What Should Be Learned?

- Chris Atkeson, CMU, 4/27/19
- I am interested in “motor” skills
- Learning from:
  - Observation
  - Practice
  - Thinking (including simulation)

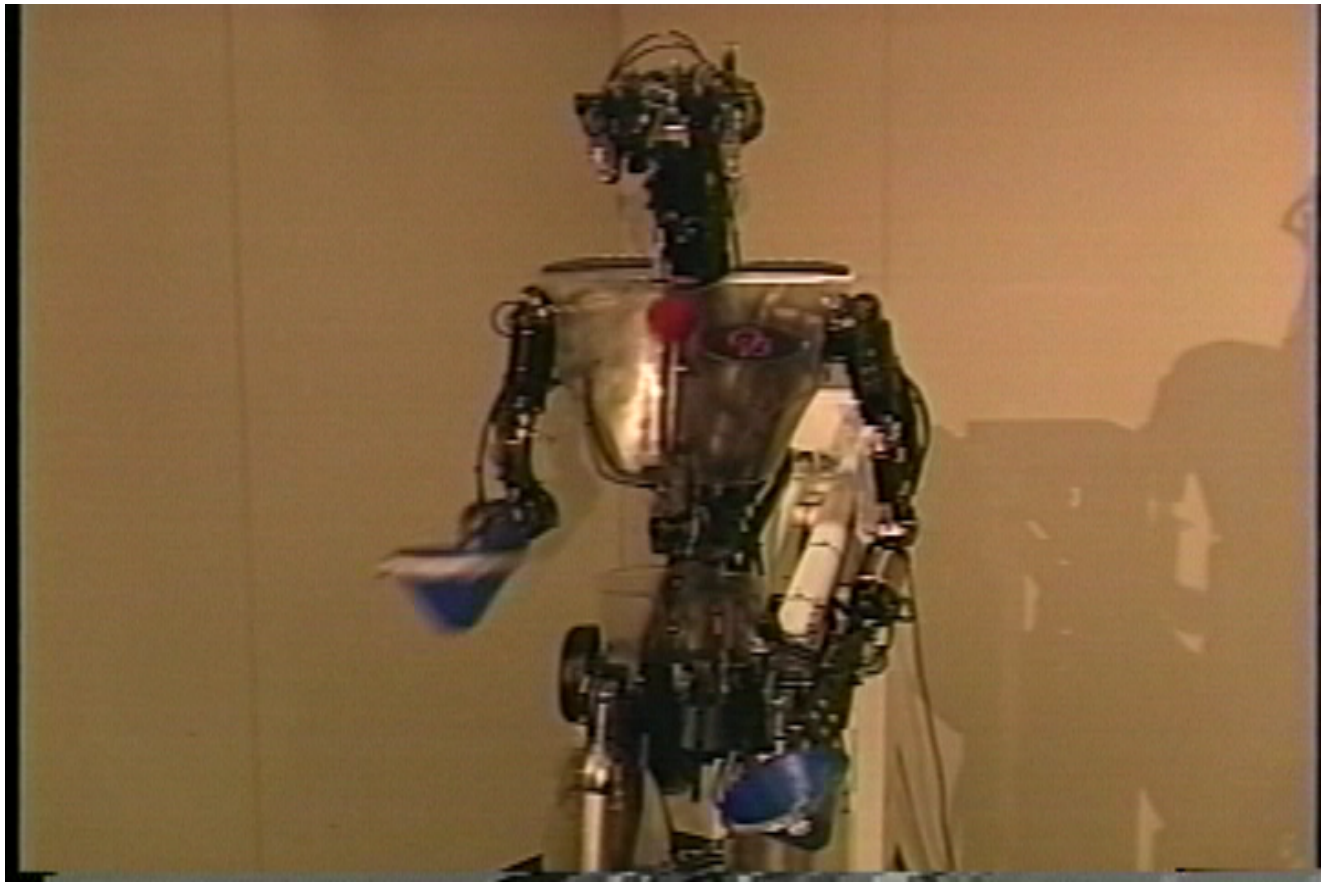

Learning From Practice

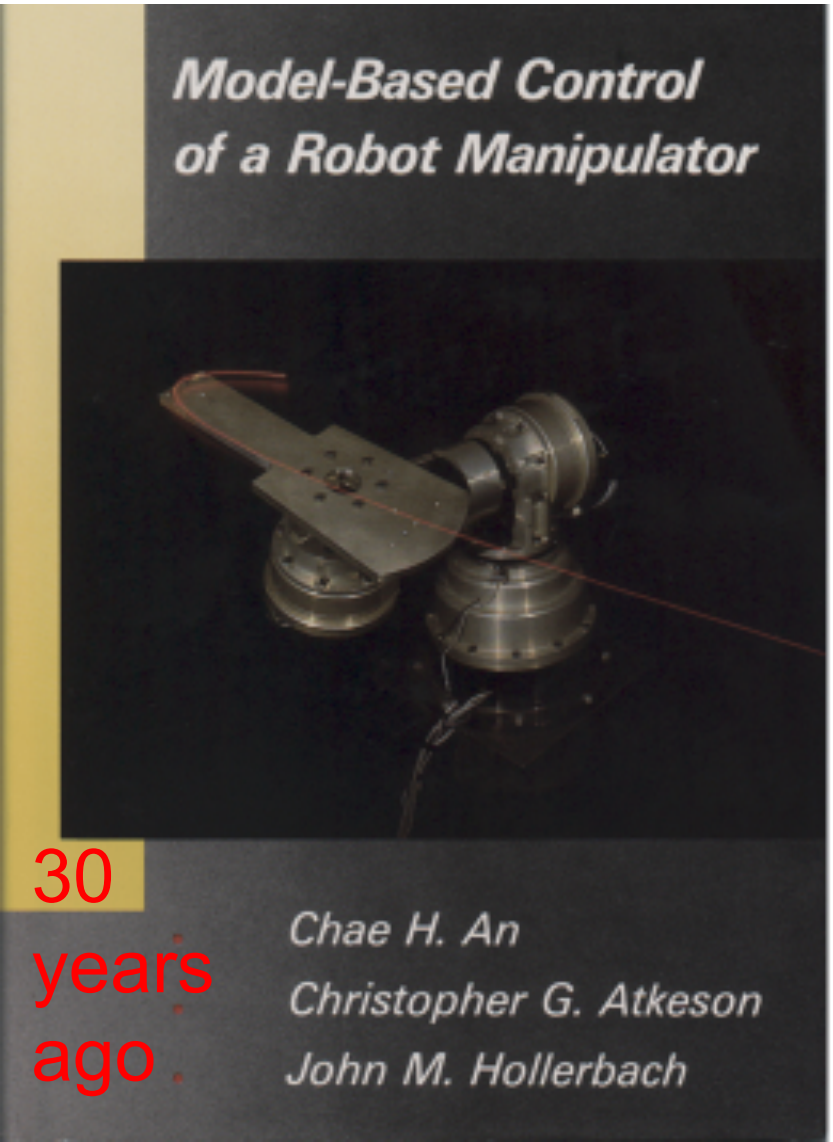

Commands

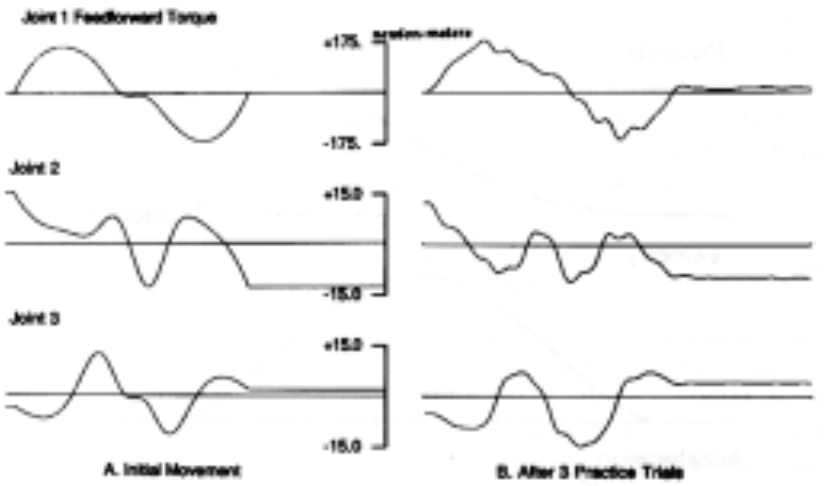

Figure 7.5: Feedforward Torques.

Before

After

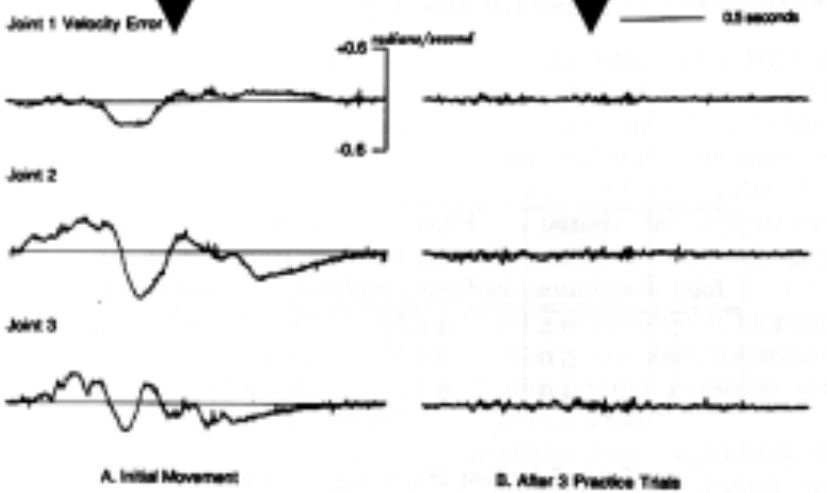

Figure 7.6: Velocity Errors.

Errors

Learning from Demonstration

Swing Up

(after learning)

25 years ago


25 years ago

Air Hockey

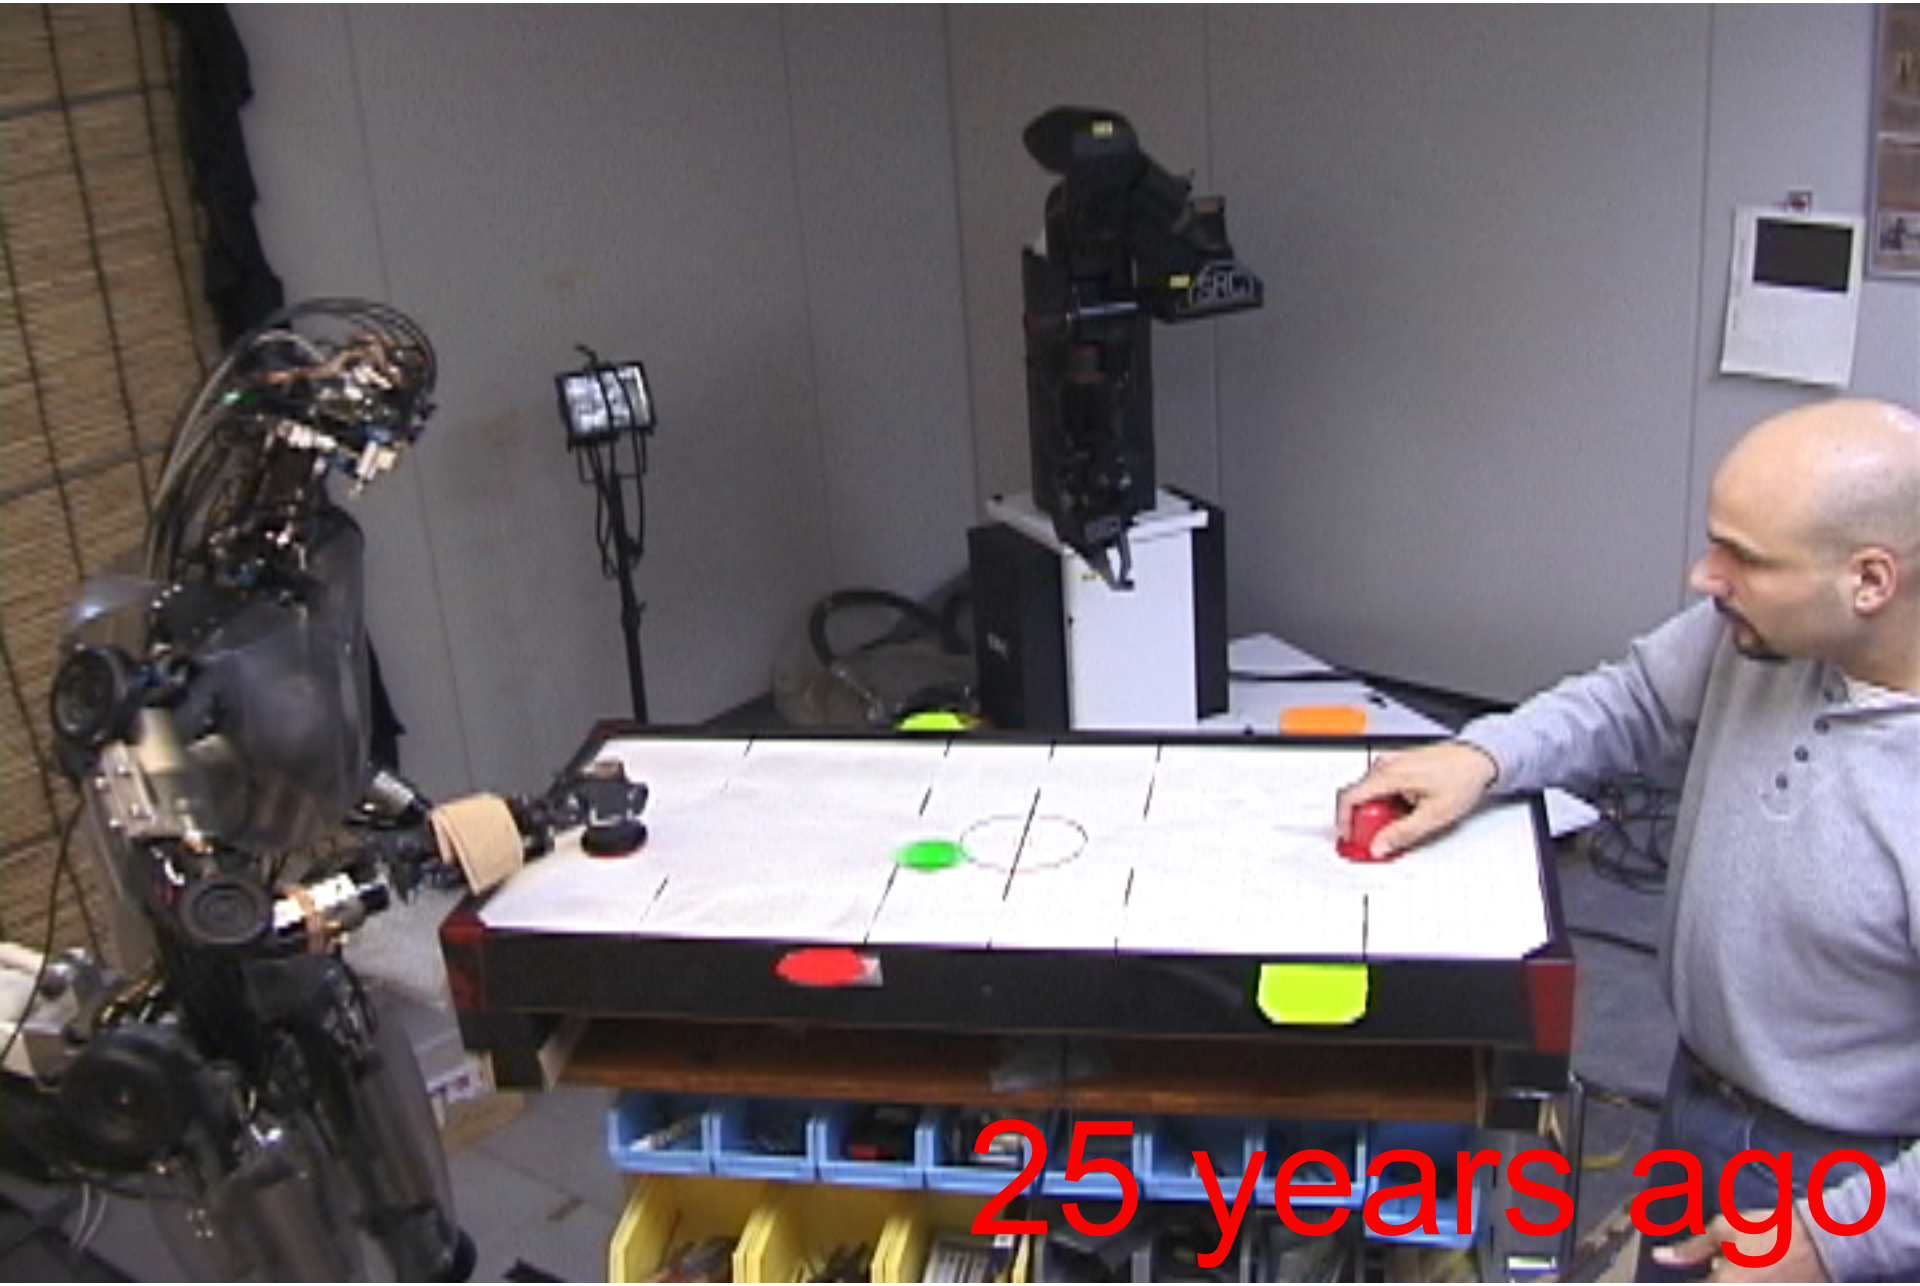


25 years ago

Learning new strategies

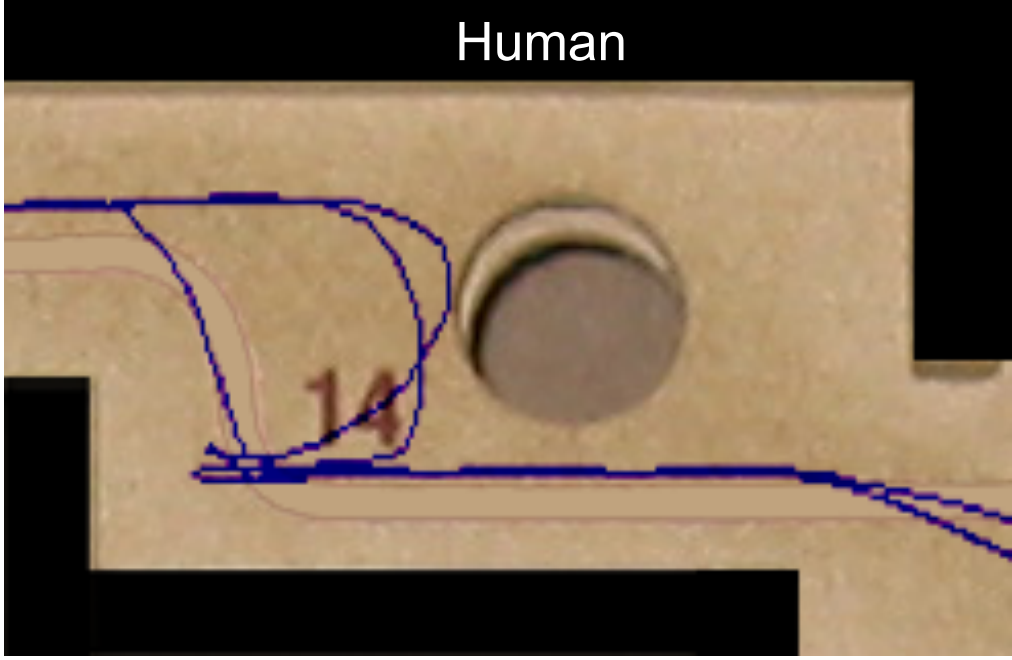

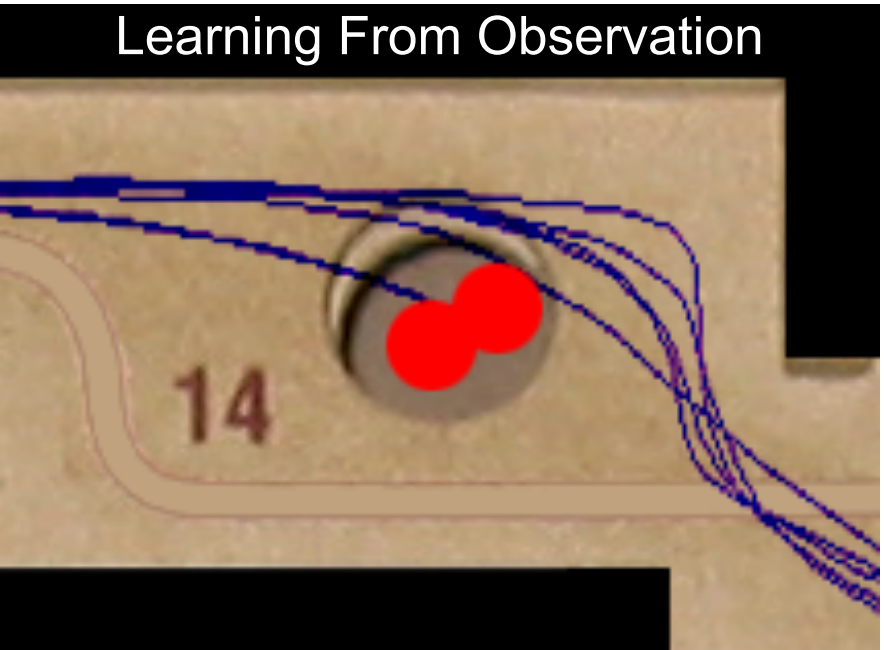

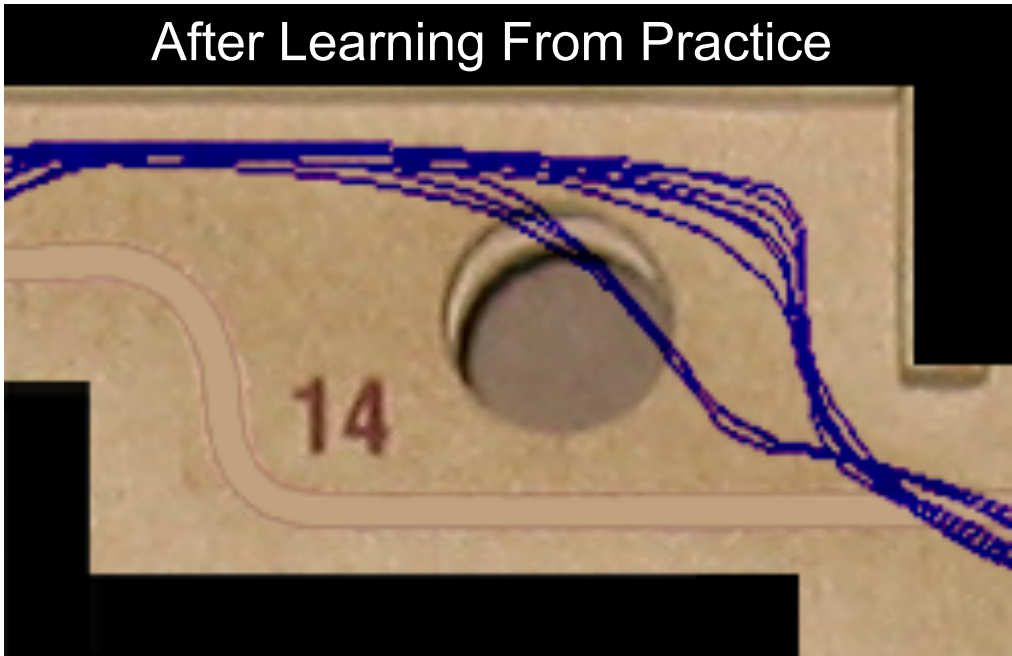

Key Points Of This Talk

- Task-level model-based learning is effective and data efficient (<100 trials). For learning to control robots, that is the standard to meet. Can handle contact, deformable objects, liquids, granular materials, and other “hard-to-model” situations.
- “Model-free” learning can play a useful role in perceptual and end-to-end learning.
- The most useful role for “model-free” learning is model-based planning which generates policies in simulation (which is what most model-free RL is doing).
- We need non-parametric learning to find better ways of doing things (better task strategies).

Learning agile and dynamic motor skills for legged robots

Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun and Marco Hutter

Science Robotics

The Paddle Juggling War
aka The Parameterization Trap

30 years ago

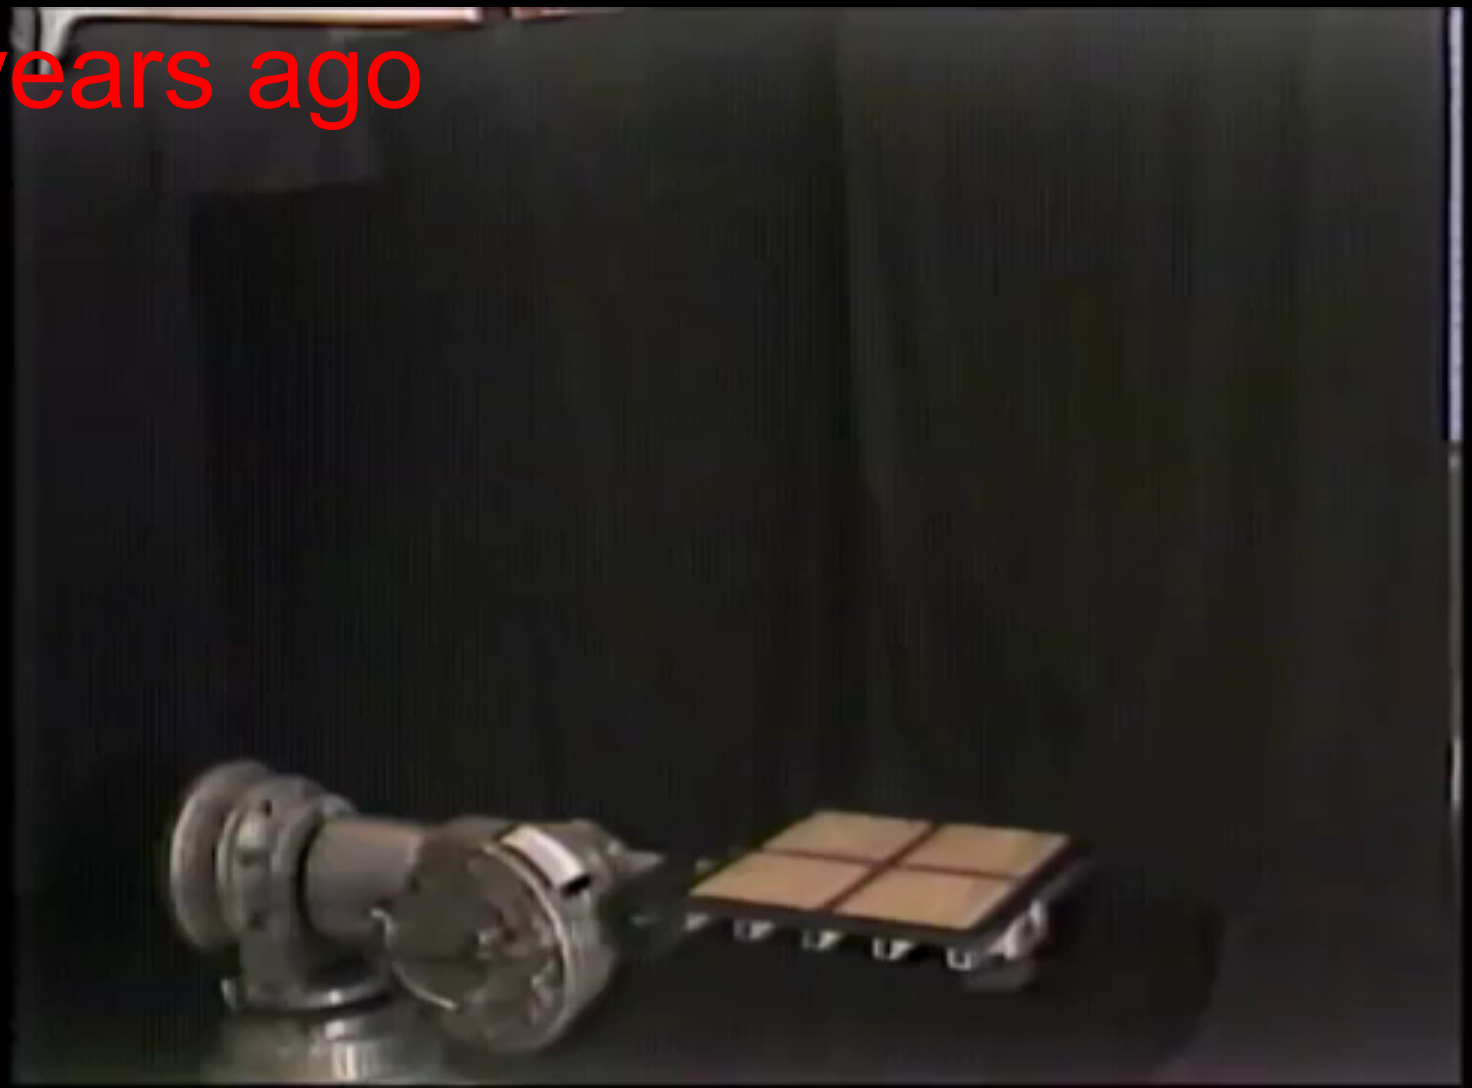

MIT Direct Drive Arm I Paddle Juggling With Vision, 1988

After Task-
Level Learning

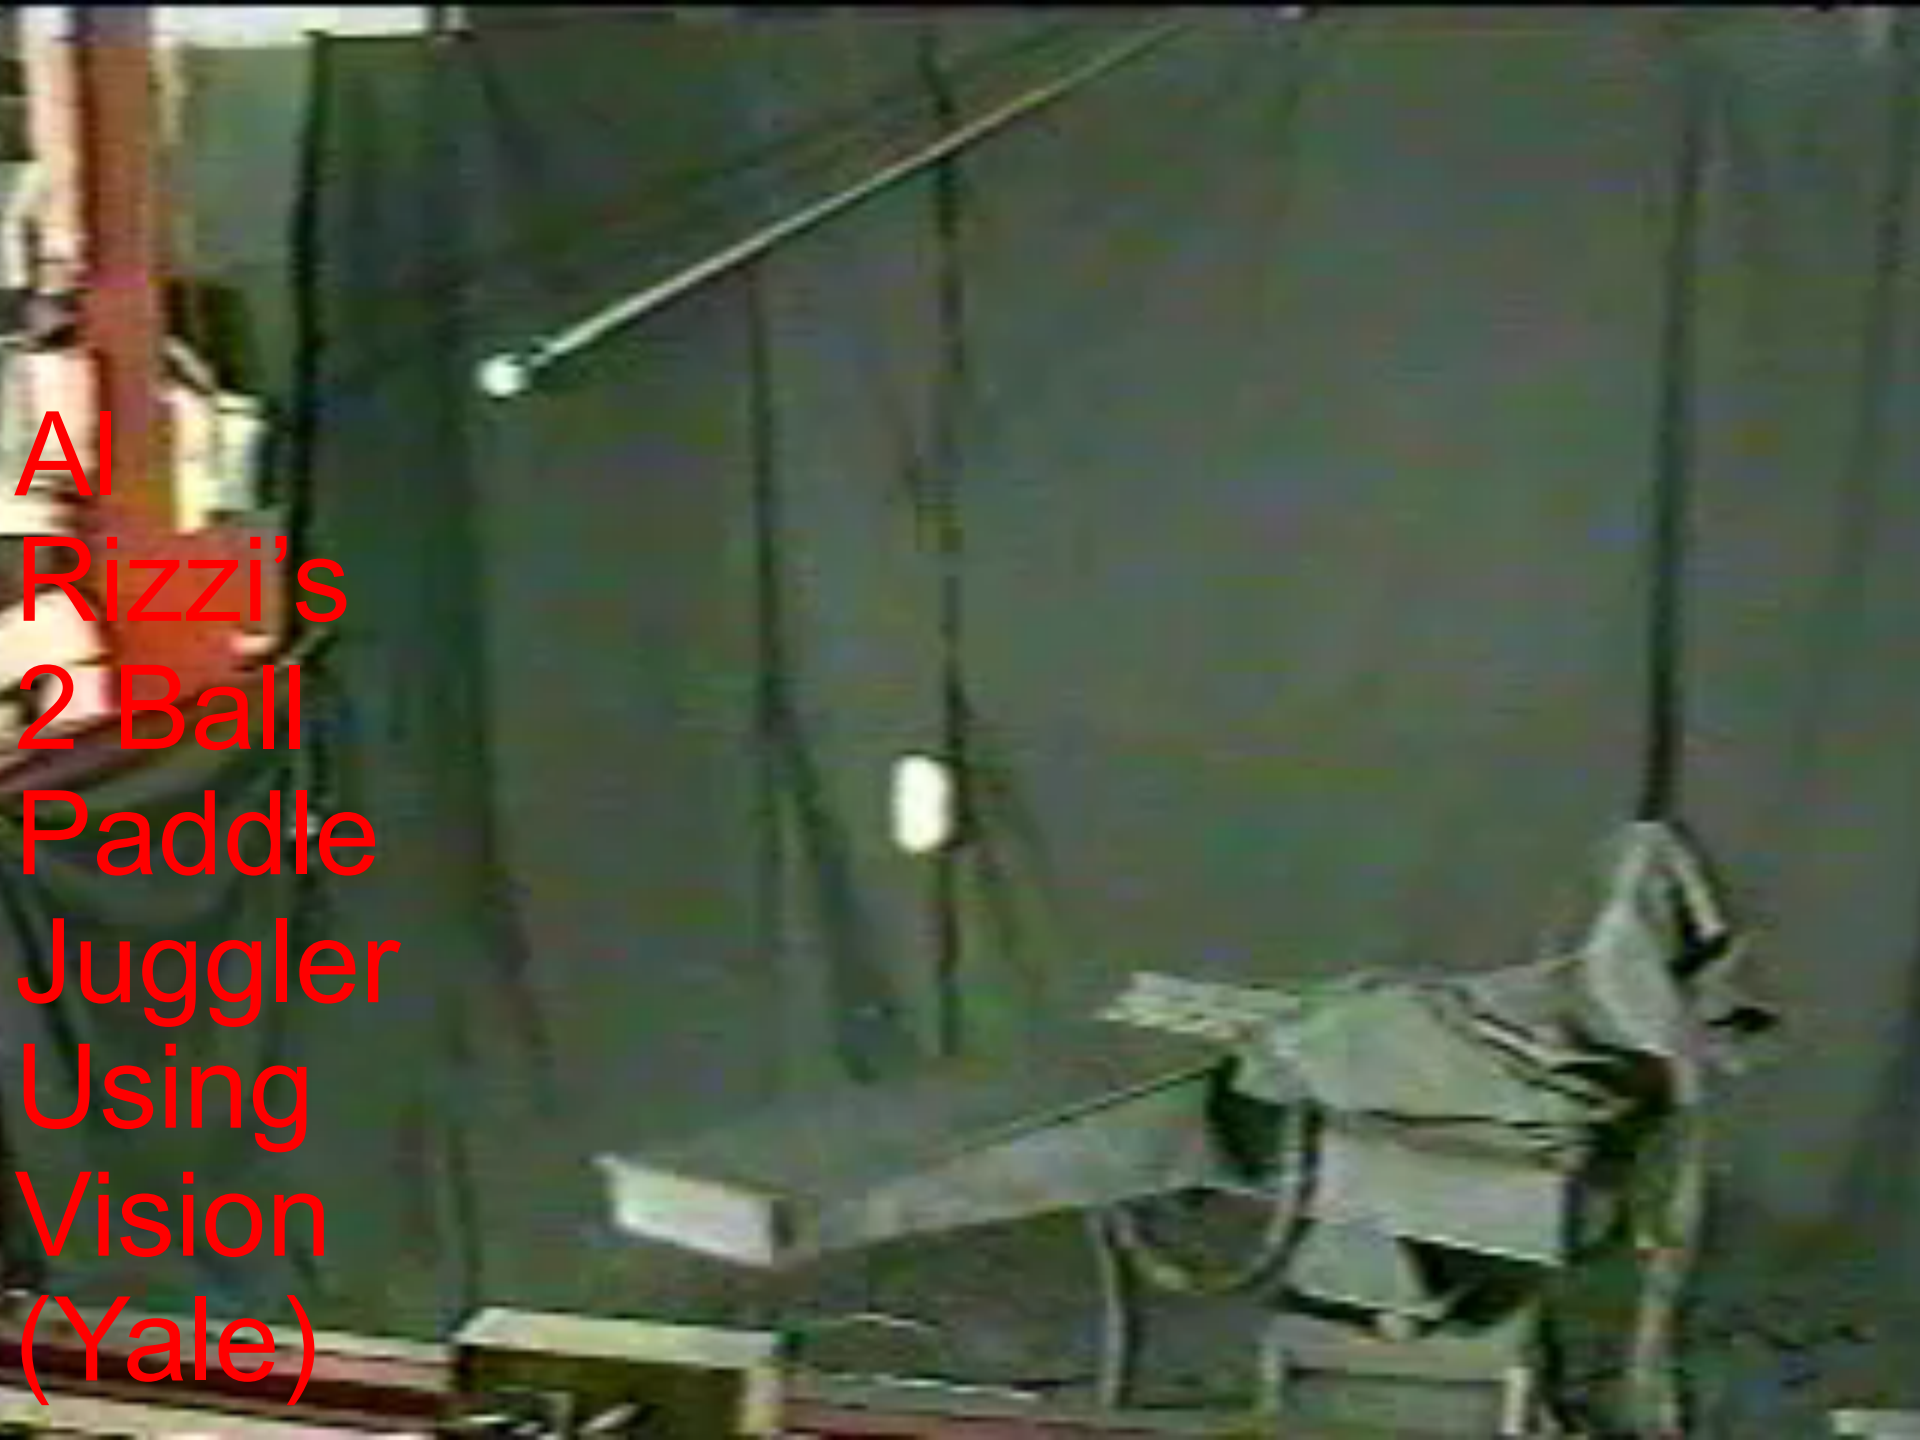
Al
Rizzi's
2 Ball
Paddle
Juggler
Using
Vision
(Yale)

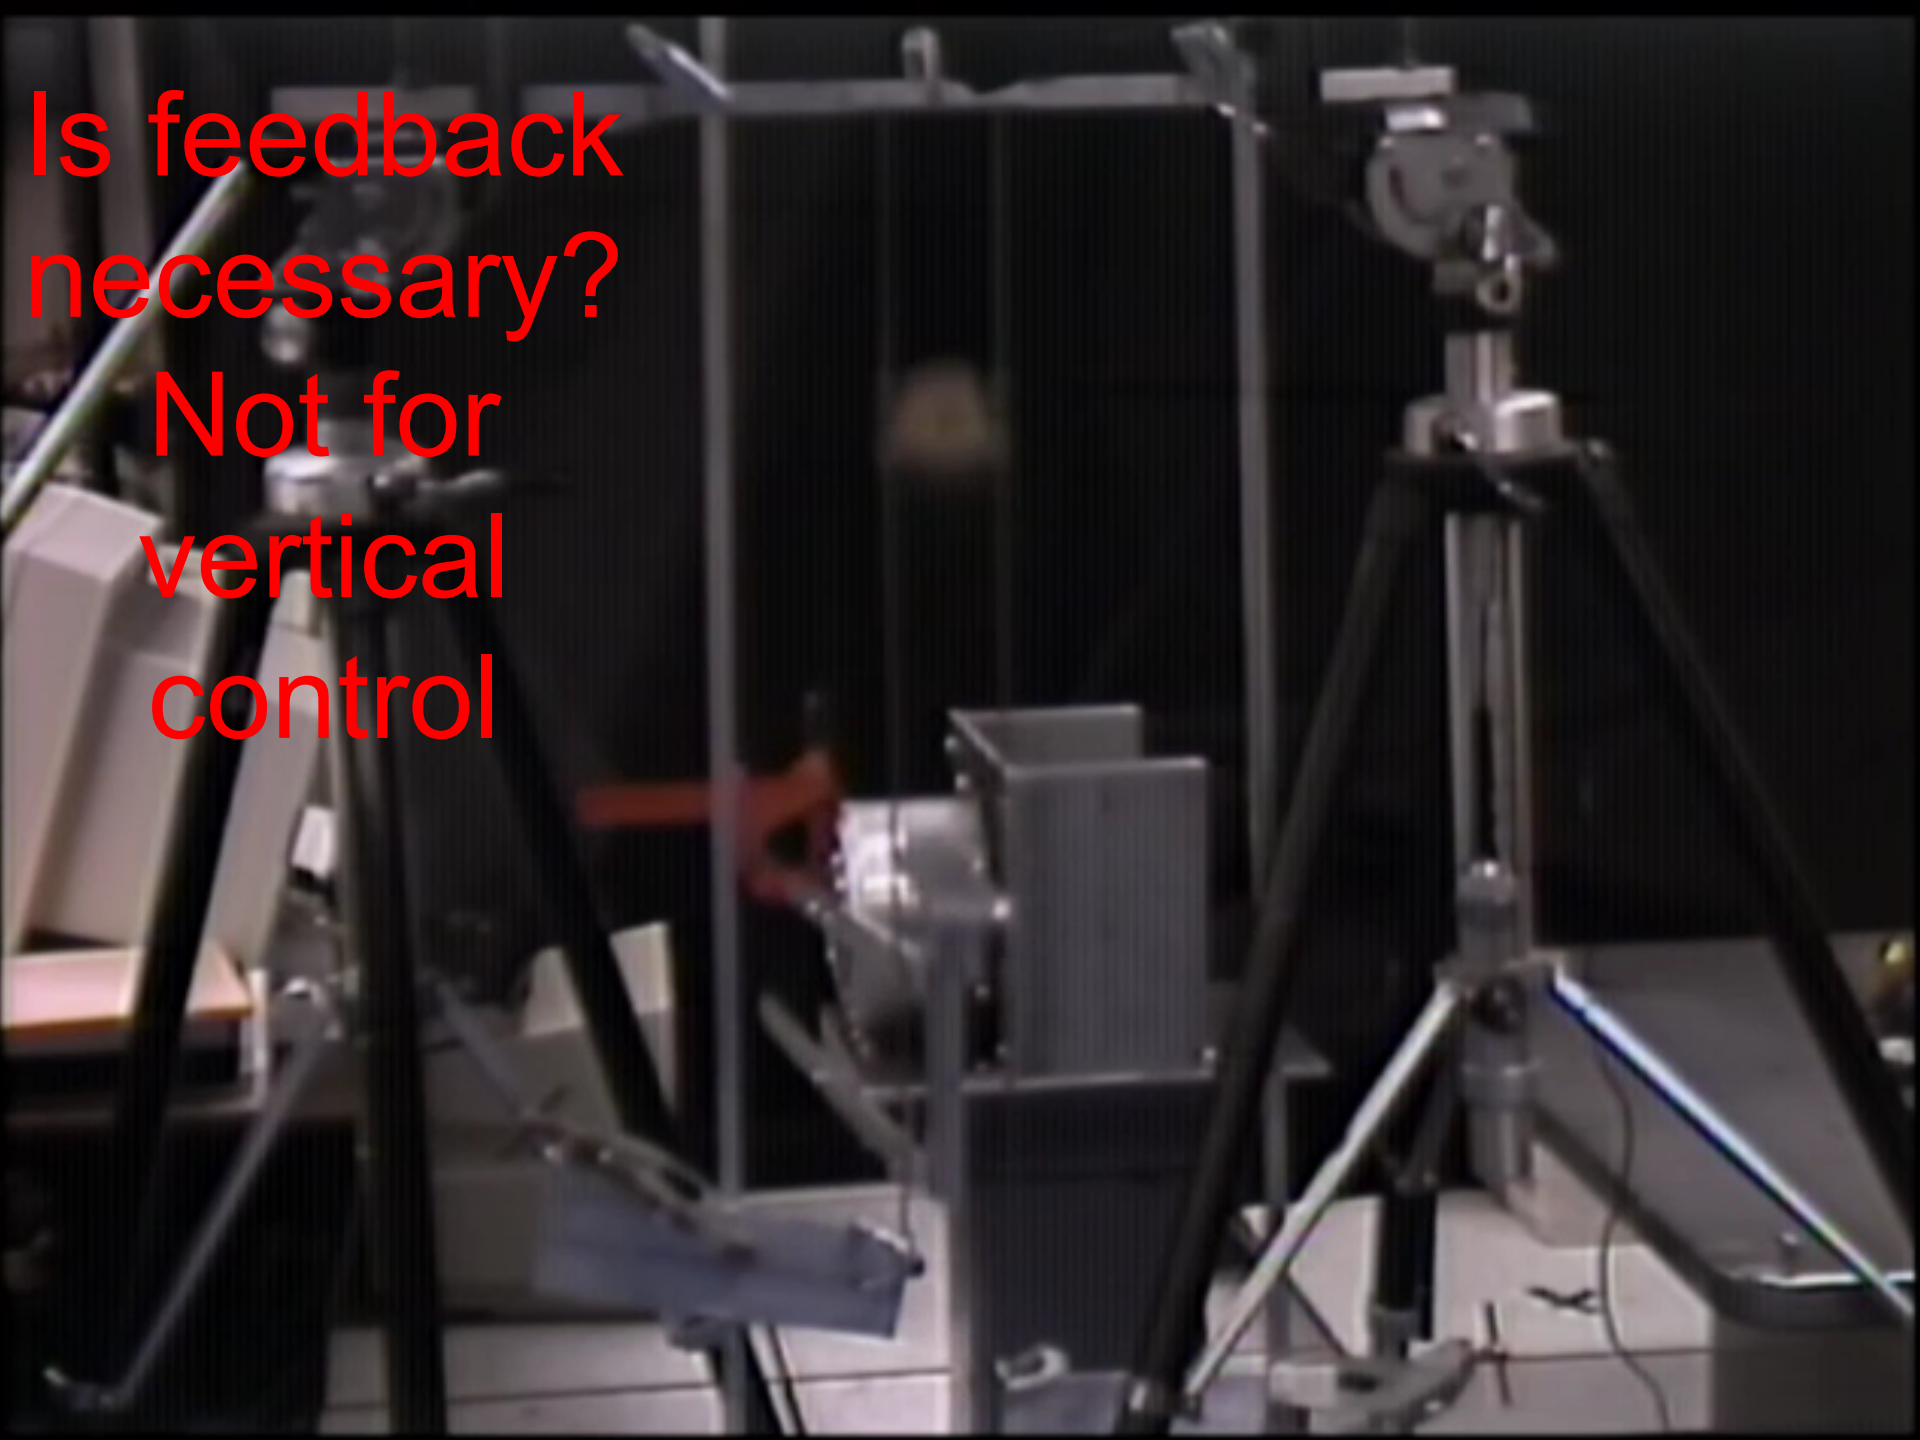
Is feedback necessary?
Not for vertical control

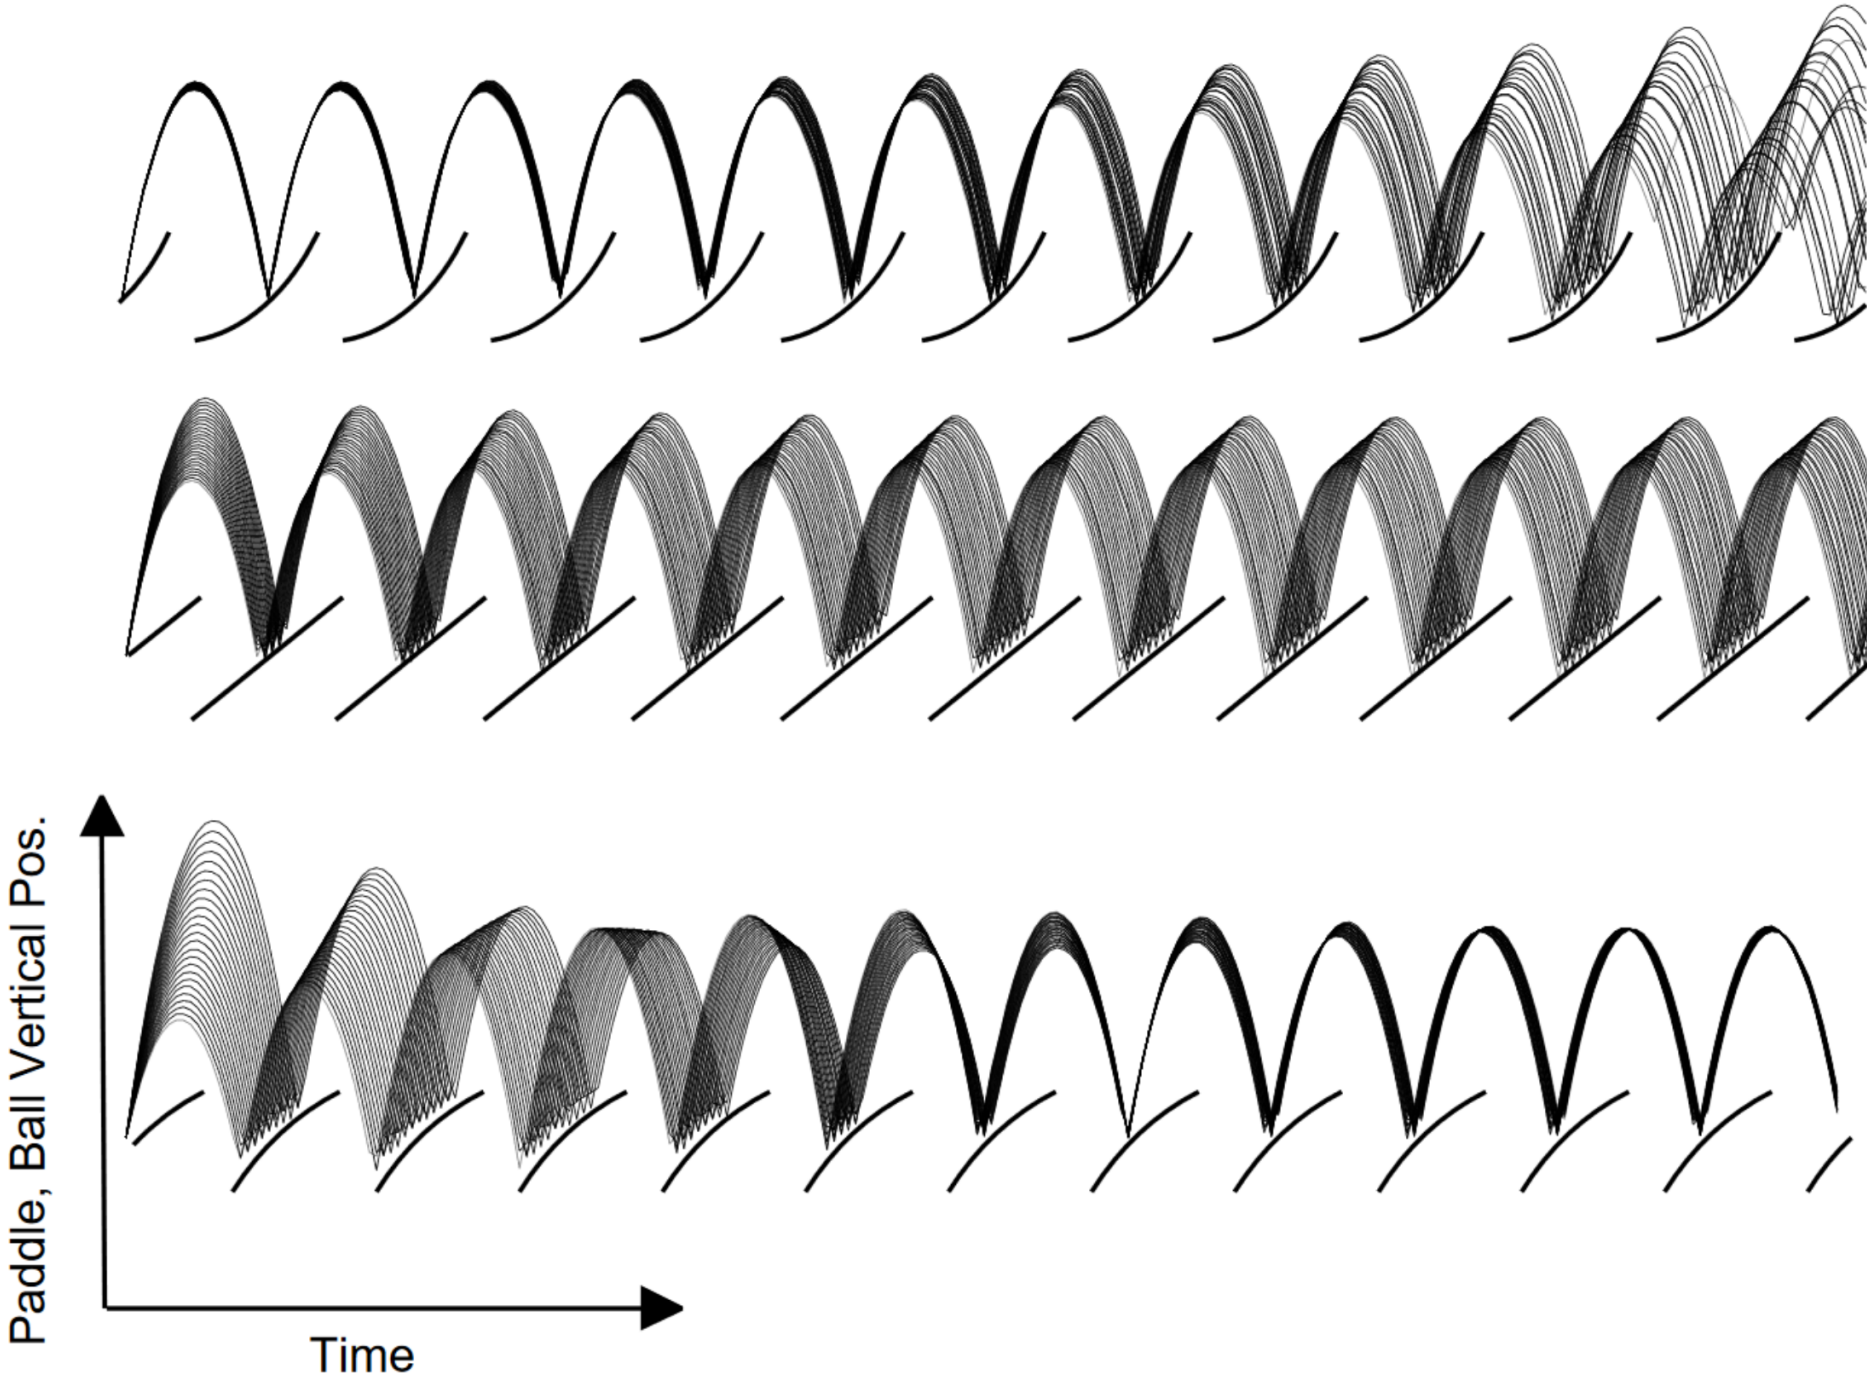
Paddle, Ball Vertical Pos.
Time

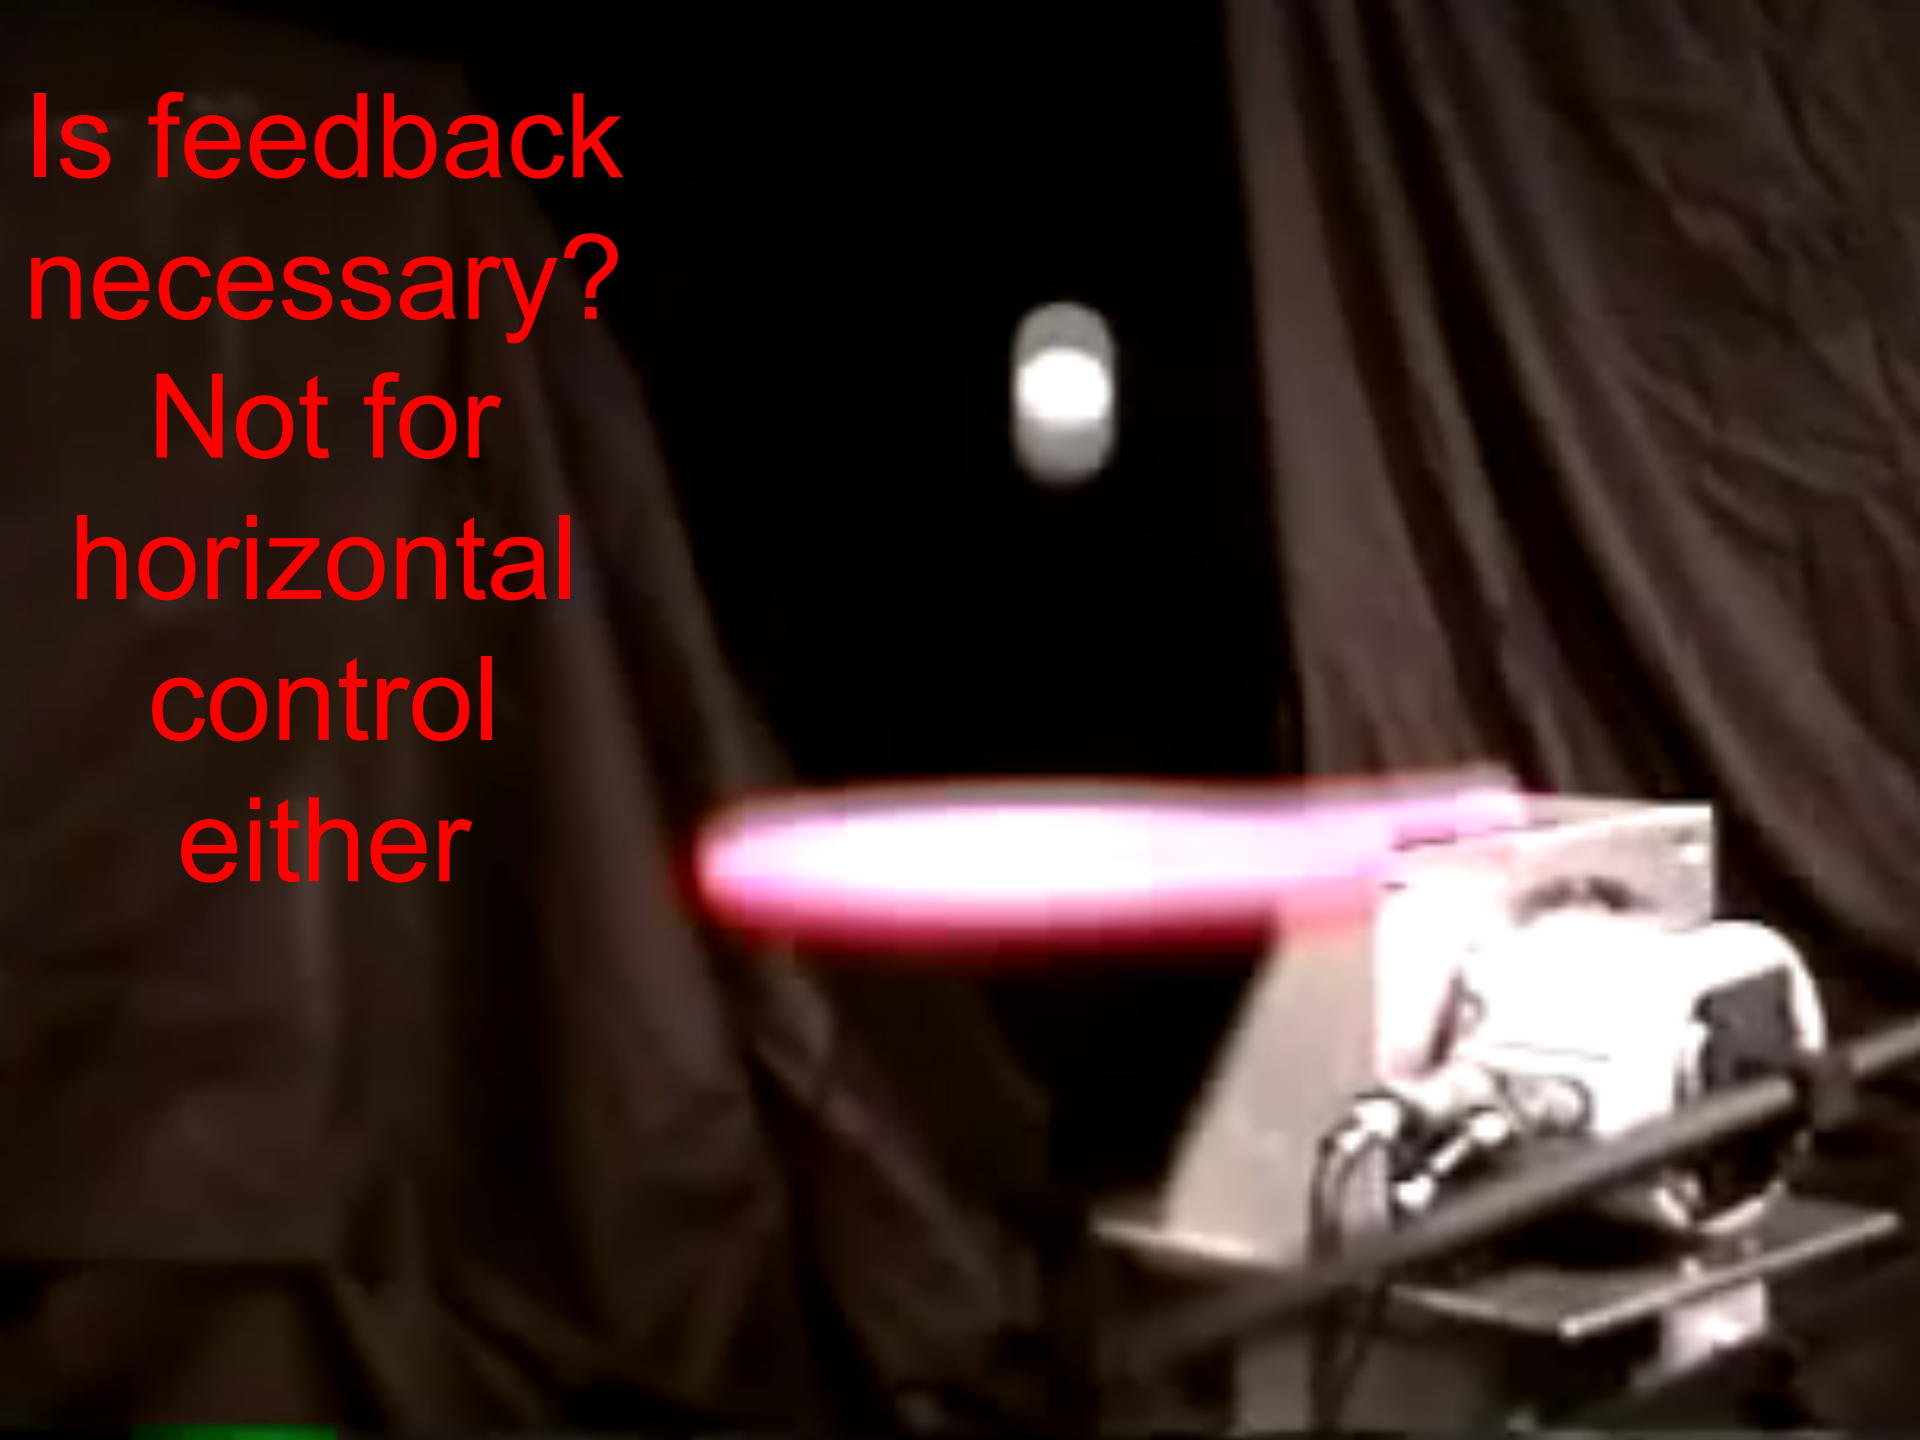
Is feedback necessary?
Not for horizontal control either

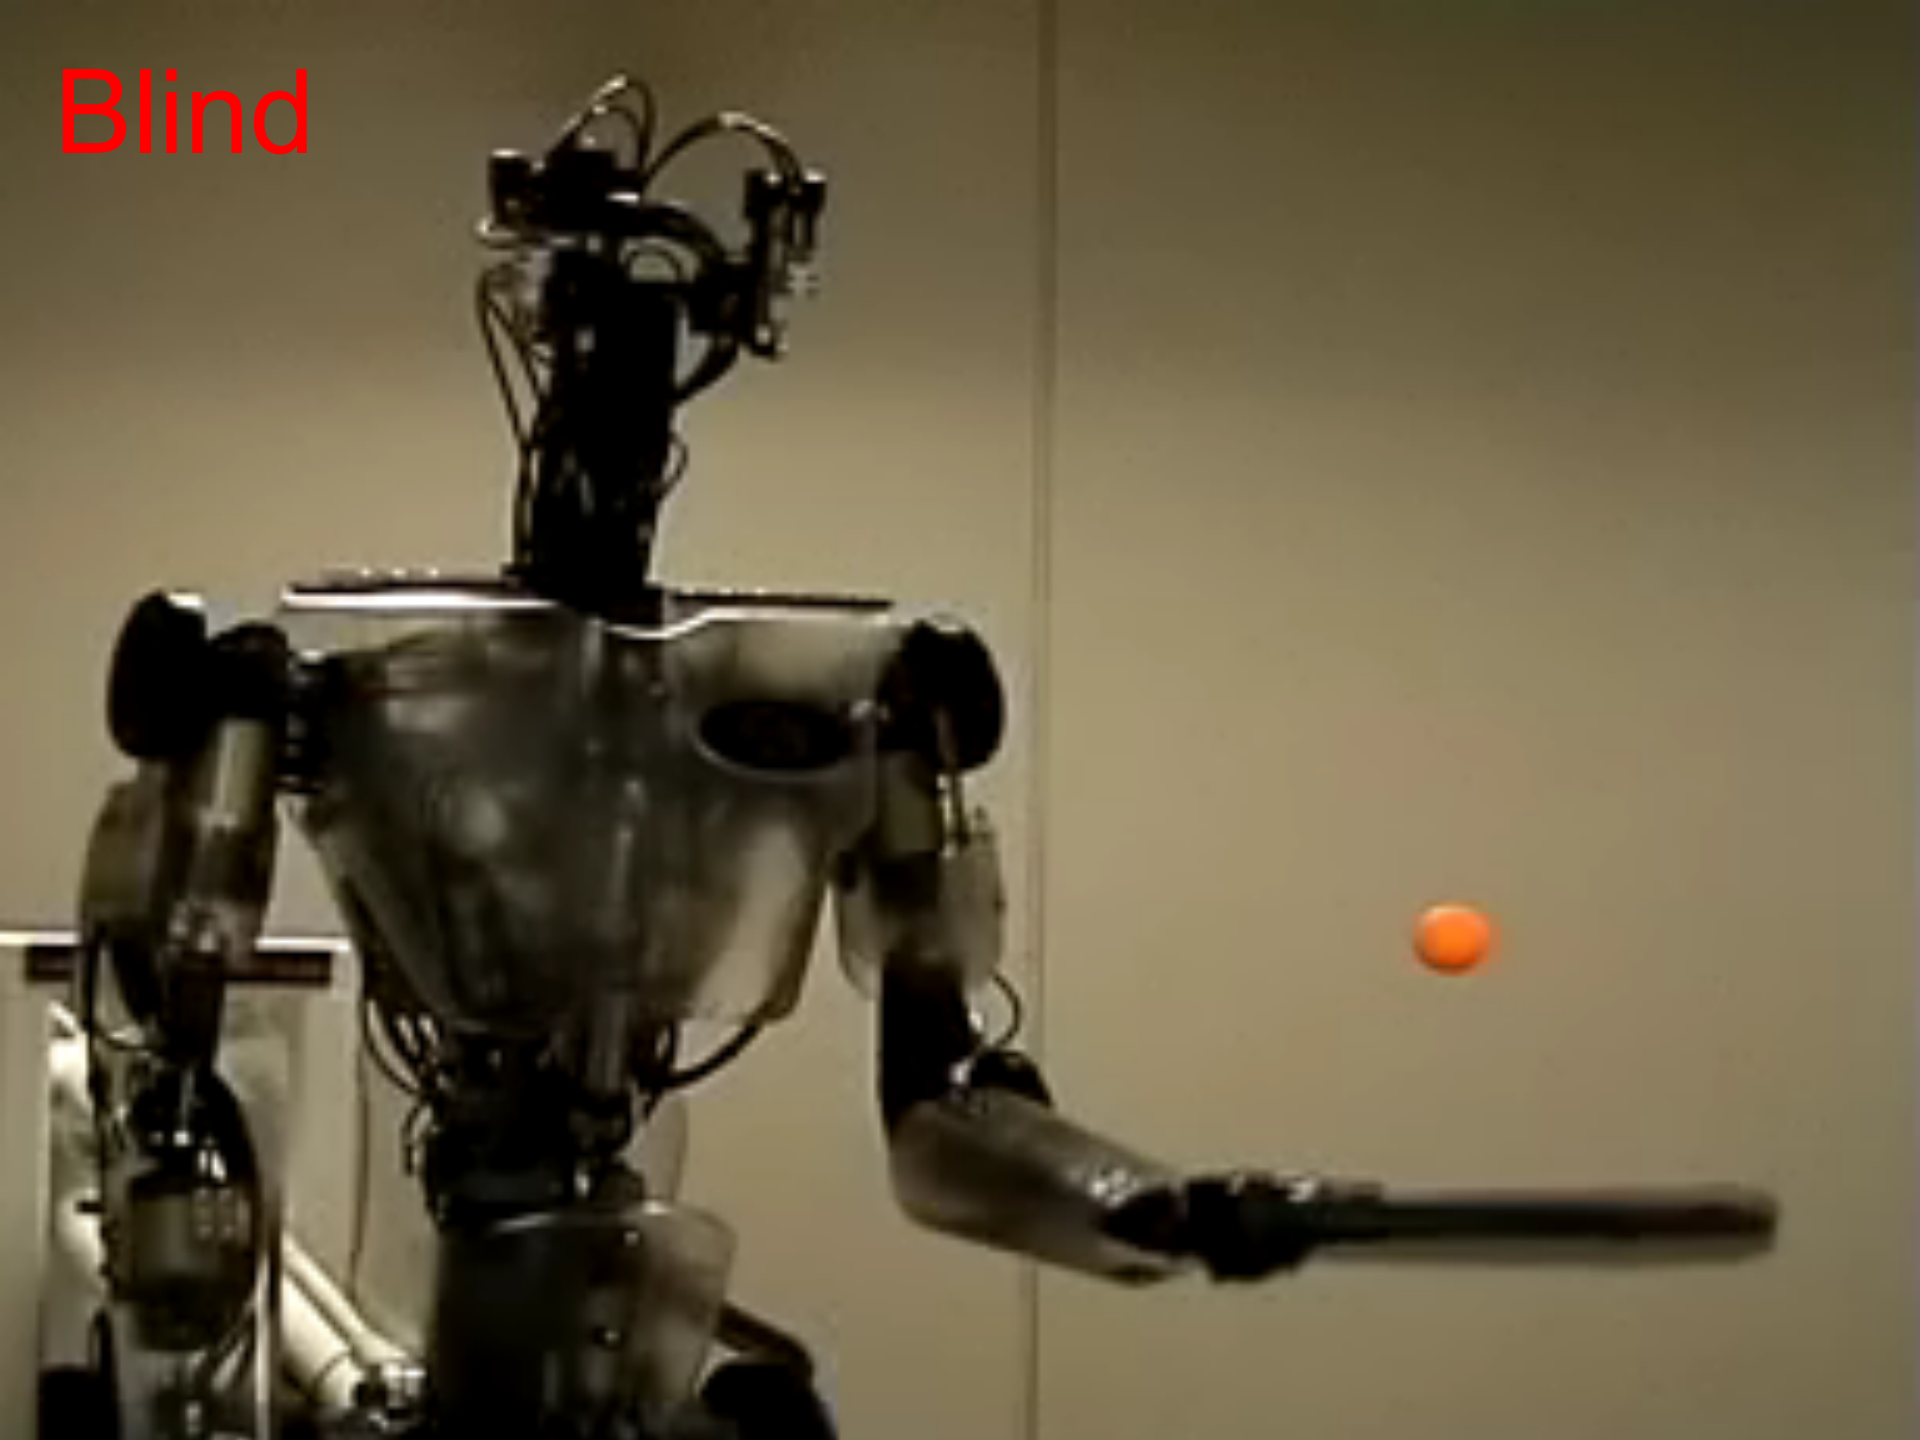
Blind

Blind Juggler Cloverleaf Paddle
IDSC, ETH Zurich

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

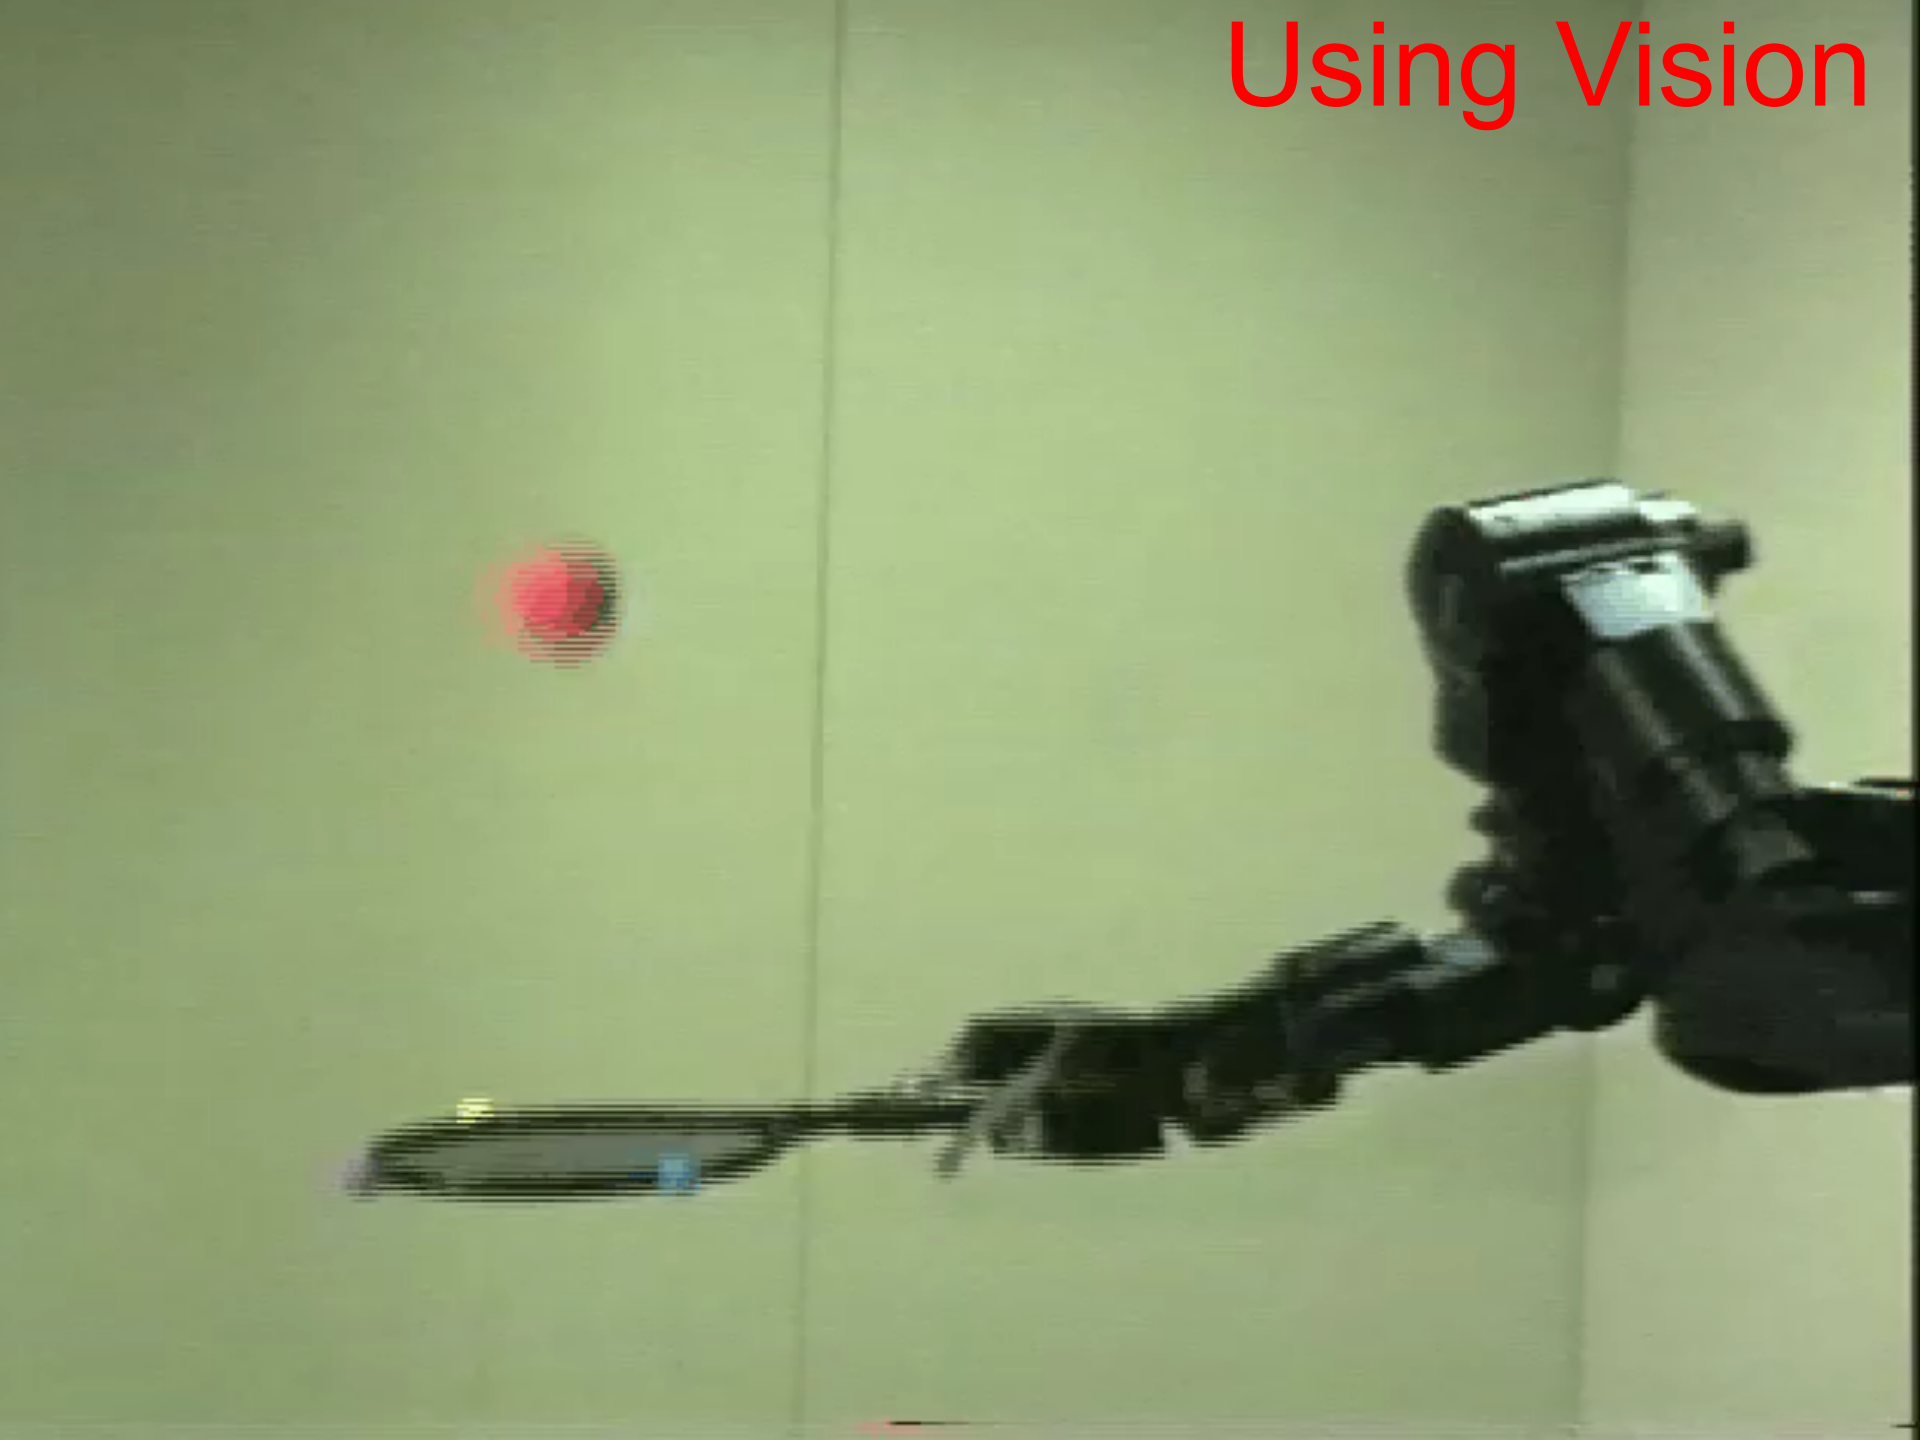
Using Vision

Model-based Reinforcement Learning of Devilsticking

Stefan Schaal & Chris Atkeson


25 years ago

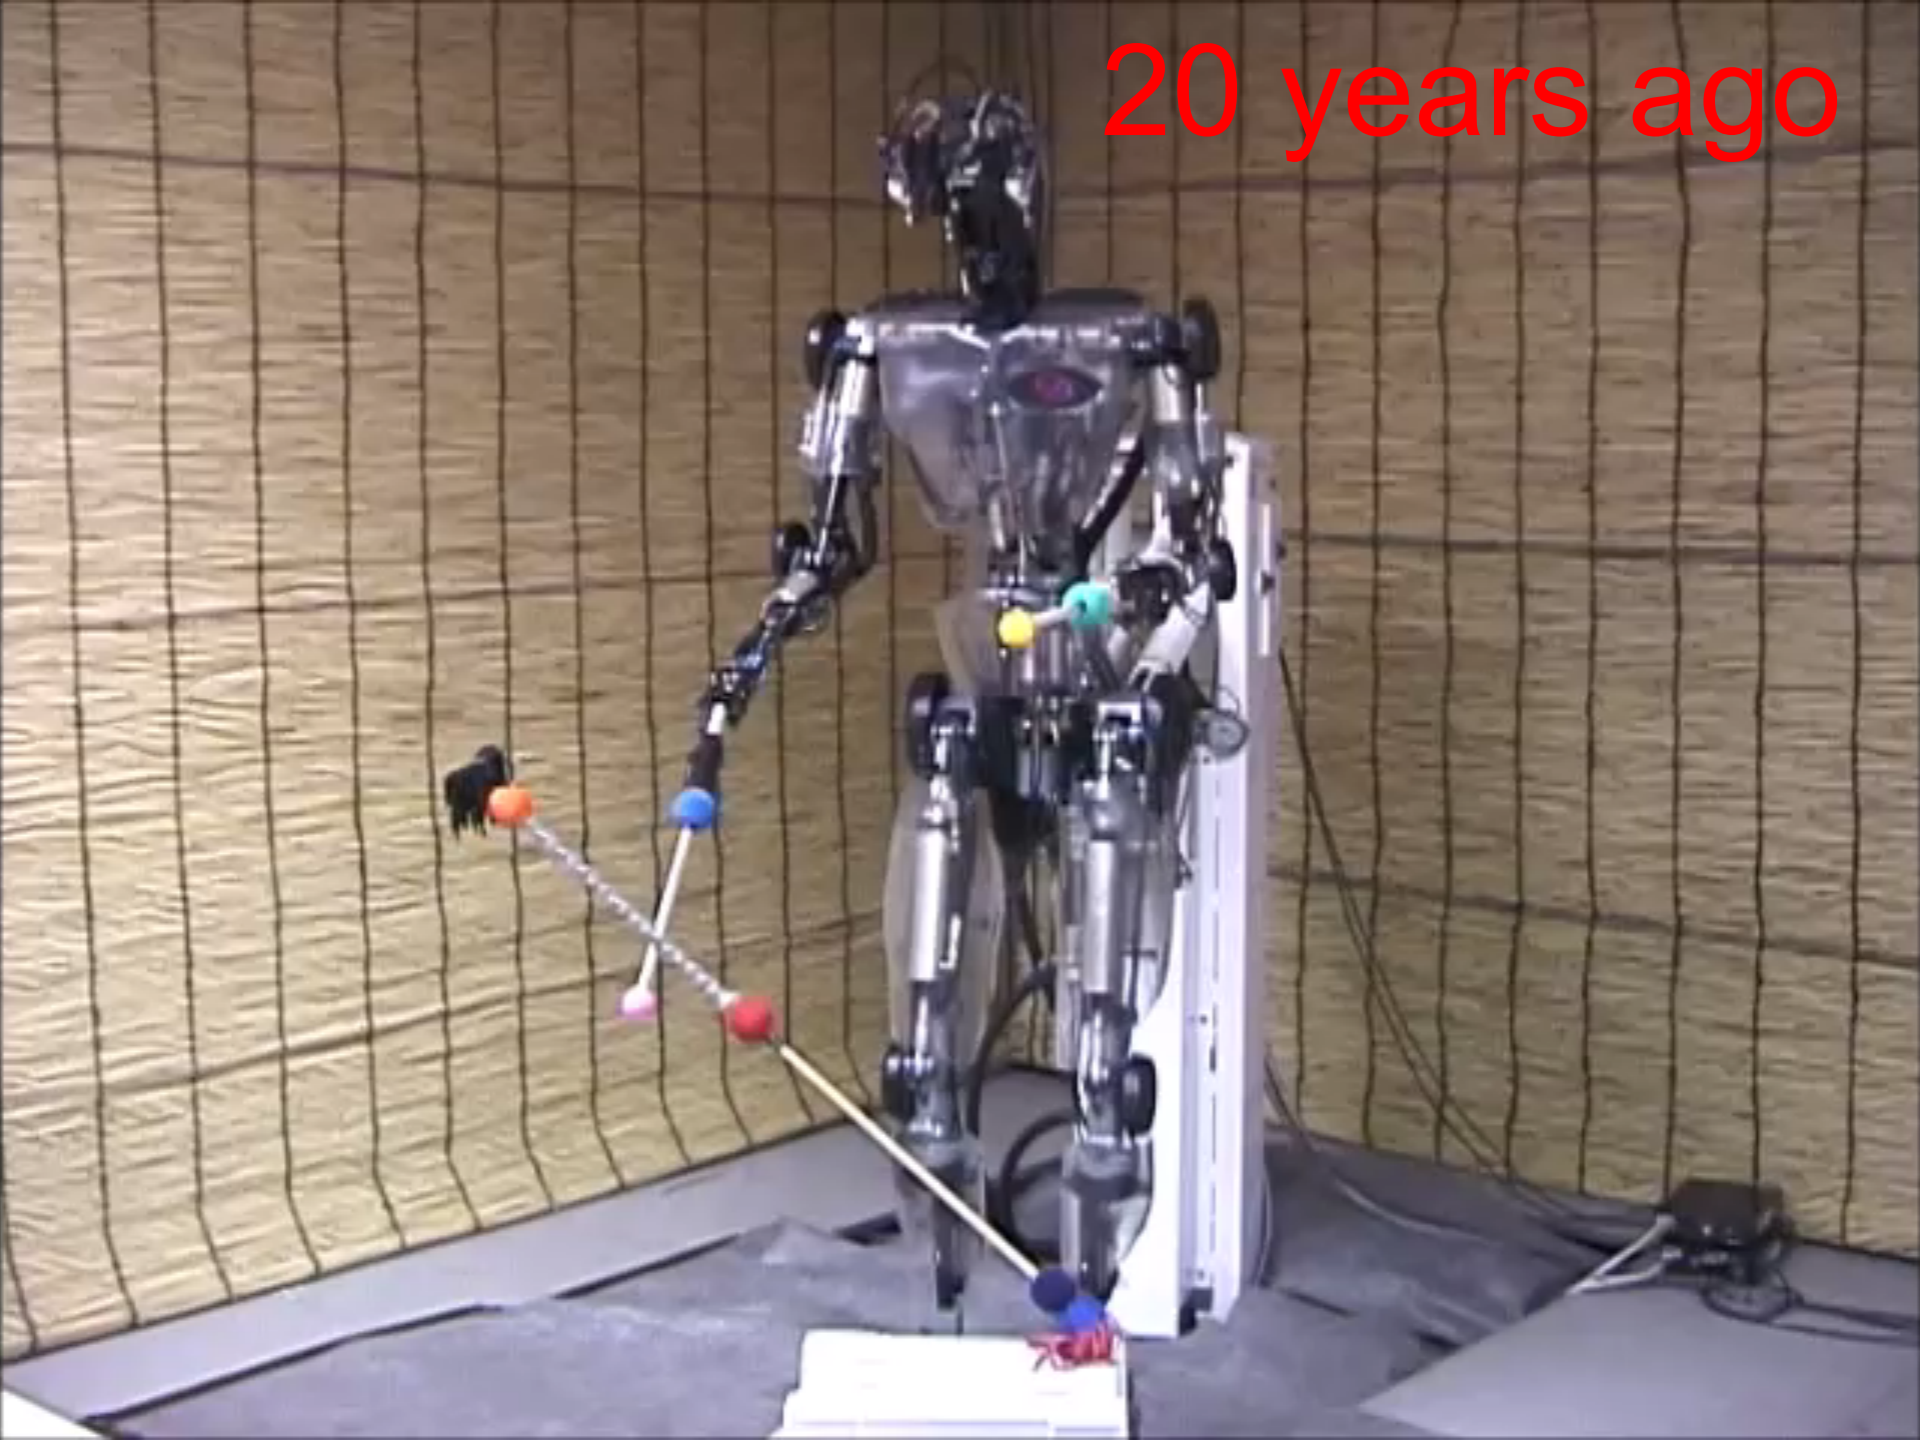
20 years ago

So why does current robot control use so little learning?

- Actually, we use lots of learning, in the form of system identification and model learning.
- Nonparametric model learning and deep models are useful if a lot of data is available (usually from simulation).
- Optimal control, a form of model-based reinforcement learning, is used to generate behaviors and policies.
- Precision and accuracy matter, especially for balance.

Push recovery

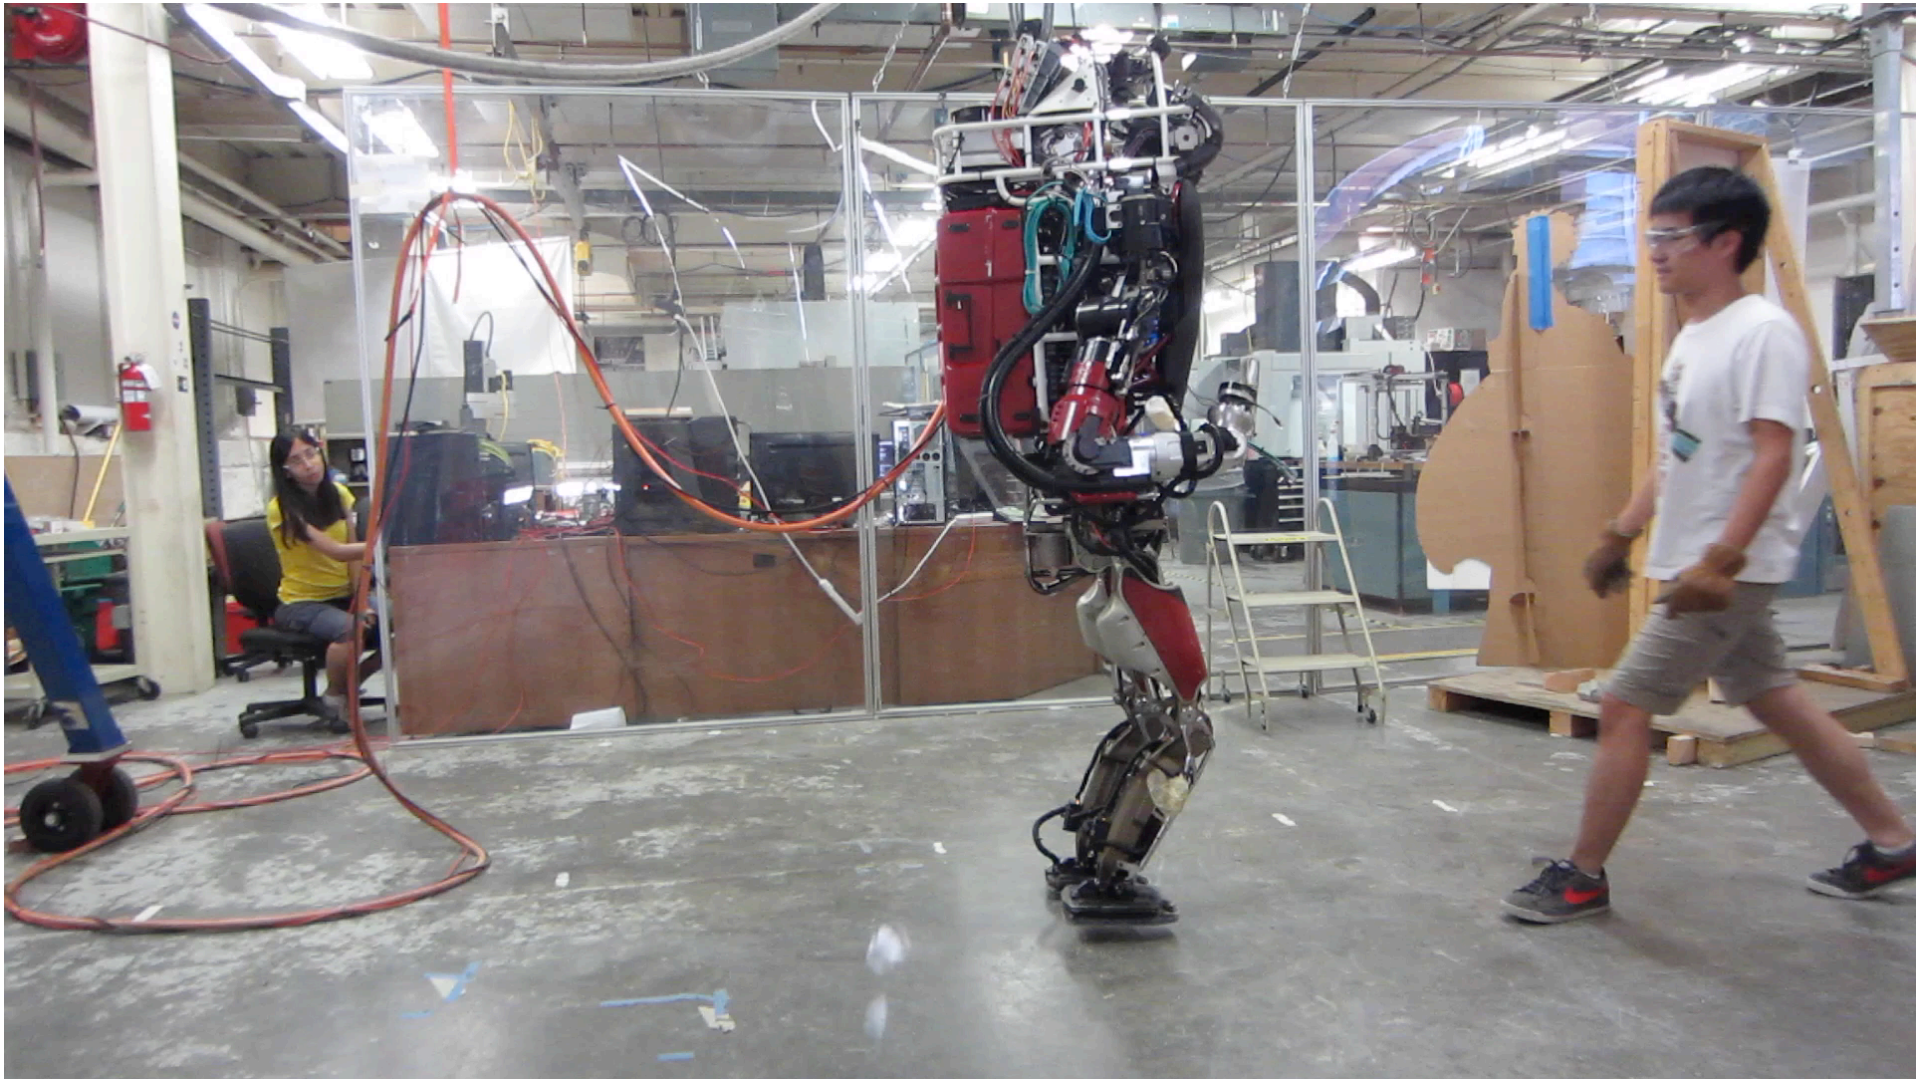

Consensus On Control: Hierarchical Receding Horizon Optimization

- Footstep Optimization (Discrete + Continuous)
- Trajectory Optimization
- Receding Horizon Control
- Optimization-Based Inverse Dynamics: Greedy continuous optimization (Quadratic Program = QP) for full body at the current instant.

Controller overview

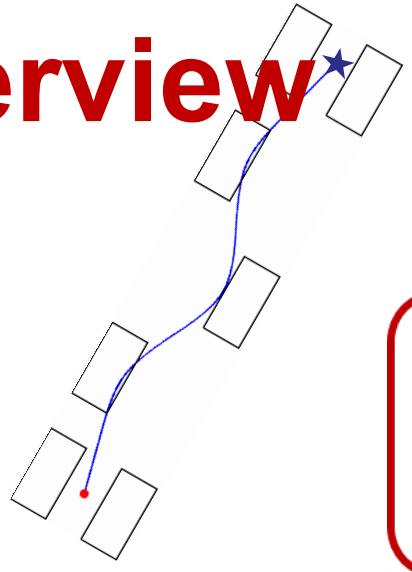

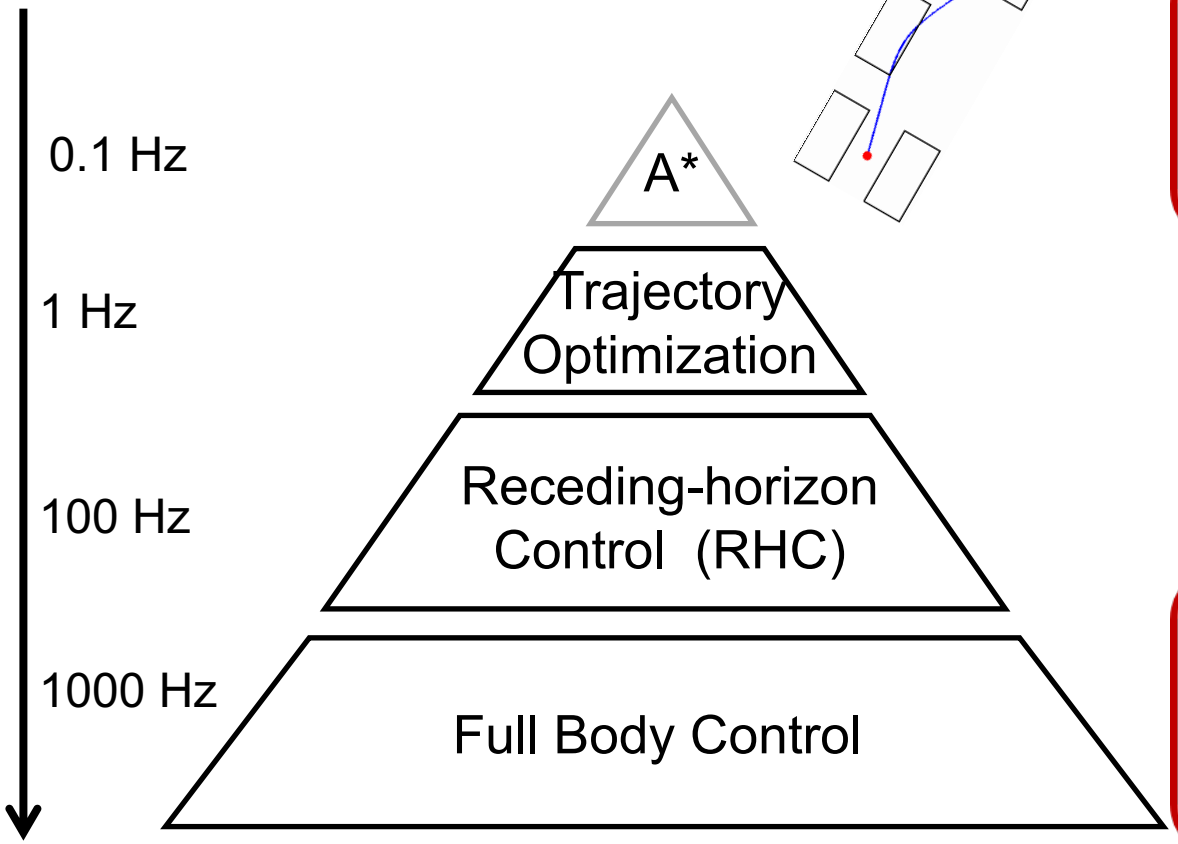

- Long horizon
- Slow computation
- Simple model

$V(x)$

- Short horizon
- Fast computation
- Complex model

$$[M(q) \quad -J^T(q) \quad -I] \begin{bmatrix} \ddot{q} \\ F \\ \tau \end{bmatrix} = -h(q, \dot{q})$$

Walking on loose rough terrain

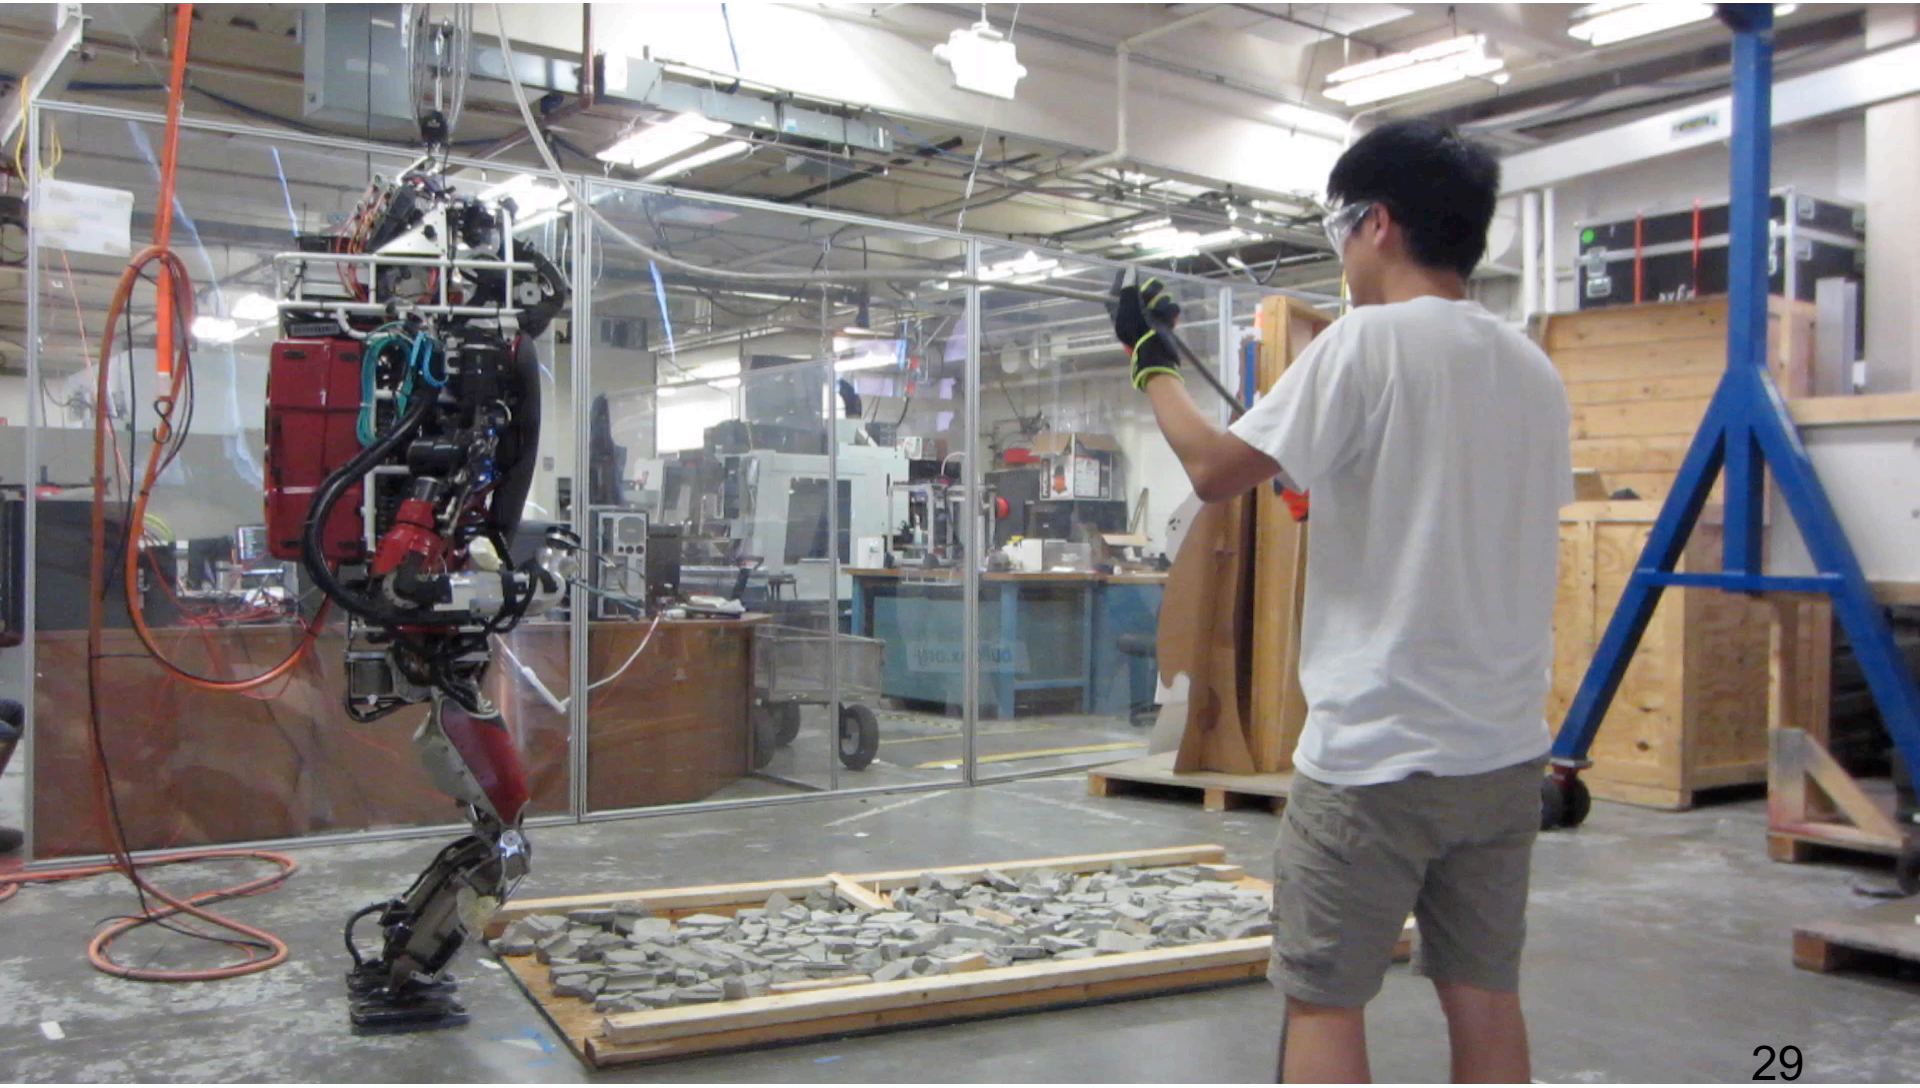

What humans can learn from:

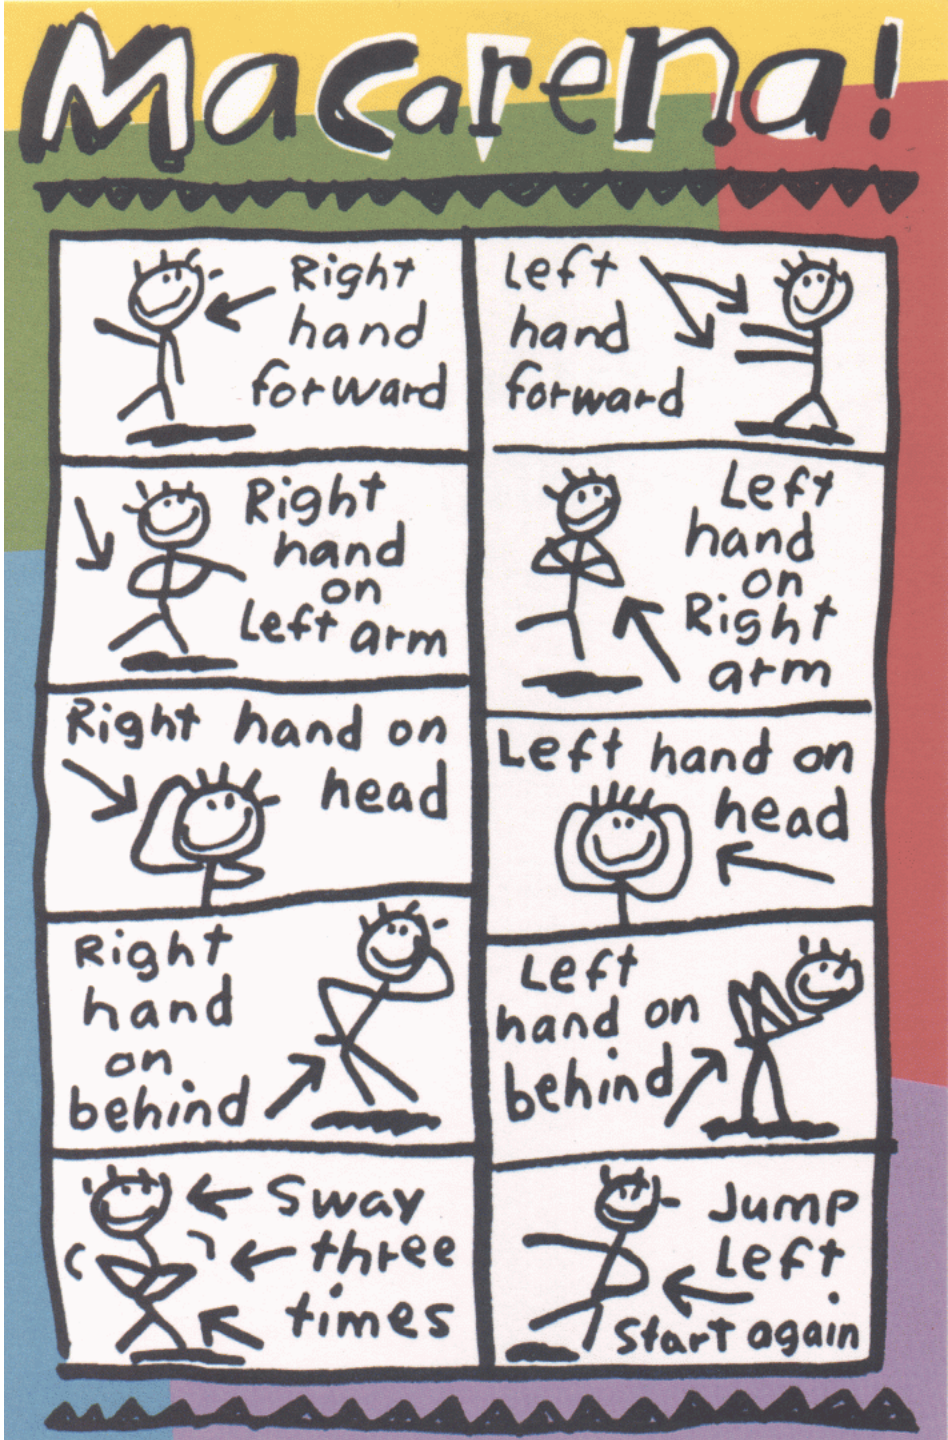

How-To Books

- are often sparse,
- incomplete,
- qualitative, and
- physically impossible.

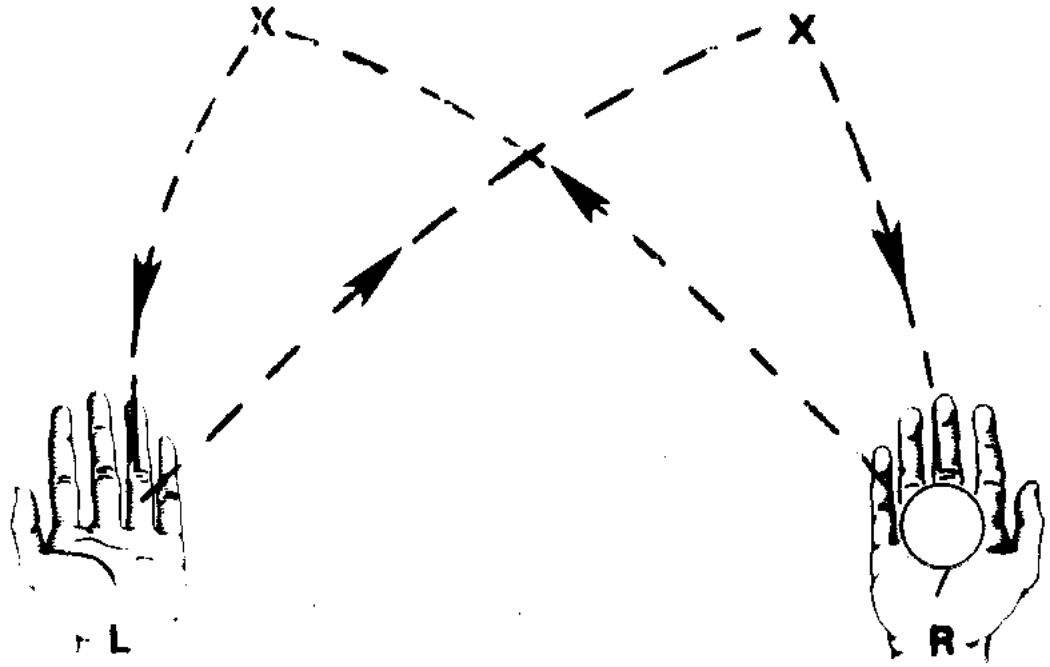

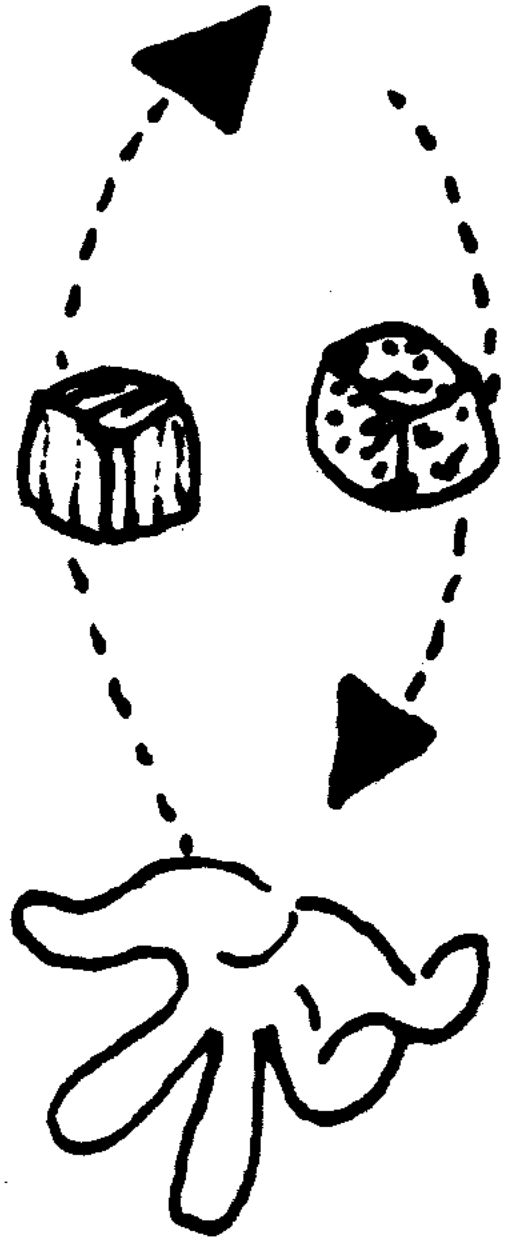

Finding Better Strategies

MacGyver

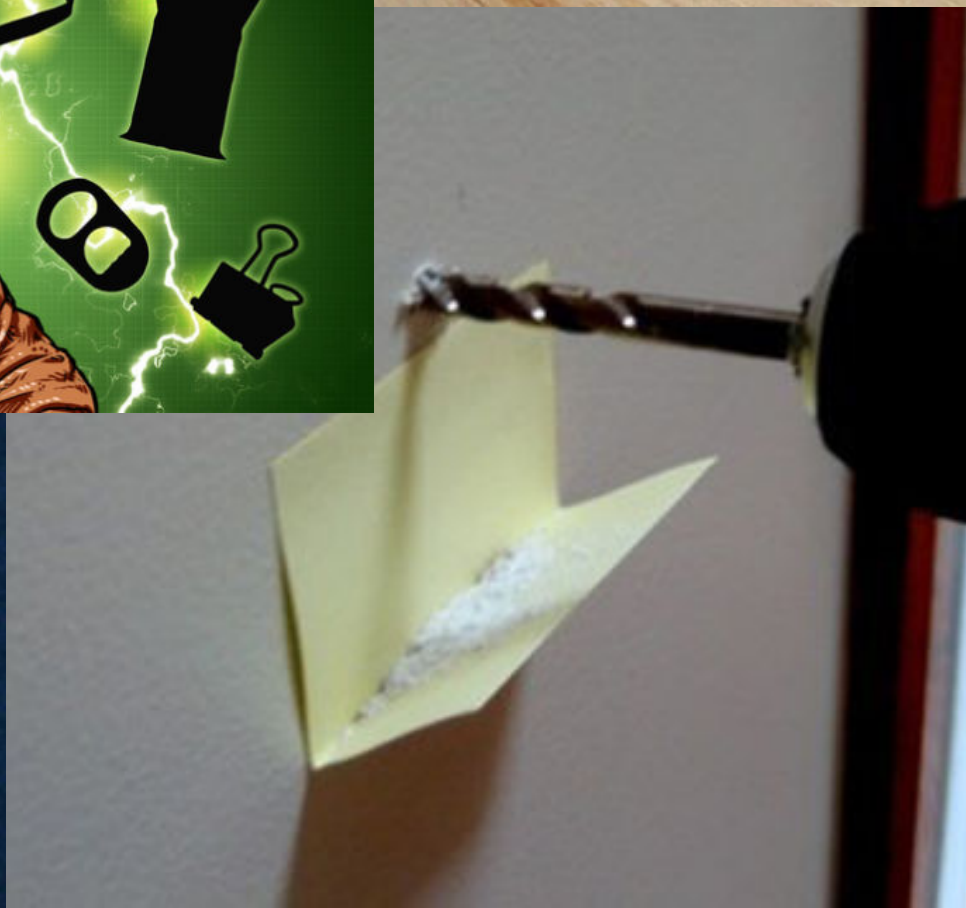

MacGyver: Pilot Episode Tricks

Disarmed missile with a paperclip
Matches & rope got rifle to shoot itself
Smashed pistol barrel to use as rocket thruster
Kicked grate to test for hidden laser
Lit cigarette with hidden laser
Smoked cigarette to find hidden laser
Used binocular eyepiece to catch laser beam
Aimed laser beam at source to "kill" it
Knotted fire hose to build up water pressure
Used said hose to lift steel beam
Tested heat on door with a stick (it caught fire)
Milk Chocolate candy stops sulfuric acid leak
Used shirt to filter gases
Sodium metal & cold capsule explosive
Flipped lights in Morse code

www.neatorama.com/2007/02/20/macgyverisms-complete-list-of-macgyver-episode-tricks/

Opening A Jar: Solved Problem?

Power grip to loosen the lid

Precision grip to unscrew it

Audience Participation: Tricks To Open A Stuck Jar

- What is the trick?
- Where did it come from?
  - You invented it?
  - You were told?
  - Learn from observation?
- How do you adapt it?
- Why does it work?


ROASTED
PEPPERS


wikiHow


Mt.Olive


Mt.Olive


Mt. Olive
wikiHow


GIFAK.NET


Sorghum - (Molass
Made By: Guenther Family
4064 Muddy Pond Road • Monterey, TN 38574
931-445-3357
1 Quart • (32 oz.)


KEWPIE


I
My
Jar
Opener


WAITROSE
REDUCED SUGAR
STRAWBERRY
JAM
EXTRA FRUIT


AS SEEN ON TV
OneTouch
www.luster-twitem.com
Automatic Jar Opener
Lid pops off
Self adjusts to any size!

The Power of Brute Force Search

Key Points Of This Talk

- Task-level model-based learning is effective and data efficient (<100 trials). For learning to control robots, that is the standard to meet. Can handle contact, deformable objects, liquids, granular materials, and other “hard-to-model” situations.
- “Model-free” learning can play a useful role in perceptual and end-to-end learning.
- The most useful role for “model-free” learning is model-based planning which generates policies in simulation (which is what most model-free RL is doing).
- We need non-parametric learning to find better ways of doing things (better task strategies).