Learning Applied to Ground Robots (LAGR)

Larry Jackel
ljackel@darpa.mil
DARPA IPTO / TTO

DARPA PR References: C/N 44224, C/N 1944 and BAA 04-25

Current State of the Art: DARPA PerceptOR Program

- 4-year program
  - concluded February, 2004
- Perception for off-road navigation
- Un-rehearsed tests at government sites
  - Ft AP Hill VA (Dec 2003)
  - Yuma, AZ (Feb 2004)

PerceptOR results

- Good performance in uncluttered environments
- Much worse than human RC operation in cluttered environments
  - No learning from mistakes (ping-pong between obstacles)
  - Significant obstacle classification errors
    - e.g. can't always tell compressible vegetation from rocks or stumps

Near-sighted sensing gives poor path-planning

How autonomous navigation is done today

1. Sense the environment, usually with LADAR
2. Create a 3-D model of the space with solid and empty volume elements
3. Identify features in the environment:

   Ditches, Grass, Water, Rocks, Trees, Etc.
4. Create a 2-D map of safe areas (black) and dangerous areas (red)
5. Run a path planning algorithm to decide on the next move toward the goal, staying in the “black” areas
6. Move the vehicle
7. Repeat

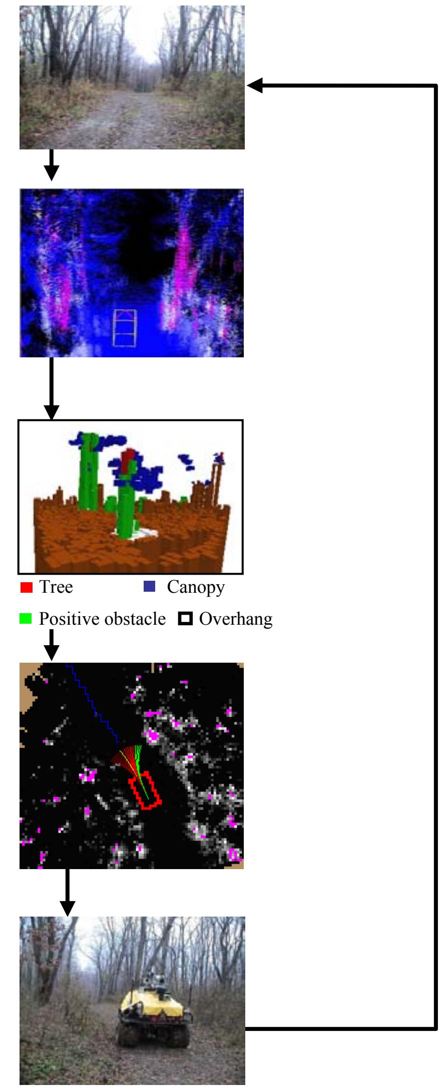

DARPA PR References: C/N 44224, C/N 1944 and BAA 04-25

Near-Sighted Behavior

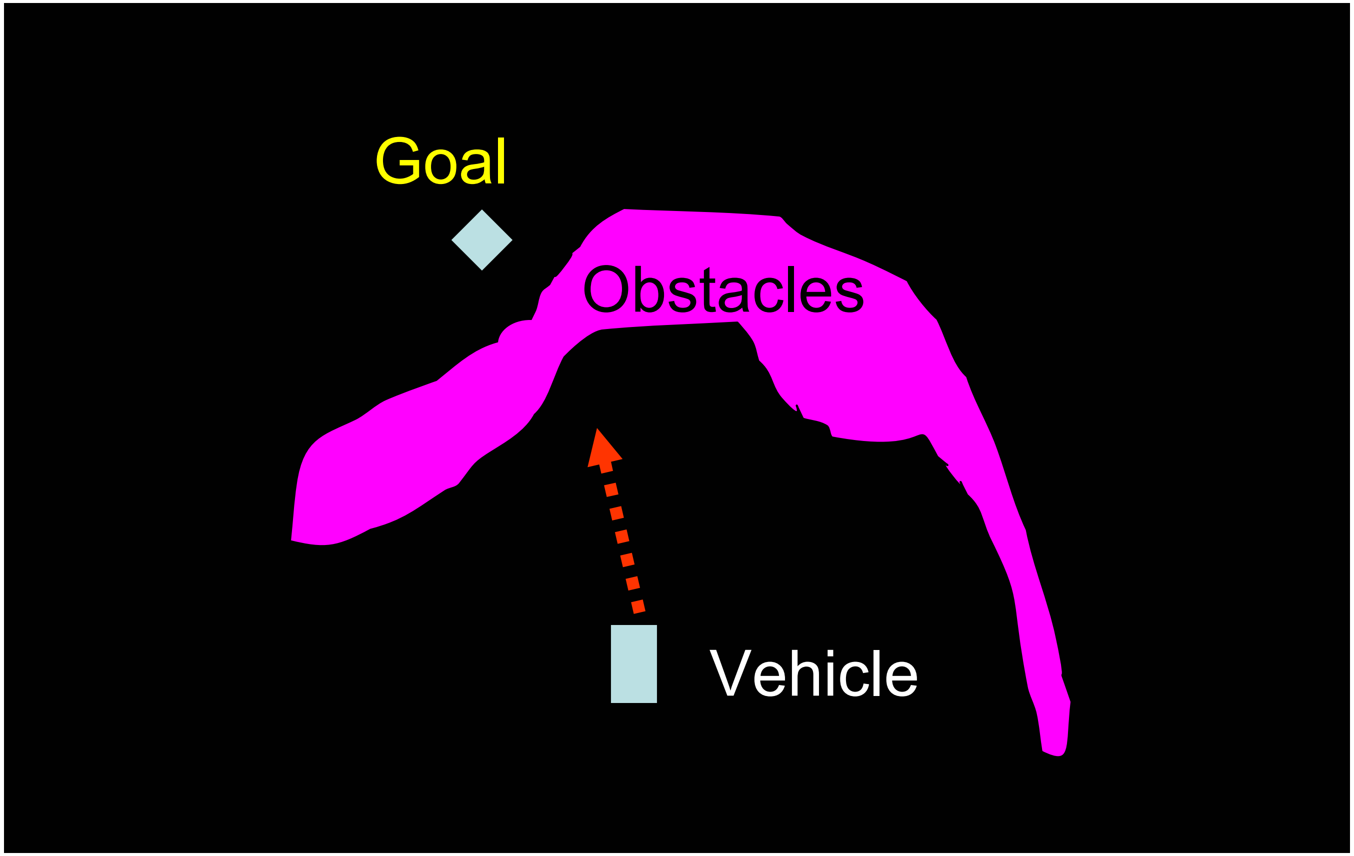

No use of LADAR

We need new approaches to autonomous navigation

- Learned navigation
  - General rules of navigation
    - Learning from example
    - Reinforcement learning
  - One shot learning
    - Don’t repeat mistakes
- Image Understanding
  - go beyond object recognition
- Brute force vehicles (E.g. CMU’s Spinner)

New DARPA program: Learning Applied to Ground Robots (LAGR)

- 3 years / 2 phases
- Hope for ~10 performers
  - Each performer supplied with a simple robot
    - About 0.5 meters on a side
    - Simple differential drive steering
    - Sensor suite includes, stereo cameras, GPS, inertial navigation unit, proximity sensor, compass
    - On-board high-end Linux computer
    - Baseline PerceptOR code
- Focus on learning

DARPA PR References: C/N 44224, C/N 1944 and BAA 04-25

LAGR Competitions

- ~Every month
  - Gov't site(s) – code uploaded to gov't robot
  - 3 runs / performer
    - Course length about 100 meters
    - Use knowledge gained in earlier runs in later runs
  - Data logs shared with all performers

- ~Every 6 months
  - New training data gathered by Gov't team
  - Performers shrink-wrap code learns new data on Gov't computer
  - New navigation algorithms tested at Gov't site

Scoring

- Three runs per performer each competition
  - Same course – but possibly changing conditions
    e.g. weather, sun angle
  - Time, Tmax allowed to run each course

- Score is average of “effective” speeds for each run

- Effective speed: S= (D/T) * F

D = shortest possible traversable distance to complete course
F = fraction of Euclidean distance to goal (= 1 for completed course)
T = time on course, or Tmax if course not completed

Runs may be aborted by govt for safety reasons or blatant navigation failure
in such cases T = Tmax by definition.

Note that if $F \neq 1$, then S = (D /T max) * F

Score Examples

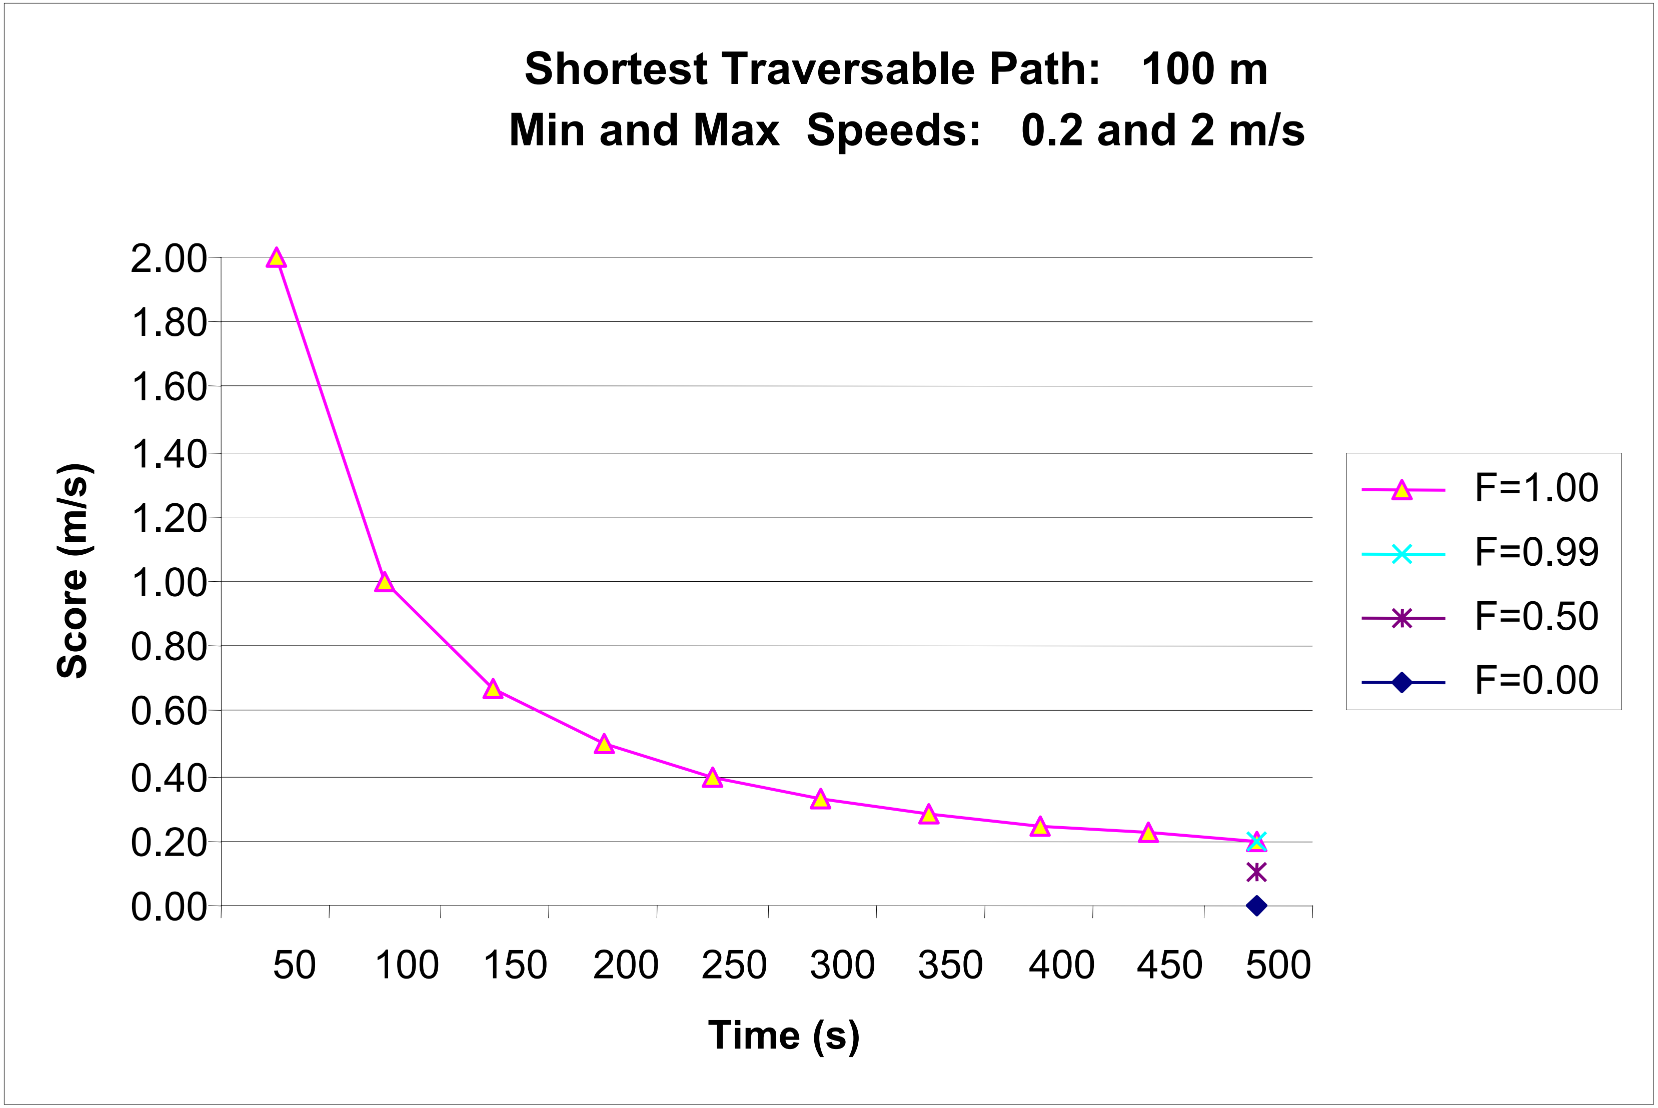

LAGR Schedule

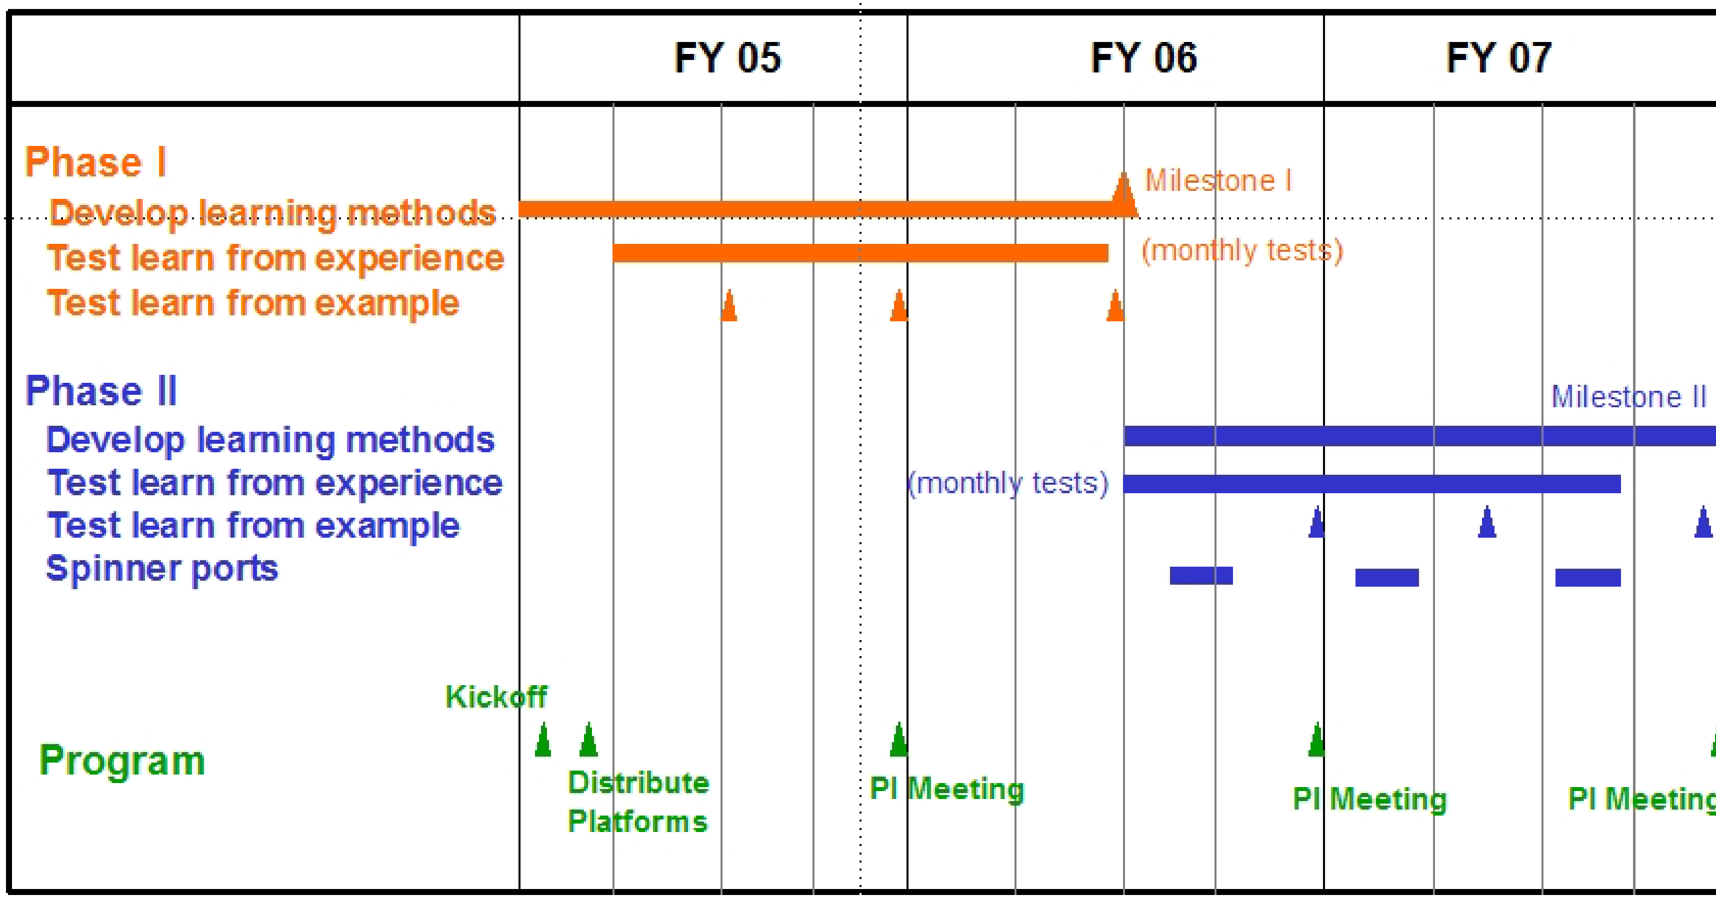

Milestone I. Learned system travels 10% faster than baseline system, same course, same vehicle

Milestone II. Learned system travels 2X faster than baseline system, same course, same vehicle

Must meet speed metrics at 18 month milestone to stay in program:

DARPA PR References: C/N 44224, C/N 1944 and BAA 04-25

Code Sharing

- Performers may share source code
  - Credit must be given where credit is due
- Object code available from each performer
  - Logs of competition runs available to performers
    - Logs can be used as training data

PI Meetings

- About 2/year

- Performers need to explain methods so that some one “skilled in the art” can reproduce results

Image Understanding

DARPA PR References: C/N 44224, C/N 1944 and BAA 04-25

Encourage Image Understanding

- Much of the course will be visible from the start but will be beyond stereo range
- There will be cul-de-sacs
- Much better performance will be possible if performers go beyond stereo and attempt to “visually” plan route
  - Possible strategy:
    Work in image plane representation, not map representation
- Some Govt competitions will only allow monocular vision
  - Enhanced range info from optical flow, motion parallax?

A Different Approach to Machine Vision

Port to Spinner

- In 2nd Phase (18 -36 months) training data from Spinner will be supplied to performers
- Best code will be tested on Spinner