CS 287 Lecture 24 (Fall 2019)
Autonomous Helicopter Flight

Pieter Abbeel
UC Berkeley EECS


STANFORD
STANFORD
STANFORD
STANFORD

Challenges in Helicopter Control

- Unstable
- Nonlinear
- Complicated dynamics
  - Air flow
  - Coupling
  - Blade dynamics
- Noisy estimates of position, orientation, velocity, angular rate (and perhaps blade and engine speed)

Success Stories: Hover and Forward Flight

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

[Ng, Coates, Tse, et al, 2004]

Alan Szabo – Sunday at the Lake

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

Stationary vs. Aggressive Flight

- Hover / stationary flight regimes:
  - Restrict attention to specific flight regime
  - Extensive data collection = collect control inputs, position, orientation, velocity, angular rate
  - Build model + model-based controller

→ Successful autonomous flight.

- Aggressive flight maneuvers --- additional challenges:
  - **Task description**: What is the target trajectory?
  - **Dynamics model**: How to obtain accurate model?

Aggressive, Non-Stationary Regimes

- Gavrilets, Martinos, Mettler and Feron, 2002
  - 3 maneuvers: split-S, snap axial roll, stall-turn
  - Key: Expert engineering of controllers after human pilot demonstrations

Sunday in Open Loop

Aggressive, Non-Stationary Regimes

- Our work:
  - Key: Automatic engineering of controllers after human pilot demonstrations through machine learning
  - Wide range of aggressive maneuvers
  - Maneuvers in rapid succession

Learning Dynamic Maneuvers

- **Learning a target trajectory**
- Learning a dynamics model
- Autonomous flight results

Target Trajectory

- Difficult to specify by hand:
  - Required format: position + orientation over time
  - Needs to satisfy helicopter dynamics
- Our solution:
  - Collect demonstrations of desired maneuvers
  - Challenge: extract a clean target trajectory from many suboptimal/noisy demonstrations

Abbeel, Coates, Ng, IJRR 2010

Expert Demonstrations

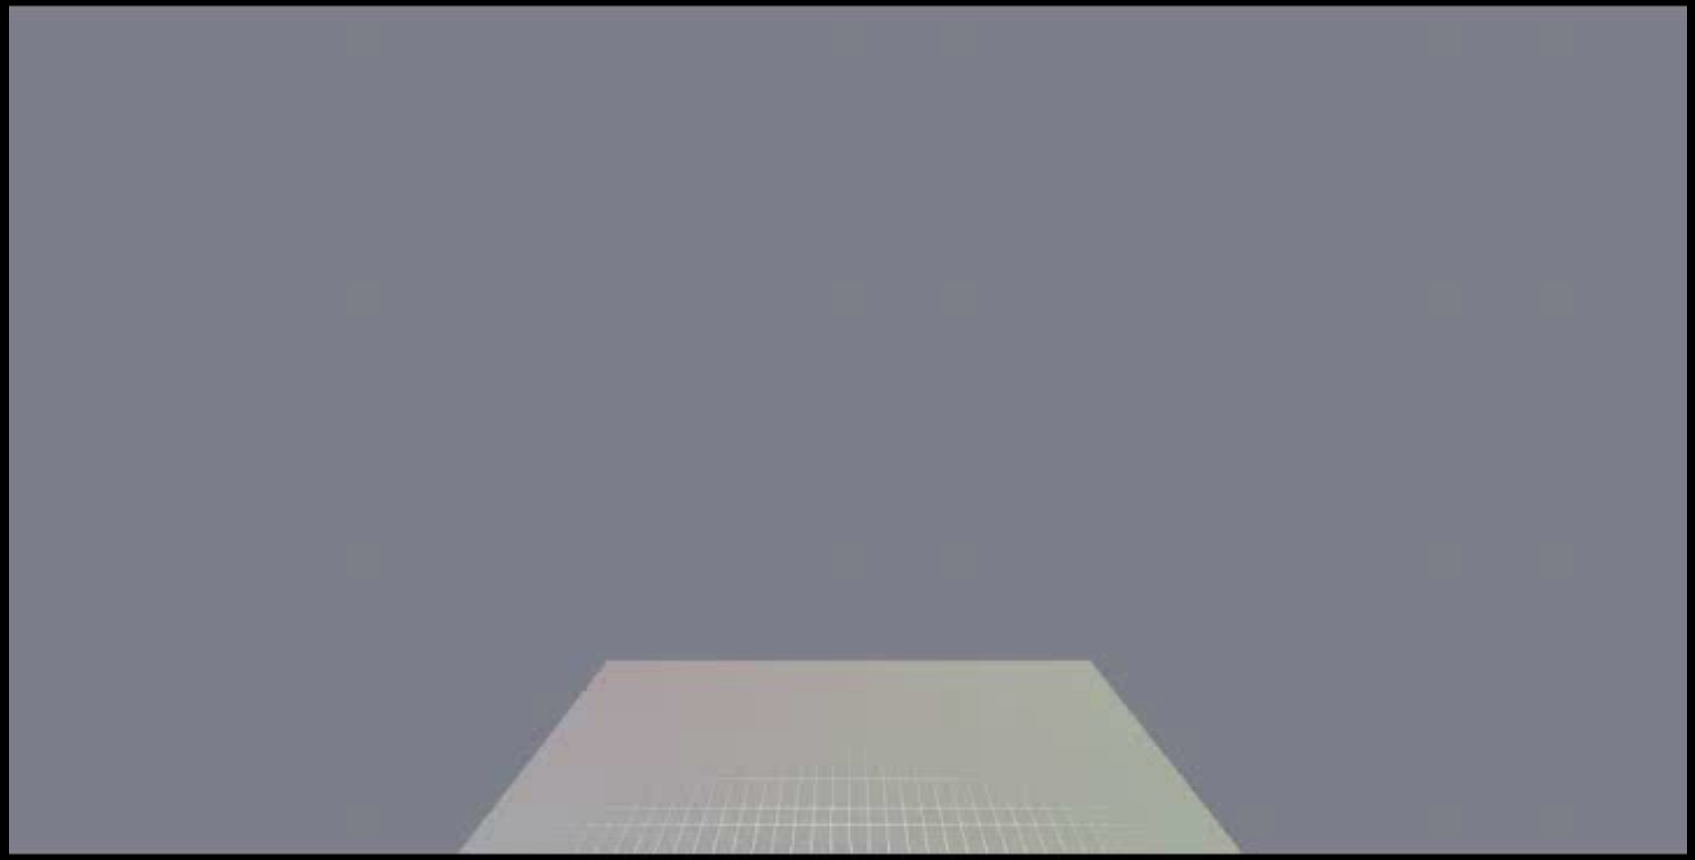

Learning a Trajectory

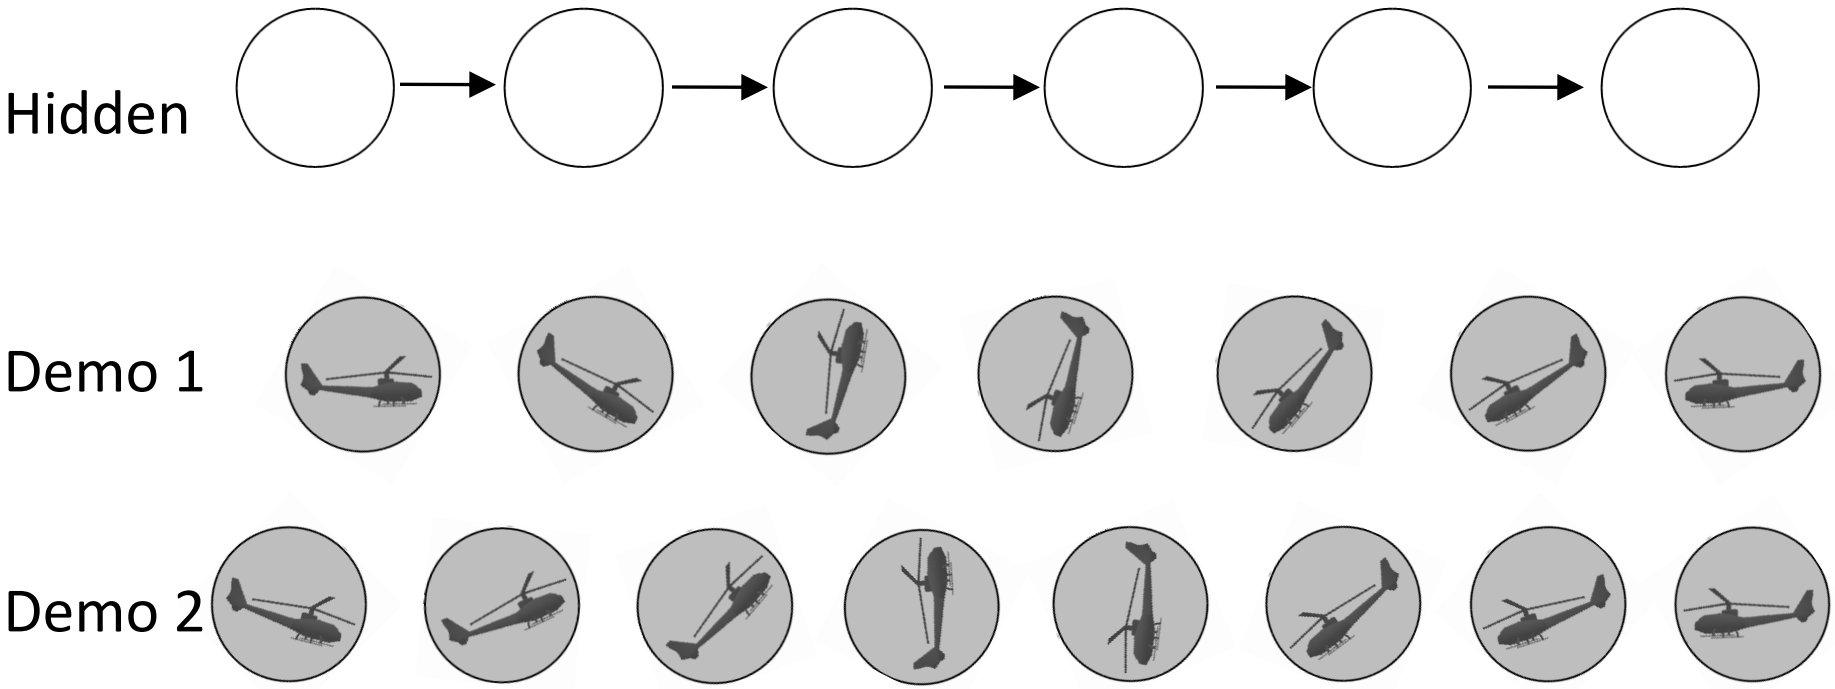

- HMM-like generative model
  - Dynamics model used as HMM transition model
  - Demos are observations of hidden trajectory
- Problem: how do we align observations to hidden trajectory?

Abbeel, Coates, Ng, IJRR 2010

Learning a Trajectory

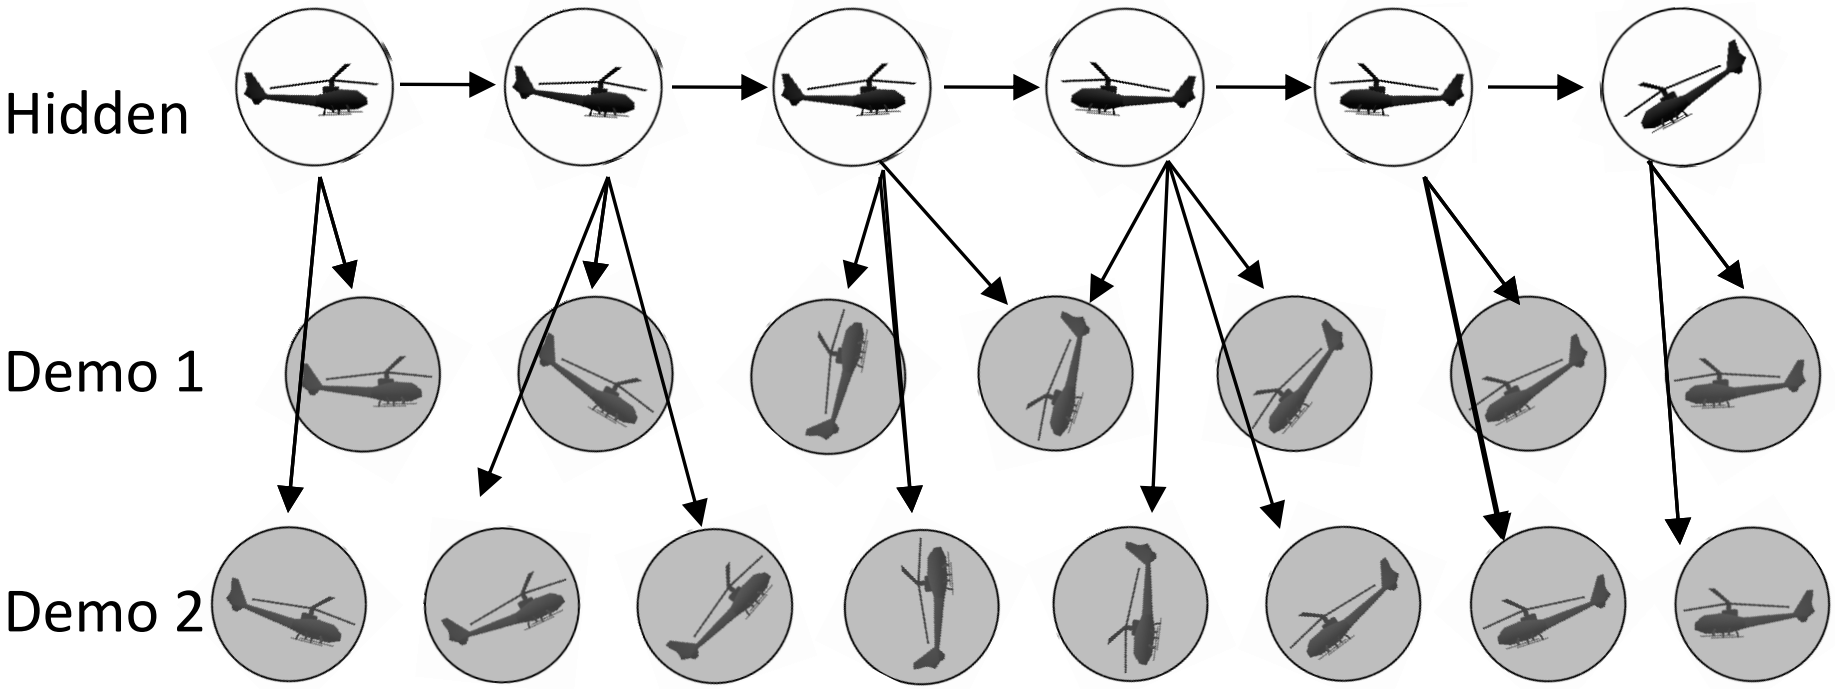

- Dynamic Time Warping (Needleman&Wunsch 1970 Sakoe&Chiba, 1978)
- Extended Kalman filter / smoother

Abbeel, Coates, Ng, IJRR 2010

Results: Time-Aligned Demonstrations

- White helicopter is inferred “intended” trajectory.

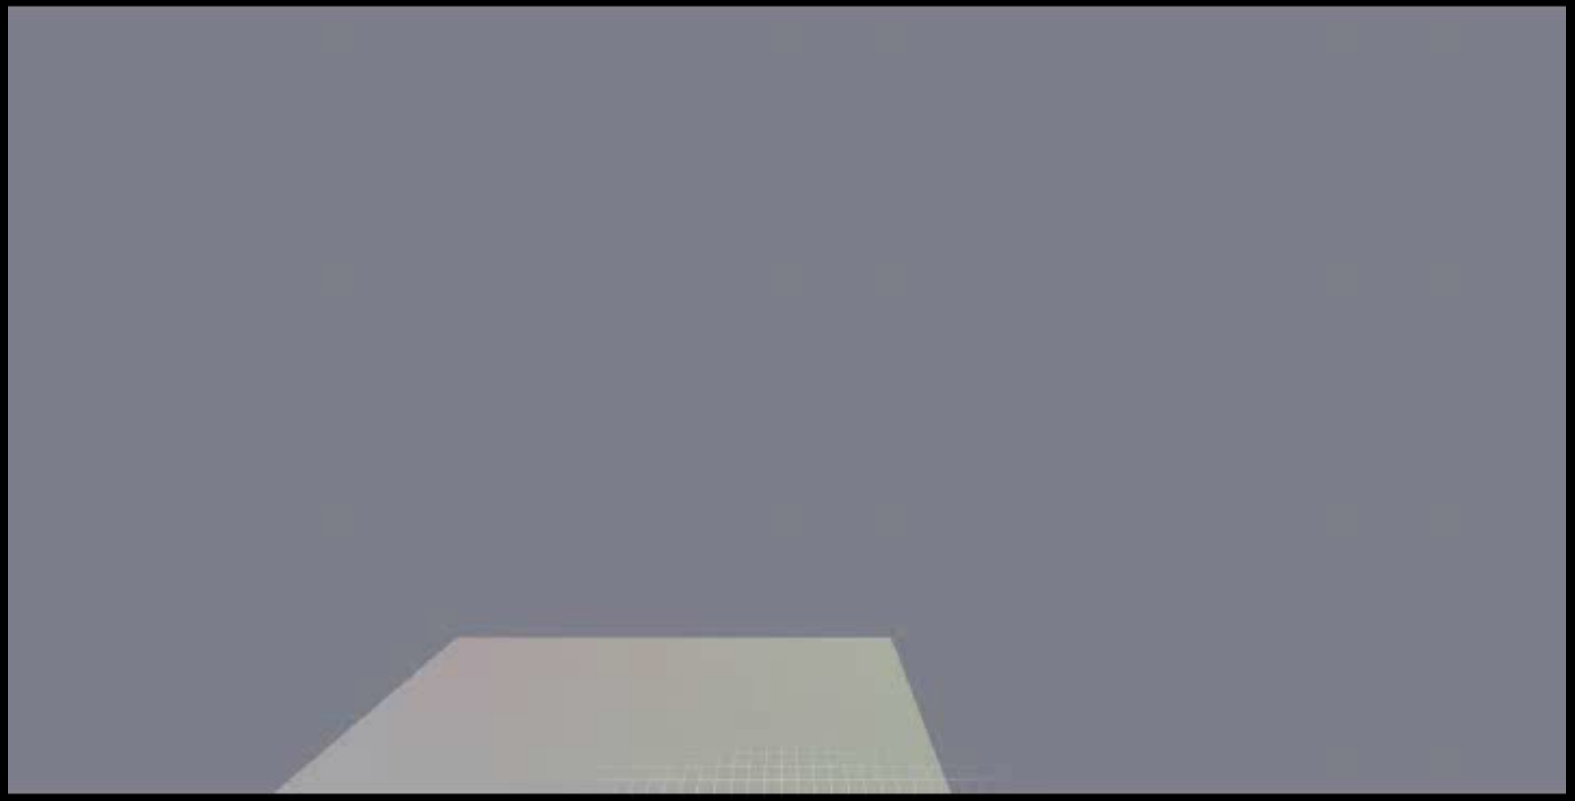

Results: Loops

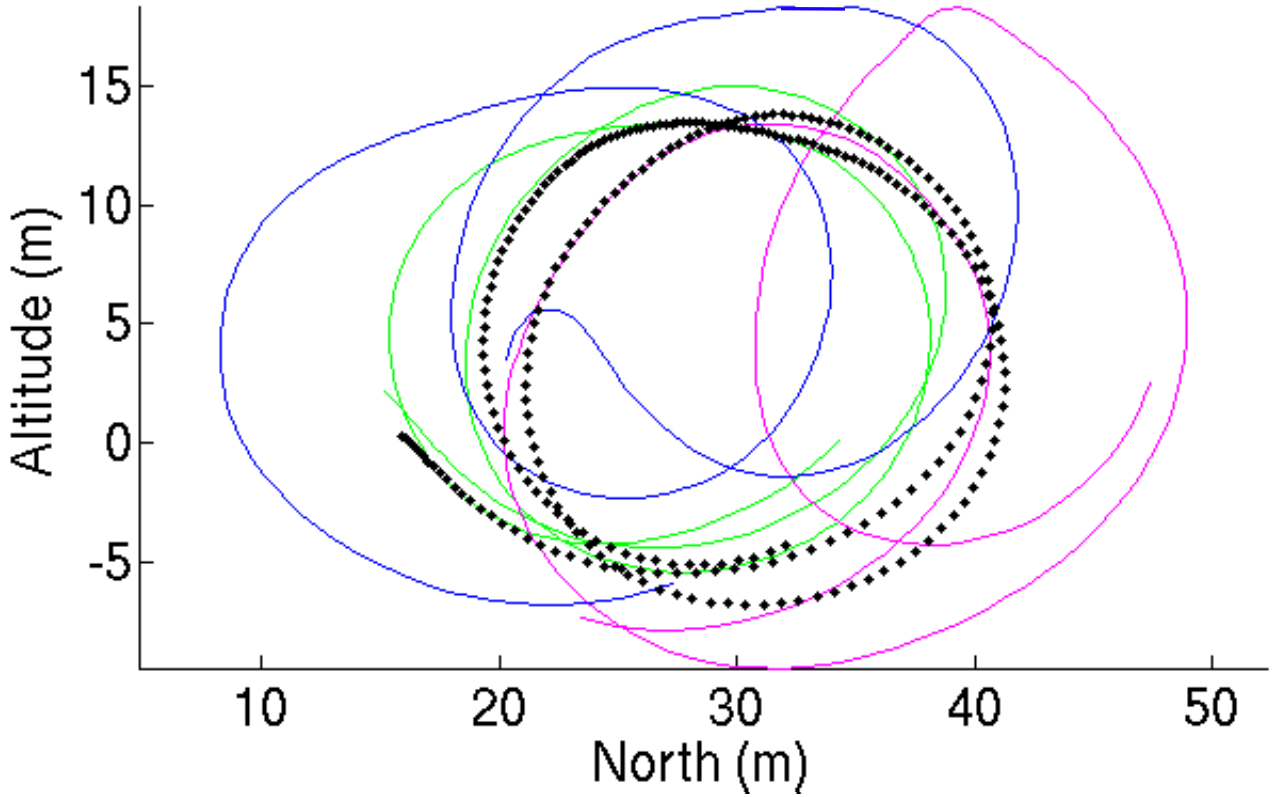

Even without prior knowledge, the inferred trajectory is much closer to an ideal loop.

Abbeel, Coates, Ng, IJRR 2010

Learning Dynamic Maneuvers

- Learning a target trajectory
- **Learning a dynamics model**
- Autonomous flight results

Standard Modeling Approach

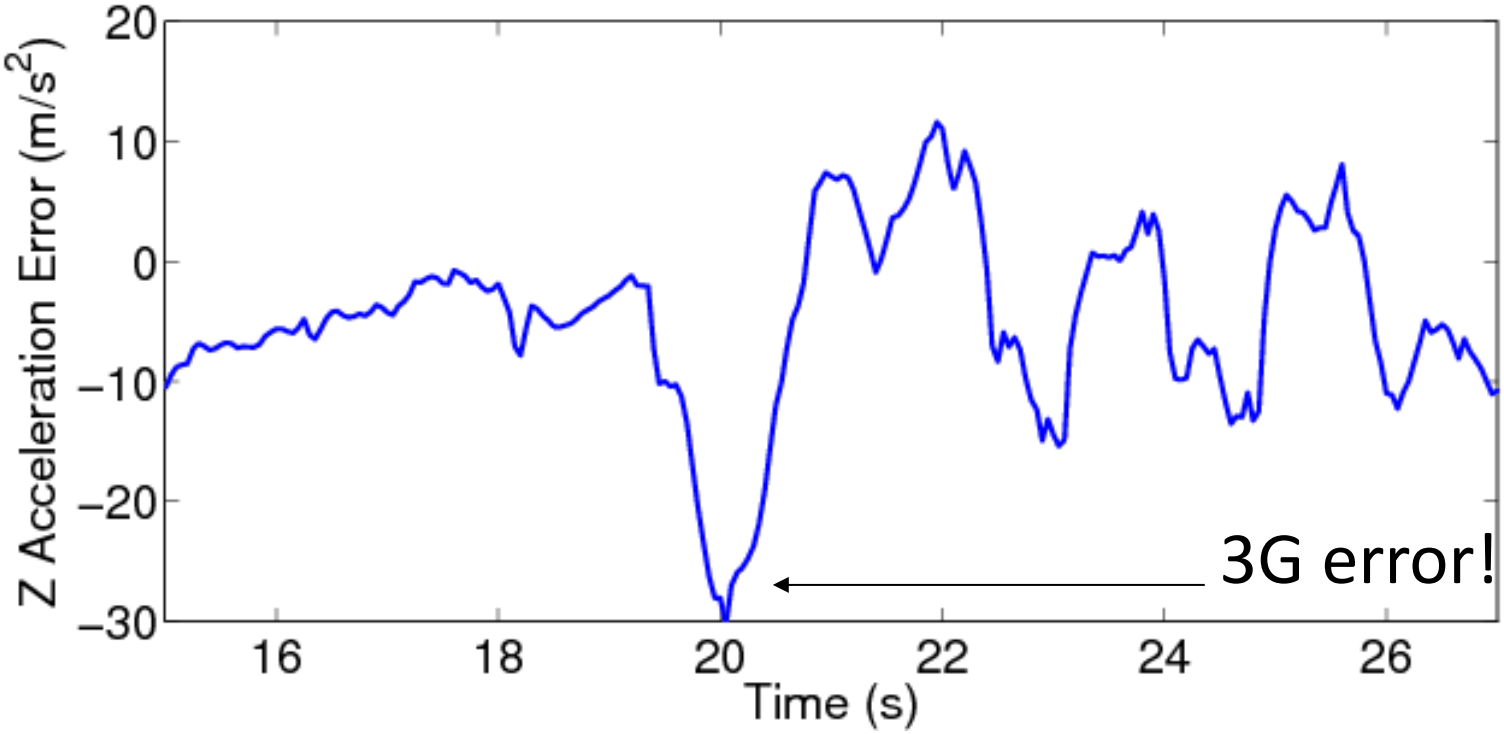

Abbeel, Coates, Ng, IJRR 2010

Key Observation

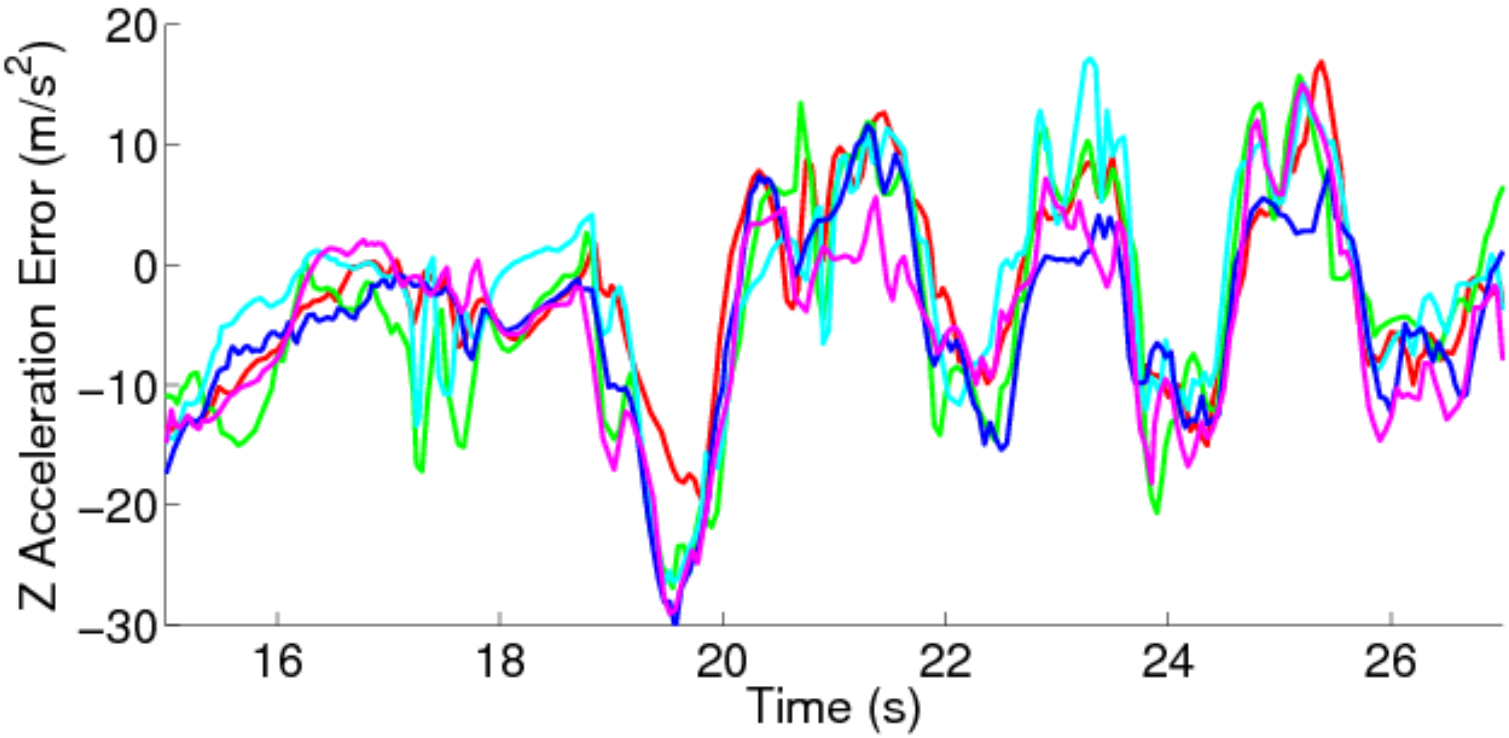

Errors observed in the “baseline” model are clearly consistent after aligning demonstrations.

Abbeel, Coates, Ng, IJRR 2010

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.

~ muscle memory for human pilots

Abbeel, Coates, Ng, IJRR 2010

Trajectory-Specific Local Models

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

Abbeel, Coates, Ng, IJRR 2010

Learning Dynamic Maneuvers

- Learning a target trajectory
- Learning a dynamics model
- **Autonomous flight results**

Abbeel, Coates, Ng, IJRR 2010

Experimental Setup

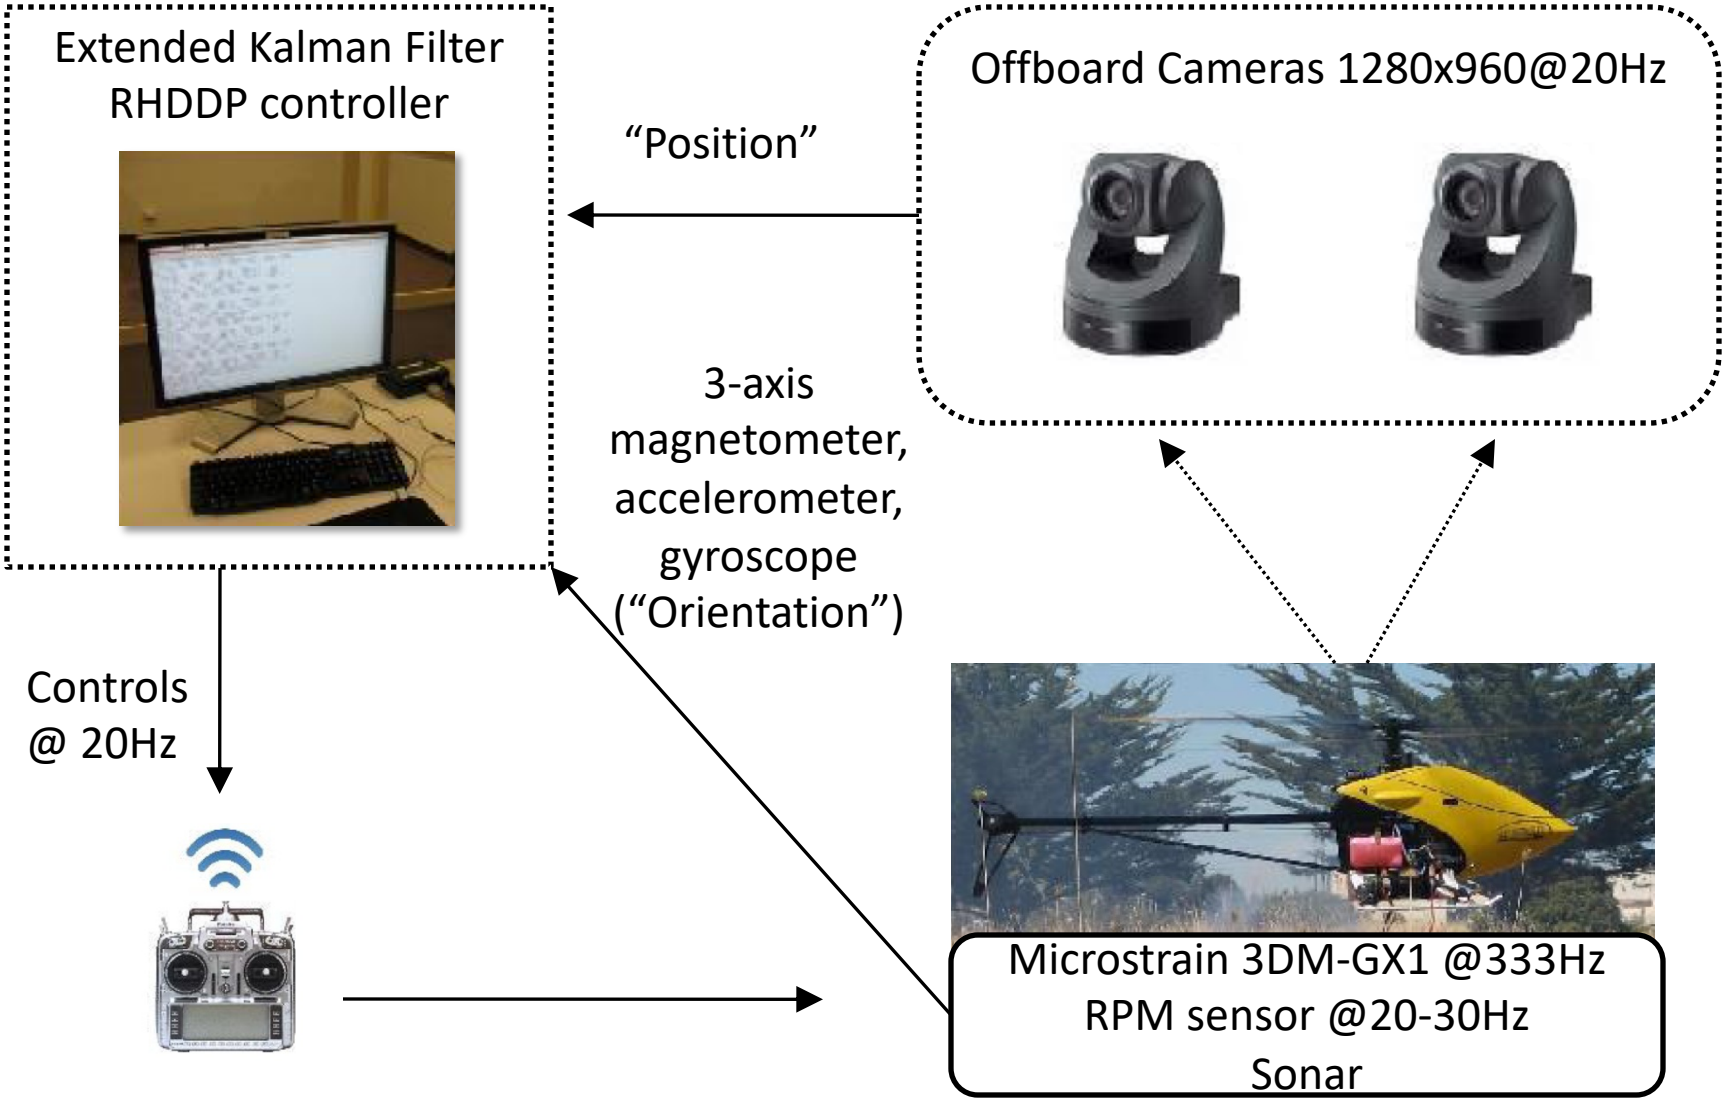

Abbeel, Coates, Quigley, Ng, NIPS 2007

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.**

Abbeel, Coates, Ng, IJRR 2010

Results: Autonomous Airshow

Results: Flight Accuracy

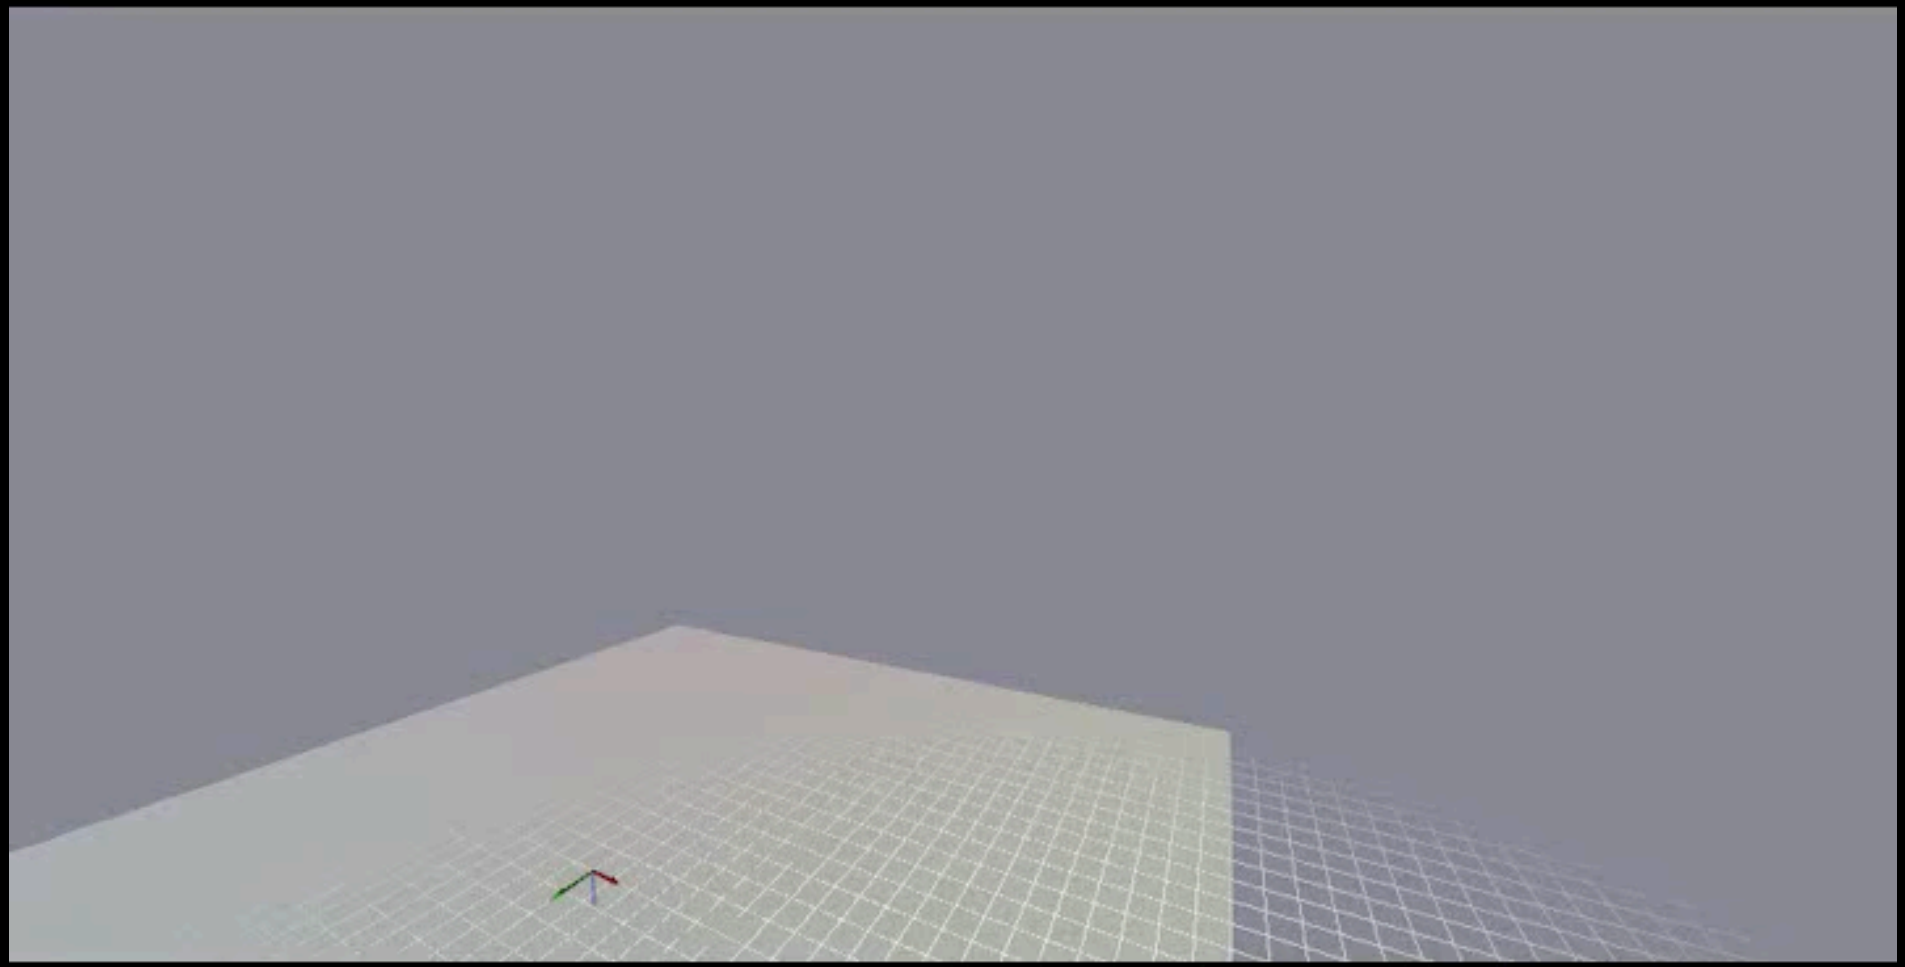

Autonomous Autorotation Flights

Abbeel, Coates, Hunter, Ng, ISER 2008

Chaos [“flip/roll” parameterized by yaw rate]

Behind the scenes

Thank you!