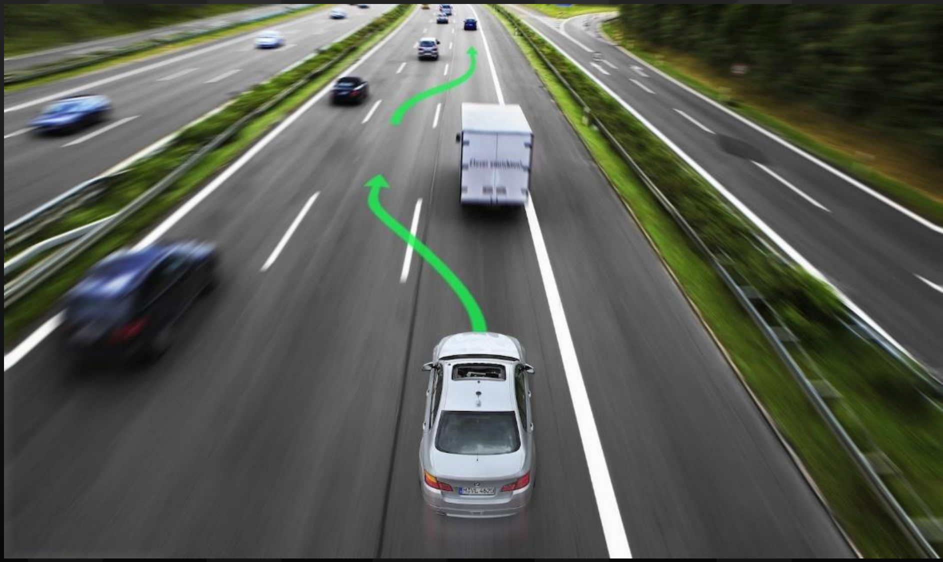

Robotic Paradigm : Primitives

- Sense
  - Gather noisy data from various sensors
  - Fuse data into a consistent model
  - Perception: semantic understanding of the world

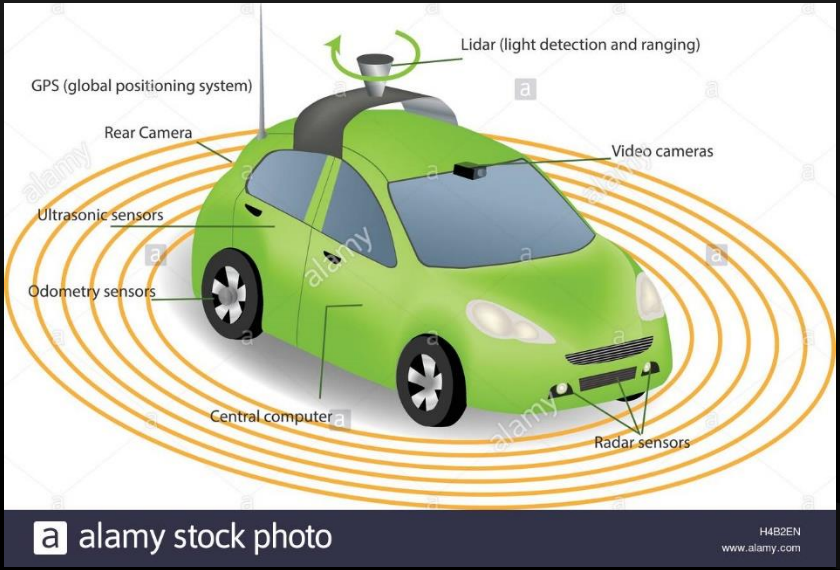

Robotic Paradigm : Primitives

- Plan
  - Different abstractions of planning
  - Higher abstraction: Knowledge based reasoning
    - "Find someone who knows about P"
    - "Go to position B"
  - Lower abstraction: Motion planning
    - Given the current setting of the robot, find a valid or optimal trajectory for the robot to reach goal B
      - Collision-free
      - Other constraints: Dynamic/ kinematic feasibility
      - Optimality criterion: shortest path, min-time, smooth etc.

Robotic Paradigm : Primitives

- Act
  - Sequence of actuator commands
  - Realizing the generated plan
  - Generates the actual motion of the robot/agent

Hierarchical Paradigm

- Traditional Paradigm
- Powerful approach for “deliberative” and complex planning

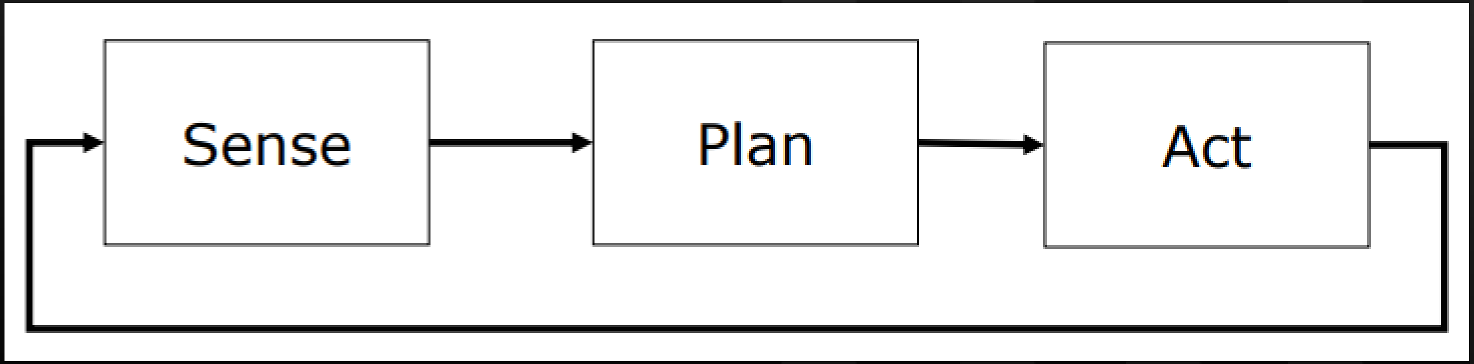

Hierarchical Paradigm

- Limitations
  - Knowledge representation
    - Closed world assumption
    - Size of the state space can explode
  - Planning can be expensive
  - No reactivity

Reactive Paradigm

- No world model; no planning
- Maps sensor input to actuator output
- Very “reactive” to sensor readings

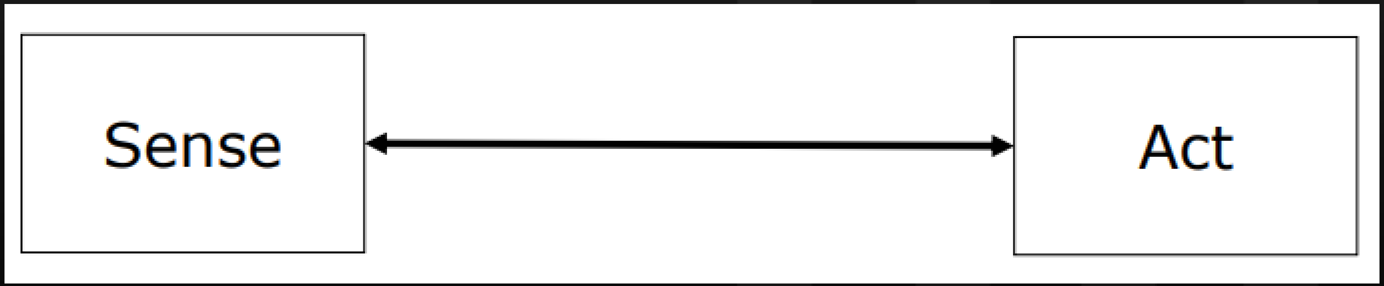

Other paradigms

- Hybrid Heirarchial / Reactive Paradigms
  - Reactive functions for low level control
  - Deliberation for higher level tasks

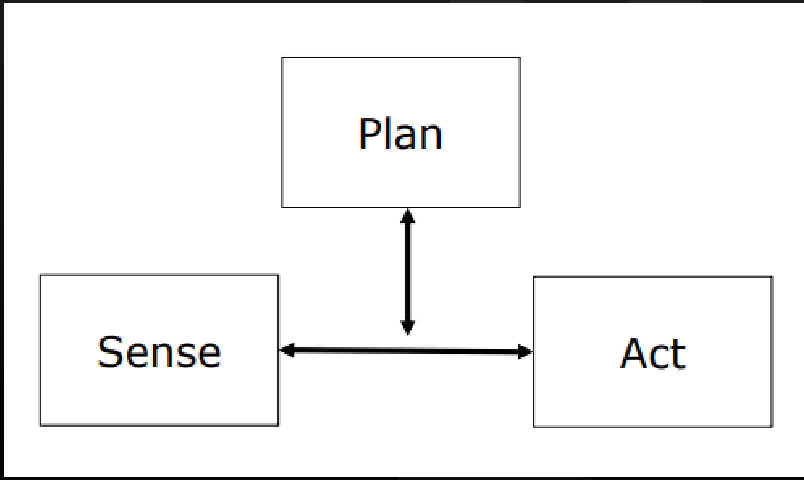

Problems to consider

- Moving obstacles
- Multiple agents
- Complex environments
- Goal is to acquire information by sensing
- Nonholonomic constraints
- Dynamic constraints
- Stability constraints
- Optimal planning
- Uncertainty in model, control and sensing
- Exploiting task mechanics (under-actuated systems)
- Integration of planning and control
- Integration with higher-level planning

Problems to consider in simulation

- Accuracy
  - Reflect real world conditions
  - Results should be transferrable to the real world

- Efficiency
  - Cost of a single timestep
  - Stability: ability to take large time steps

- Robustness

Structure

- Introduction
- Course details
- Background
- Multi-agent simulation
    - Crowd simulation
    - Pedestrian tracking
- Autonomous Driving

Multi-agent simulation

- Study of agents planning in a shared environment
- Environment
  - Static and Dynamic obstacles
- Goals
  - Generate optimal and feasible plans for all agents with respect to give constraints.
- Complexity
  - Linear in the number of robots
  - Exponential in the dimensionality of the configuration space

Multi-agent simulation

- Centralized vs Distributed Planning
  - Centralized
    - Planning is centralized, execution is distributed
  - Distributed
    - Both planning and execution are distributed

Multi-agent simulation

- Coordinated vs Independent Planning
  - Coordinated
    - Explicit communication and coordination between agents
  - Independent
    - Implicit communication (observations) and no explicit coordination between agents

Crowd Simulation

- Study of how pedestrians flow through a shared environment
- Goals:
  - Understanding Human Crowd Behavior
  - Predicting / Replicating pedestrian behavior
  - Design and Plan with Pedestrians in mind
- Multiple approaches
  - Agent Based (Distributed and Independent)
  - Fluid-Dynamic or Continuum (Centralized)
  - Event Based

Crowd Simulation

- Agents have:
  - Independent sensing
  - Independent Goals
  - Independent Planning
  - No implicit Communication
- Modeling pedestrians
  - Simple 2D shapes: circles (or ellipses)
  - Some high level constraints to generate human-like motion
    - Range of motion, dynamic stability, limb acceleration etc

Crowd Simulation Framework: Menge

- Menge is a modular, pluggable framework for crowd simulation developed at UNC.
- Menge is Open-Source and publicly available.
- Pluggable components:
  - Behaviors
    - State transitions
  - High level planning: goal selection
  - Motion planning
- Easy to create and simulate complex scenarios with 1000’s of agents.

Menge: Applications

- Modeling physiological and psychological factors that effect density in crowds

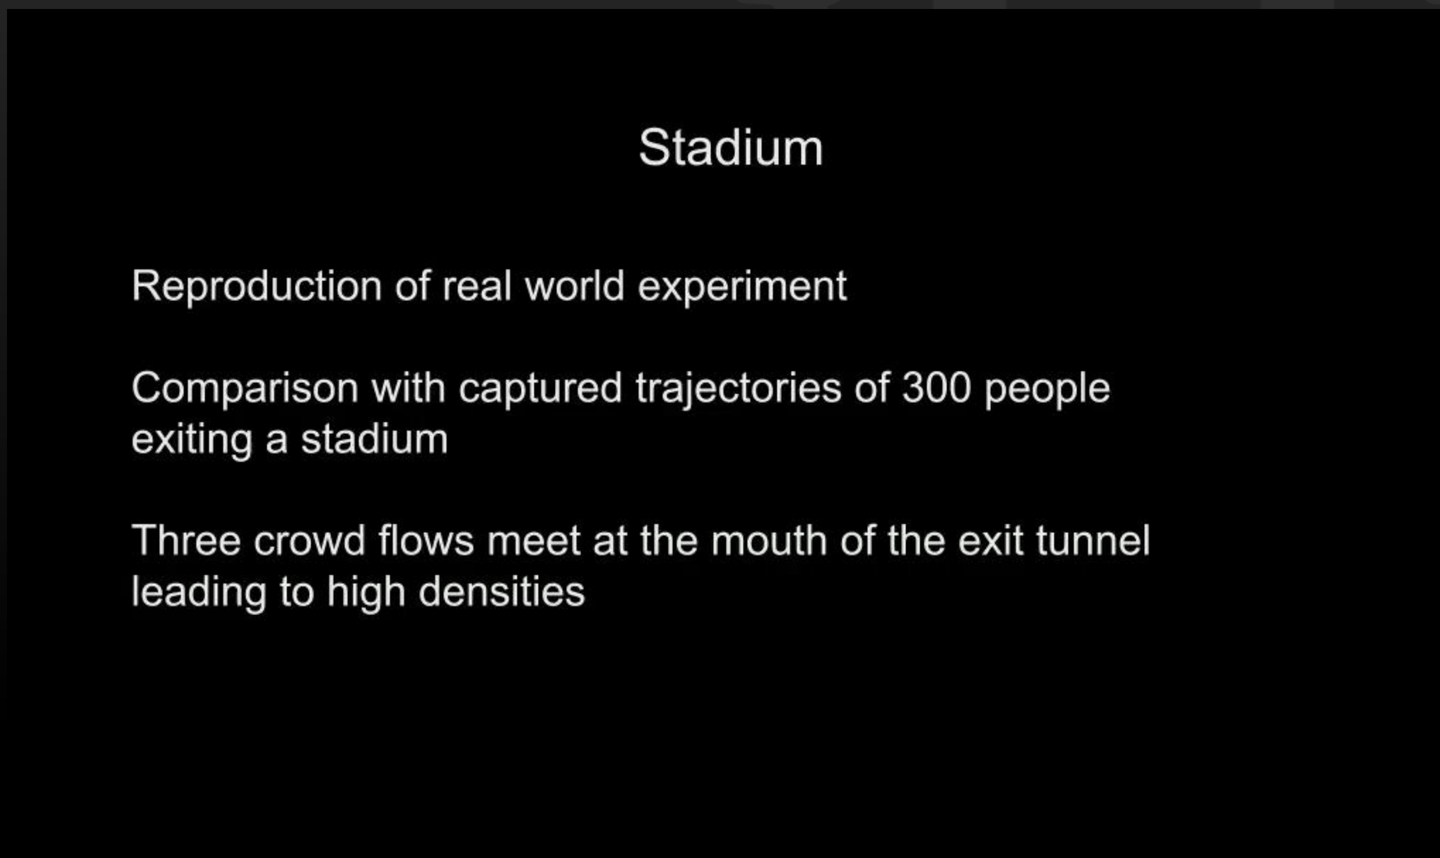

Menge: Applications

- Loading a Boeing aircraft

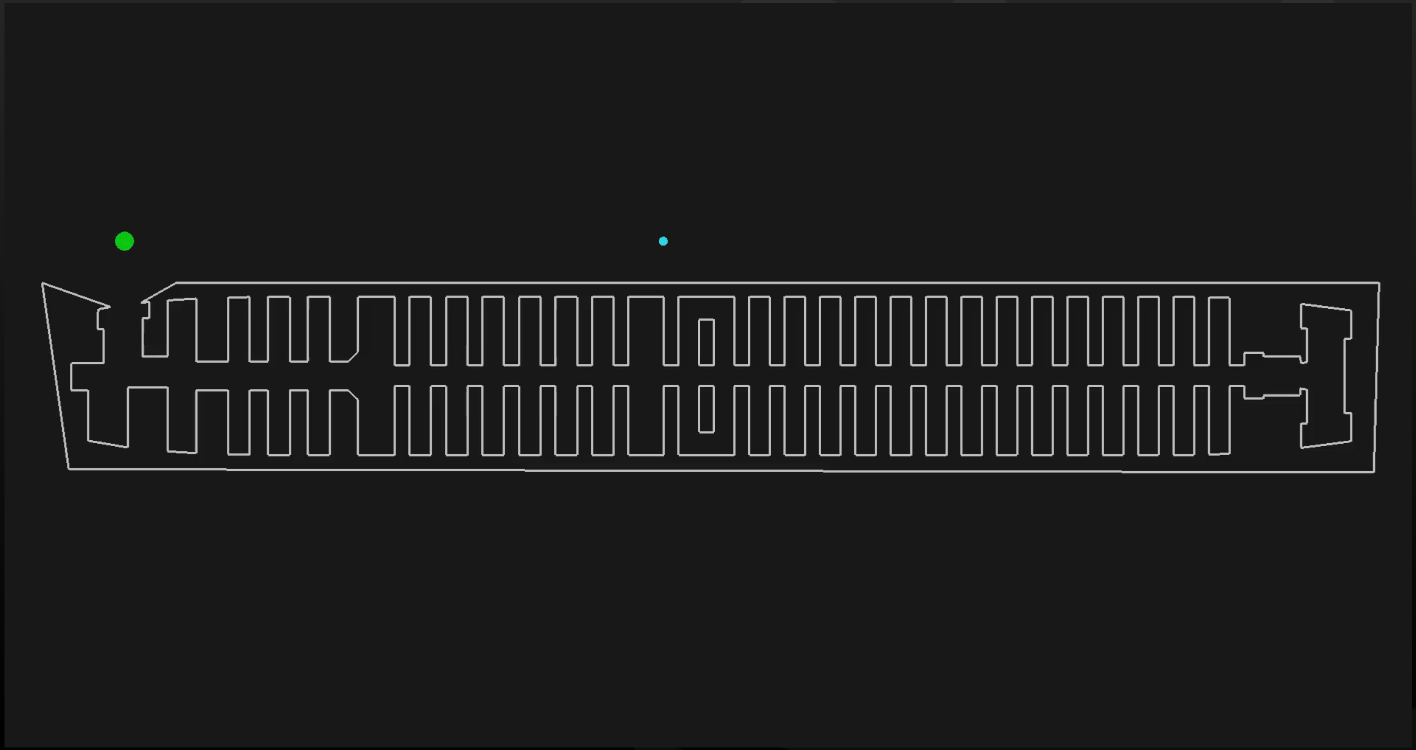

Menge: Applications

- Unloading a Boeing aircraft

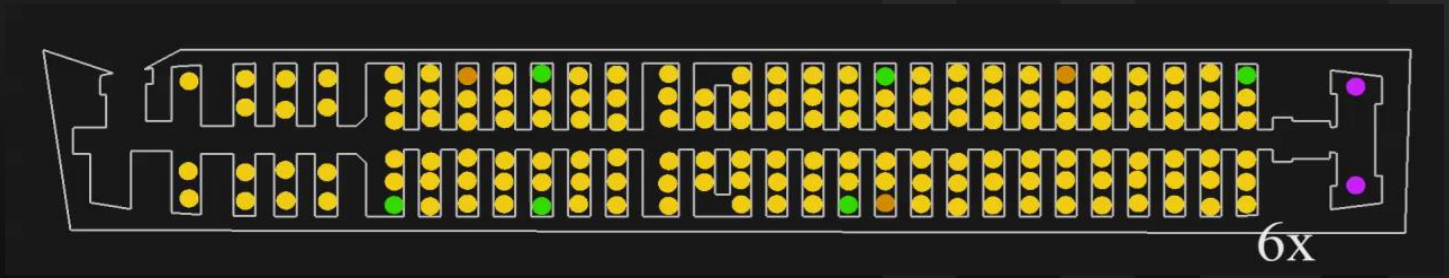

Menge: Applications

- Modeling human motion constraints

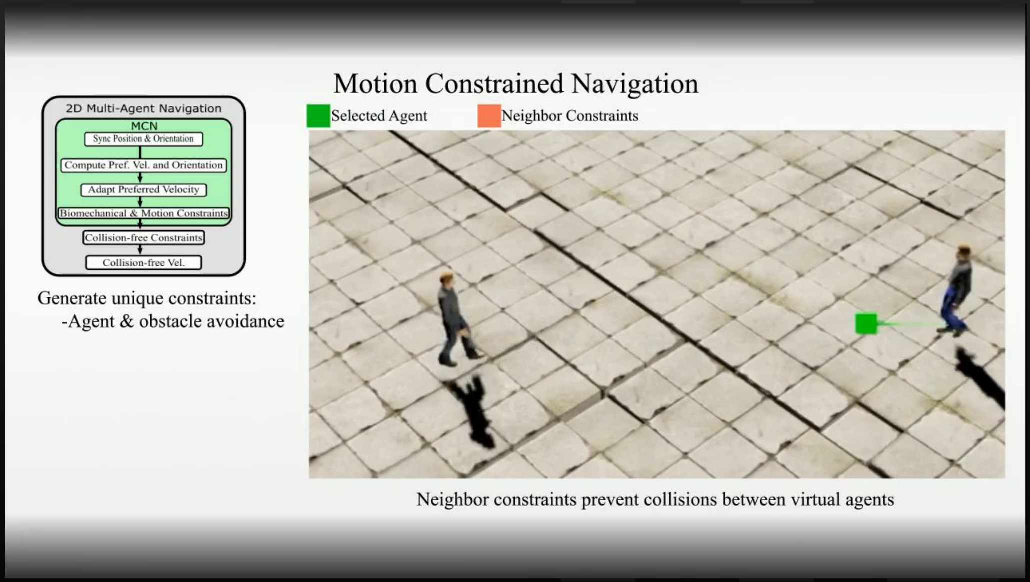

Menge: Applications

- User – agent interactions in VR

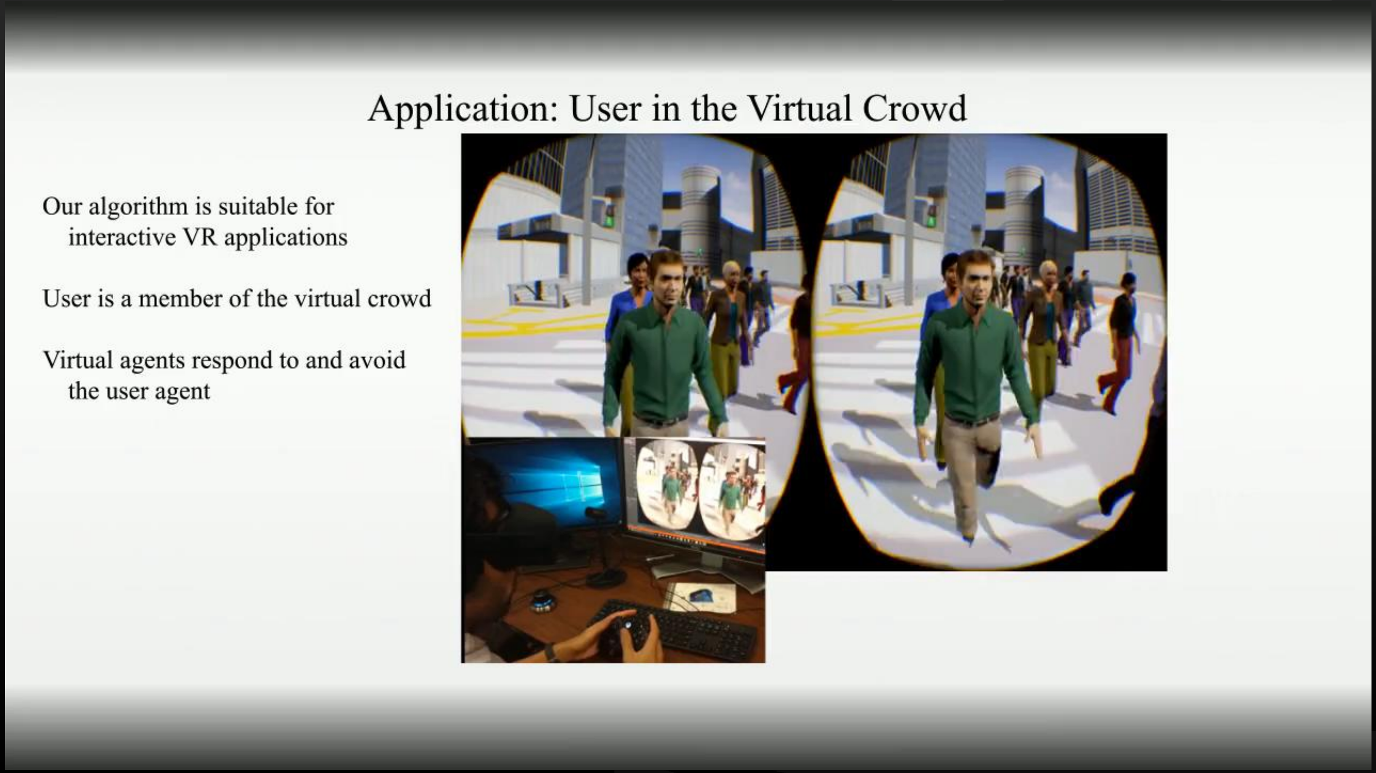

Structure

- Introduction
- Course details
- Background
- Multi-agent simulation
  - Crowd simulation
  - Pedestrian tracking
- Autonomous Driving

Pedestrian Tracking

- Locating a pedestrian (or pedestrians) along a window of time in a video.
- Tracking corresponds to computing the projected trajectory on a 2D plane assuming that the pedestrian is represented as a small circle.

NPLC-2

Medium Density, MOTP - 71.28%

Pedestrian Tracking

Stable multi-target tracking in real-time surveillance video – *Benfold et al. (2011)*

Tracking with Local Spatio-Temporal Motion Patterns in Extremely Crowded Scenes - *Kratz et al. (2012)*

People tracking with human motion predictions from social forces - *Luber et al. (2010)*

Tracking multiple people using laser and vision– *Cui et al. (2005)*

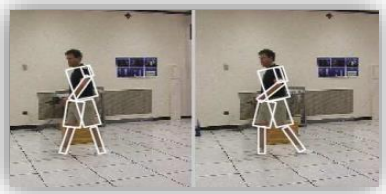

Tracking people by learning their appearance – *Ramanan et al. (2007)*

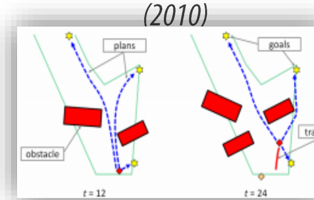

Multi-hypothesis motion planning for visual object tracking – *Gong et al. (2011)*

+ Realtime
- Low density

+ Dense videos
- Offline

+ Realtime
- Poor accuracy

Pedestrian Prediction

- Determining future pedestrian positions and velocities based on past data.
- Short term prediction as future pedestrian positions for 1–2 seconds and long term prediction as future positions for 5 or more seconds.

Pedestrian Prediction

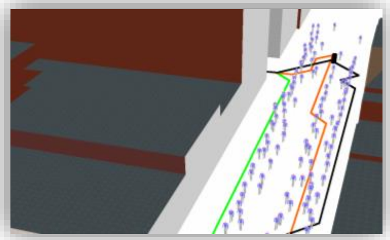

Learning to Navigate Through Crowded Environments – *Henry et al. (2010)*

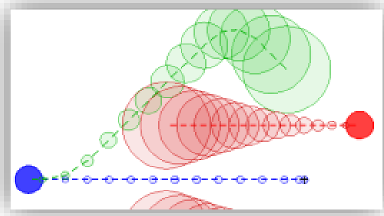

Robotic motion planning in dynamic, cluttered, uncertain environments – *Toit et al. (2010)*

+ Realtime
- Low density

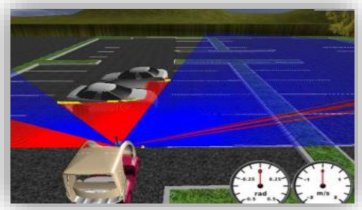

Dynamic obstacle avoidance in uncertain environment combining pvos and occupancy grid – *Fulgenzi et al. (2007)*

Learning behavior patterns from video – *Zhong et al. (2015)*

+ Accurate
- Costly

Trajectory Analysis and Prediction for improved Pedestrian Safety – *Møgelmose et al. (2015)*

Feature-based prediction of trajectories for socially compliant navigation– *Kuderer et al. (2012)*

+ Realtime
- Scene dependent

Pedestrian Behavior Learning

- We compute personality personalities based on based on Eysenck Personality Theory, a well-known psychology trait theory work.

Pedestrian Behavior Learning

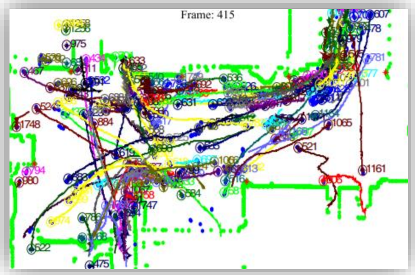

A Fully Online and Unsupervised System for Large and High Density Area Surveillance – *Song et al. (2013)*

+ Realtime/Online
- Low density

Coherent filtering: Detecting coherent motions from crowd clutters – *Zhou et al. (2013)*

+ Dense crowds
- Offline

Identifying behaviors in crowd scenes using stability analysis for dynamical systems– *Solmaz et al. (2012)*

+ Online
- Specific patterns

Crowd Density

Low Density
(<1 pedestrians/m^2)

Medium Density
(1-3 pedestrians/m^2)

High Density
(>3 pedestrians/m^2)

Pedestrian Tracking - Challenges

- Change in illumination
- Change in appearance
- From certain camera angles, pedestrians look alike
- Occlusions
- Rapid change in velocity

Long-term Pedestrian Path Prediction

Path Prediction - Issues

- Most prior work limited to local interactions between pedestrians.
- Long-term predictions prone to error.
- Scene specific and limited to pre-learnt behaviors.

Behavior Learning - Challenges

- Most prior work on behavior learning is offline.
- No prior work on automatically classifying pedestrian personality.

Personality

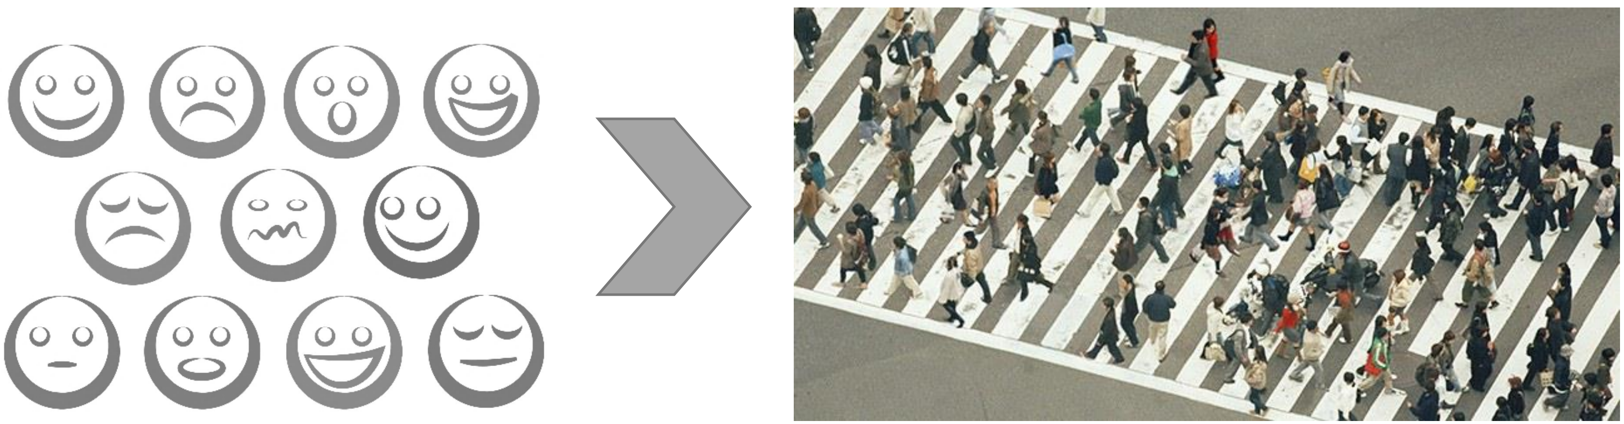

Crowd behavior is not a sole product of the crowd itself; rather, it is defined by the individual pedestrians in that crowd.

Personality Traits

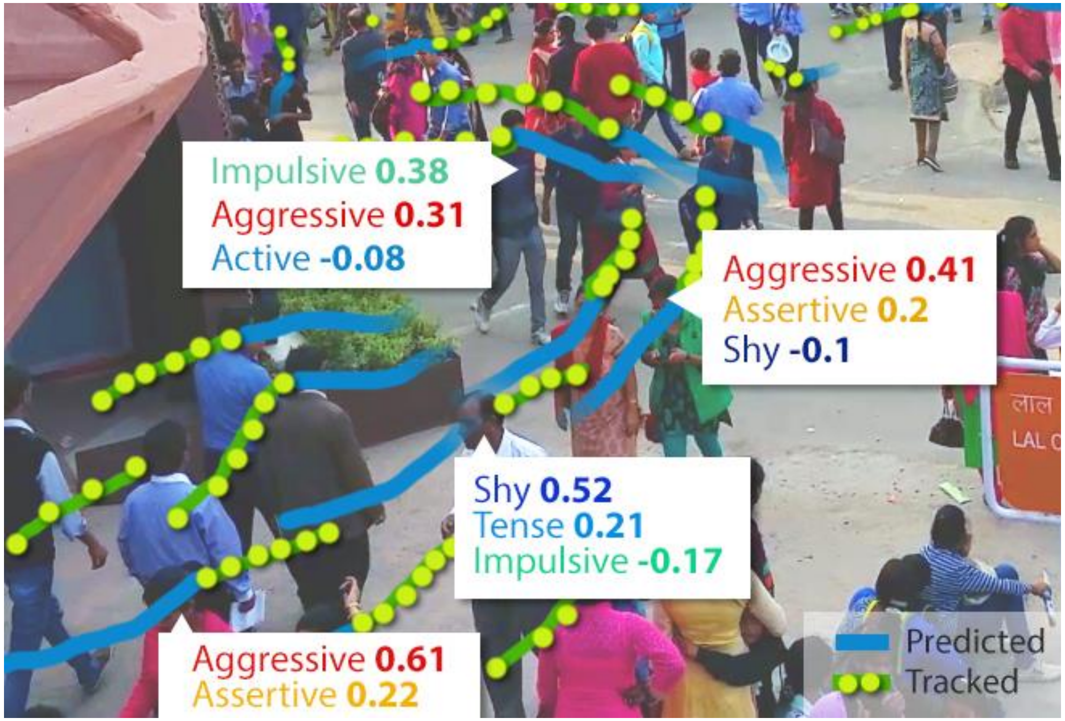

Video: International Trade Fair, New Delhi 2016

Anomaly Detection - Issues

- Most prior work offline.
- Requires precomputation and apriori learning.
- Limited to sparse crowds.

Overview

Overview

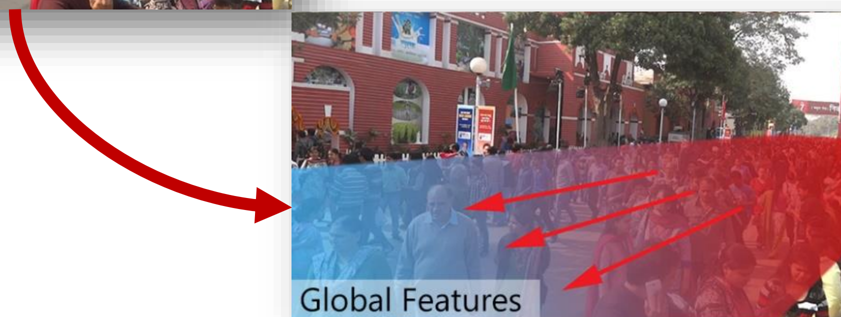

Overview

GLMP

Realtime Pedestrian Path Prediction using Global and Local Movement Patterns

ICRA 2016 Submission

Supplementary Video

Structure

- Introduction
- Course details
- Background
- Multi-agent simulation
  - Crowd simulation
  - Pedestrian tracking
- Autonomous Driving

Autonomous Driving

- Autonomous vehicle:

Tesla

Waymo

Nutonomy

Autonomous Driving

- Autonomous vehicle: a motor vehicle that uses artificial intelligence, sensors and global positioning system coordinates to drive itself without the active intervention of a human operator
- Focus of enormous investment [$1b+ in 2015]

Tesla

Waymo

Nutonomy

Autonomous Driving

- Levels of Autonomy
  - 0: Standard Car
  - 1: Assist in some part of driving
    - Cruise control
  - 2: Perform some part of driving
    - Adaptive CC + lane keeping
  - 3: Self-driving under ideal conditions
    - Human must remain fully aware
  - 4: Self-driving under near-ideal conditions
    - Human need not remain constantly aware
  - 5: Outperforms human in all circumstances

Autonomous Driving

- Cutting Edge of numerous disciplines
  - Robotics
  - Sensor and signal analysis
  - Computer-vision
  - Motion-planning
  - Human-factors psychology
  - Civil engineering
  - Digital Ethics
  - Economics

Autonomous Driving Challenges

- Recall primitive: Sense, Plan, Act
- Sensing Challenges
  - Sensor Uncertainty
  - Sensor Configuration
  - Weather / Environment

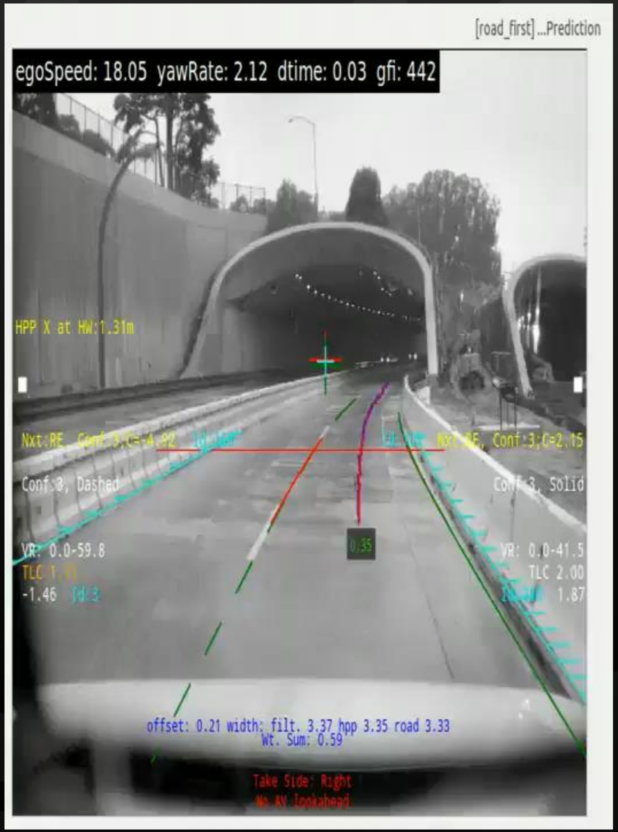

Autonomous Driving Challenges

- Sensor Misclassification
  - "When is a cyclist not a cyclist?"
  - When is a sign a stop sign?
  - Whether a semi or a cloud?

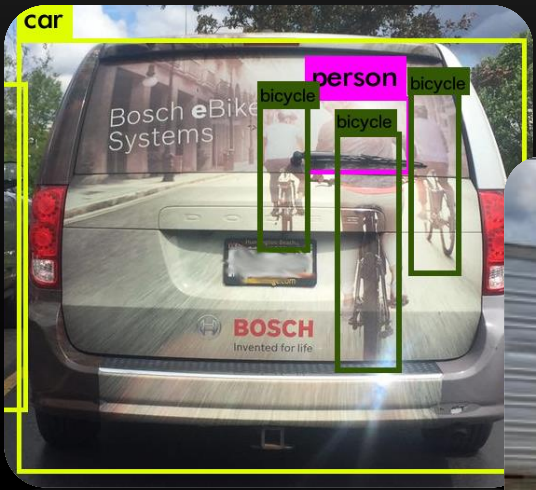

Autonomous Driving

- Planning challenges
  - Behavior of others
  - Reliance on Implicit knowledge / norms
  - Weather / Environment

Autonomous Driving

- Behavior of others
  - Humans are notoriously hard to predict
  - Cyclists operate as vehicles and pedestrians

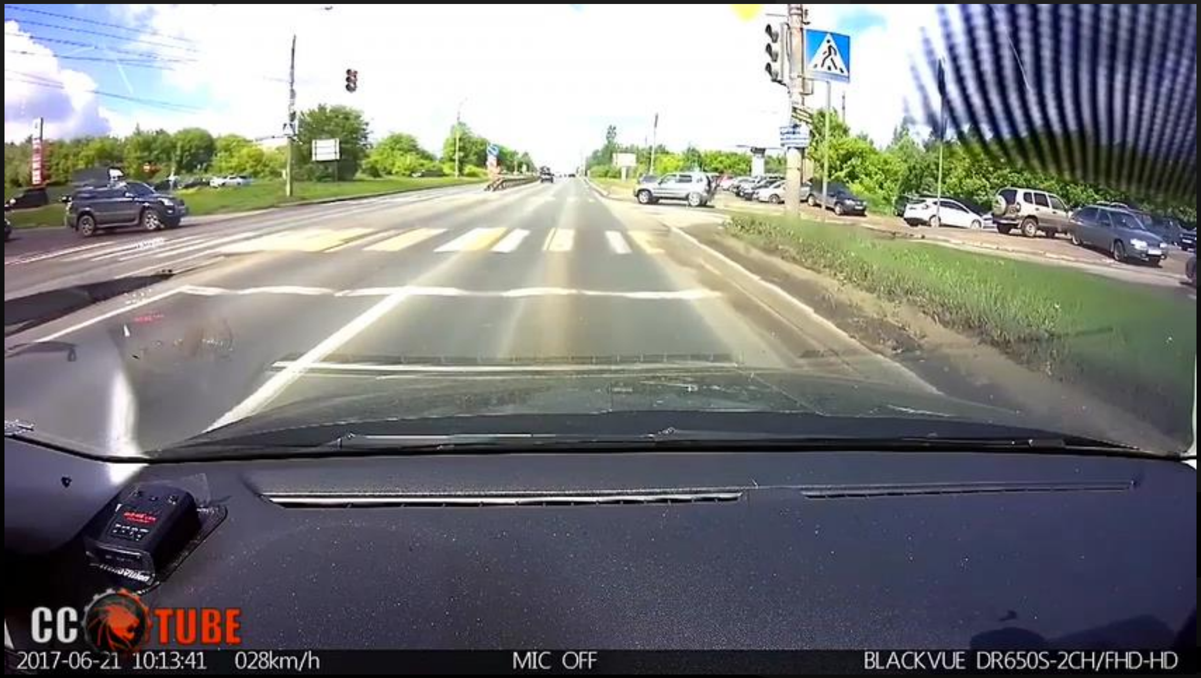

Autonomous Driving

- "Act" challenges
  - Vehicle dynamics complex and uncertain
  - Weather / Environment!

Autonomous Driving

- Vehicle Dynamics modelling
  - Tire properties change with speed
    - Traction
    - Pressure
    - *Shape*
  - Tread level difficult to predict
  - Forward simulation expensive considering forces
    - Load transfer
    - Slip equations

Autonomous Driving

- Other challenges:
  - Communication
  - Coordination
  - Ethical Issues
    - Trolley Problem

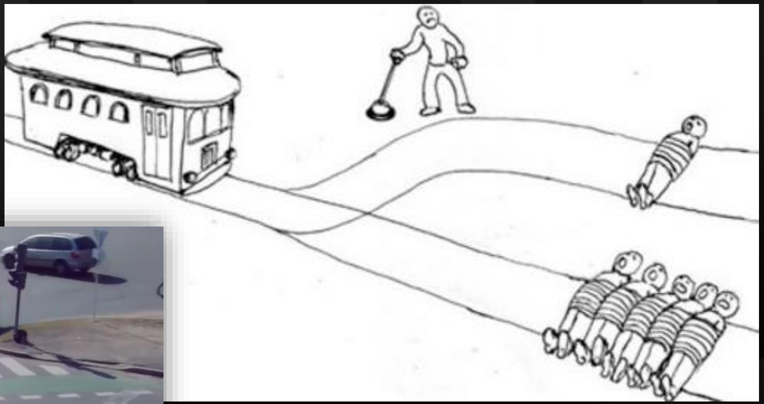

Autonomous Driving

- Other challenges:
  - MIT “Moral Machine” [https://goo.gl/RL4pr5]

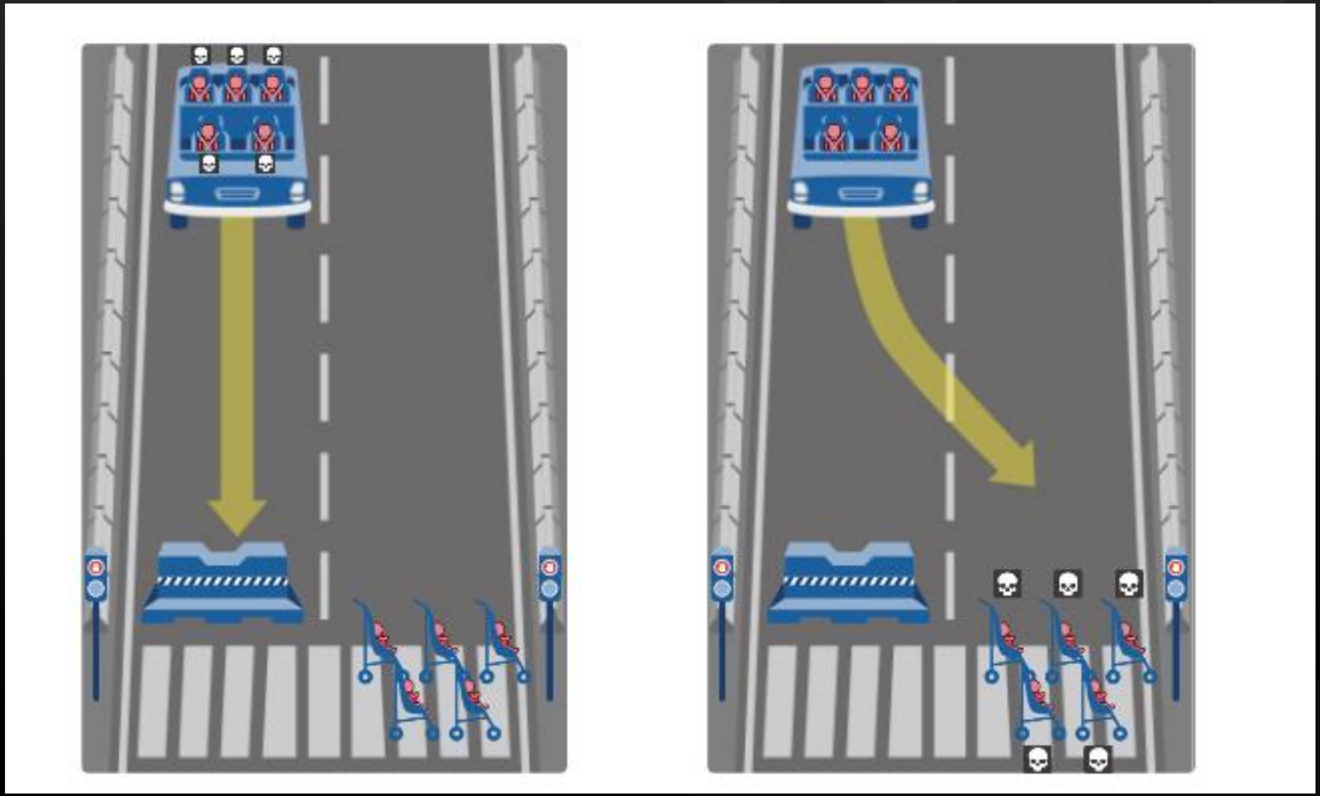

MIT Moral Machine

Autonomous Driving

- Civil Engineering / Ethics
  - Traffic impacts?
    - Pro: Vehicles should respond appropriately to traffic reducing jams
    - Con: Many more vehicles per person possible
  - People may not own cars?
    - Pro: Less emission? Less Traffic?
    - Con: Less access?

Autonomous Driving SOA

- Lidar Visualization

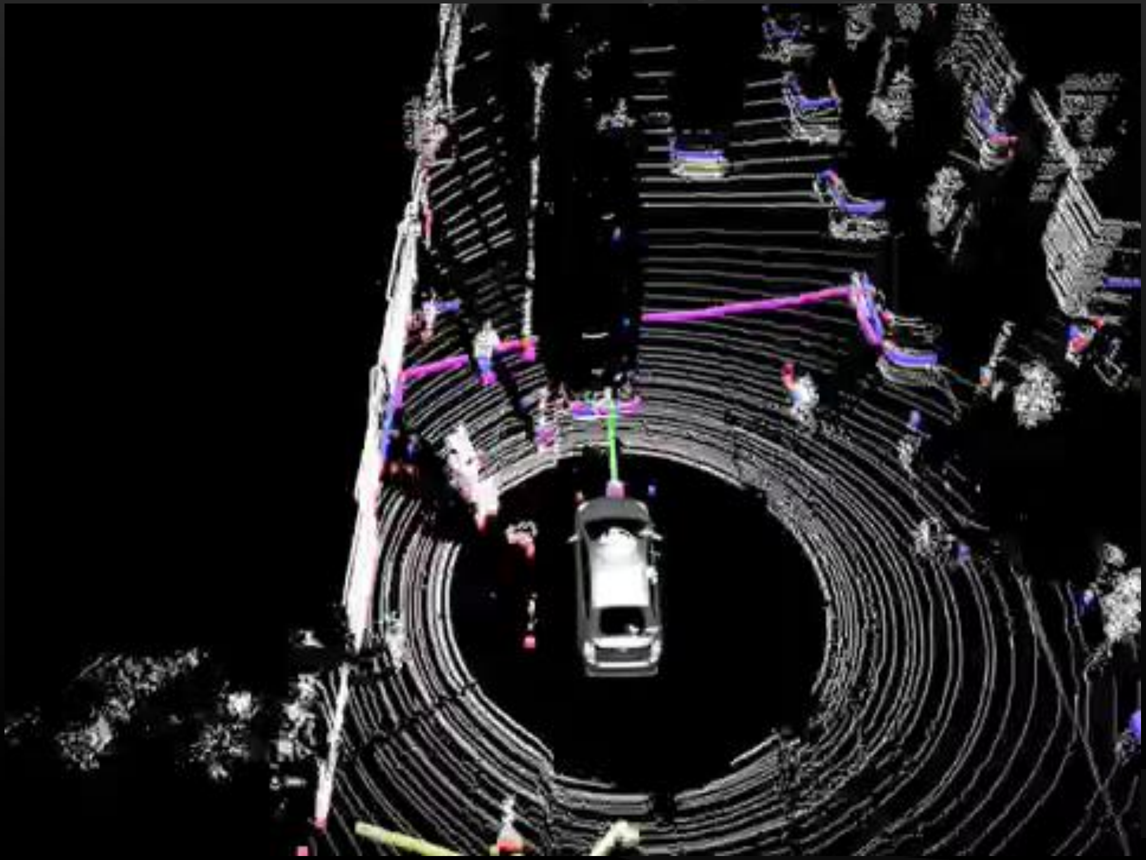

Autonomous Driving SOA

- CMU Boss

Autonomous Driving SOA

- Waymo

Autonomous Driving SOA

- Multiple approaches demonstrated
- Nvidia Pilotnet

Pilotnet

Autonomous Driving SOA

- AutonoVi-Sim

Jaywalking Pedestrian

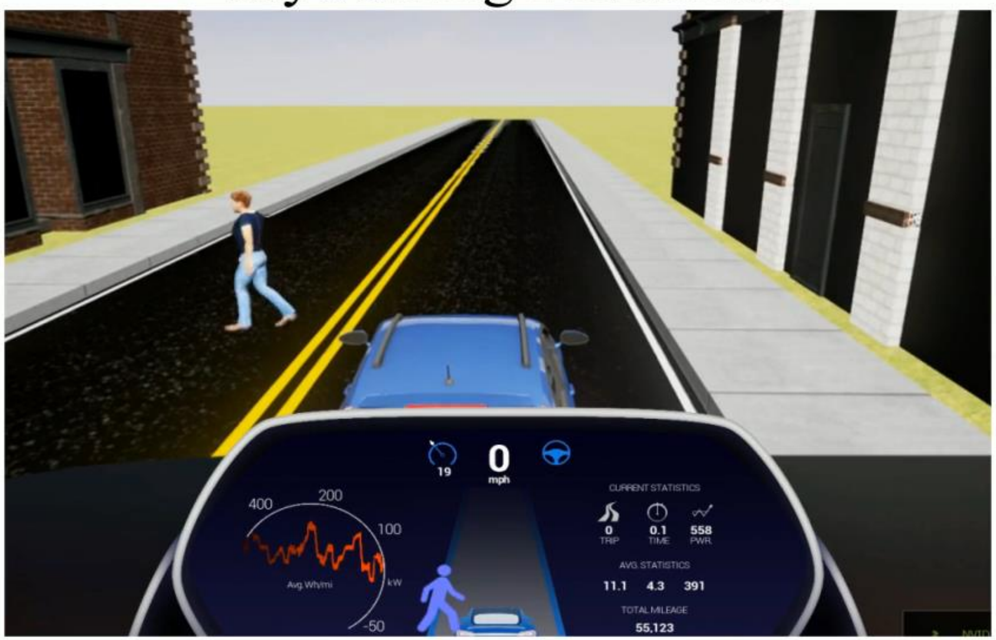

The vehicle respects pedestrians and slows until they have safely crossed the road

Multi-agent Simulation @ UNC

- Crowd and Multi-agent Simulation
  - http://gamma.web.unc.edu/research/crowds/
  - http://gamma.cs.unc.edu/menge/
- Autonomous Driving
  - http://gamma.cs.unc.edu/AutonoVi/
- Motion and Path Planning
  - http://gamma.web.unc.edu/research/robotics/