Opportunistic Perception

Paul Fitzpatrick

machine perception

training data + learning = perception

machine perception

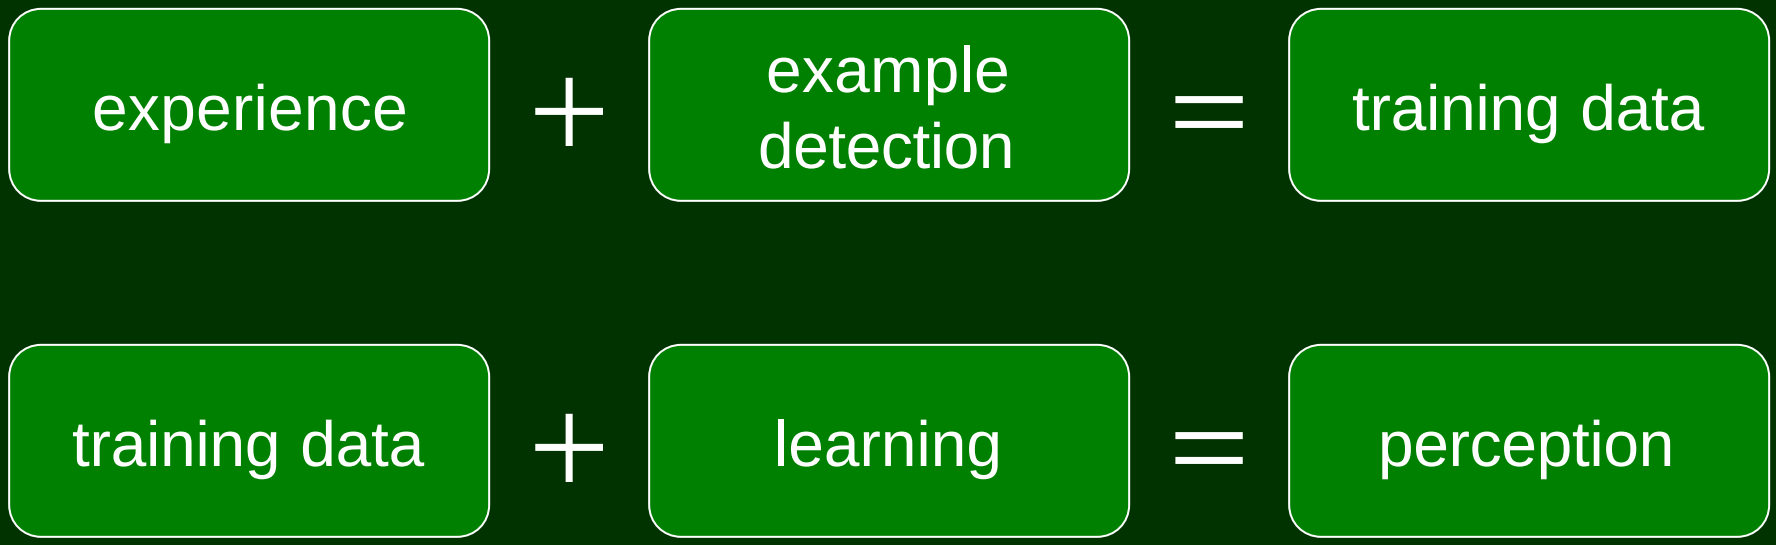

machine perception

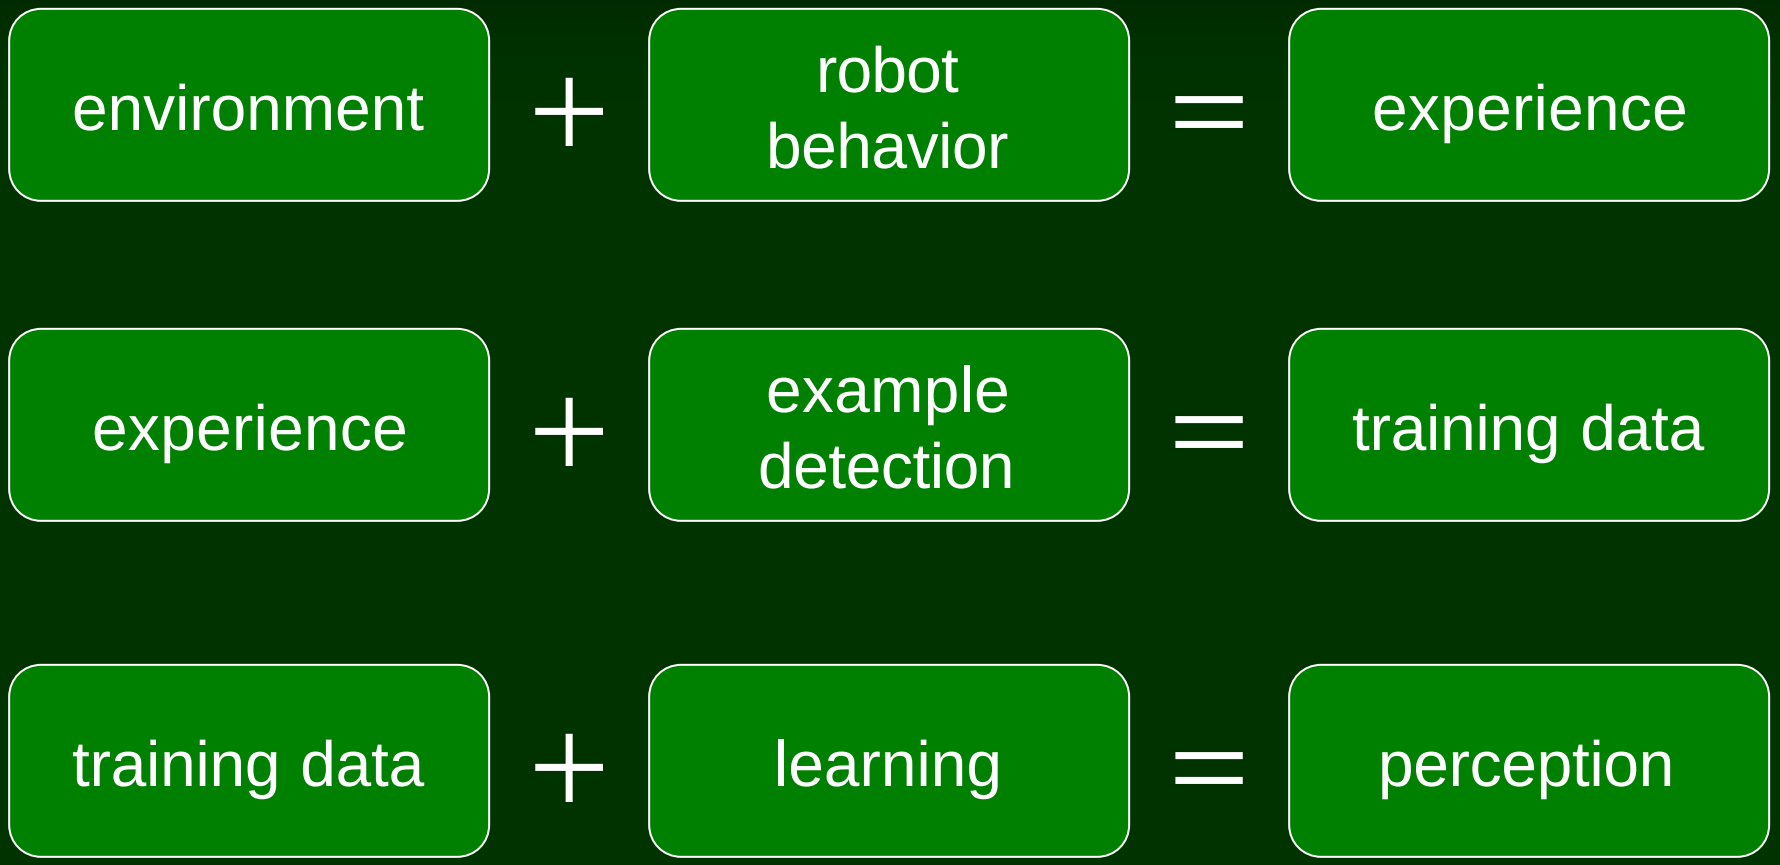

Active Segmentation:

an active vision approach to object segmentation

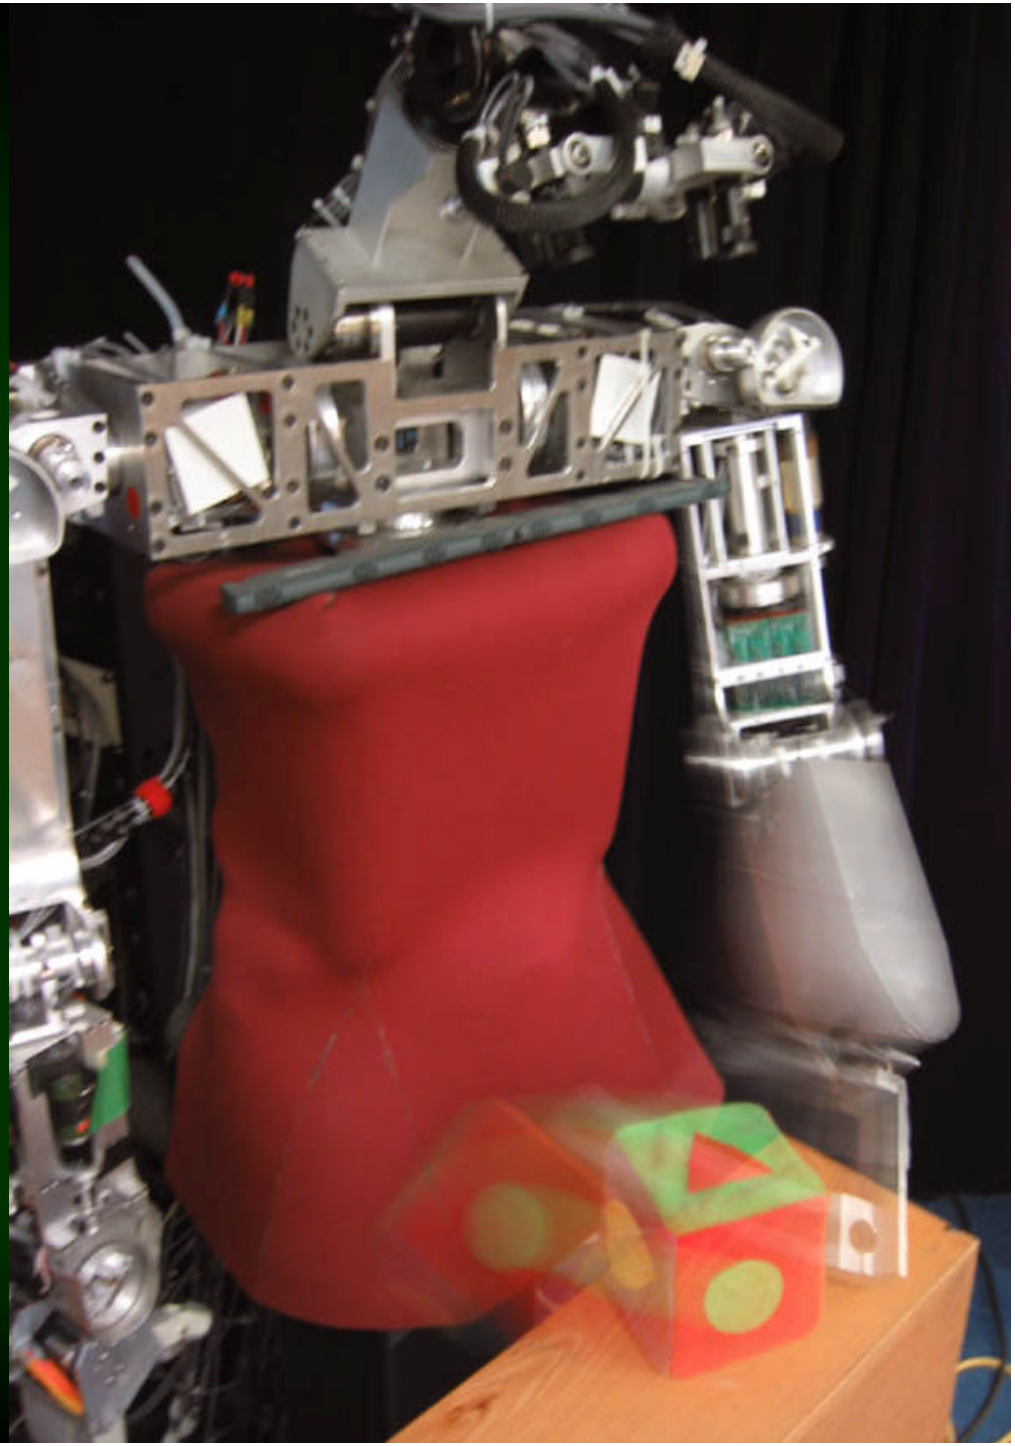

active segmentation

- Object boundaries are not always easy to detect visually
- Solution: Robot sweeps arm through ambiguous area
- Any resulting object motion helps segmentation
- Robot can learn to recognize and segment object without further contact

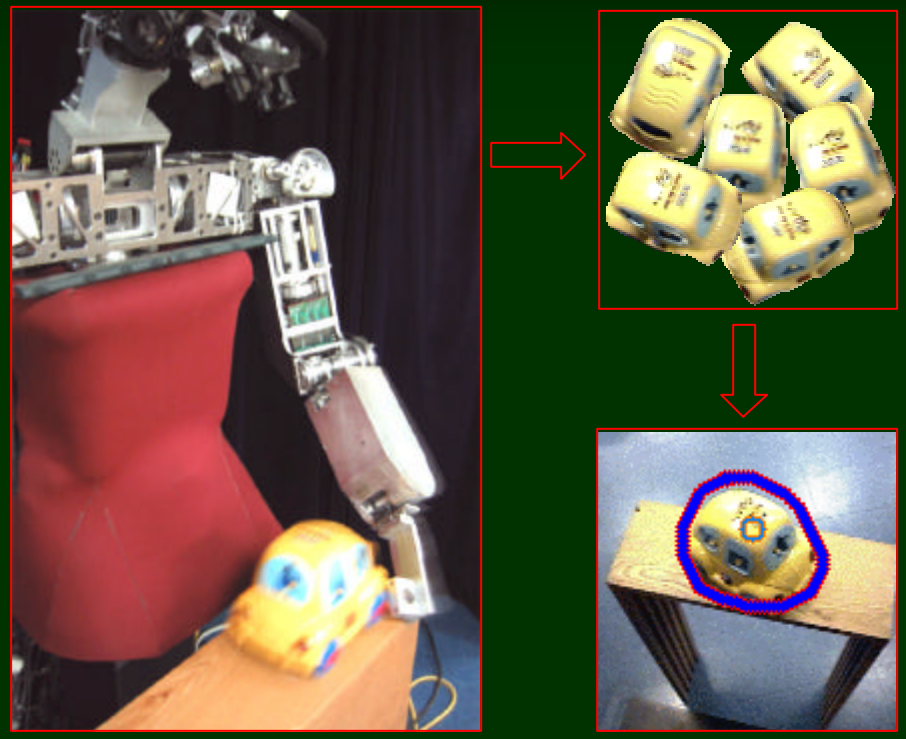

segmentation example

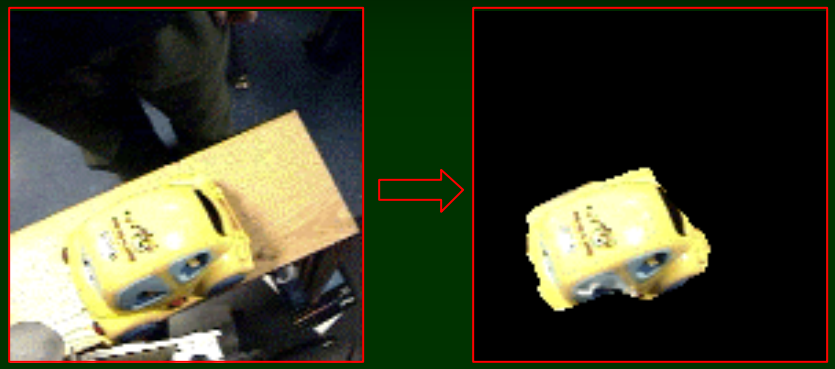

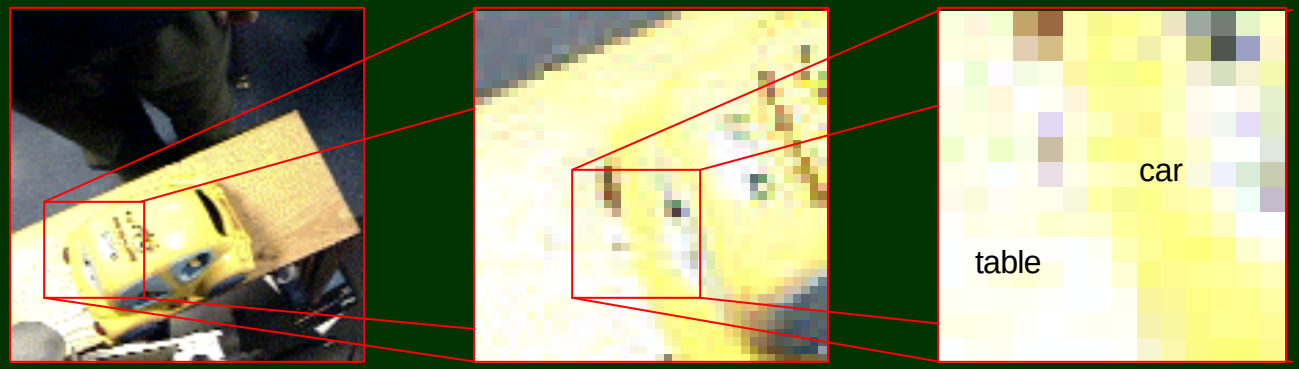

what is it good for?

- Not always practical!
- No good for objects the robot can view but not touch
- No good for very big or very small objects
- But fine for objects the robot is expected to manipulate

Head segmentation the hard way!

learning about and exploiting affordances

with
Giorgio Metta

a bottle
it rolls along its side

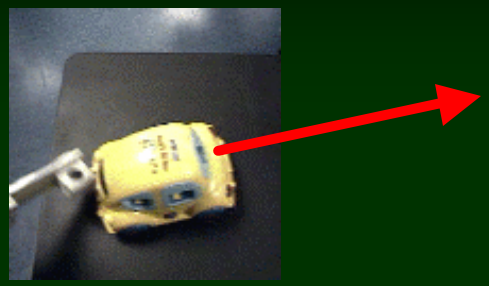

a toy car
it rolls forward

a toy cube
it doesn't roll easily

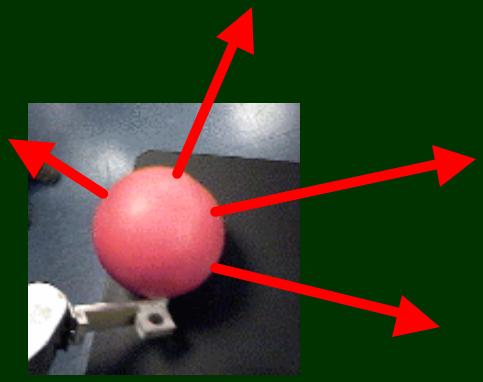

a ball
it rolls in
any direction

Feel the Beat: using amodal cues for object perception

with
Artur Arsenio

amodal versus modal cues

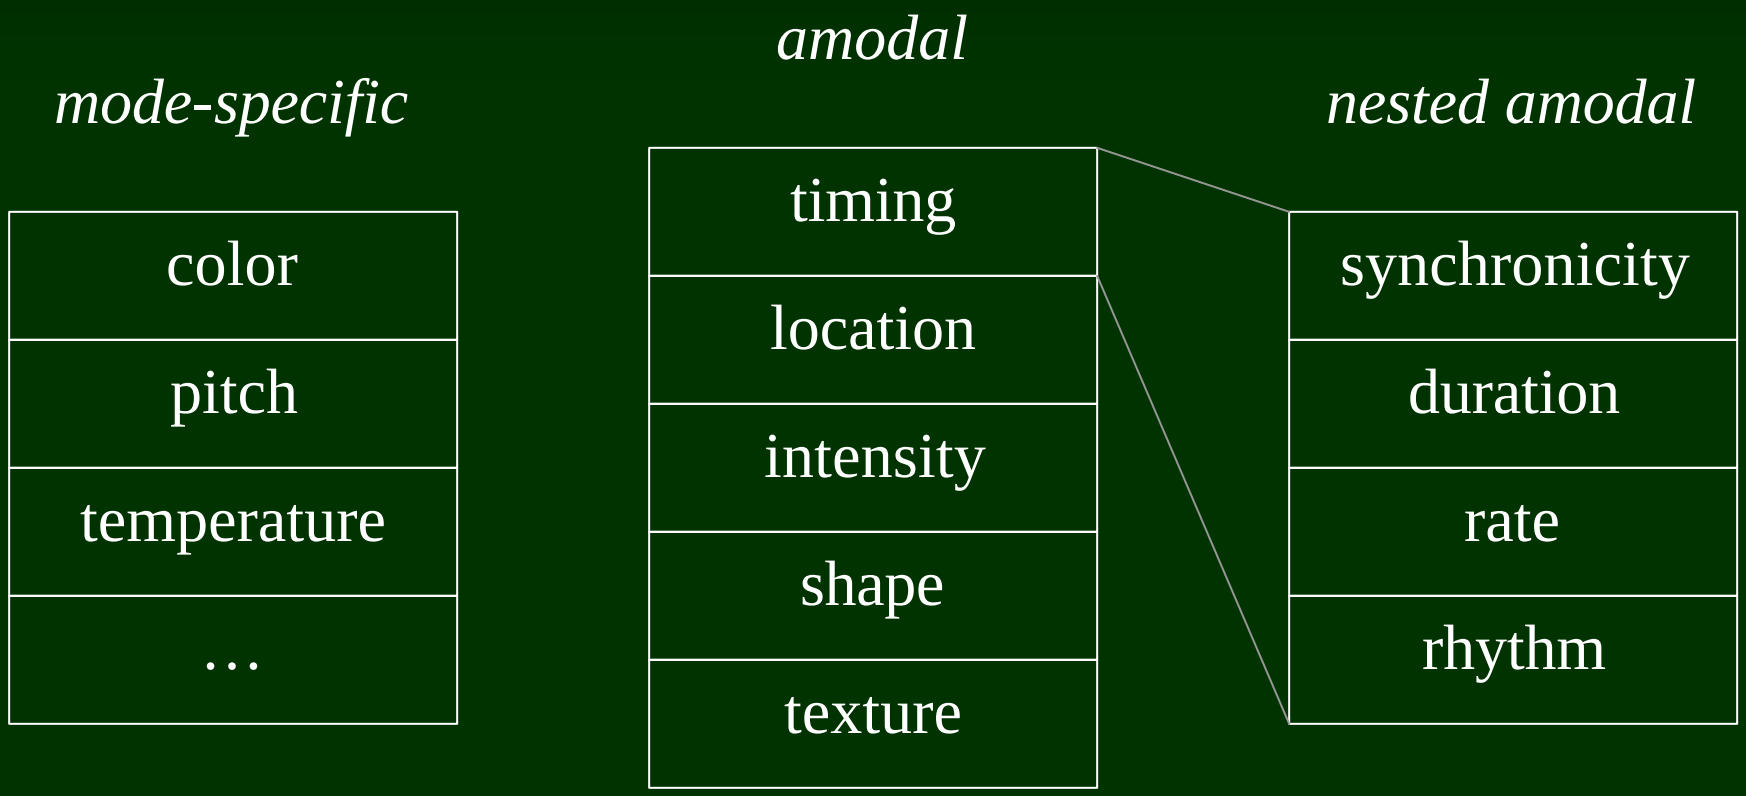

matching sound and vision

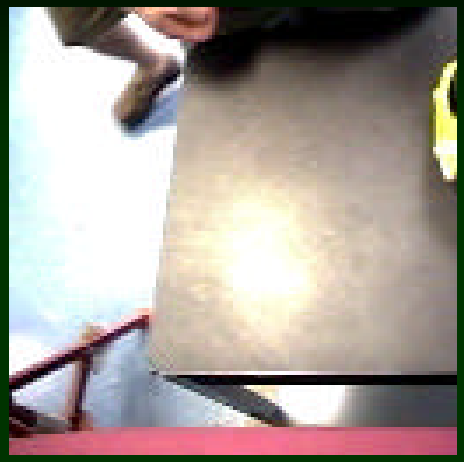

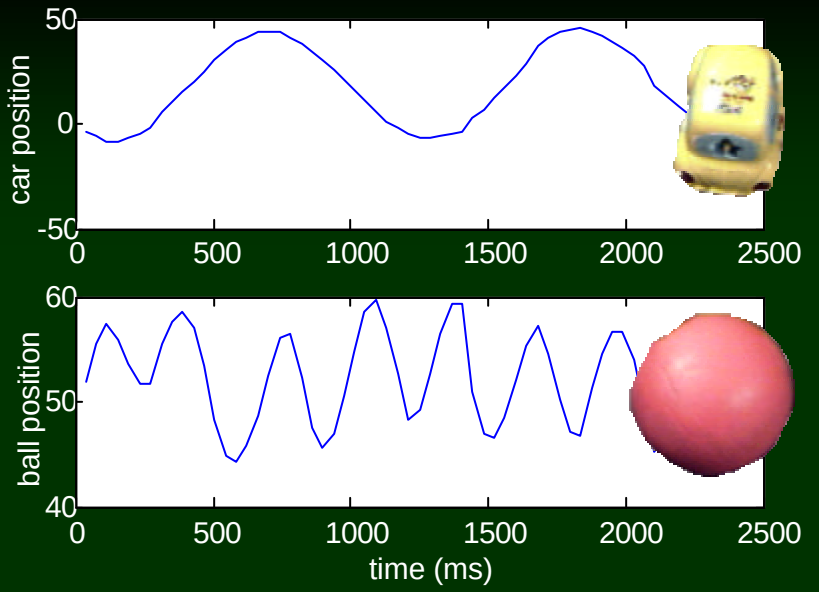

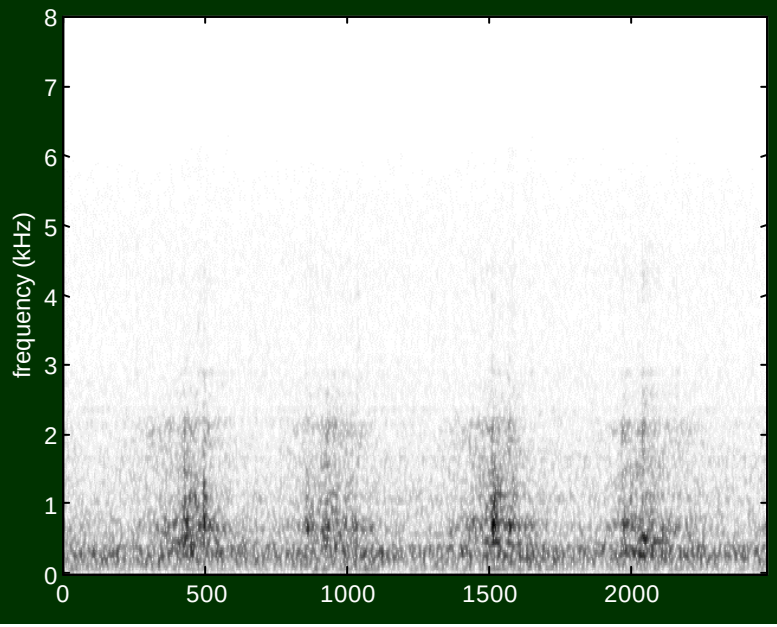

- One object (the car) making noise
- Another object (the ball) in view
  - Problem: which object goes with the sound?
  - Solution: Match using amodal cues (period) and intermodal cues (relative phase)

Cross-modal object recognition

Causes sound when changing direction after striking object; quiet when changing direction to strike again

Causes sound while moving rapidly with wheels spinning; quiet when changing direction

Causes sound when changing direction, often quiet during remainder of trajectory (although bells vary)

Cross-modal object recognition

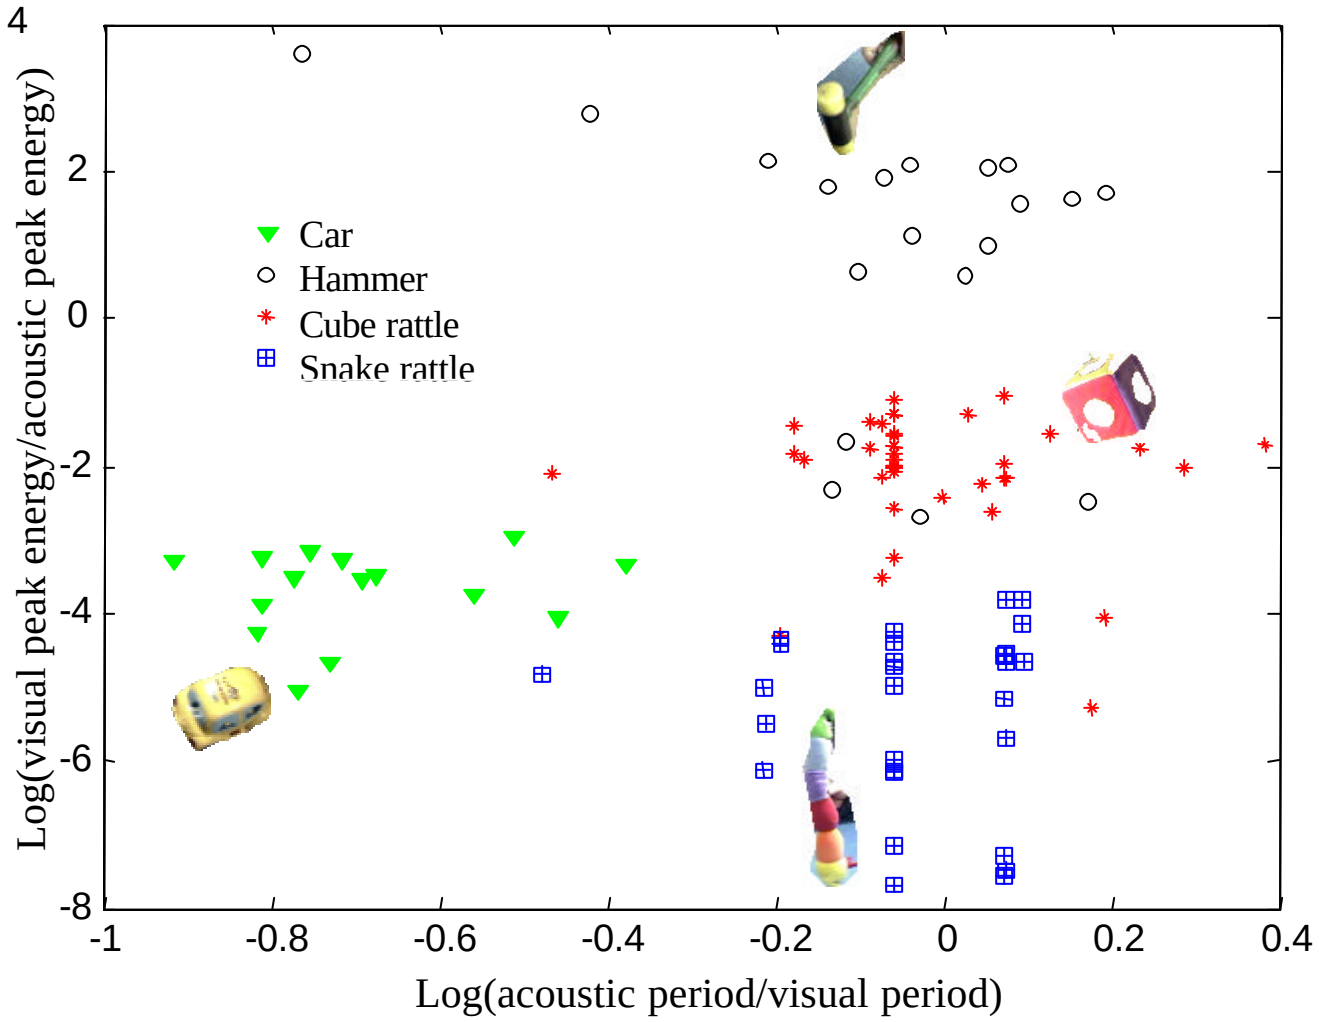

Cross-modal recognition rate: 82:1%

recognizing the body

appearance, sound, and action of the arm all bound together

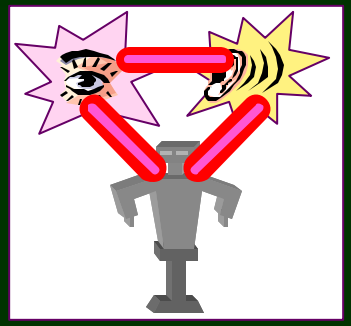

robot is looking towards its arm as human moves it

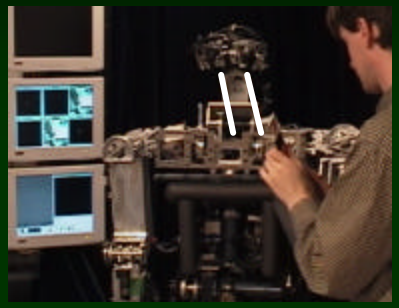

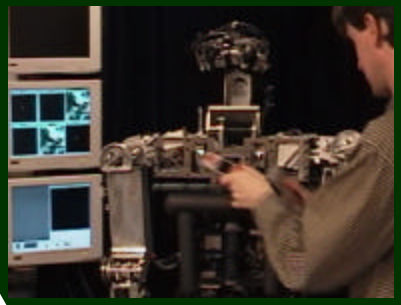

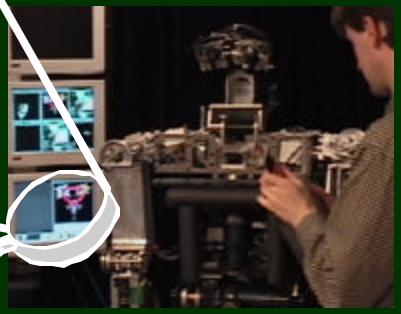

sound detected and bound to the motion of the arm

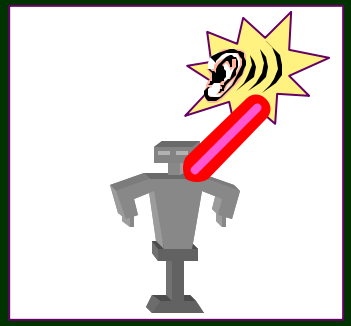

robot is looking away from its arm as human moves it

Shadowy Contacts: Time to contact from shadows

with

Eduardo Torres-Jara

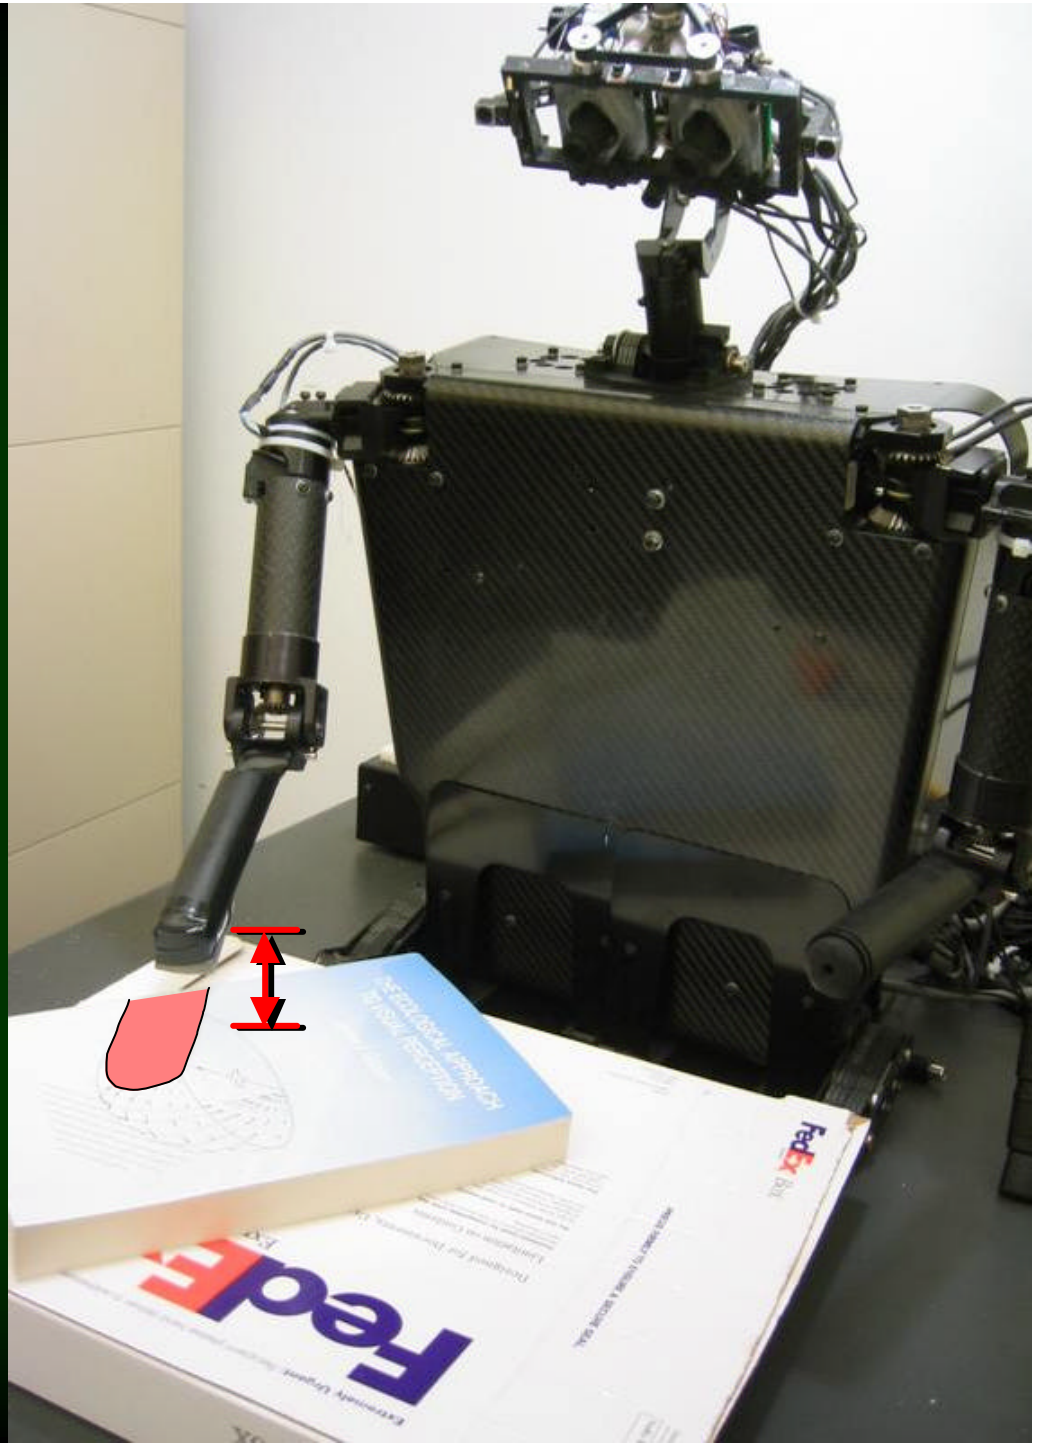

visually-guided touching using shadows

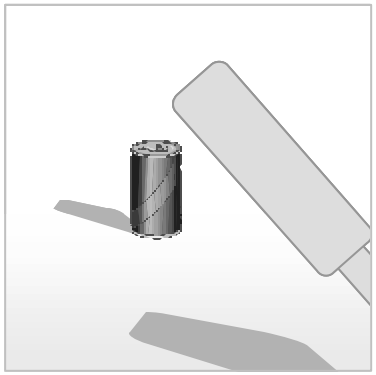

Robot sees target, arm, and arm's shadow

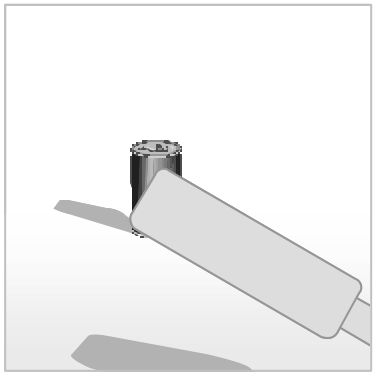

Robot moves to reduce visual error between arm and target

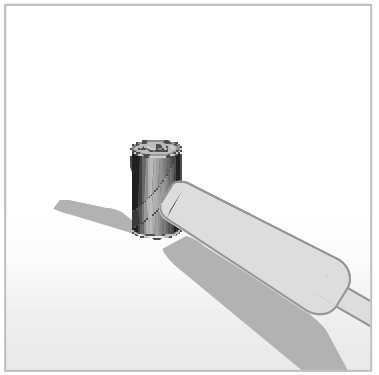

Robot moves to reduce visual error between arm's shadow and target

shadow cast by weak ambient light

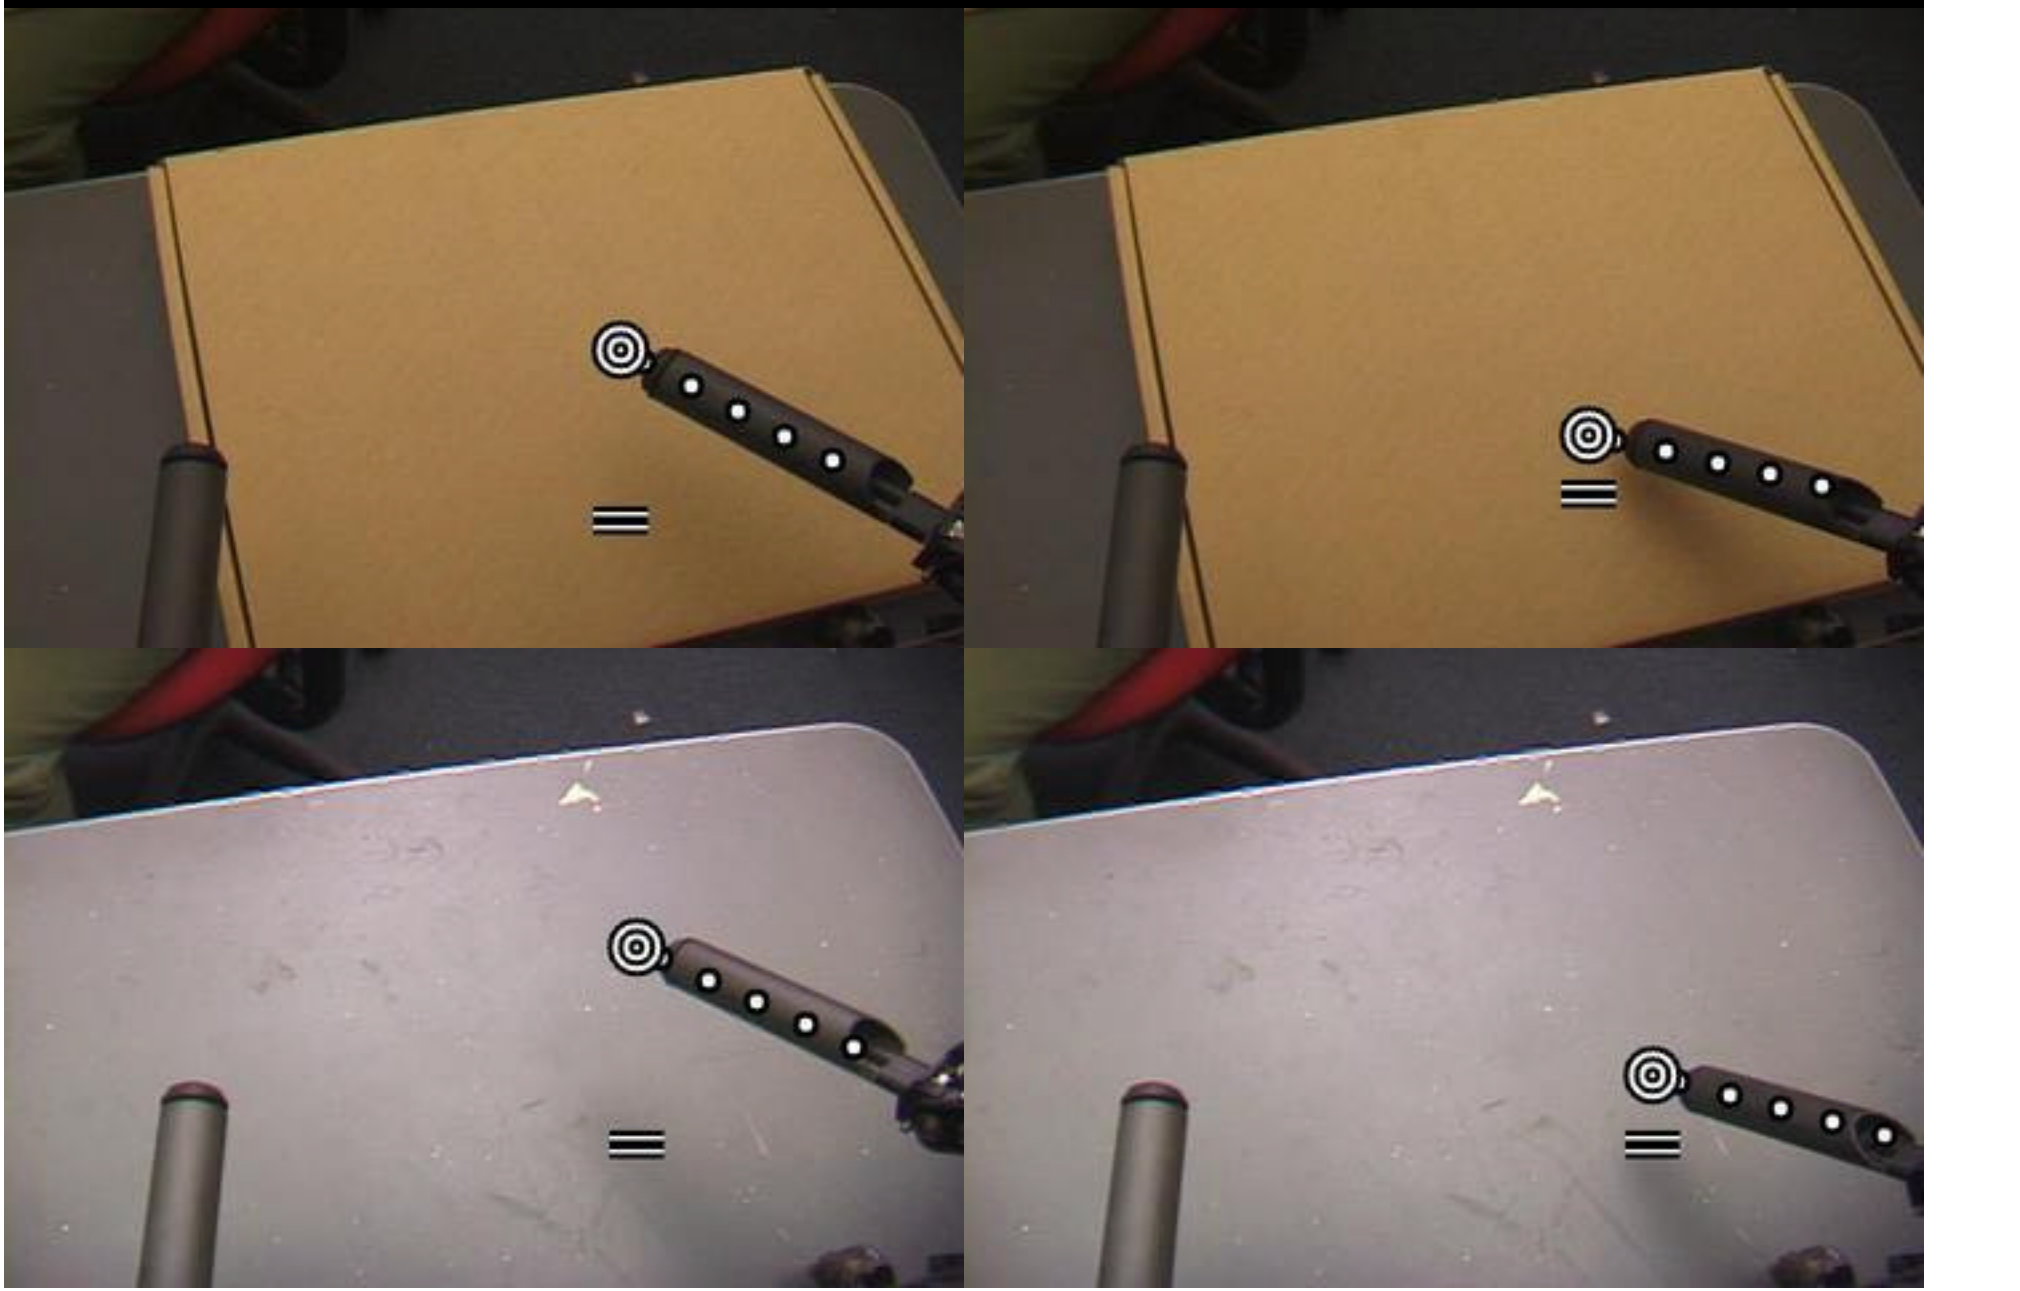

shadow cast by strong directional light

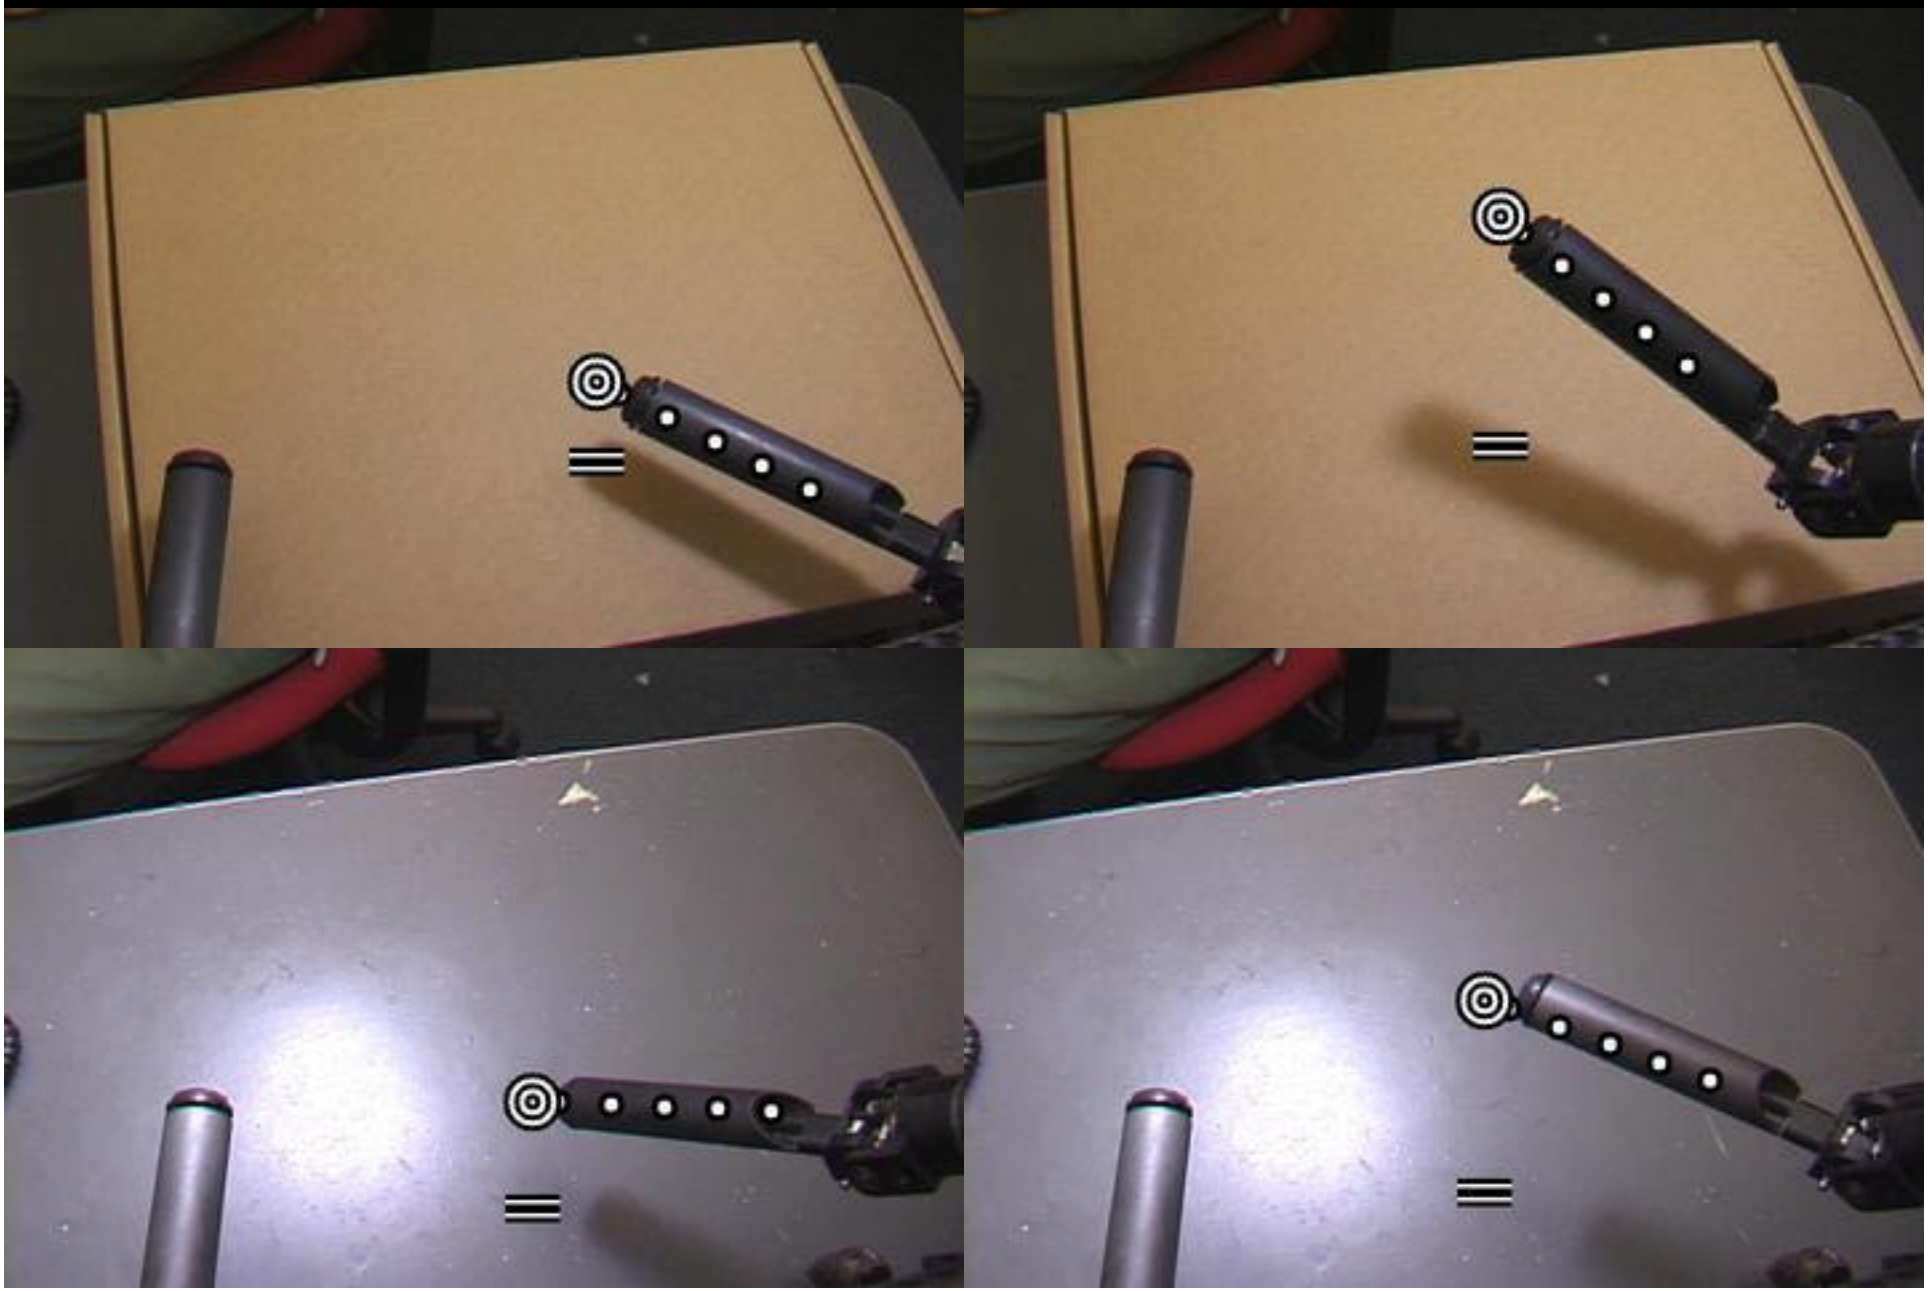

time to contact estimation

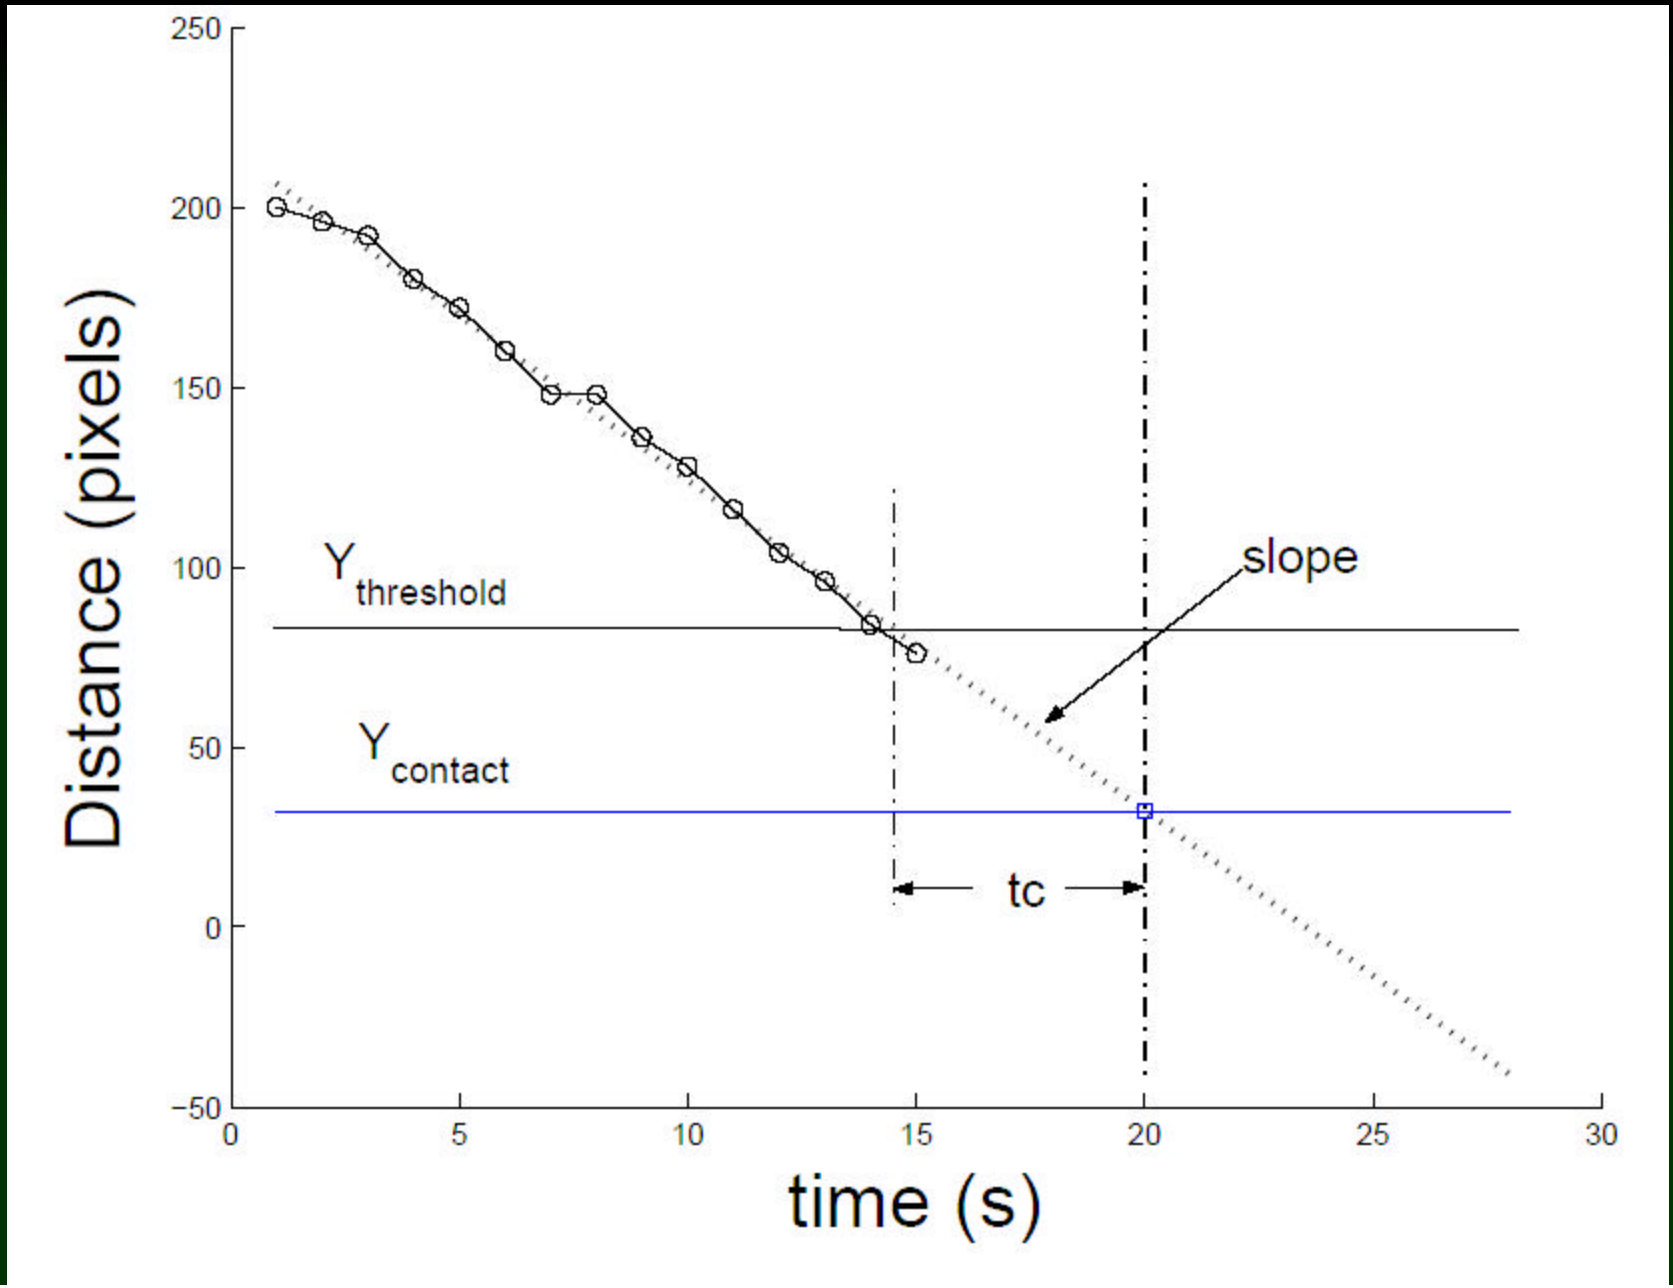

reflection of arm in mirror

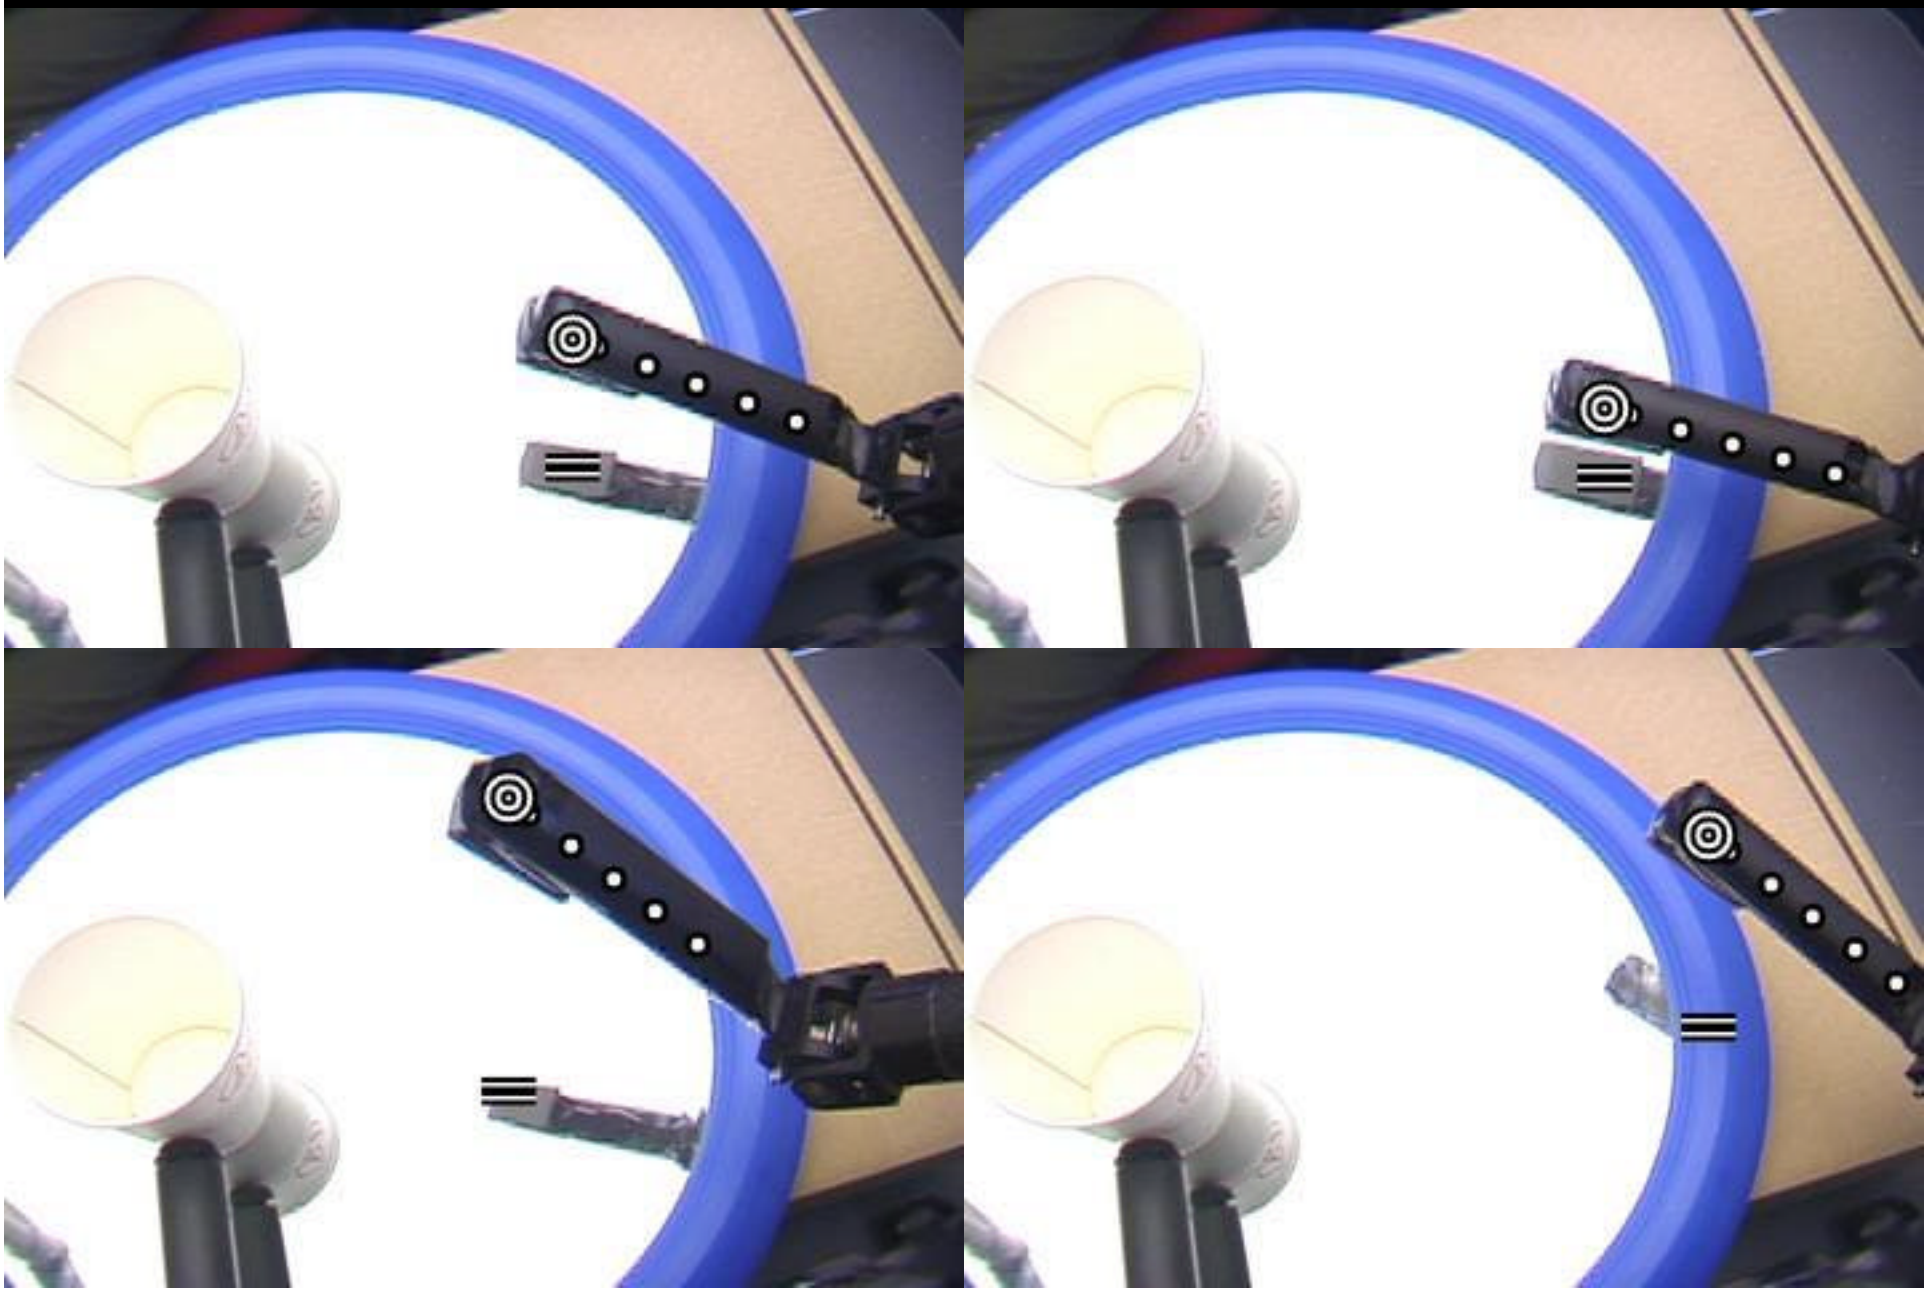

reflection of arm in water

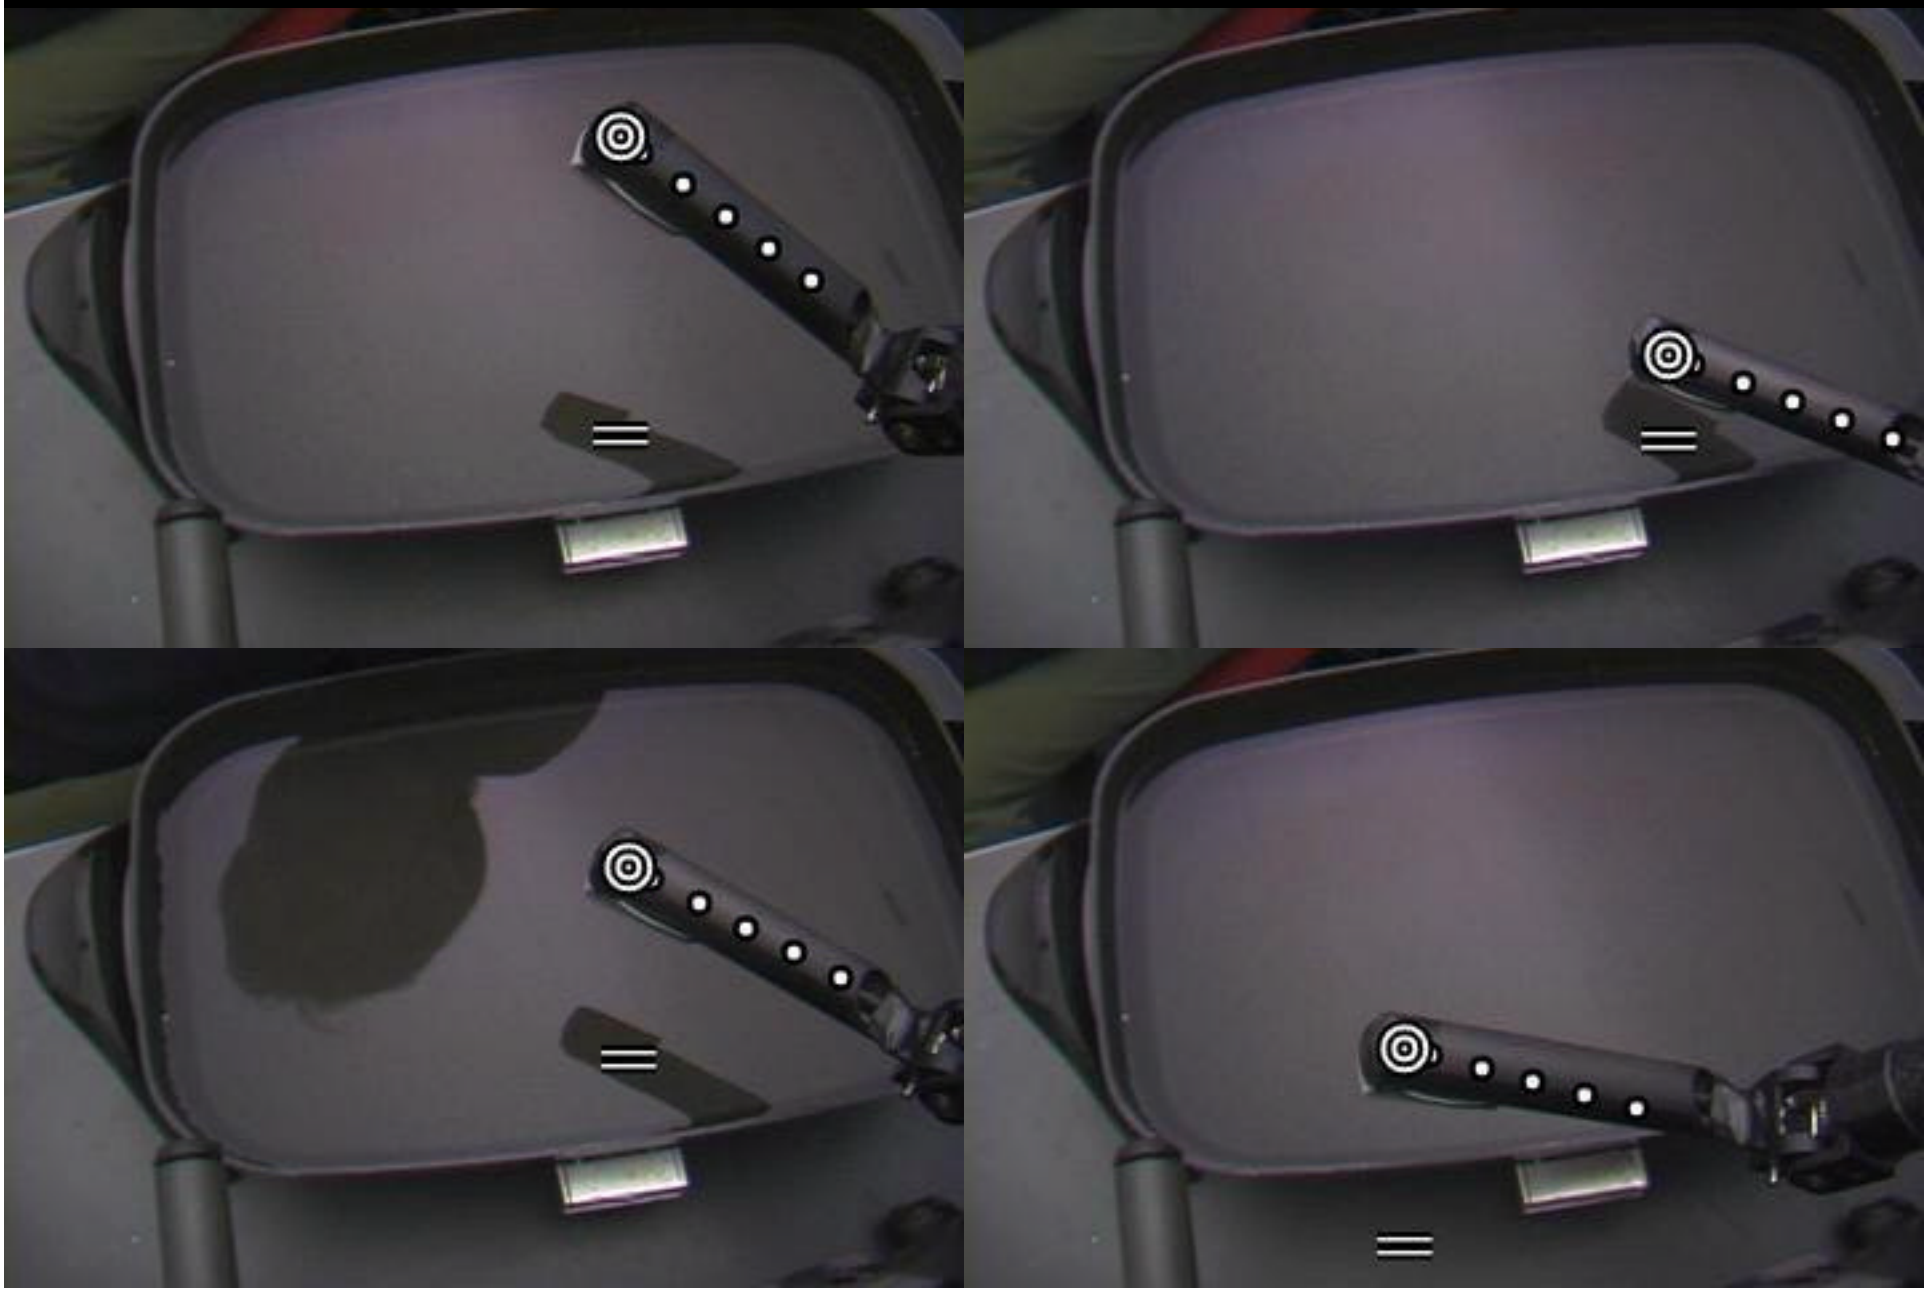

reflection of arm on acrylic

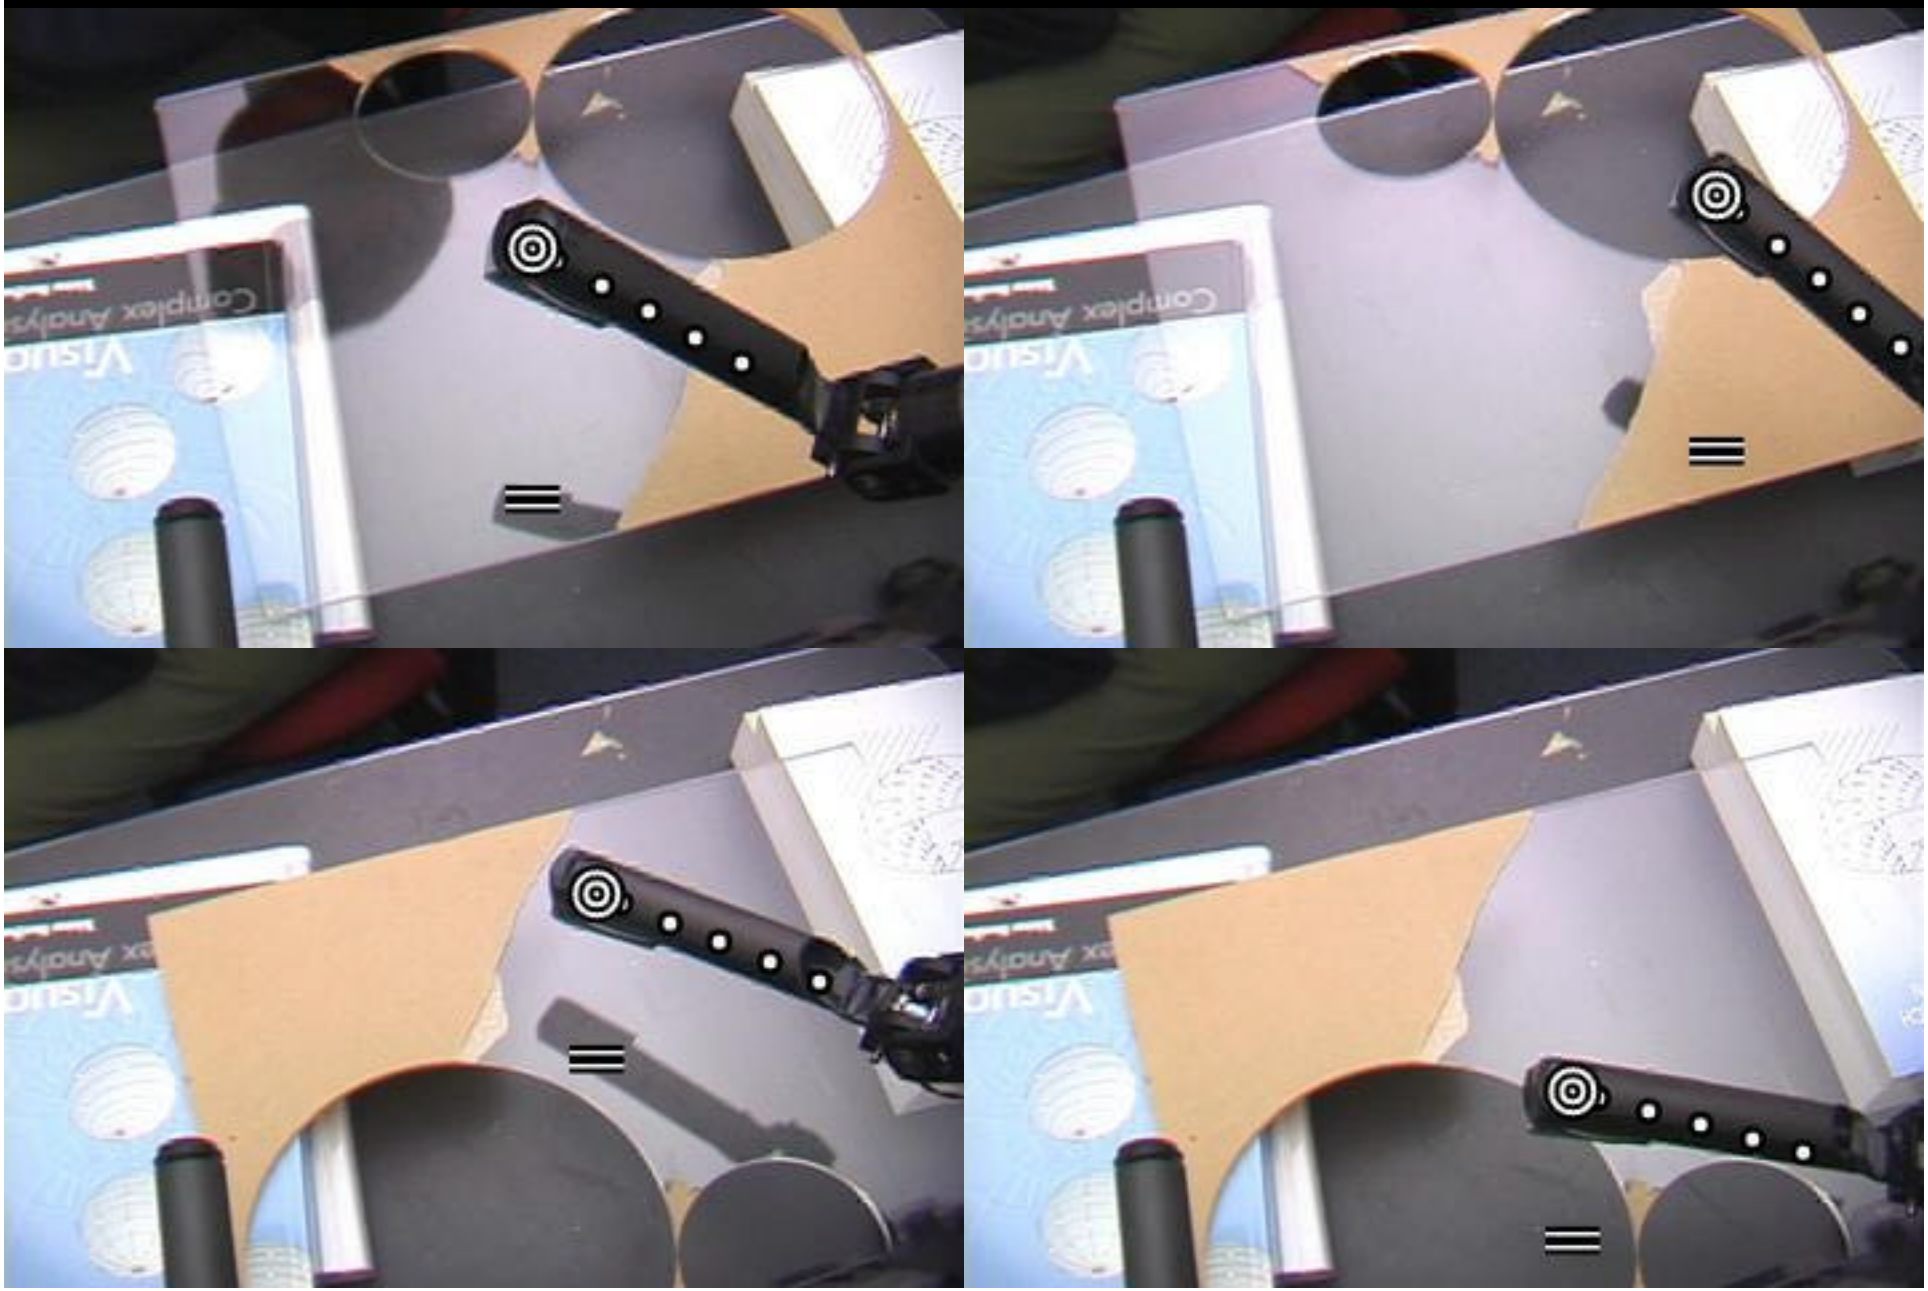

detecting object shadows

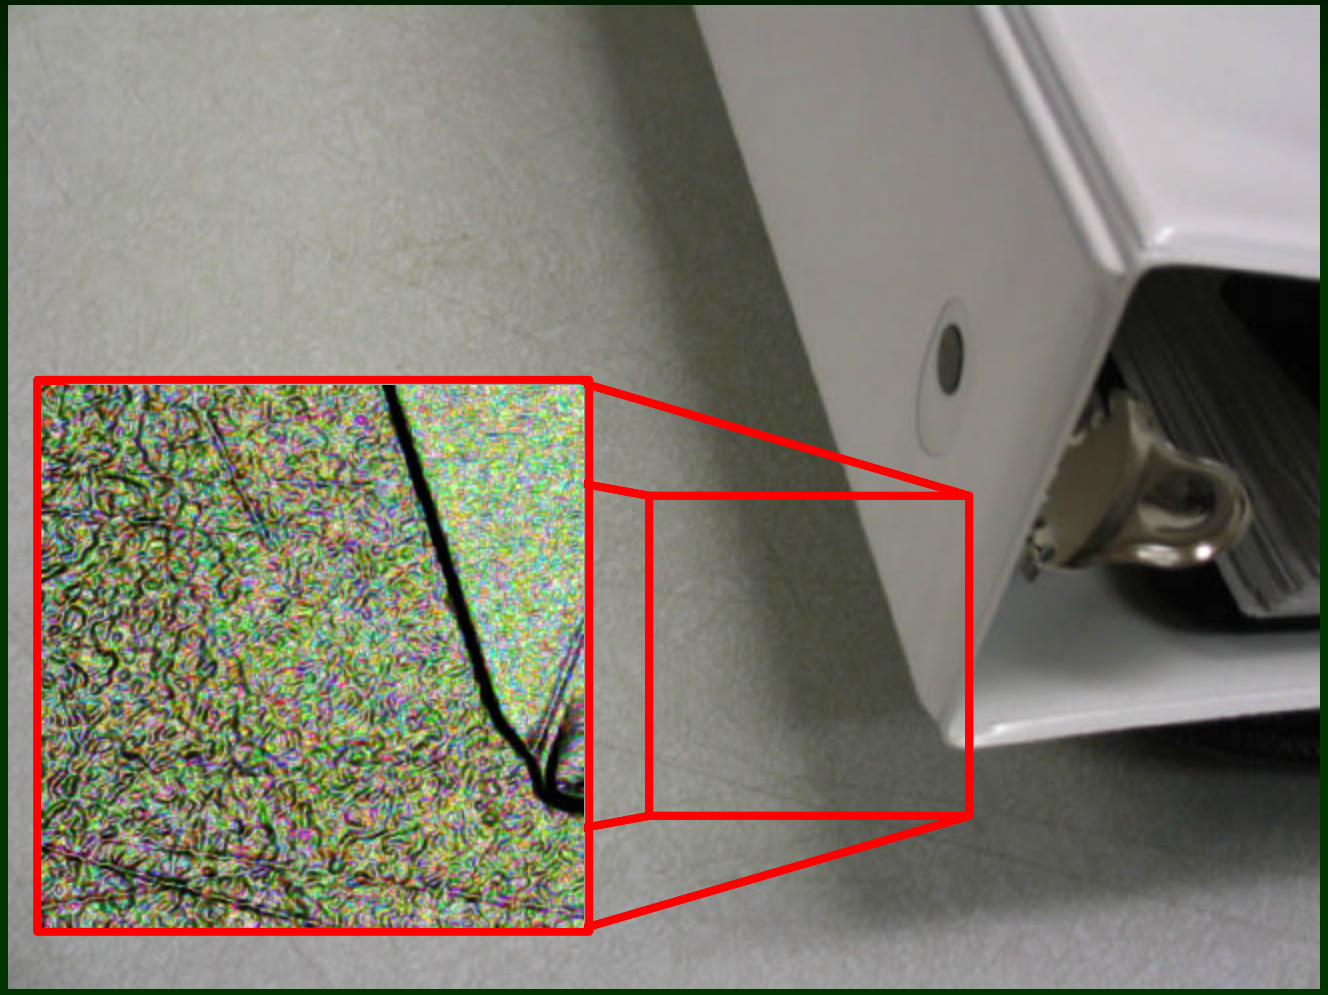

Platform Shoe: shoes as a platform for vision

with
Charlie Kemp

view from a shoe

detecting when the foot is planted

- darker image
- motion blur
- large time derivative

- lighter image
- motion blur
- large time derivative

- average image
- no motion blur
- small time derivative

the features

Image brightness

$$I_0 = \frac{1}{N}\sum_{x,y} I(x,y), \qquad N = \sum_{x,y} 1$$

Temporal derivative

$$\Delta I_t = \frac{1}{I_0 N}\sum_{x,y} |I(x,y,t) - I(x,y,t-1)|$$

Spatial derivative

$$\Delta I_x = \frac{1}{I_0 N}\sum_{x,y} |I(x,y,t) - I(x-1,y,t)|$$

Combined & Filtered

$$s = \alpha \Delta I_t - \beta \Delta I_x - \gamma I_0$$

gait analysis

spatial derivative

temporal derivative

image brightness

combined & filtered

swing/planted detection

orientation

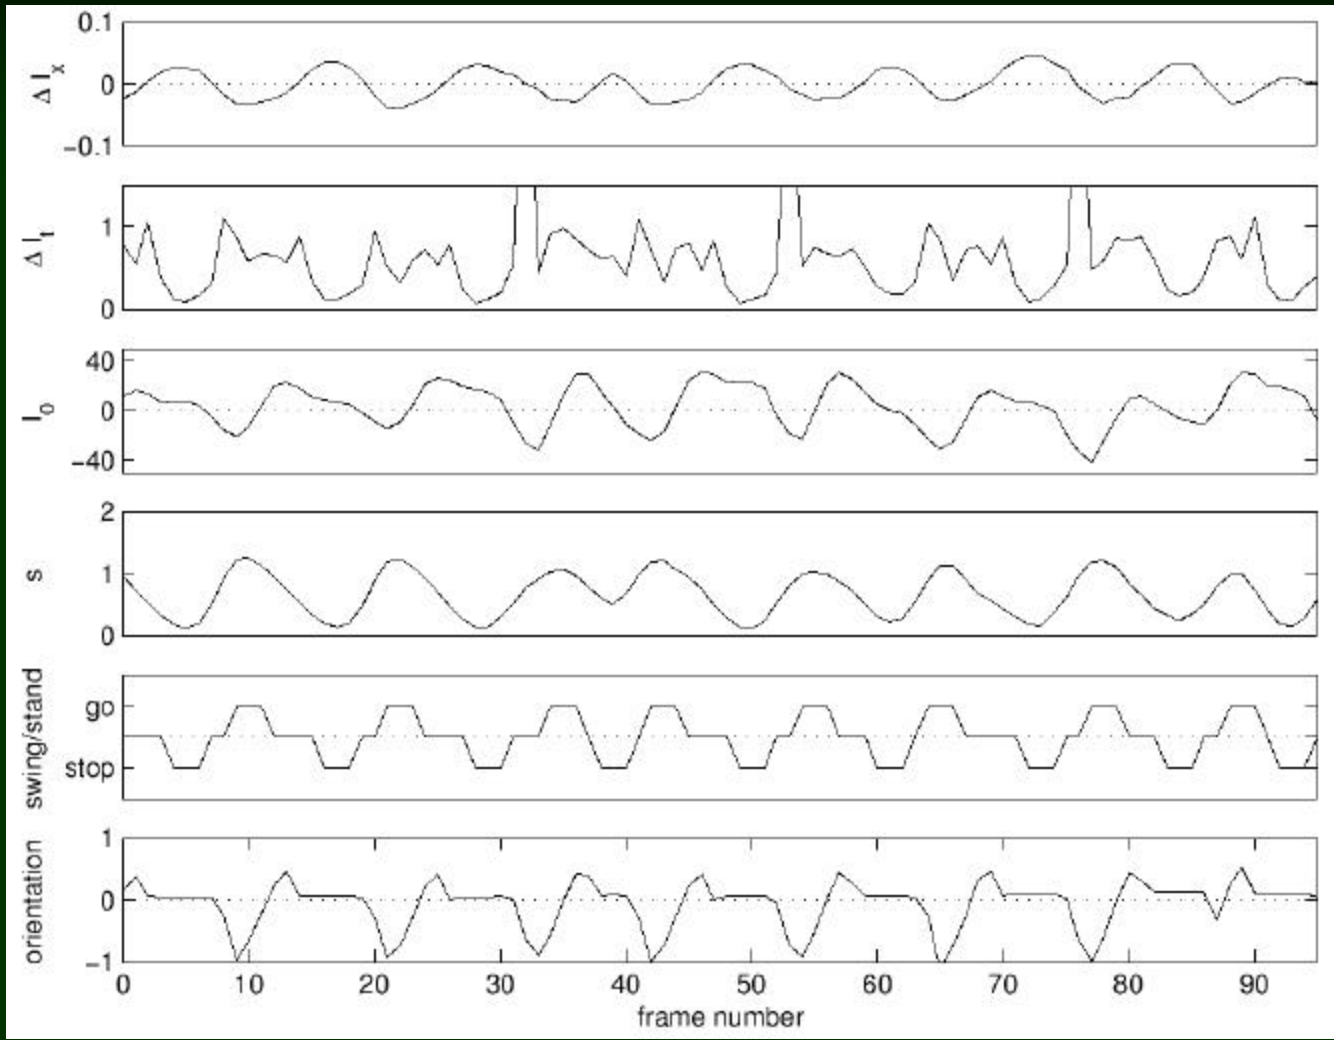

ground segmentation

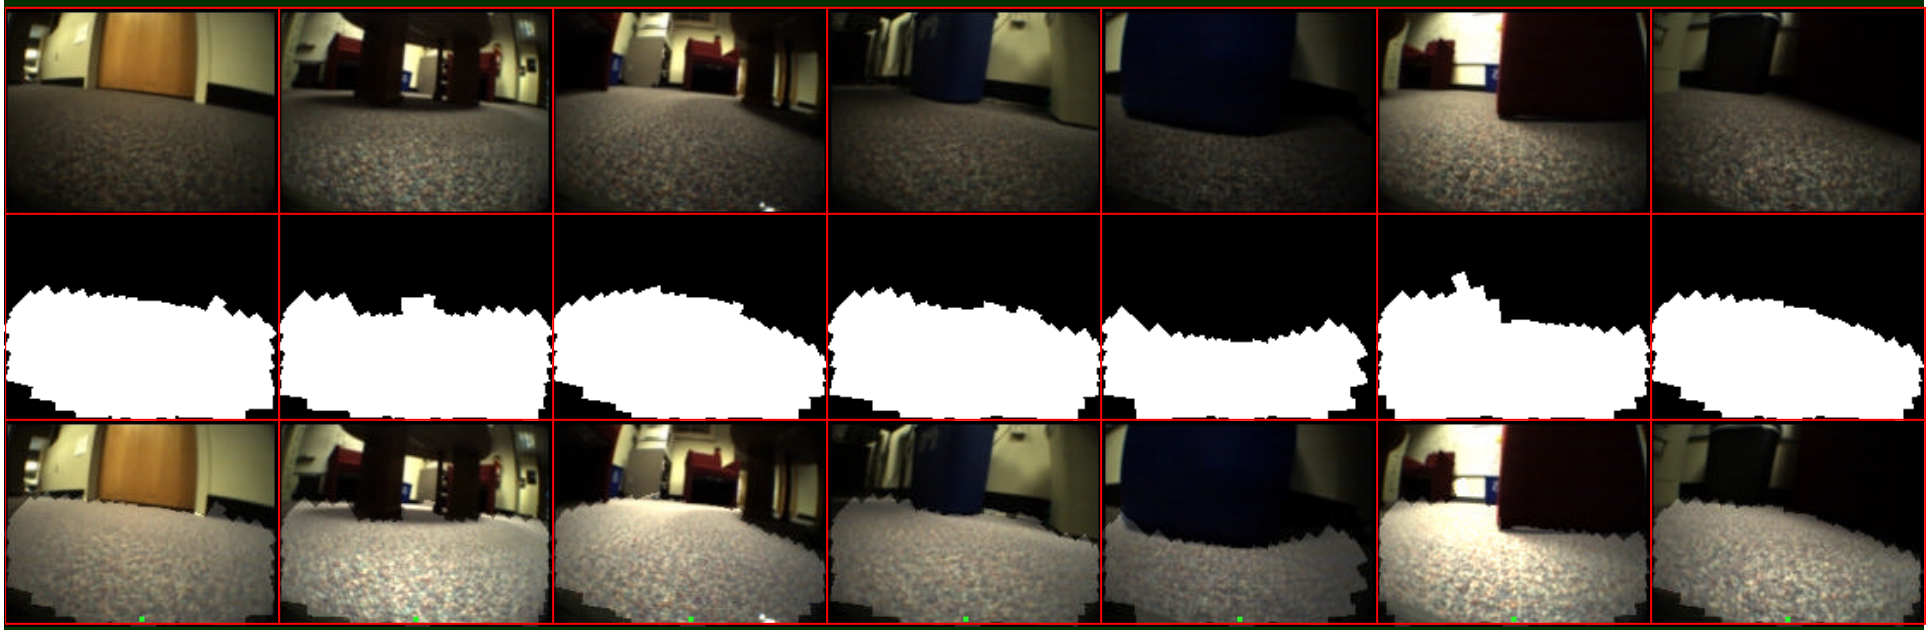

extract stable views for recognition