Evaluating the Performance of Unmanned Ground Vehicle Water Detection

Performance Metrics for Intelligent Systems Workshop (PerMIS)

September 30, 2010

Dr. Arturo Rankin

Jet Propulsion Laboratory, California Institute of Technology

This research was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under the Army Research Laboratory (ARL) Robotics Collaborative Technology Alliances program, through an agreement with NASA.

RCTA Program

- ARL Robotics Collaborative Technology Alliances (RCTA) Program
  - Technology development in areas of importance to future Army
  - Consortium of industrial, academic, and government lab partners:
    - GDRS, CMU, ASI, JPL, Alion, BAE, Sarnoff, SRI, FAMU, UMD, PercepTek, Robotic Research, SSC, Howard U., NCA&T, UPenn, Skeyes Unlimited
  - 5 year base + 3 year extension (May 31, 2001 - Dec 31, 2009)
  - JPL's Advanced Perception tasks:
    - Improvements to stereo vision
    - Terrain classification
    - Pedestrian detection (plus vehicle detection)

Stereo ObsDet evaluation platform

GDRS Instrumented Train

Pedestrian detection evaluation platforms

GDRS Suburban GDRS Escape

Water detection evaluation platform

GDRS XUV

Topics Addressed

- Motivation
- Brief summary of water detection methods
- Methods for evaluating water **detection**
  - Ground truthing water in 2D (image) space
  - Image space evaluation
  - Map space evaluation (after temporal filtering)
- Method for evaluation water **localization**
  - Ground truthing water in 3D (map) space
  - Map space evaluation
- Conclusions

Motivation

- **Issue:** multiple image-space water detection algorithms exist
  - Color variation based water detection
  - Sky reflection based water detection
  - Stereo reflection based water detection
  - Multi-cue water detection (color, texture, stereo)
- **Need:** method for evaluating the performance of each detector

- **Issue**: localization errors can lead to poor autonomous navigation performance
- **Need:** method for evaluating localization accuracy

Water Detection Evaluation: Was it detected?

- Water detection based on color variation
  - Estimate the horizon line
  - Segment low texture regions below the horizon, growing them as long as intensity gradient is low
  - Perform connected components
  - Threshold size and average delta variance across blob boundary
  - Perform ellipse fit of remaining blobs
  - Threshold ellipse width, length, and density (assuming blob is in the ground plane)
  - Threshold bri/sat line fit from leading to trailing edge

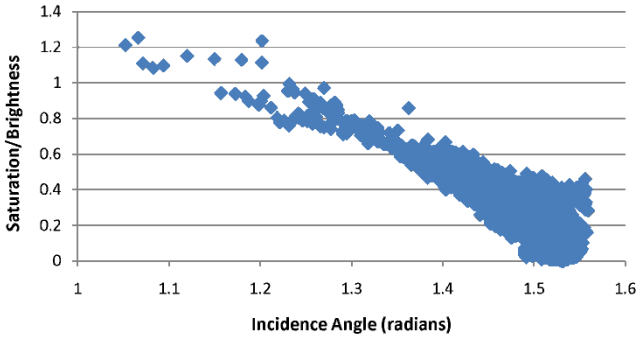

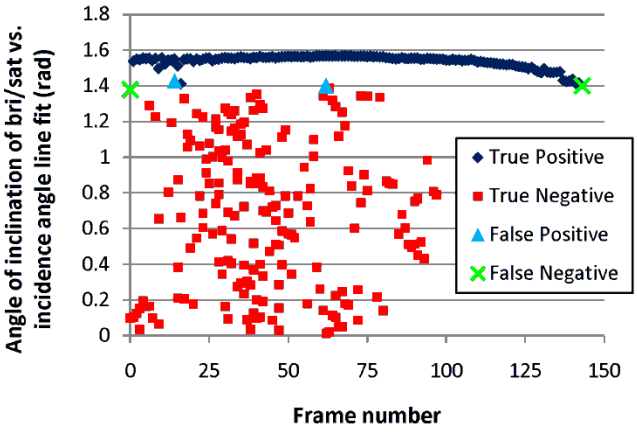

Water detection result

Water Detection Evaluation: Was it detected?

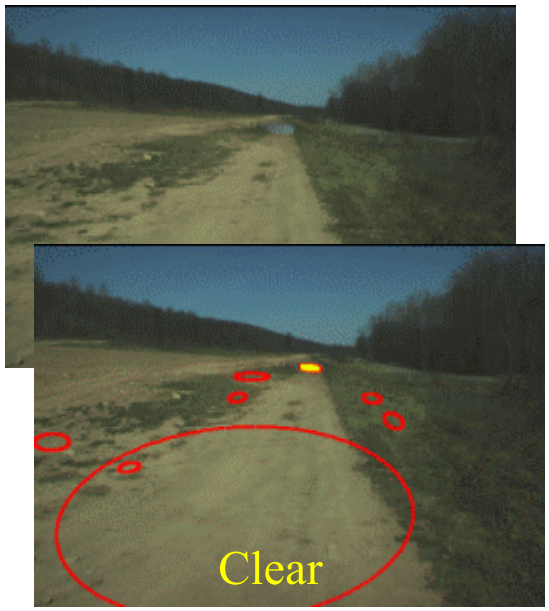

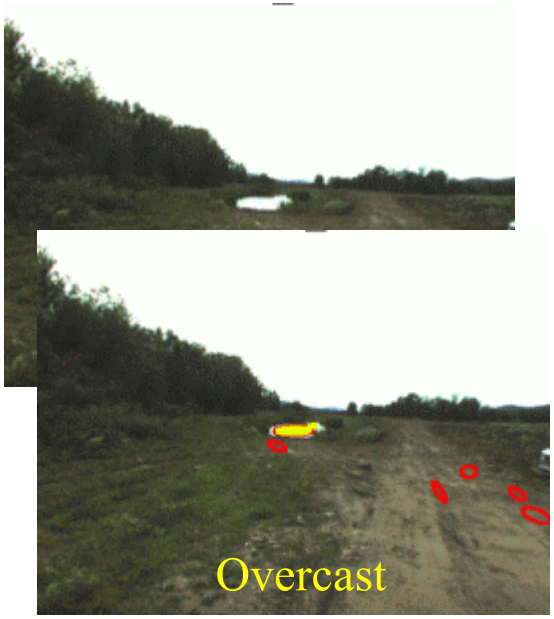

Water detection based on color variation results

Scene	1st frame distance to water	Num frames	True positive detection	False positive detection	Avg frame time
Clear	64 meters	237	227 (95.76%)	1 (0.45%)	128ms
Over-cast	35 meters	334	323 (96.71%)	2 (0.60%)	76ms
Cloudy	47 meters	162	160 (98.77%)	1 (0.62%)	54ms

Water Detection Evaluation: Was it detected?

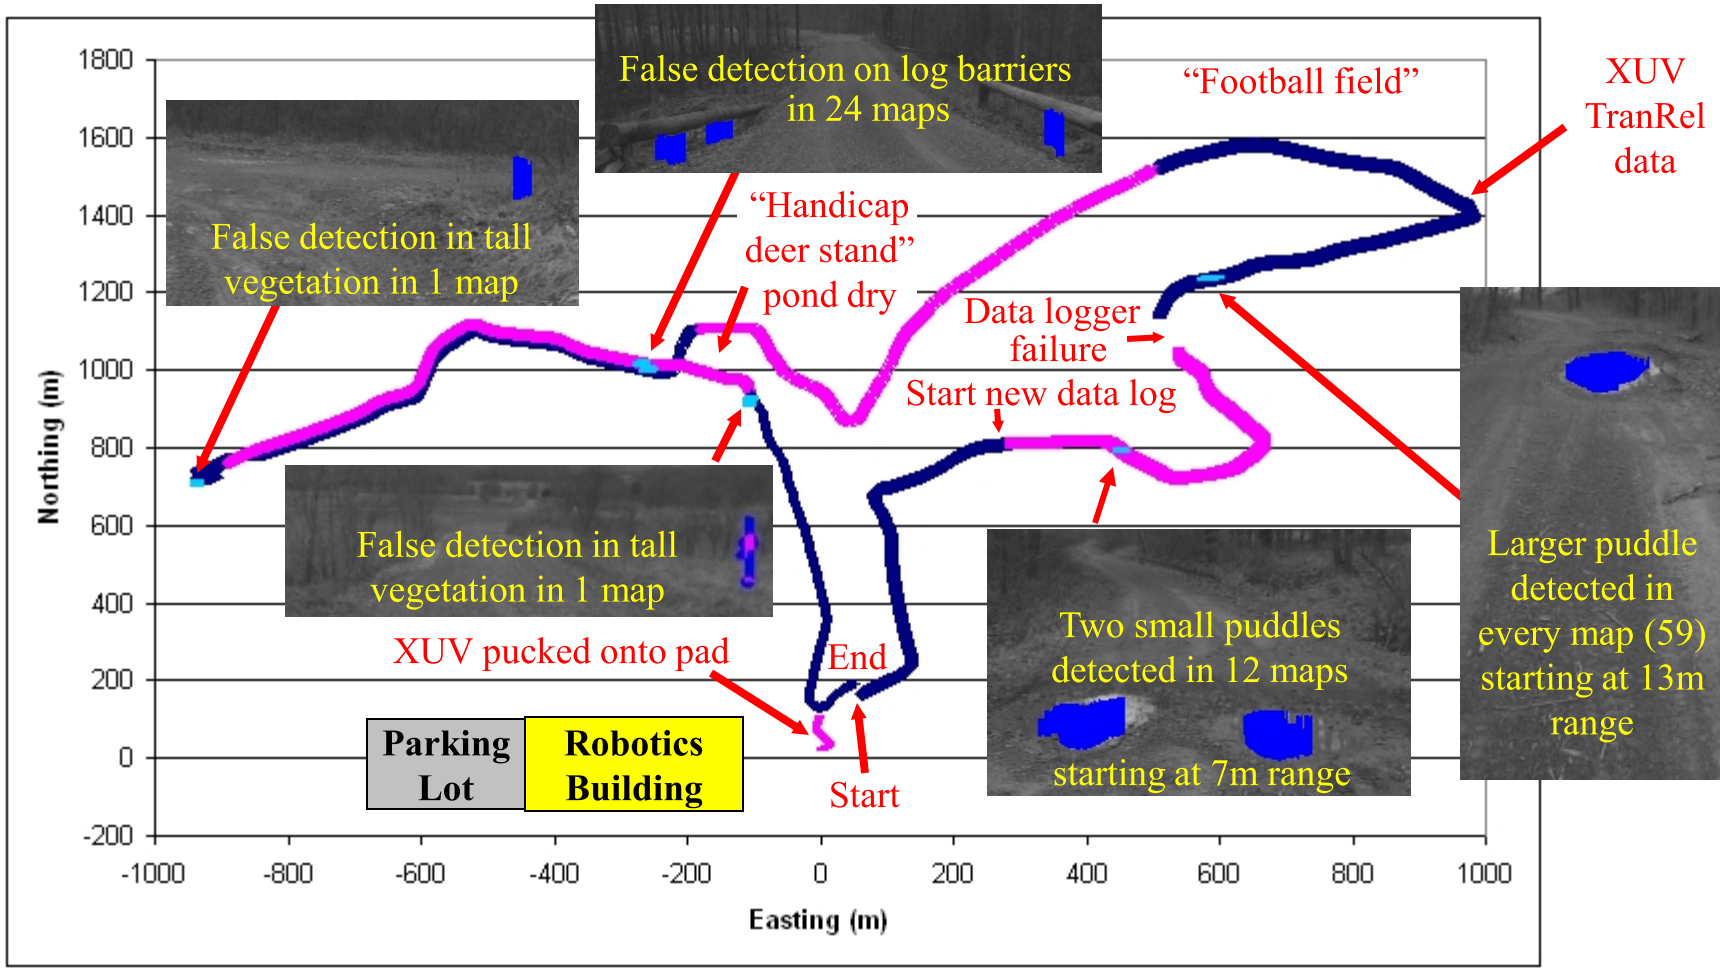

- 6.9km course (12,265 stereo image pairs processed @ 512x384 resolution)
- All water bodies detected (world map temporal filtering N=2)
- 0.2% false pos detection rate (26/12,265 world maps contain false pos detection)

Water Detection Evaluation: How much of it was detected?

- Developed OpenGL tool to ground truth water bodies
- Select vertices around water body perimeter, connect them w/ lines
- Use stereo to find the avg vertex elevation and vertex 3D coordinates
- From frame-to-frame, update vertex 2D image coordinates by performing 3D to 2D mapping using camera model
- From frame-to-frame, adjust vertices as needed

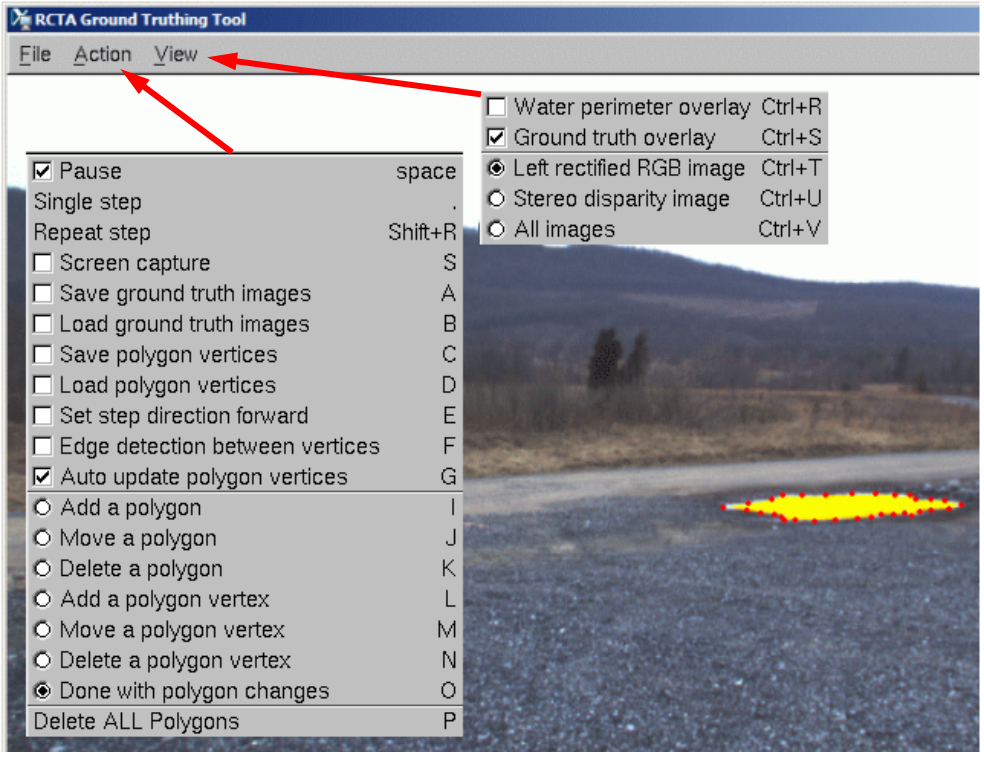

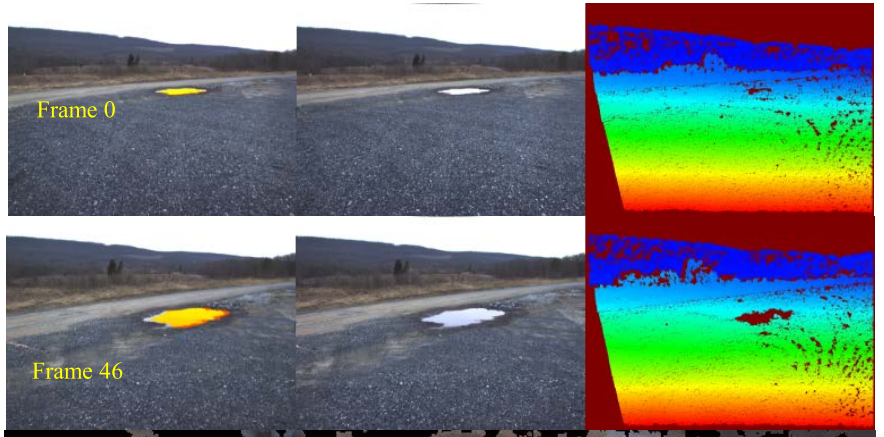

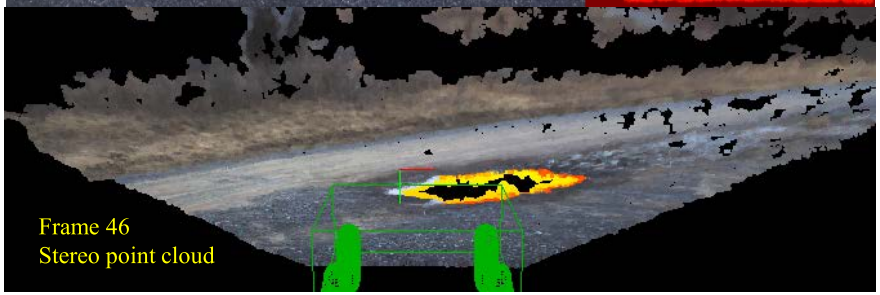

Water Detection Evaluation: How much of it was detected?

- Some ways to minimize ground truthing error:
  - Allow the user to step through the sequence, pausing at every frame and verifying accurate labeling and alter vertices.
  - Provide the option of processing the sequence in reverse order from the final image to the first image.
  - Perform non-linear segmentation between vertices to improve the modeling of the water body perimeter.

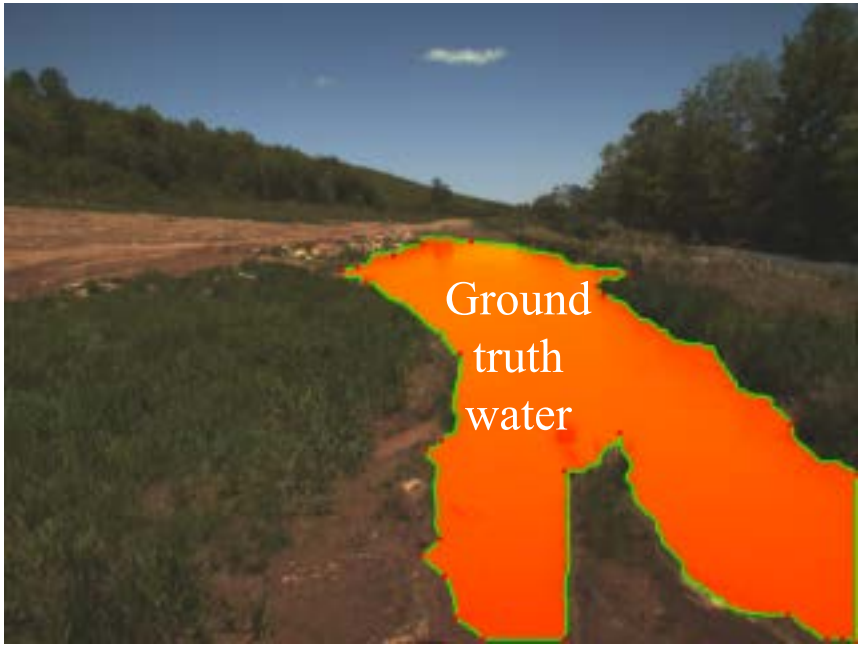

Water Detection Evaluation: How much of it was detected?

- Non-linear segmentation:
  - The open-source code for **intelligent scissors**, available under GNU Image Manipulation Program (GIMP), was integrated into the water body ground truthing tool.
  - This algorithm attempts to find the most grayscale contrast closed-loop boundary (Laplacian zero-crossing) while keeping the boundary edge smooth (gradient direction) and the texture around the boundary consistent (gradient magnitude).
  - An optimal graph search called *live-wire boundary* is performed based on Dijkstra's path finding algorithm to find a minimal cost path via dynamic programming.

A portion of a sample water body

Edge detection is performed between selected perimeter vertices

Water Detection Evaluation: How much of it was detected?

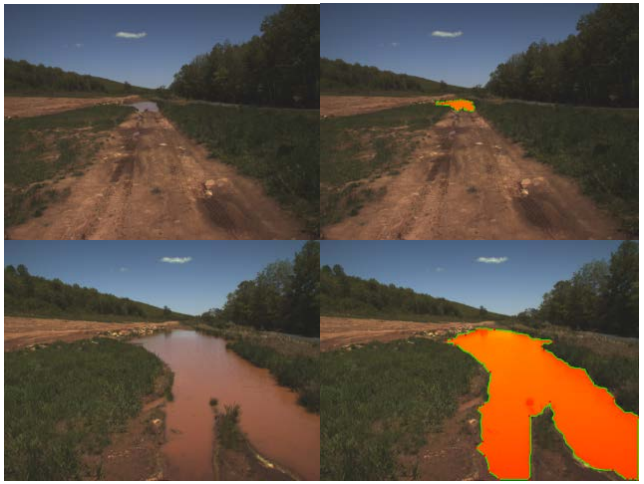

Ground Truth

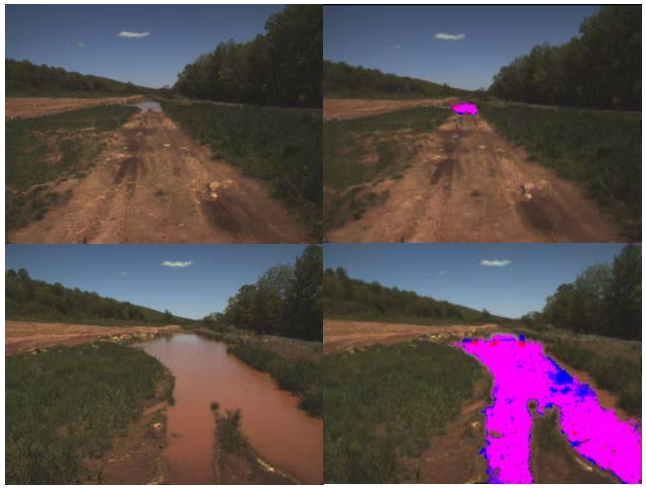

Water Detection

- One cue
- Two cues
- Three cues

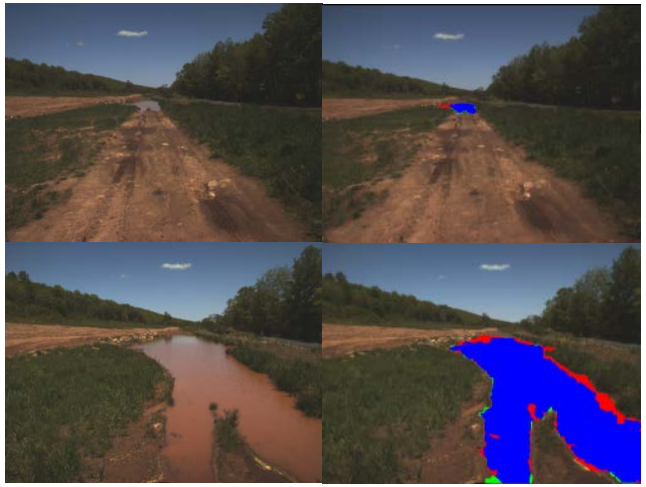

Comparison

- True positive
- False negative
- False Positive

Water Detection Evaluation: How much of it was detected?

- The true positive detection rate ranged from 68% (at a range of 28 meters to the leading edge) to 90% (at a range of 4 meters to the leading edge).
- The false positive detection rate was 3.3% in one frame but remained below 0.8% in the rest of the frames.

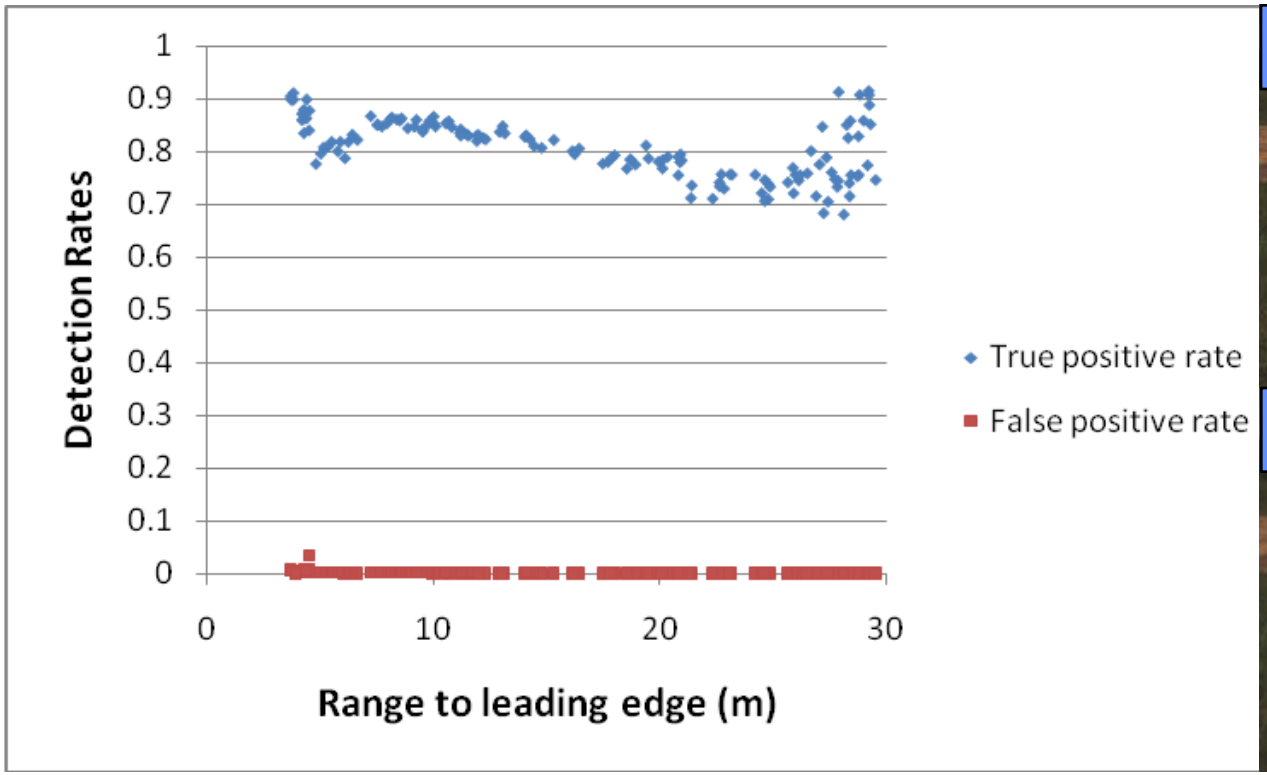

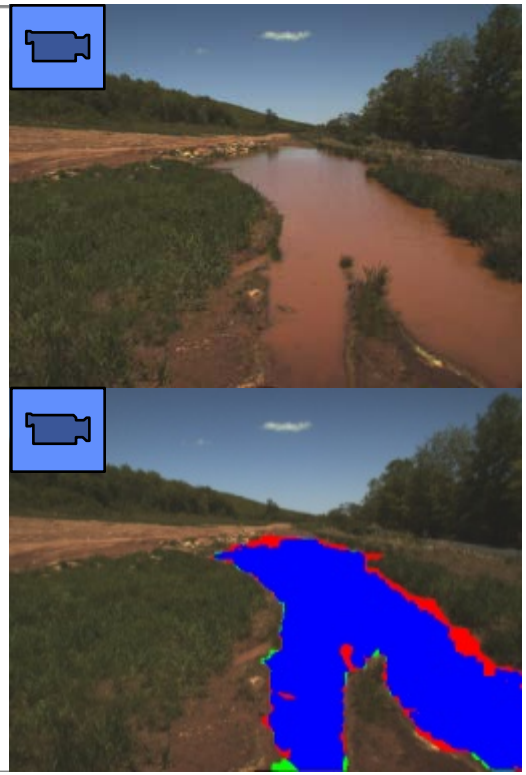

Water Detection Evaluation: Was it detected in the right place?

- Test course
  - GDRS instrumented train
    - XUV autonomous mobility computer/software
    - IMU/GPS Kalman filtered nav solution (accurate to 0.5% of distance traveled)
    - General Electric 24 volt DC motor
    - Ogura Fail-Safe Brake
    - Honda EU 1000i Generator
  - Two man-made water bodies
    - Corners ground truthed w/ DPGS (1cm + 1ppm)
    - DGPS spot checked w/ surveying instrument

Water Detection Evaluation: Was it detected in the right place?

- Thus far, GDRS has released one ground truth position to JPL.
- Using a 20cm resolution map, the JPL water detection software localized the corner of the second man-made water body within 16cm of the ground truth position.

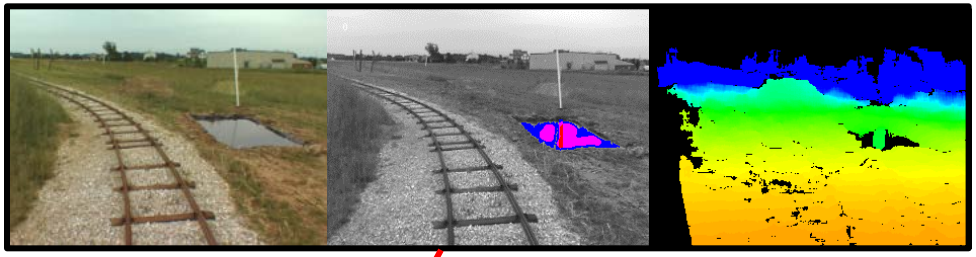

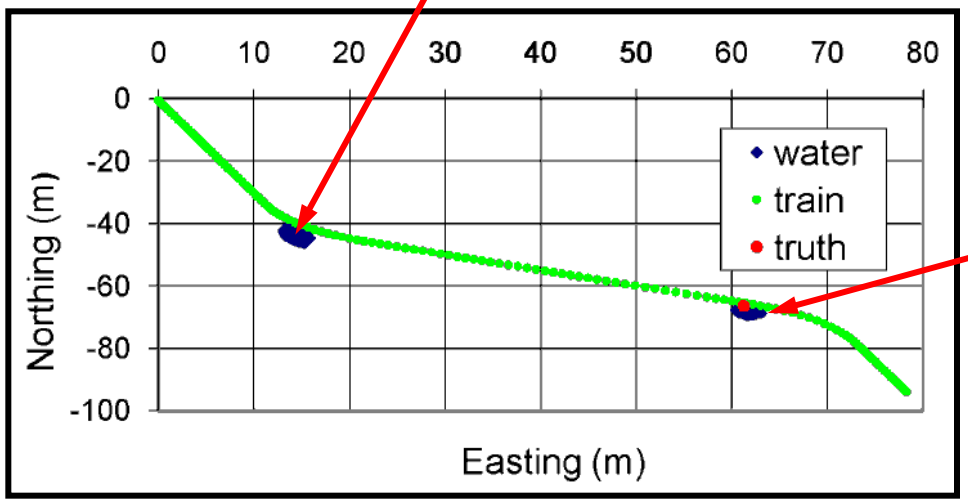

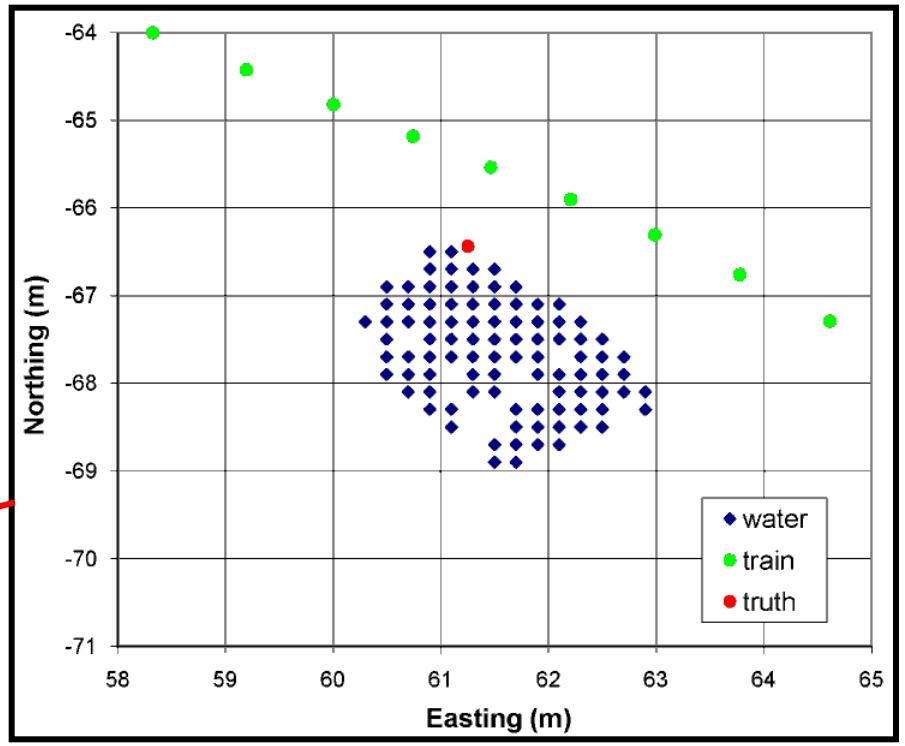

Water Detection Evaluation: Was it detected in the right place?

- Some useful detection metrics
  - Range of detection
  - Strength of detection
- Other potential measures of accuracy:
  - Difference in the detected and ground truth water body centroid (units: meters).
  - Percentage of the detected water body within the ground truth water body.
  - Percentage of the detected water body outside of the ground truth water body.
  - Percentage of the ground truth water body detected as water.
  - Maximum distance the detected water body perimeter strays from the ground truth water body perimeter (units: meters).

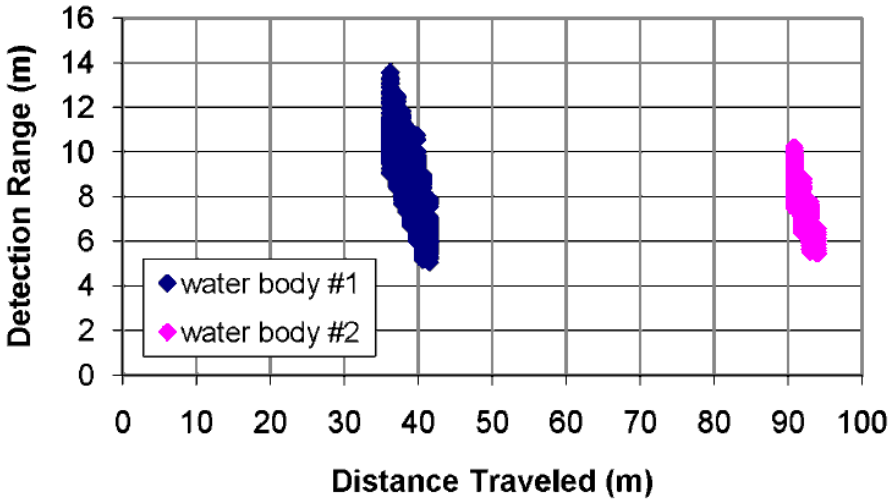

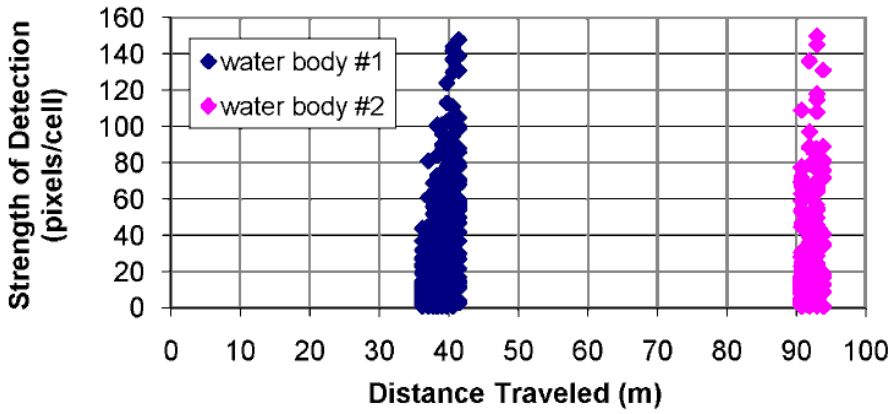

Conclusions

- JPL has developed a software tool for ground truthing water bodies in stereo image sequences.
- In the 143 frame sequence, the water body was detected in every frame.
  - The true positive detection rate ranged from 68% (at a range of 28 meters to the leading edge) to 90% (at a range of 4 meters to the leading edge).
  - The false positive detection rate was 3.3% in one frame but remained below 0.8% in the rest of the frames.
- The corner of a man-made water body was localized to within 16cm of ground truth.
- We outlined several measures of accuracy for comparing the 3D coordinates of a water body localized with JPL's water detection software with 3D ground truth water body perimeter measurements.
- More work is needed to determine the usefulness of these measures.

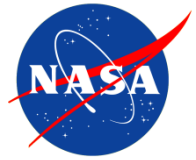

Questions?