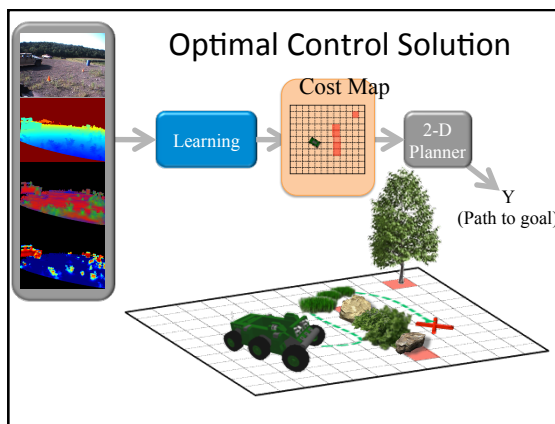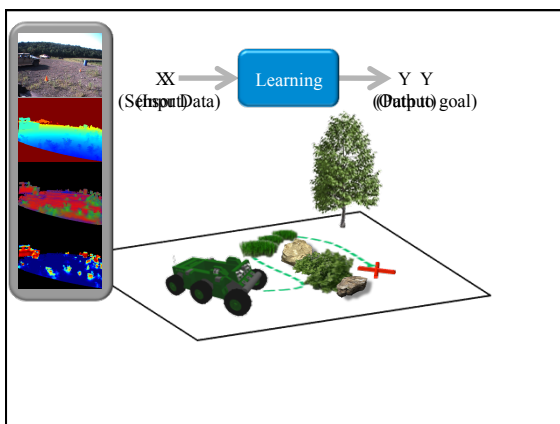## CS 473: Intro to Artificial Intelligence
Dieter Fox

Inverse Optimal Control
(Inverse Reinforcement Learning)

Most slides by
Drew Bagnell / Brian Ziebart
Carnegie Mellon University

Optimal Control Solution

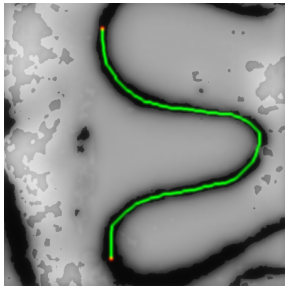Mode 1: Training example

Mode 1: Training example

Mode 1: Learned behavior

Mode 1: Learned behavior

Mode 1: Learned cost map

Mode 2: Training example

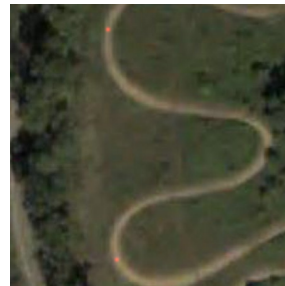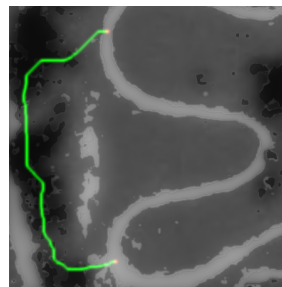Mode 2: Training example

Mode 2: Learned behavior

Mode 2: Learned behavior

Mode 2: Learned cost map

Feature vector
Cost =w'F
Weighting vector

Ratliff, Bagnell, Zinkevich 2005
Ratliff, Bradley, Bagnell, Chestnutt, NIPS 2006
Silver, Bagnell, Stentz, RSS 2008

(, High Cost)
(, Low Cost)
Learn F_1

Ratliff, Bagnell, Zinkevich, ICML 2006
Ratliff, Bradley, Bagnell, Chestnutt, NIPS 2006
Silver, Bagnell, Stentz, RSS 2008

$[F_1]$
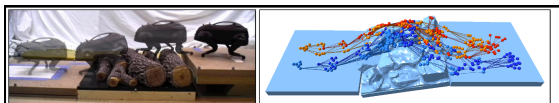(, High Cost)
(, Low Cost)
Learn F_2

Ratliff, Bagnell, Zinkevich, ICML 2006
Ratliff, Bradley, Bagnell, Chestnutt, NIPS 2006
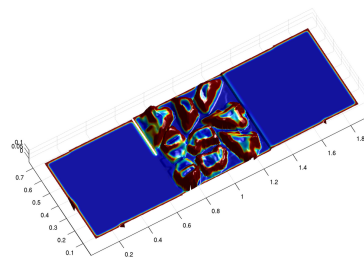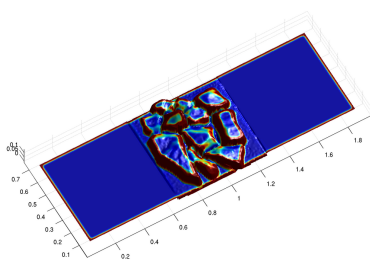Silver, Bagnell, Stentz, RSS 2008

example path

Ratliff, Bradley, Chesnutt, Bagnell 06

Zucker, Ratliff, Stolle, Chesnutt, Bagnell, Atkeson, Kuffner 09

Learned Cost Function Examples

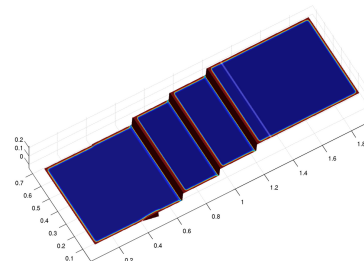
Learned Cost Function Examples

Learned Cost Function Examples

Mimicry implies Prediction

Pedestrian Trajectory Prediction

Kitchen

Secretary Desk

Lounge

Pedestrian Trajectory Prediction

Pedestrian Trajectory Prediction

Pedestrian Trajectory Prediction

(Using O(1) Bayes' Rule for Goals)

Pedestrian Trajectory Prediction

Staying out of People's Path
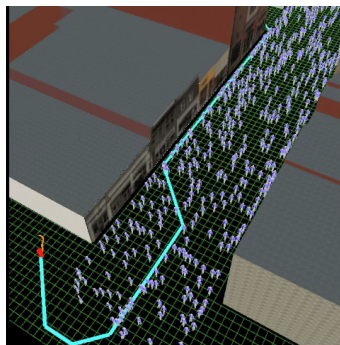
Car Prediction
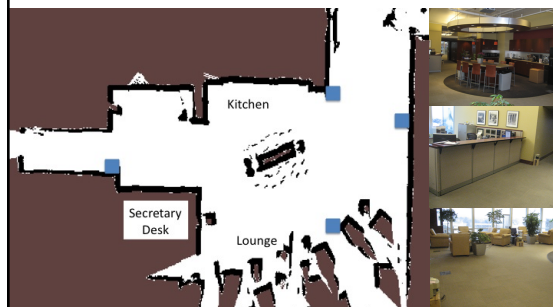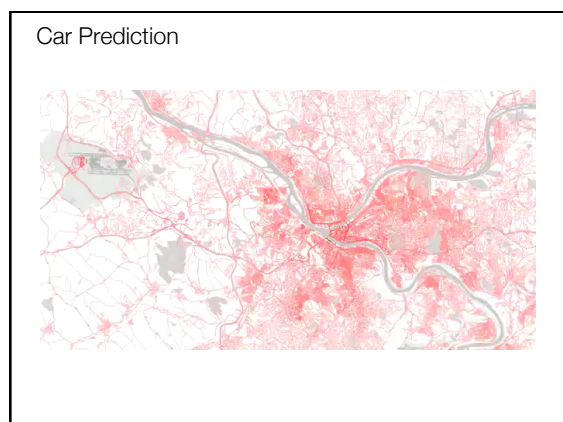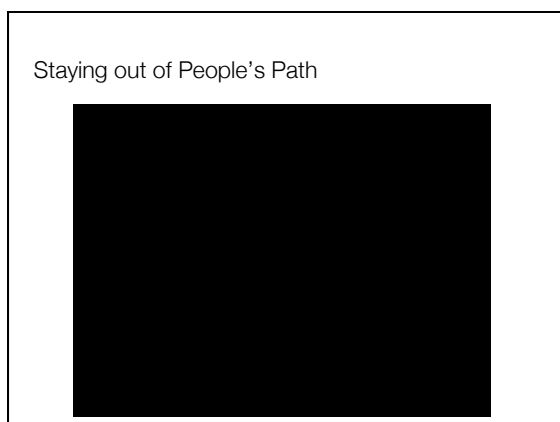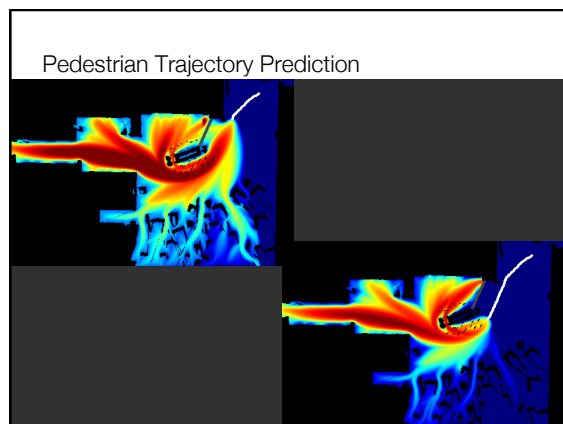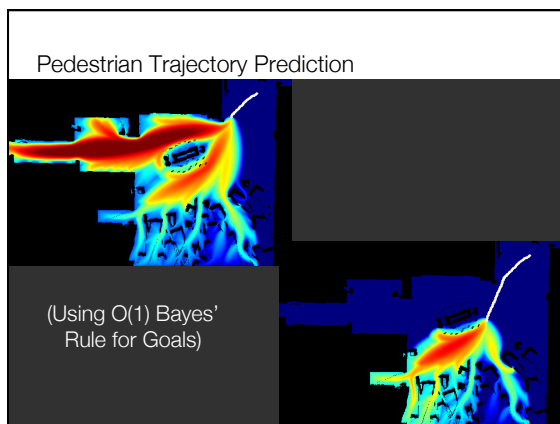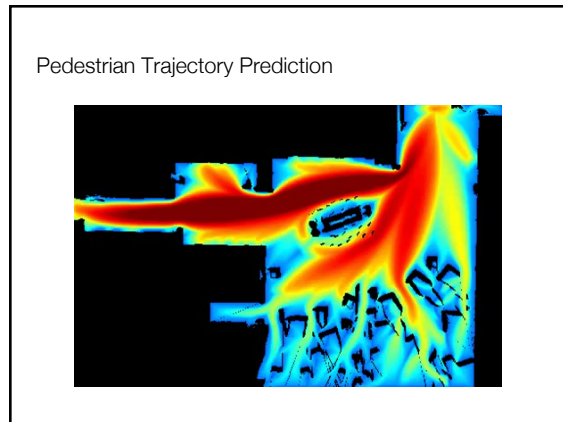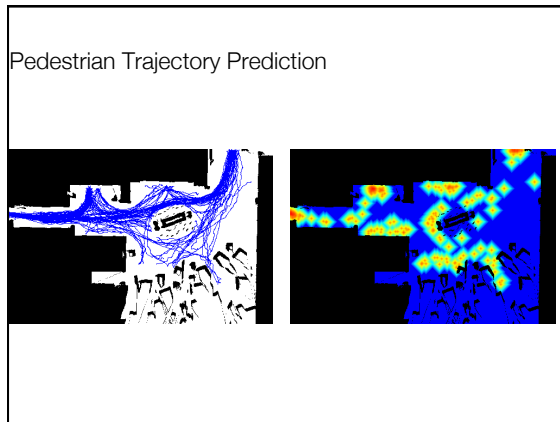
Personal Navigation System

Learning to Navigate Through Crowded Environments (Henry et al)

Motivation

- Robots should move naturally and predictably within crowded environments
  - Move amongst people in a socially transparent way
  - More efficient and safer motion
- Humans trade off various factors
  - To move with the flow
  - To avoid high density areas
  - To walk on the left/right side
  - To reach the goal

Mall Scenario

Lane Formation