Apprenticeship Learning for Autonomous Flight and Surgical Robotics

Pieter Abbeel
UC Berkeley EECS

Collaborators: Adam Coates, Andrew Ng, Jie Tang, Arjun Singh, Nimbus Goehausen; Jur van den Berg, Stephen Miller, Ken Goldberg, Humphrey Hu, Daniel Duckworth

Objective

- Autonomous execution of trajectory-based tasks for systems with complicated dynamics
- E.g.,

- Challenges:
  - How to specify the trajectory?
  - How to build a controller?

Outline

- **Learning a target trajectory**
- Learning a dynamics model
- Autonomous flight results
- Surgical robotics
- Conclusion

Target trajectory

- Difficult to specify by hand:
  - Required format: position + orientation over time
  - Needs to satisfy dynamics

- Our solution:
  - Collect demonstrations of desired maneuvers
  - Challenge: extract a clean target trajectory from many suboptimal/noisy demonstrations

Expert demonstrations: Airshow

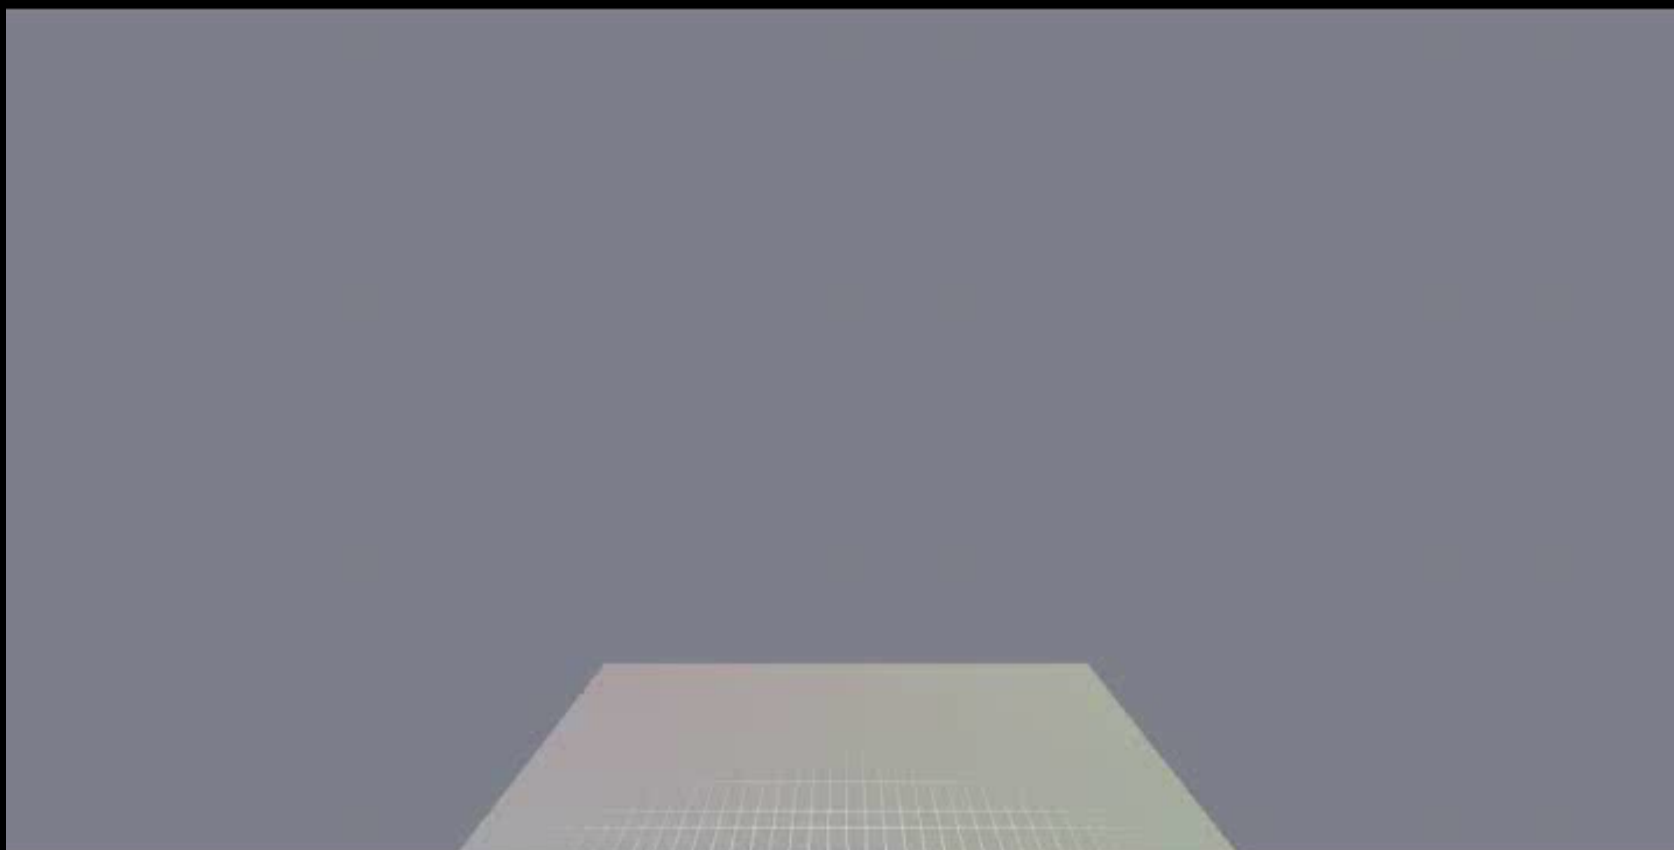

Learning Trajectory

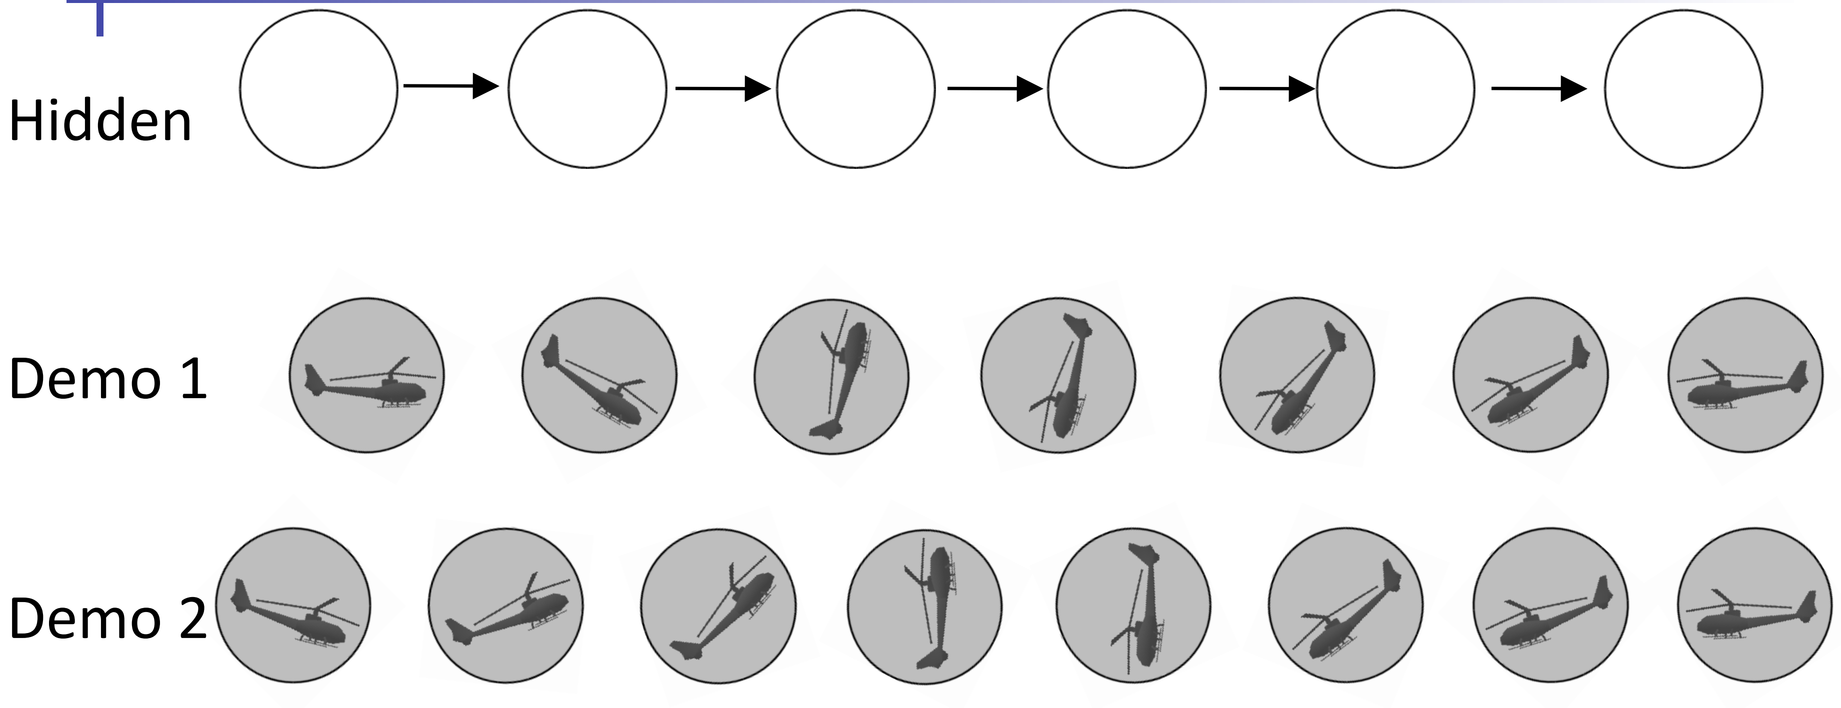

- HMM-like generative model
  - Dynamics model used as HMM transition model
  - Demos are observations of hidden trajectory
- Problem: how do we align observations to hidden trajectory?

Learning Trajectory

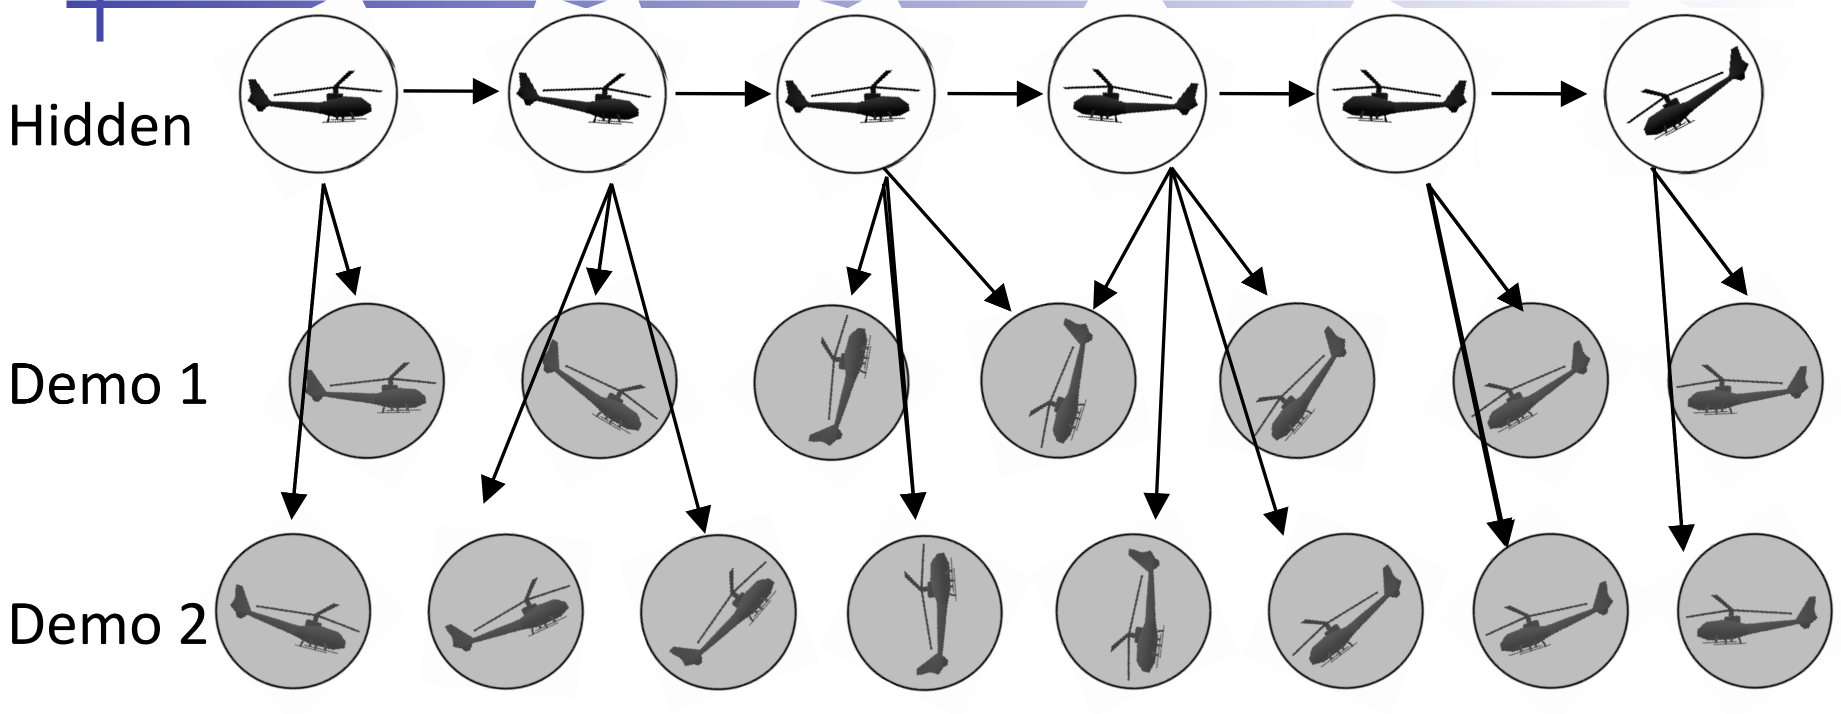

- Dynamic Time Warping (Needleman&Wunsch 1970, Sakoe&Chiba, 1978)
- Extended Kalman filter / smoother

Results: Time-aligned demonstrations

- White helicopter is inferred "intended" trajectory.

Results: Loops

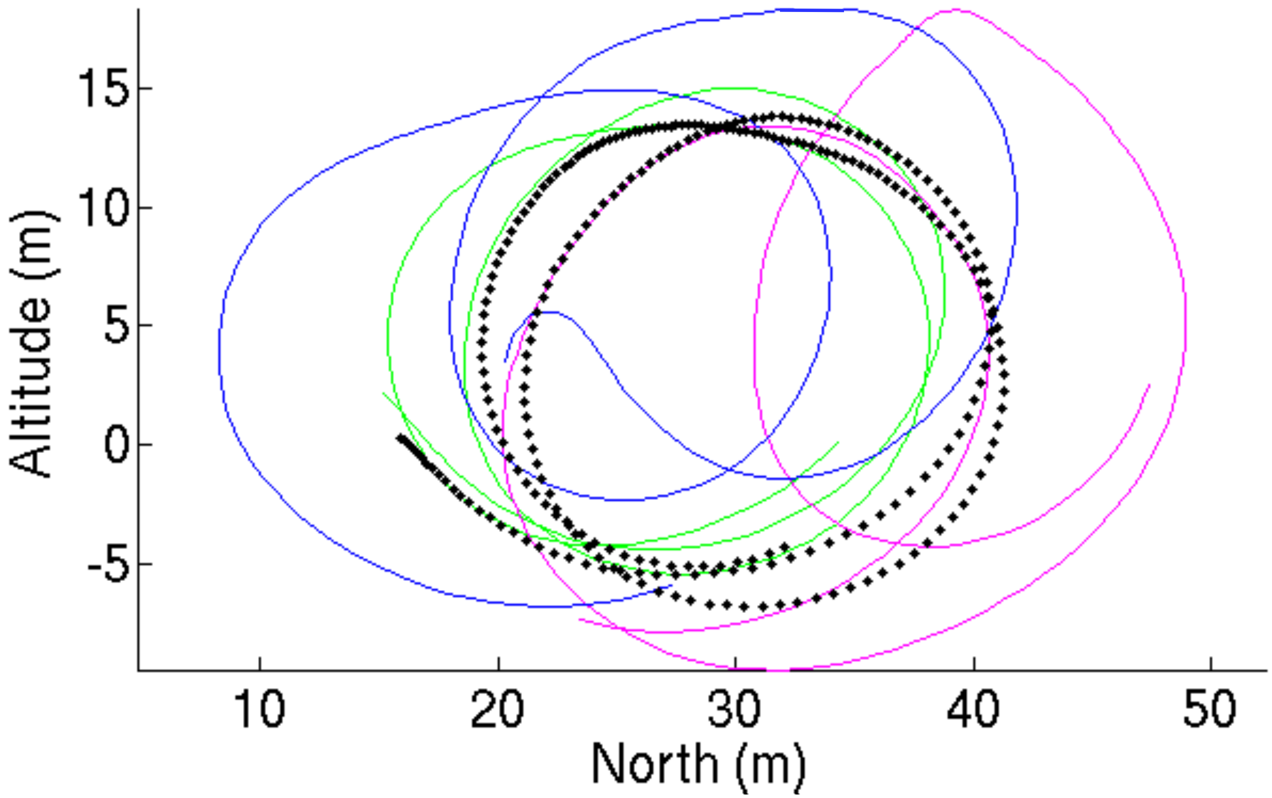

- Even without prior knowledge, the inferred trajectory is much closer to an ideal loop.
- If desired, can incorporate prior knowledge as prior.

Outline

- Learning a target trajectory
- **Learning a dynamics model**
- Autonomous flight results
- Surgical robotics
- Conclusion

Empirical evaluation of standard modeling approach

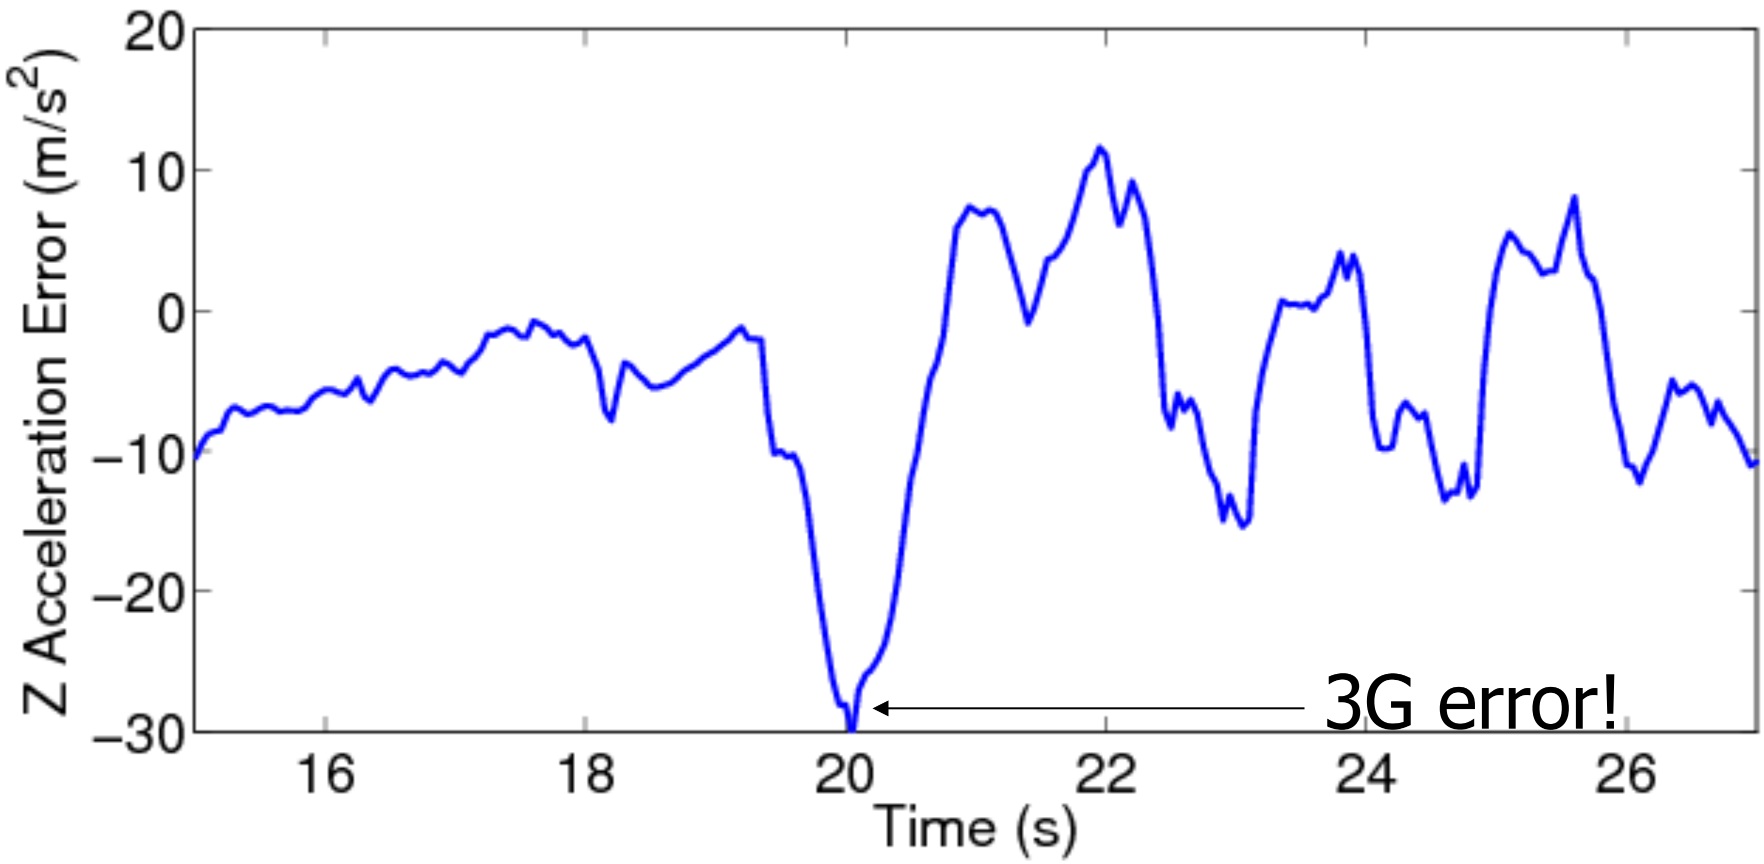

Key observation

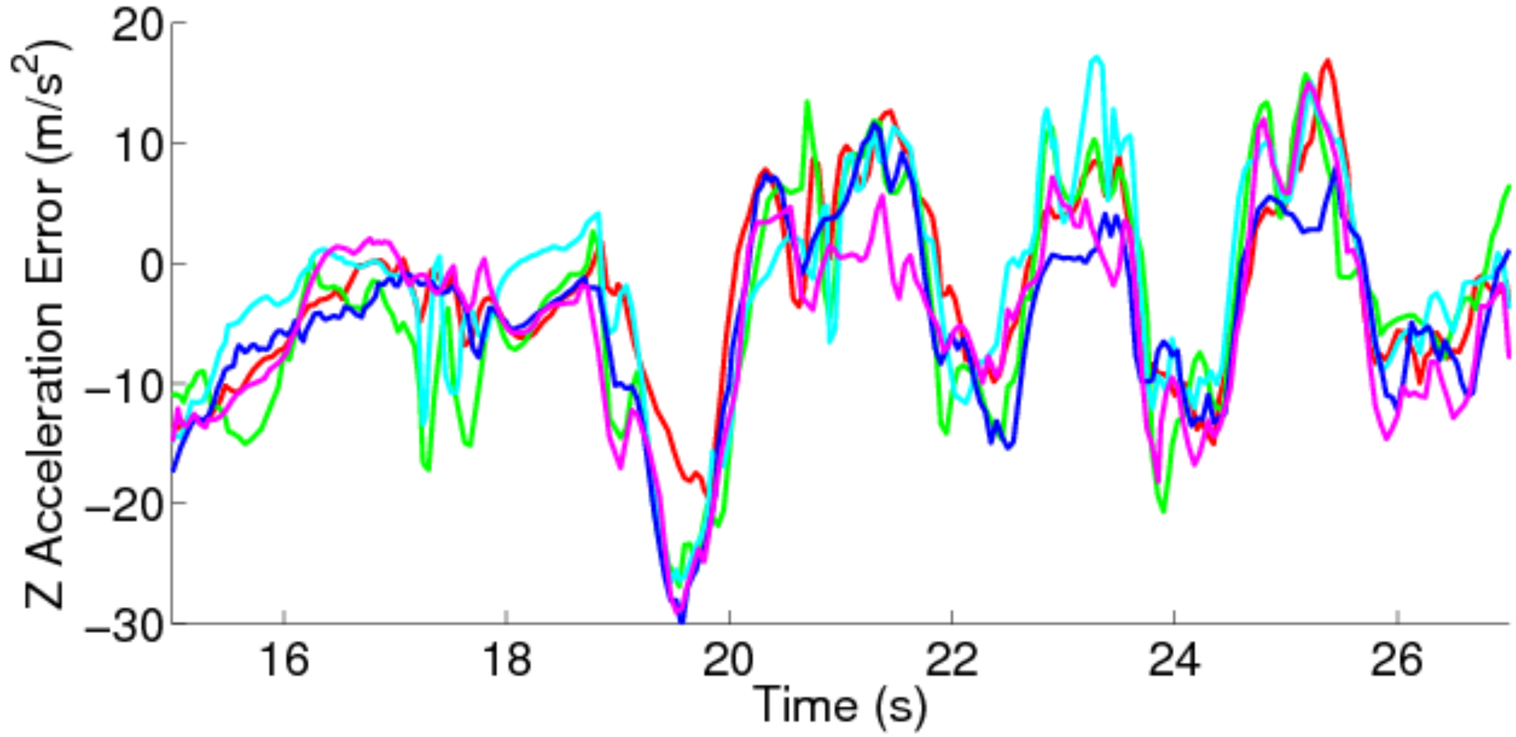

- Errors observed in the “baseline” model are clearly consistent after aligning demonstrations.

Key observation

- If we fly the same trajectory repeatedly, errors are consistent over time once we align the data.

  - There are many unmodeled variables that we can't expect our model to capture accurately.
    - Air (!), actuator delays, etc.
  - If we fly the same trajectory repeatedly, the hidden variables tend to be the same each time.

cf. muscle memory for humans

Trajectory-specific local models

- Learn locally-weighted model from aligned demonstration data
  - Since data is aligned in time, we can weight by *time* to exploit repeatability of unmodeled variables.
  - For model at time t: $W(t') = \exp(-(t - t')^2 / \sigma^2)$
  - Obtain a model for each time t into the maneuver by running weighted regression for each time t

Outline

- Learning a target trajectory
- Learning a dynamics model
- **Autonomous flight results**
- Surgical robotics
- Conclusion

Many success stories in hover and forward flight regime

- Just a few examples: Bagnell & Schneider, 2001; LaCivita, Papageorgiou, Messner & Kanade, 2002; Ng, Kim, Jordan & Sastry 2004a (2001); Roberts, Corke & Buskey, 2003; Saripalli, Montgomery & Sukhatme, 2003; Shim, Chung, Kim & Sastry, 2003; Doherty et al., 2004; Gavrilets, Martinos, Mettler and Feron, 2002; Ng et al., 2004b.
- Varying control techniques: inner/outer loop PID with hand or automatic tuning, H1, LQR, …
- Very few results outside of stationary regimes --- exception: Gavrilets, Martinos, Mettler, Feron 2002

One of our first attempts at autonomous flips [using similar methods to what worked for ihover]

Target trajectory: meticulously hand-engineered
Model: from (commonly used) frequency sweeps data

Experimental Setup

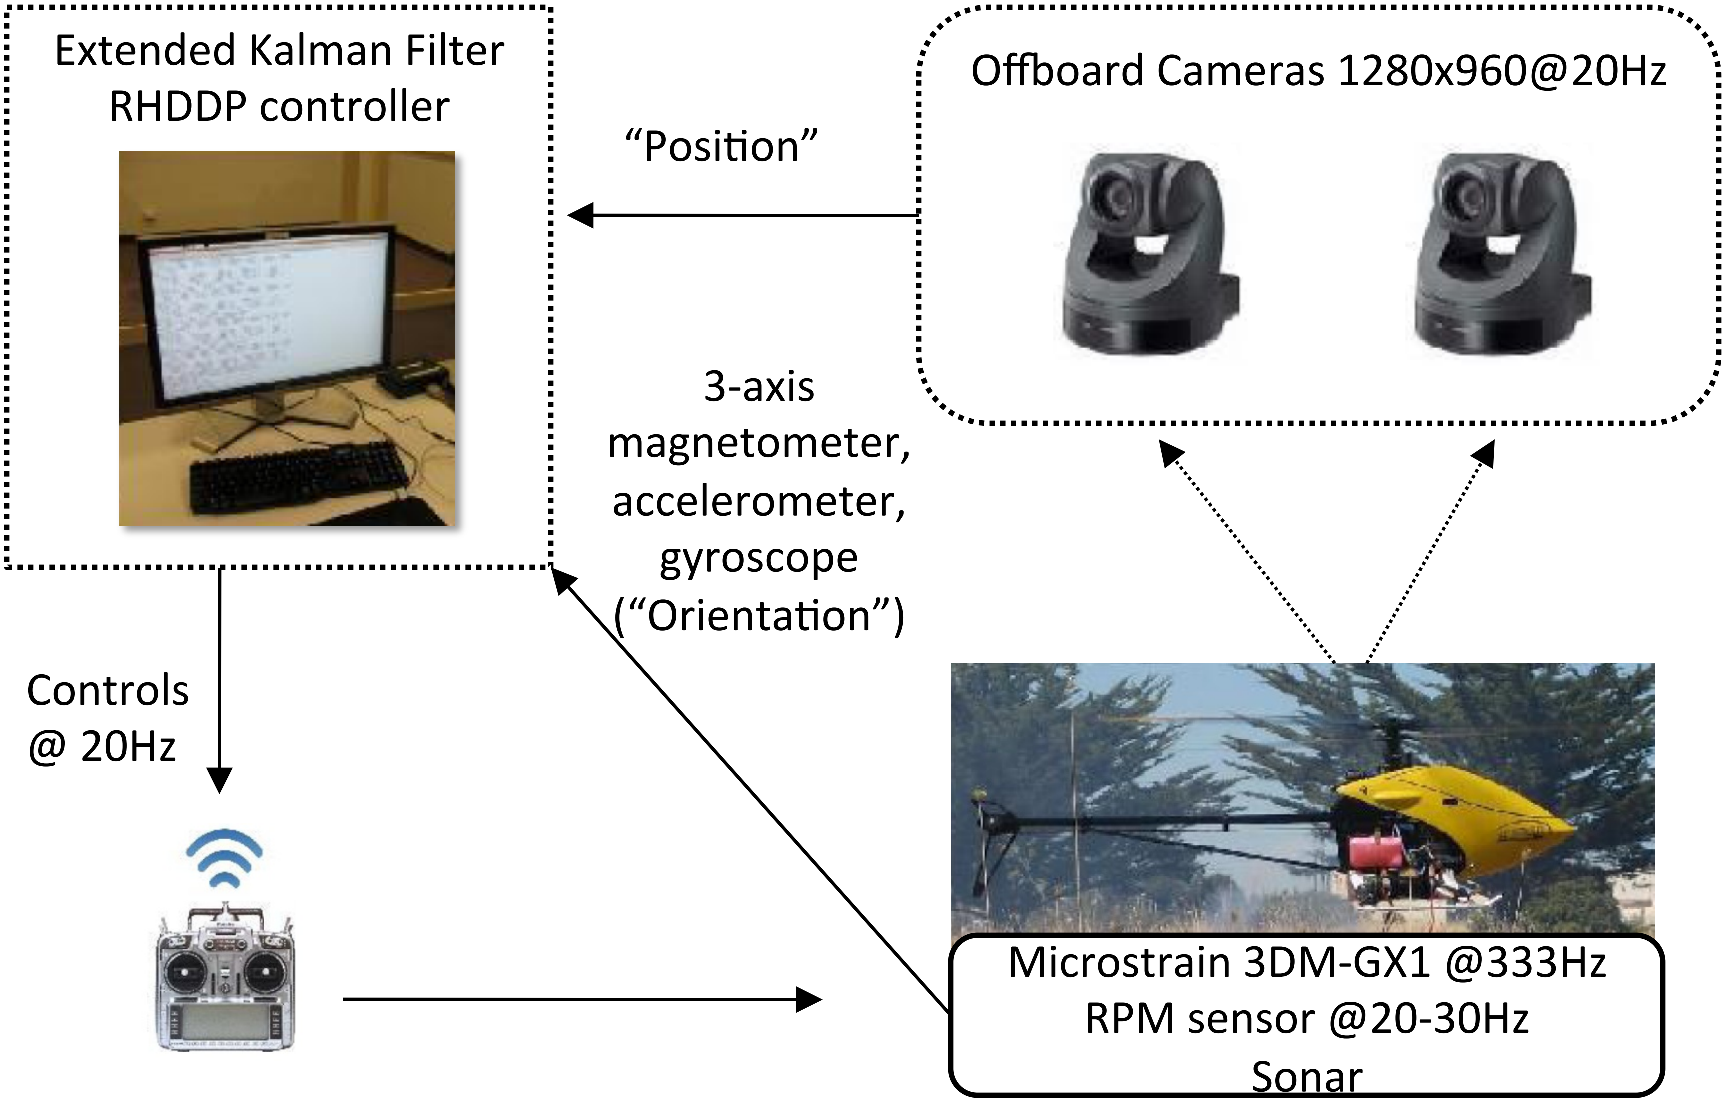

Experimental procedure

1. Collect sweeps to build a baseline dynamics model
2. Our expert pilot demonstrates the airshow several times.

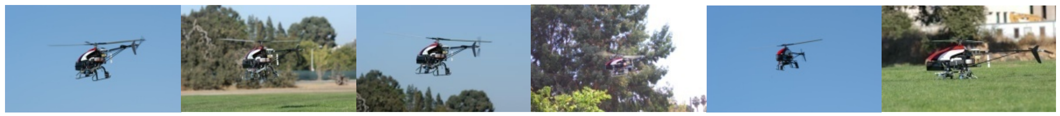

3. Learn a target trajectory.
4. Learn a dynamics model.
5. Find the optimal control policy for learned target and dynamics model.
6. Autonomously fly the airshow

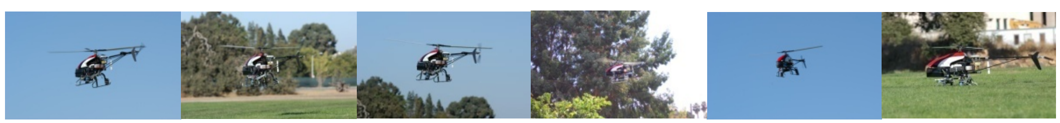

7. Learn an improved dynamics model. Go back to step 4.

→ **Learn to fly new maneuvers in < 1hour.**

Results: Autonomous airshow

Results: Flight accuracy

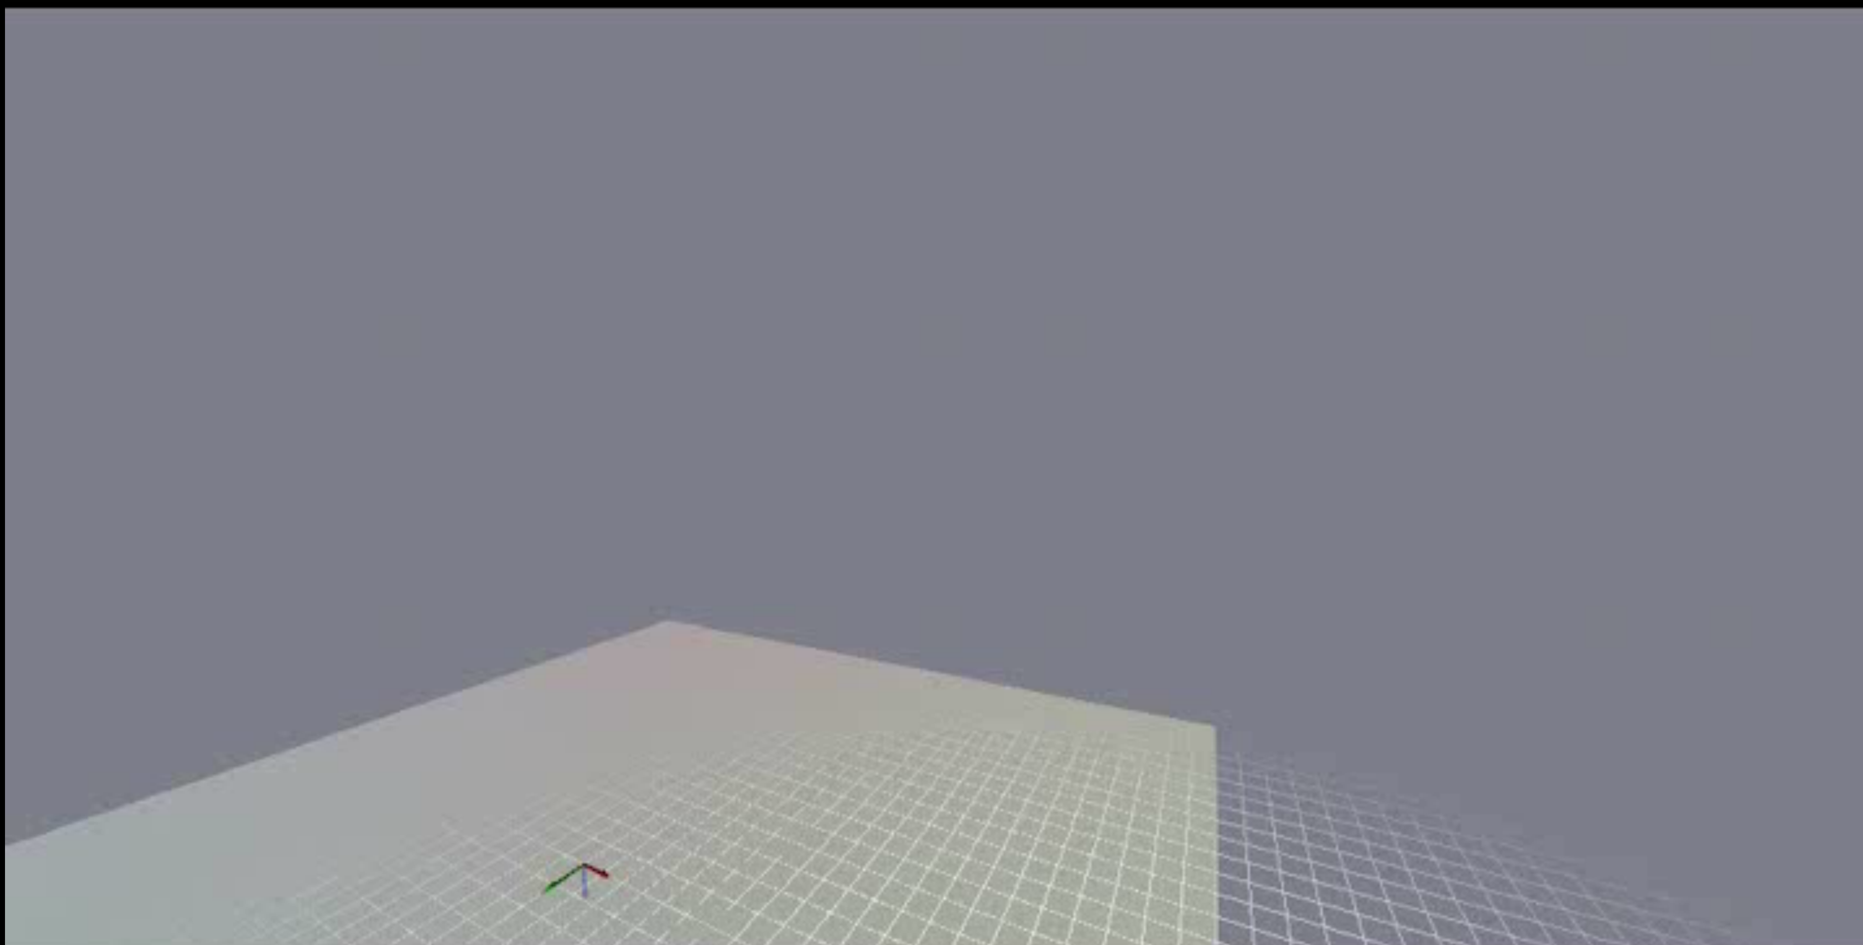

Outline

- Learning a target trajectory
- Learning a dynamics model
- Autonomous flight results
- **Surgical robotics**
- Conclusion

Objective

- Autonomous surgical assistants
  - Surgeons still perform major work
  - Robots autonomously perform menial tasks
    - Enhance surgeon performance
    - Reduce tedium and medical errors
    - Reduce operation time and costs
    - Improve patient health

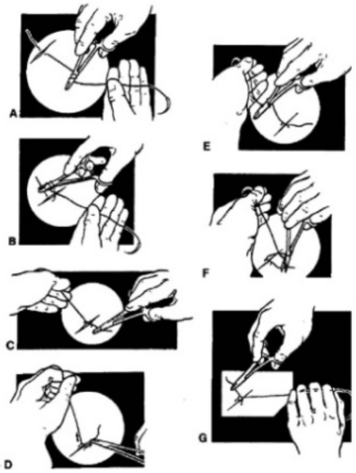

Our Robotic Setup

- Two surgical robots
  - 6 degrees of freedom
  - Integration with Da Vinci end-effector
  - Cavusoglu et al., 1999

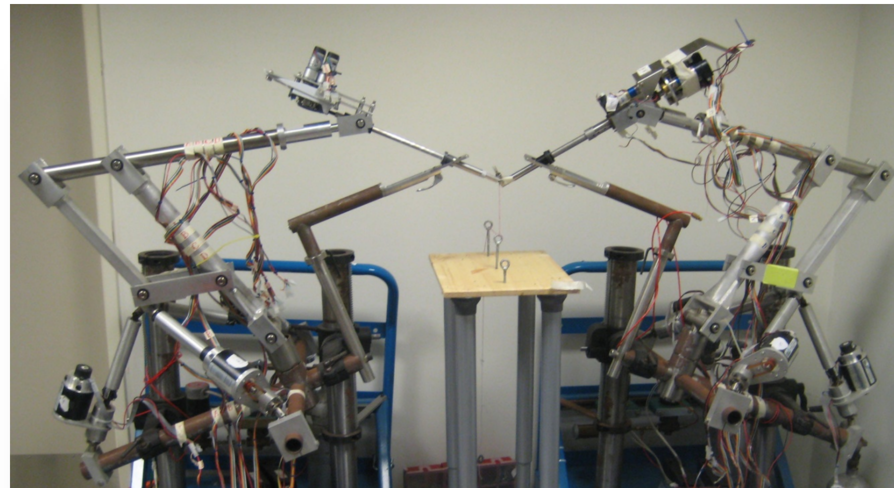

- Two input devices for human control
  - Same degrees of freedom

Results: Knot tie

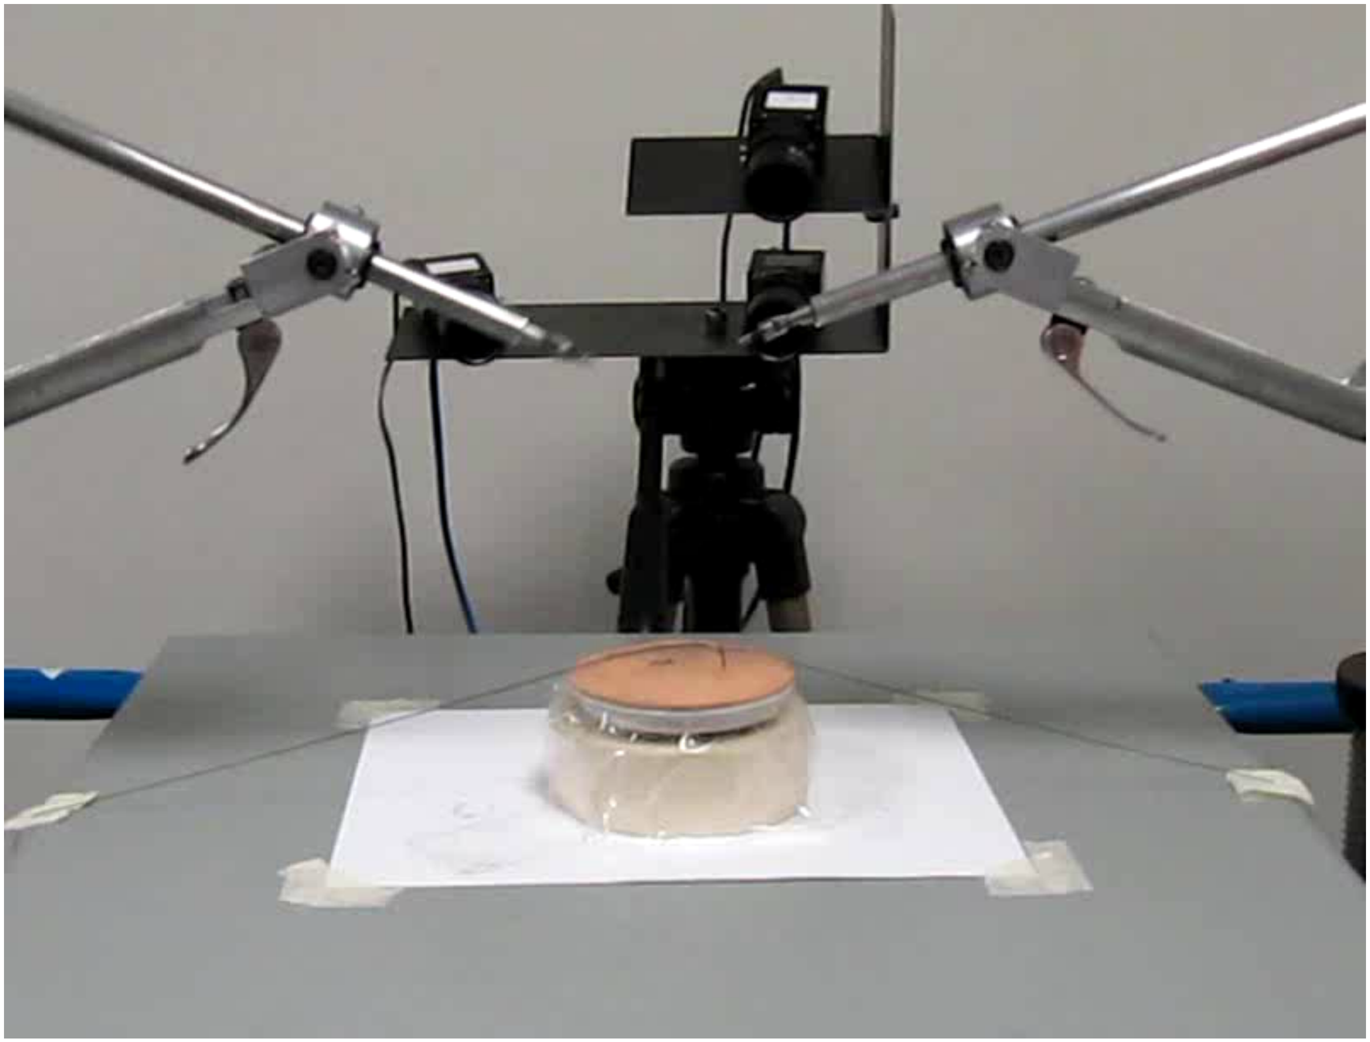

Surgical sub-skills: discussion

- Current Limitations
  - Specific to initial conditions
  - Robots are "blind"
  - Hardware limitations

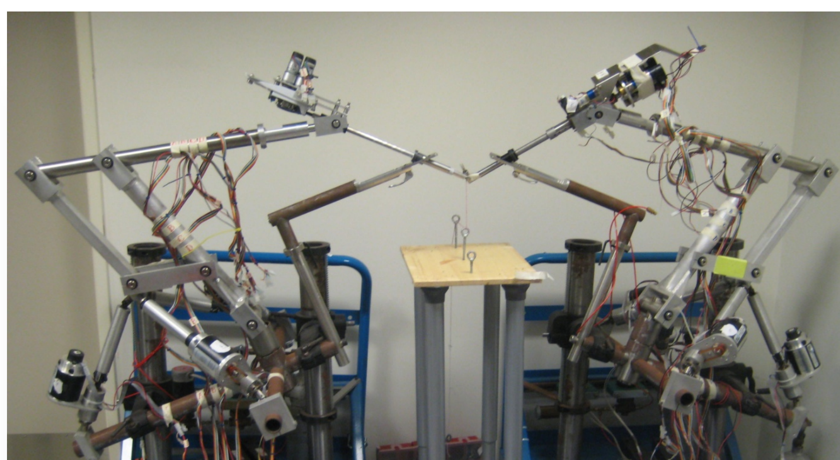

Conclusion

- For systems with complicated dynamics hard to obtain
  - Task trajectory specification
  - Dynamics model
- Our approach uses multiple expert demonstrations to learn:
  - Task trajectory
  - Dynamics models along the trajectory for control.
- Enabled robotic abilities beyond the prior state of the art
- Current directions:
  - Parameterize trajectories
  - Adapt to environment