A 3D Laser Targeting System

Master Thesis

Roman Stanchak

rs7@cec.wustl.edu

Department of Computer Science and Engineering

Washington University in St. Louis

Co-advised by Robert Pless and Bill Smart

Overall Goal

Aim laser at a point in the environment using observations from a stereo camera

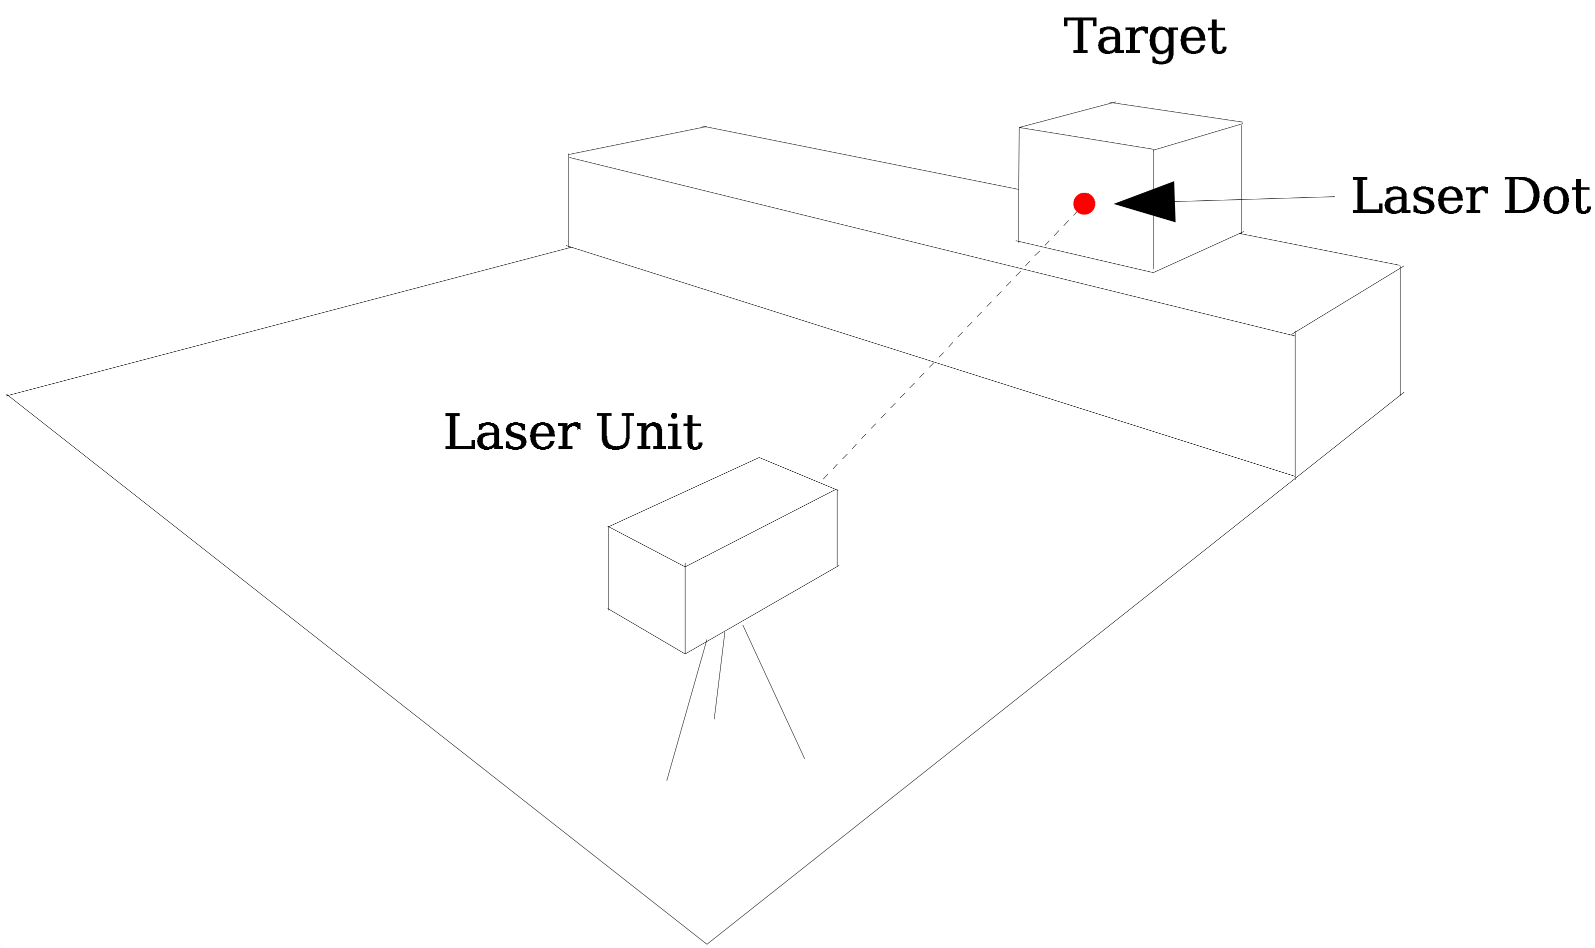

Contribution

Two methods of calibrating camera observations to laser controls

- Theoretical justification
- Implementation
- Experimental analysis
  - Overall system accuracy
    - depth of target
    - position of target
  - Accuracy of different calibration methods

Outline

- Related Work & Motivations
- Background
- Calibration Algorithms
  - Theory
  - Experimental Results
- Improving Calibration
  - Automatic Detection
  - Point Selection
  - More Experimental Results
- Concluding Remarks

Related Work: Visual Servoing

Iterative method to control robotic manipulator using camera observations

- use error gradient to pick action that minimizes difference between target and observed position

Advantages

- Analytic relationship not required
- Can dynamically adapt to to observed errors

Problems

- convergent method, never *exactly* on
- Requires consistent knowledge of laser dot position
  - Laser dot detection not robust

Current Method

Solve for transformation between laser and one plane in space.

- requires only one camera
- allows direct aiming of the laser
- calibration possible with 4 corresponding points between laser & image

Problems

- Doesn’t model full 3D geometry
- targeting outside depth plane is inaccurate
- must recalibrate to change it

New Approach

- Stereo camera measures depth
- Exact transformation allows direct aiming of laser

Two calibration methods

- Direct (3D -> laser)
- Epipolar (2D x 2 -> laser)

Background: Laser

We model the laser as a black box:

- Two inputs (u, v) control direction $\vec{X}_L$ of the laser.
- Fixed origin
- Direction $\vec{X}_L$ linear with $x_L = (u, v)$.

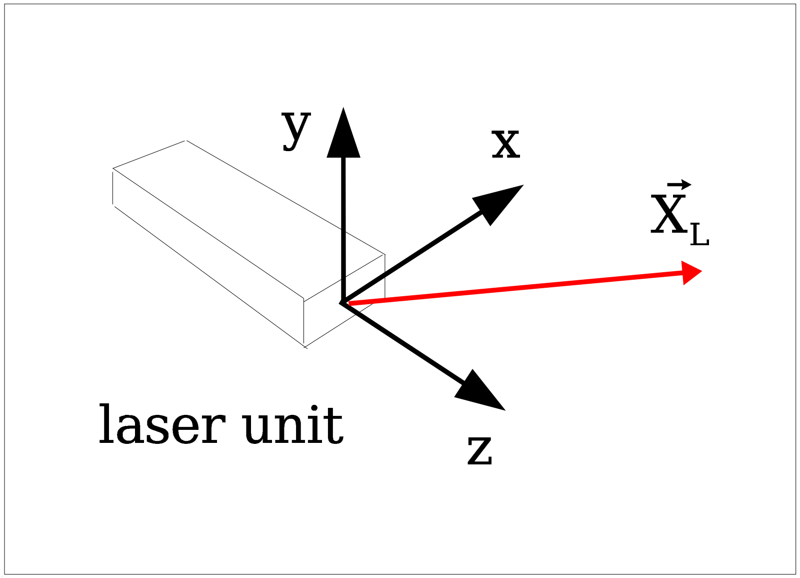

Background: Laser

Direction $\vec{X}_L$ linear with x_L.

$$w\mathbf{x_L} = \mathbf{A_L}\vec{X}_L$$

Where

- w is a scale factor
- $\mathbf{A}$ is a 3×3 laser projection matrix.

x_L projects on a line of 3D points.

Background: Depth Sensor

Requirements:

- can sense laser dot
- can report position relative to some 3D coordinate system

Tyzx Stereo Camera

- dot is visible in dim lighting
- report location relative to left camera center

Coordinate system relationship

Camera and laser 3D coordinate systems are related by a rotation and translation.

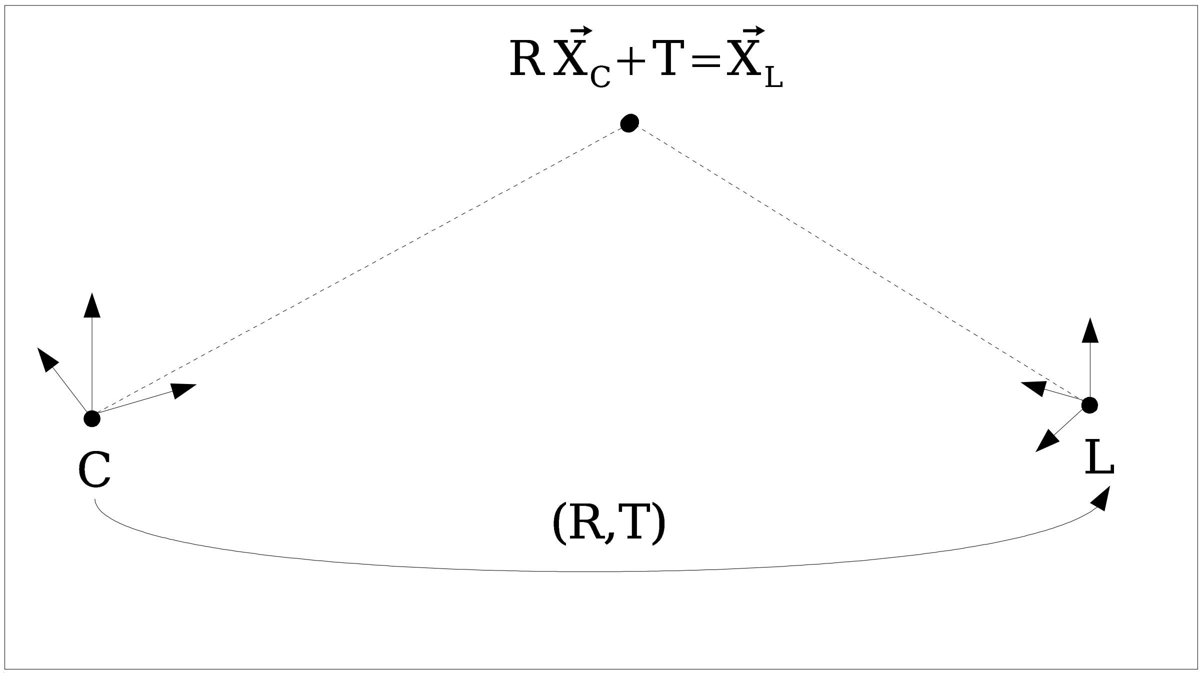

- $\mathbf{R}$ is a 3×3 rotation matrix
- $\mathbf{T}$ is a 3×1 translation vector.

Coordinate system relationship

Laser control and Camera coordinate related by

$$\mathbf{H}\mathbf{X_C} = \mathbf{x_L}$$

- Where $\mathbf{H} = \mathbf{A_L}[\mathbf{R}|\mathbf{T}]$
  - $\mathbf{A_L}$ is laser projection matrix
  - $[\mathbf{R}|\mathbf{T}]$ is 3×4 augmented matrix of rotation and translation
- Calibrate laser by solving for $\mathbf{H}$
- Control laser by multiplying $\mathbf{H}$ and the desired target $\mathbf{X_C}$

Direct Calibration

- Observe correspondance between laser, 3d coordinate of laser in camera image
- Each correspondance provides three linear constraints on $\mathbf{H}$:

$$
\begin{aligned}
Xh_1 + Yh_2 + Zh_3 + h_4 &= wu \\
Xh_5 + Yh_6 + Zh_7 + h_8 &= wv \\
Xh_9 + Yh_{10} + Zh_{11} + h_{12} &= w
\end{aligned}
$$

- Where h_i are the components of the matrix $\mathbf{H}$

Direct Calibration

- Eliminating w maintains two linear constraints

$$
\begin{aligned}
Xh_1 + Yh_2 + Zh_3 + h_4 &= u(Xh_9 + Yh_{10} + Zh_{11} + h_{12}) \\
Xh_5 + Yh_6 + Zh_7 + h_8 &= v(Xh_9 + Yh_{10} + Zh_{11} + h_{12})
\end{aligned}
$$

- Need 6 or more correspondences to solve for 12 degrees of freedom of $\mathbf{H}$ using linear least squares

Deriving Laser Controls with H

Given $\mathbf{H}$:

- Define 3D coordinate $\vec{X_C}$ of target using Tyzx Stereo camera
- Product $\mathbf{HX_C}$ = $\begin{bmatrix} wv \\ wu \\ w \end{bmatrix}$.
- Solve for laser controls (u, v) by dividing out w.

Results to come . . .

Epipolar Calibration

- 3D sensor not required
- Requires two or more conventional cameras
- Cameras can be uncalibrated

Background: Camera

Pinhole Perspective projection model.

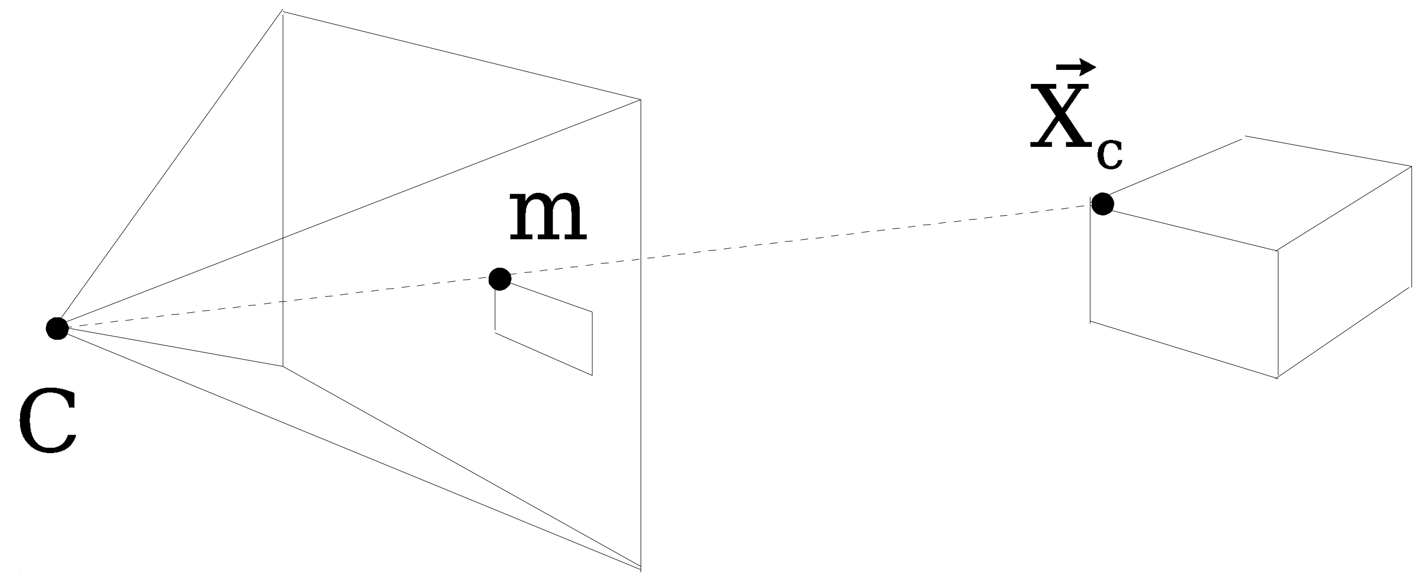

- C = center of projection
- $\vec{X}_C$ = 3D point relative to C
- $\mathbf{m}$ = projection of $\vec{X}_C$ on 2D image plane

Camera Projection Equation

$$s\mathbf{m} = \mathbf{A_C}\vec{X}_C$$

Where:

- $\mathbf{m}$ = homogeneous 2D image coordinate $\begin{bmatrix} x & y & 1 \end{bmatrix}^T$
- $\vec{X}_C$ = 3D point relative to camera center
- $\mathbf{A_C} = 3 \times 3$ camera calibration matrix encoding *intrinsic* parameters
- s = the projective depth

Background: Stereo

Cameras related by rotation and translation

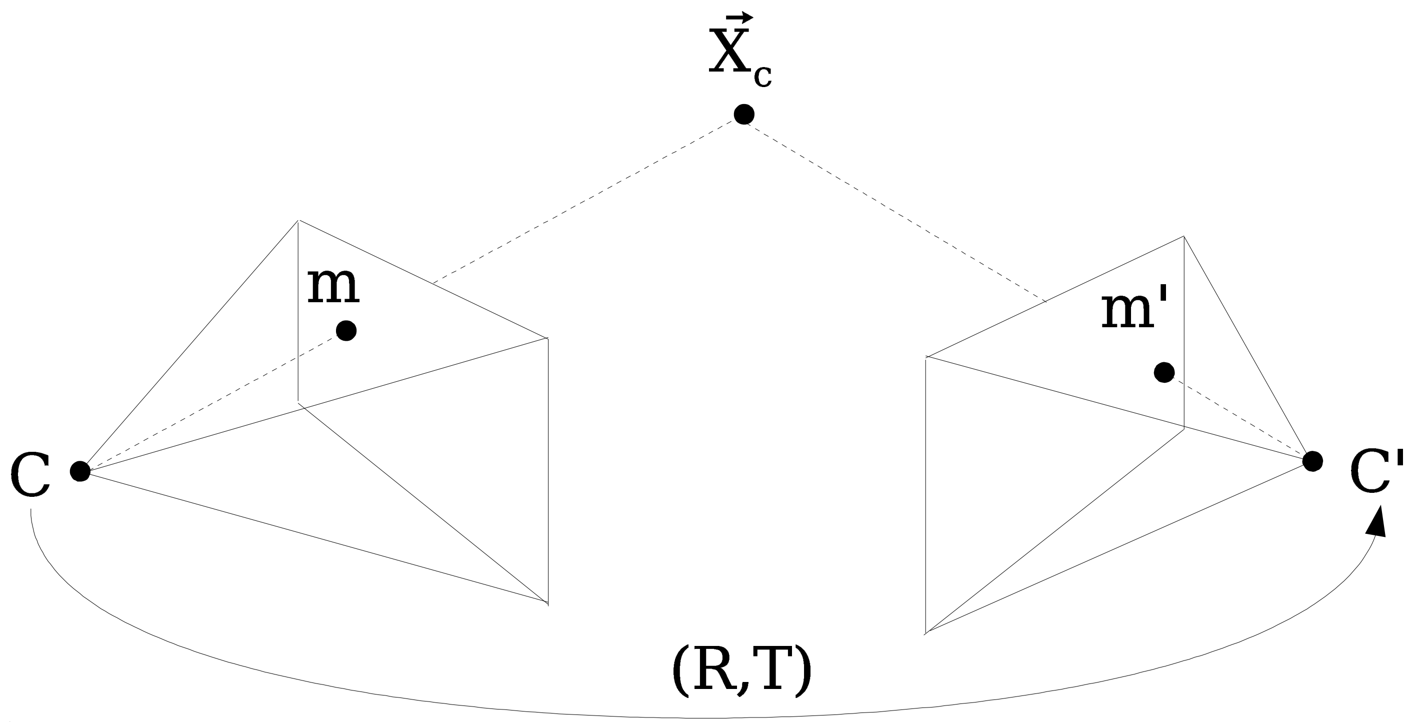

Background: Epipolar Line

Point on camera image 1 constrained to lie on a line in camera image 2 (and vice versa).

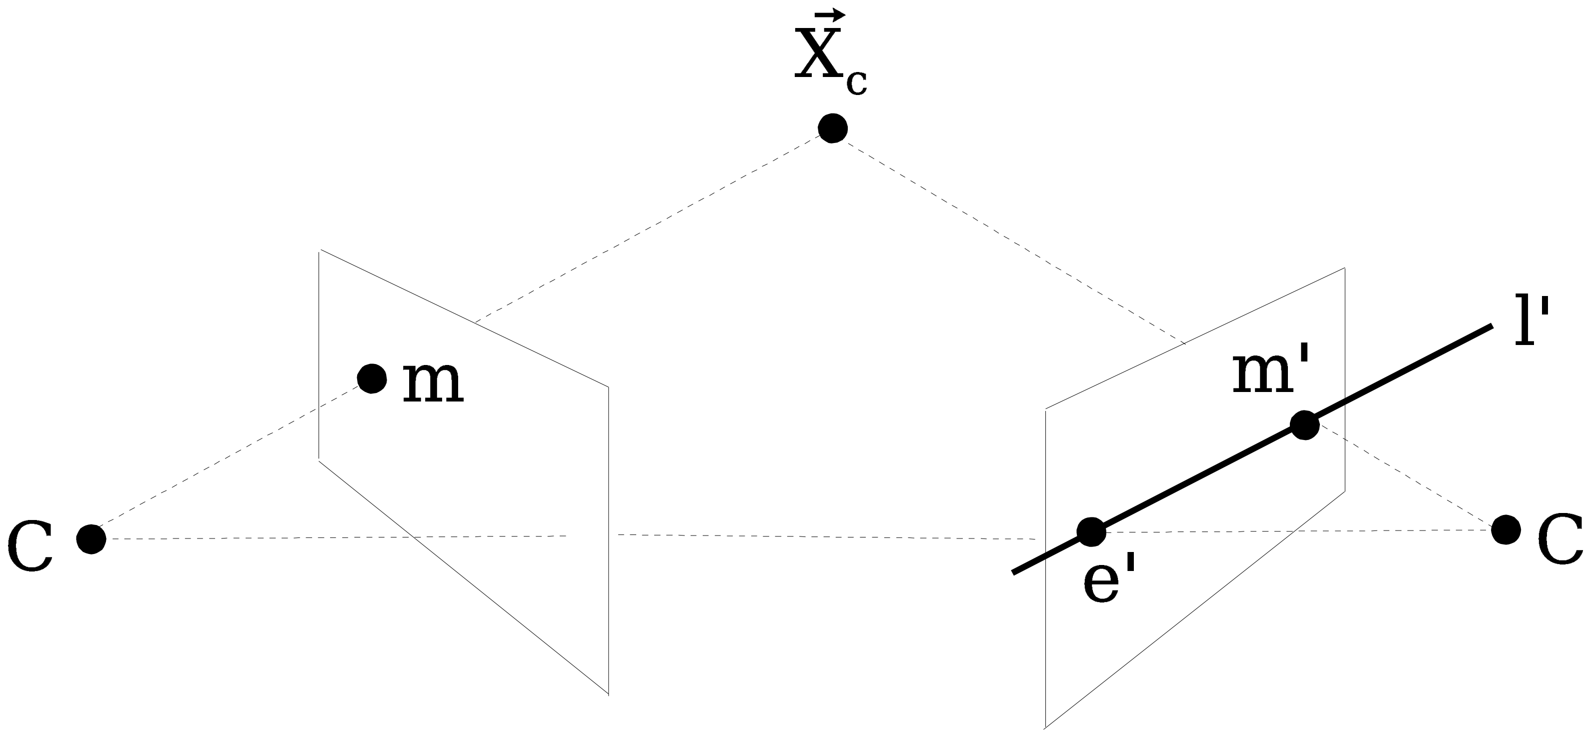

Background: Fundamental Matrix

Epipolar geometry encoded in the Fundamental Matrix:

$$\mathbf{m}^{\mathbf{T}}\mathbf{F}\mathbf{m}' = \mathbf{0}$$

- $\mathbf{F}$ is a 3×3 matrix.
- Well studied in vision literature.
- Given examples of corresponding $\mathbf{m}$,$\mathbf{m}'$, many techniques to solve for $\mathbf{F}$.

Epipolar Calibration

Key intuition: Laser is an inverted camera

- Emits light instead of absorbing it
- (u,v) laser controls congruent to (x,y) image coordinates.
- Same linear relationship.

$$\underbrace{s\mathbf{m} = \mathbf{A_C}\vec{X}_C}_{Camera} \quad \underbrace{w\mathbf{x_L} = \mathbf{A_L}\vec{X}_L}_{Laser}$$

Epipolar Calibration

One camera constrains laser control to a particular line in (u, v) space.

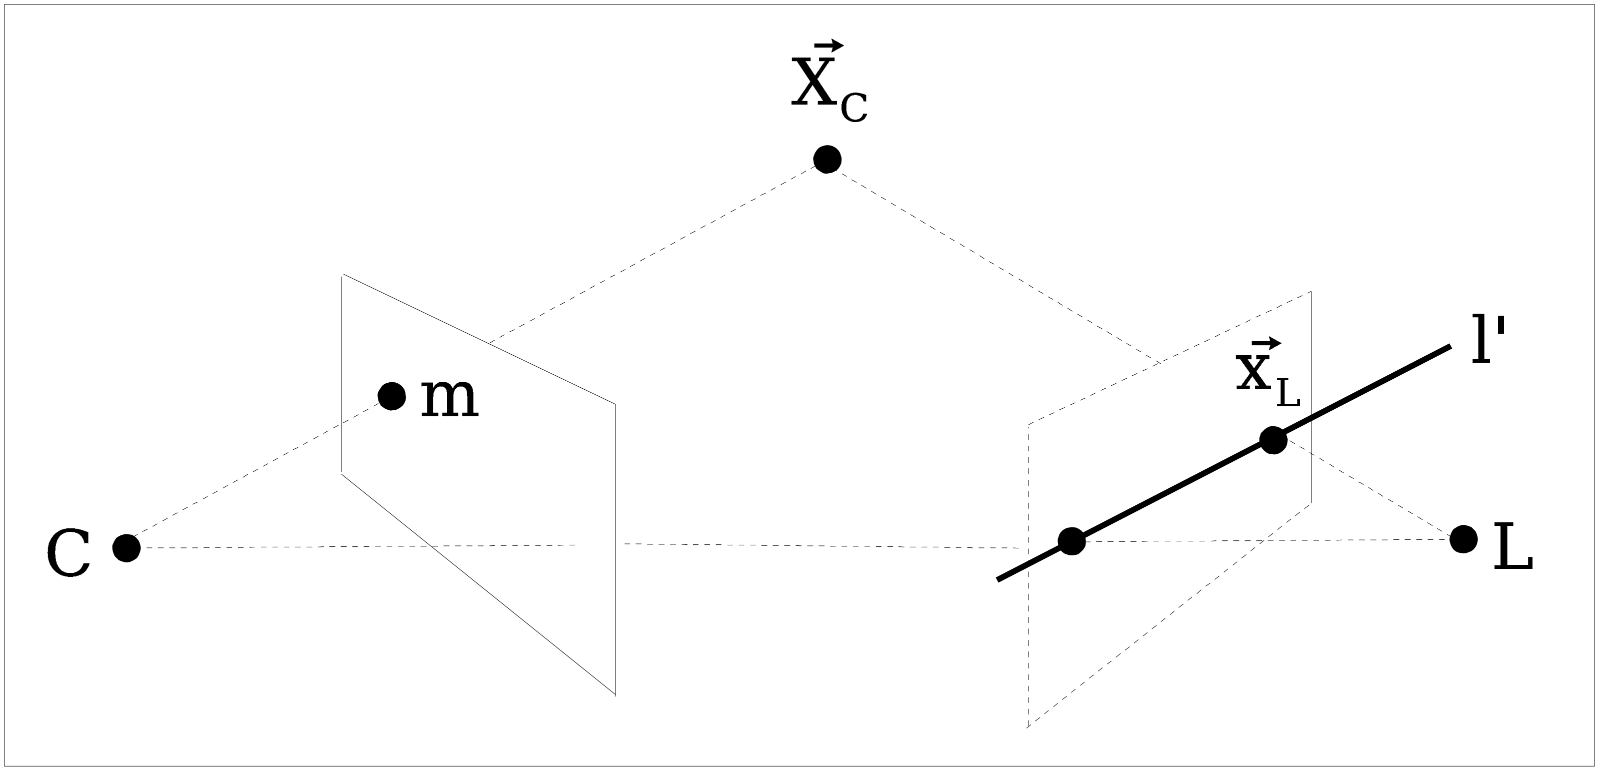

Epipolar Calibration

Two cameras constrain laser control to the intersection of epipolar lines in (u, v) space.

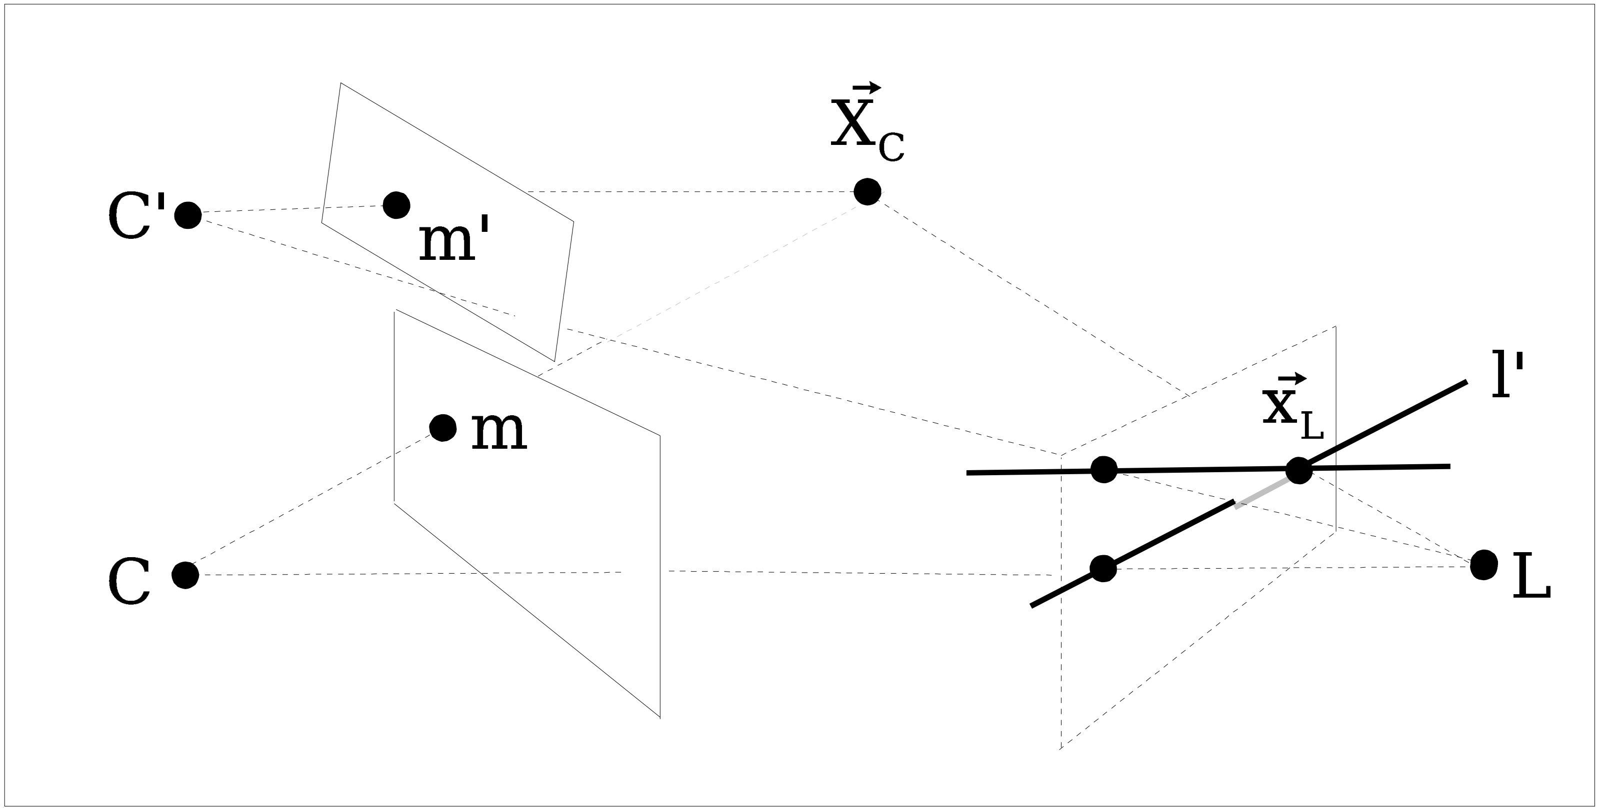

Epipolar Calibration

- Fundamental matrix F encodes this geometric relationship
- Each correspondance provides 1 constraint on F
- Utilize Hartley's Normalized 8 point algorithm to solve for F
- Need to solve for two F's:
  - Camera 1 and Laser
  - Camera 2 and Laser

Deriving Laser Control

Requires:

- Two fundamental matrices acquired during calibration
- Image coordinates of the target in each camera

Plugging these in yields:

- Two linear constraints (one for each camera)
- Two Unknowns (u, v)

Solve directly for laser control (u, v).

Experimental Procedure

Calibration

- Move laser to an arbitrary (u, v) coordinate
- Click on laser position in camera image
- Laser position, clicked image position define corresponding points.
- Laser moved in regular grid along image
- Repeated at several different depth planes

Experimental Procedure

Targeting

- Targets are the 4 extreme corners on a chessboard
- Error is difference between actual position and target in mm
- Test at 3 positions

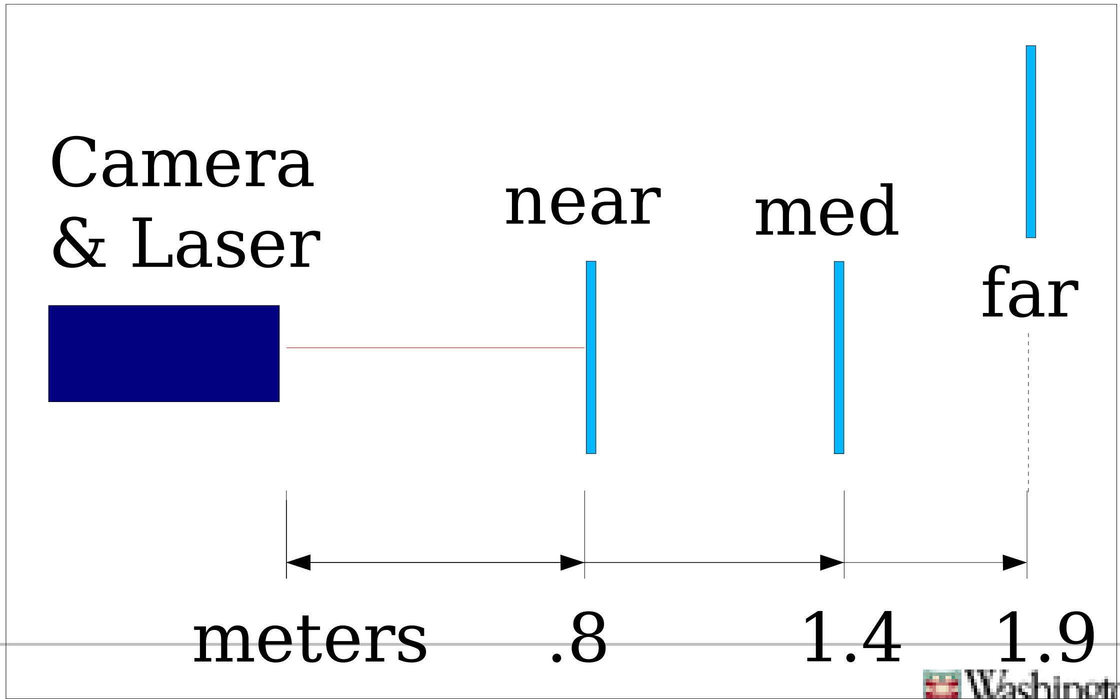

Parameter optimization

- Number of calibration planes
- Number of calibration points/plane
- Maximum angle of laser

See paper for details.

Results

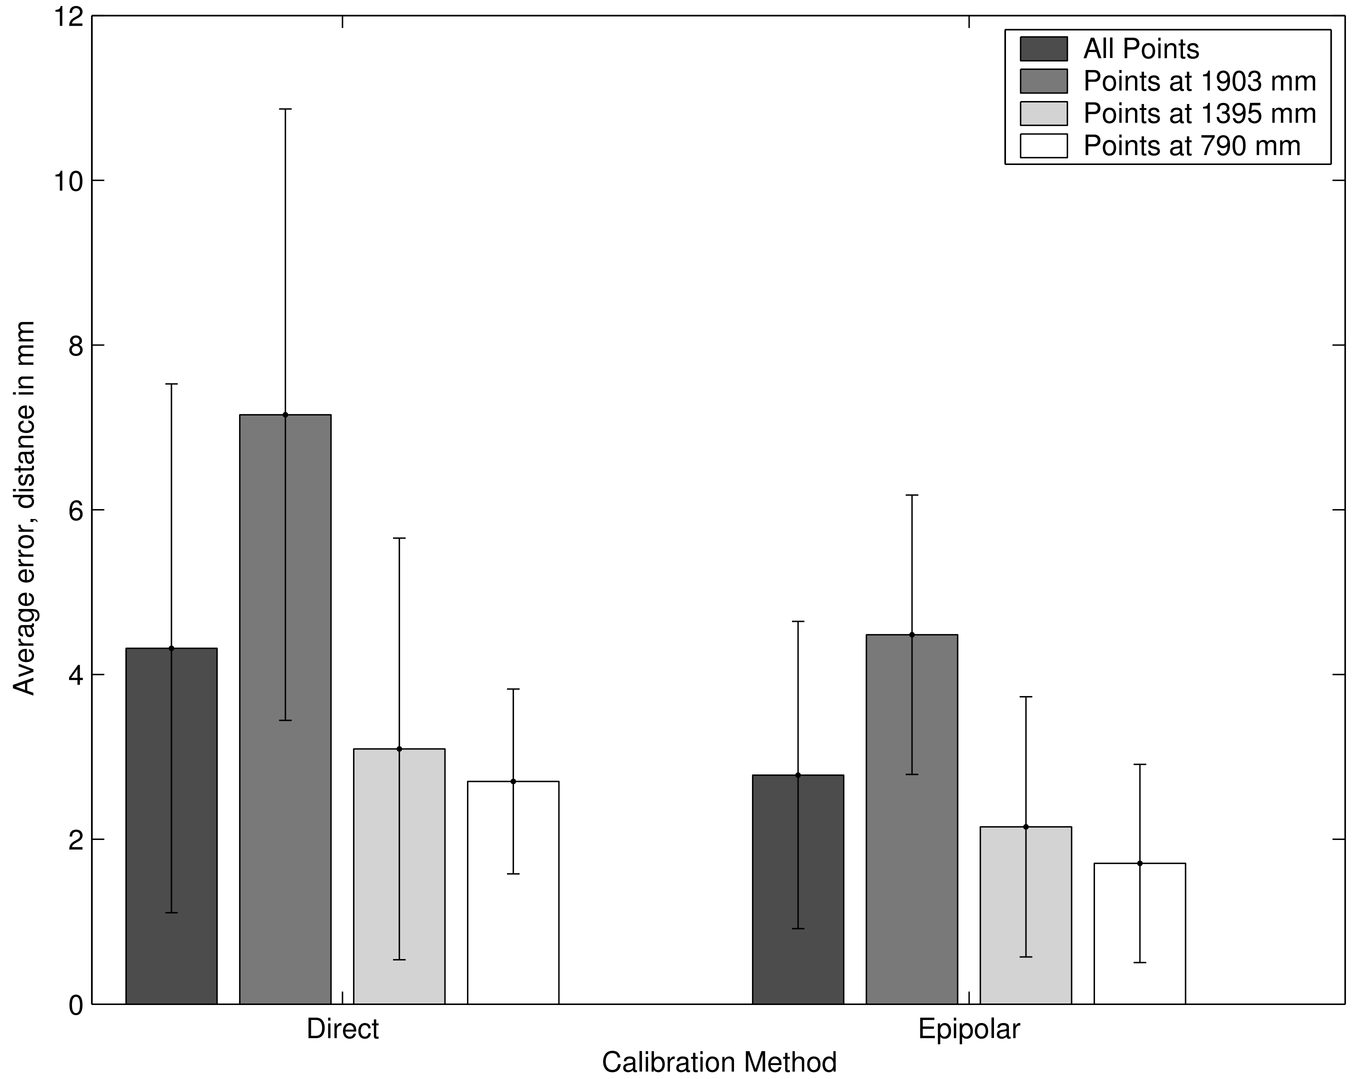

All Points
Points at 1903 mm
Points at 1395 mm
Points at 790 mm
Average error, distance in mm
0
2
4
6
8
10
12
Direct
Epipolar
Calibration Method

Discussion

- Both methods accurate to within $3-4$ mm on average
- Epipolar method slightly better at all depths
- Why?
  - Maturity of fundamental matrix solution method.
  - Noise in 3D sensor (epipolar method uses image coordinates directly)

Automatic calibration

- Mouse clicking is tiresome and prone to inaccuracies
- Automatic detection must consider laser artifacts in camera image:

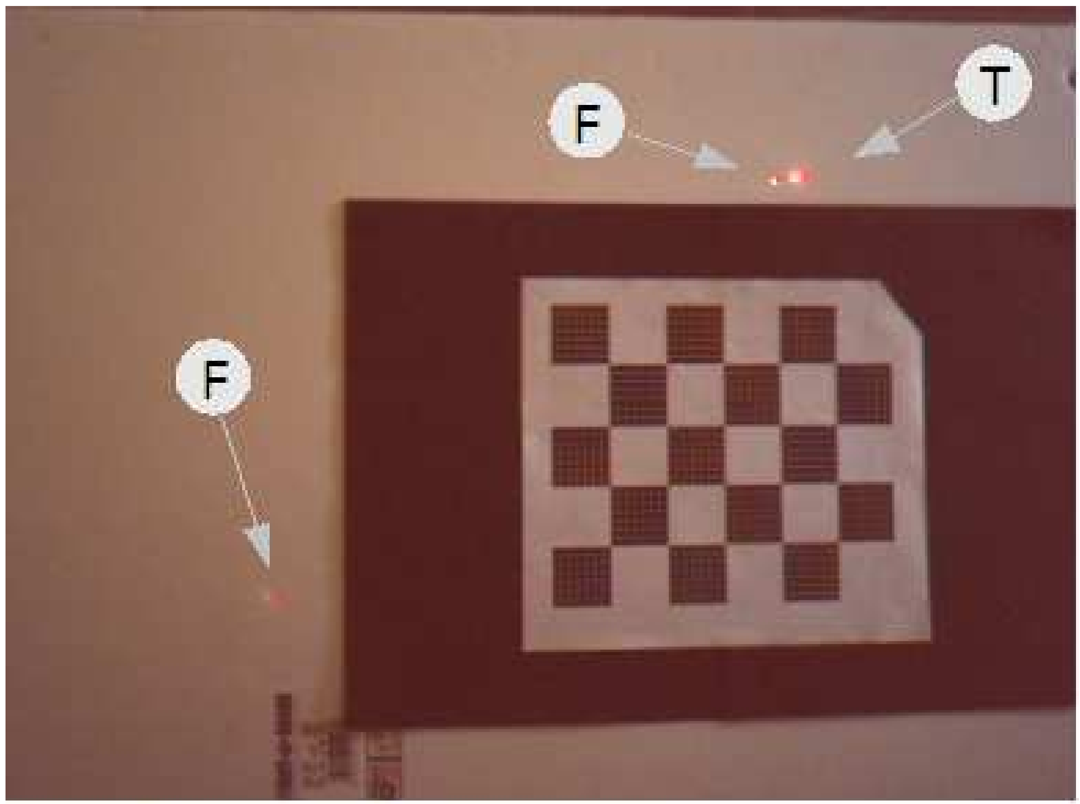

Red dot detection algorithm

- Capture background image (without laser)
- Capture image with laser, subtract out background image
- Keep red color channel only
- Threshold pixels
- Compute weighted center of mass (x, y) over entire image
- Recompute using a window around (x, y)

Results

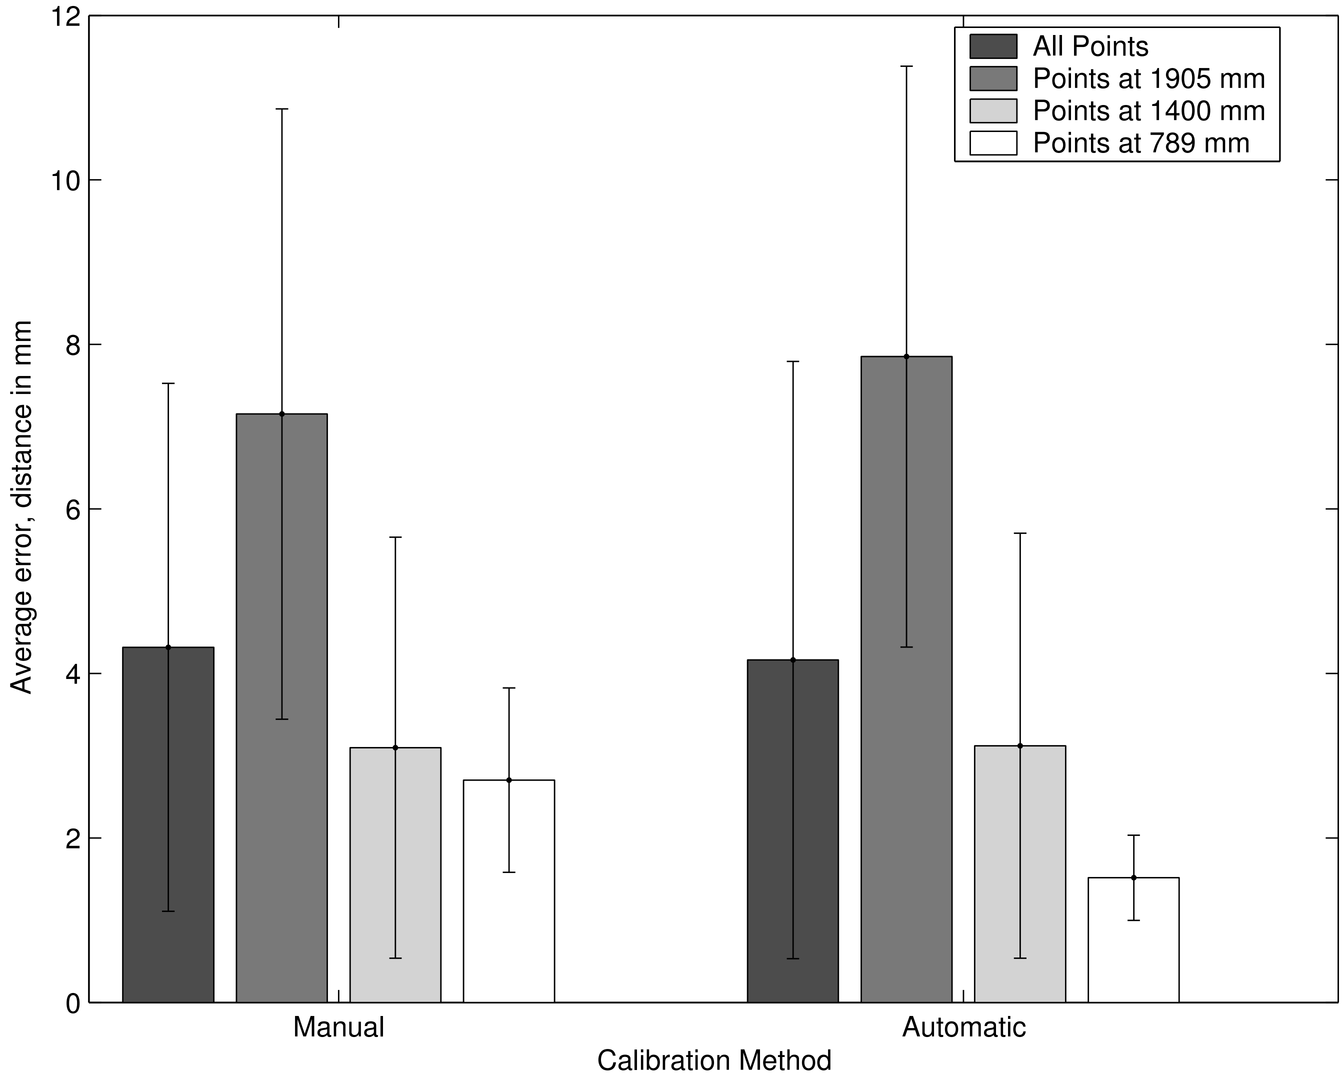

All Points
Points at 1905 mm
Points at 1400 mm
Points at 789 mm
Average error, distance in mm
0
2
4
6
8
10
12
Manual
Automatic
Calibration Method

Point Selection

- Currently specify laser coordinates
  - choose/detect corresponding image coordinate
- Stereo camera only provides sparse depth
  - Points without depth are thrown away during calibration
- Can we specify image coordinates, then move laser to match?
  - manual control laborious
  - automatic control (chicken and egg problem?)

Image Point Selection

- Algorithm:
  1. Measure distance between laser & target
  2. Move $\alpha \cdot \mathbf{x.distance}, \beta \cdot \mathbf{y.distance}$
  3. Repeat until distance = 0
  4. α, β are constants determined empirically to minimize distance
- Will probably only work if coordinate systems are roughly aligned.
- Highly unsophisticated instance of visual servoing methodolgy.
  - could be easily improved to be more robust

Experimental Procedure

- Use chessboard corners as calibration points
  - Take advantage of automatic corner detection
- Repeat for 4 positions
  - Use 8 points at each position

Results

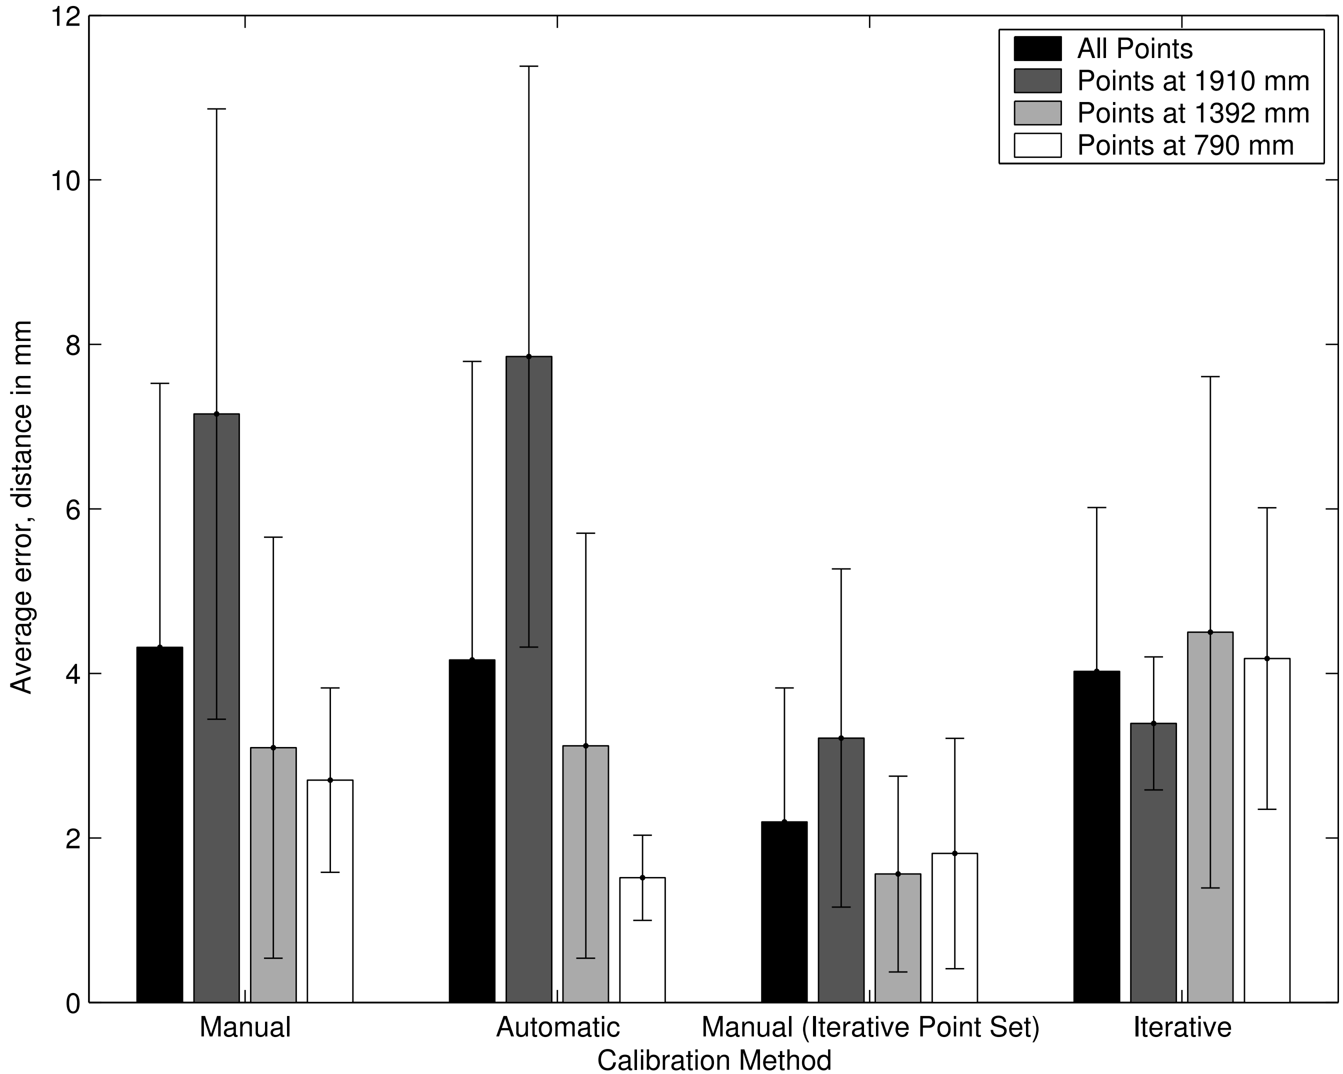

All Points
Points at 1910 mm
Points at 1392 mm
Points at 790 mm
Average error, distance in mm
0
2
4
6
8
10
12
Manual
Automatic
Manual (Iterative Point Set)
Iterative
Calibration Method

Discussion

- Both manual and automatic calibration show improvement with inverse selection.
- Point selection is more important than number of calibration points.
- Improvement possibly due to sub-pixel accuracy of corner detection.

Overall Discussion

- The best overall average accuracy achieved is around 2.5 mm.
- Good, but not perfect – bias.

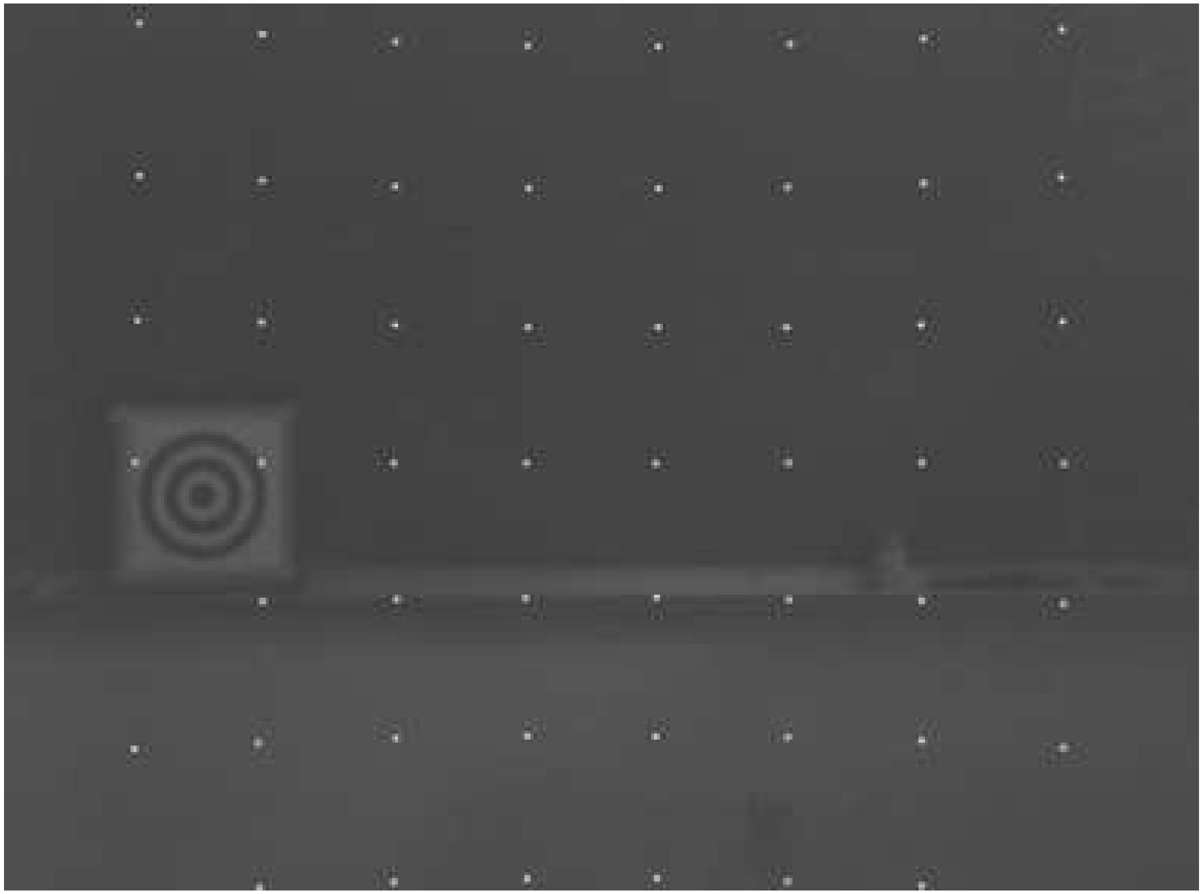

Future Work

- Identify and model non-linearity in laser unit.
- Evaluate in comparison to visual servoing as an alternative targeting approach.

Conclusion

- Two calibration methods
  - Verified by experimental results to 3-4 mm accuracy
- Automatic laser point detection
- Image point correspondence
  - Verified by experimental results to 2.5 mm accuracy

Acknowledgements

- Technical Support Working Group (TSWG) for funding the project encompassing this work.
- Drs. Bill Smart and Robert Pless for advising.
- Michael Dixon for red dot detection algorithm.
- Rachel Whipple for data entry.