Learning Parameterized Maneuvers for Autonomous Helicopter Aerobatics from Expert Demonstration

Jie Tang Arjun Singh Nimbus Goehausen Pieter Abbeel
{jietang@eecs, arjunsingh@, nimbus@, pabbeel@cs}.berkeley.edu
EECS Department, University of California, Berkeley

Introduction

- Robotics tasks often involve specifying complex trajectories
  - Ex: flying a stall turn with a helicopter
- Goal: Learn representations of difficult maneuvers that we can query to obtain novel trajectories.

Challenges

- Hard to specify maneuvers in high dimensional spaces
- Many robotics platforms cannot be accurately modeled through all operating regimes.

Solution Idea

- Leverage (suboptimal) expert demonstrations to learn target trajectories
- Learn dynamics model locally tuned for specific maneuver.
- Extend current state of the art in helicopter aerobatics: A. Coates, P. Abbeel, A. Ng, ICML 2008 [1].
  - Can fly aggressive helicopter maneuvers with few demonstrations.
- Specify waypoints to generate novel instances of maneuver

Our Approach

Input: *Maneuver type, target waypoints*
Output: *Maneuver that passes through waypoints*
1. Gather demonstrations
2. Initialize target trajectory
3. Repeat:
3a. Time-alignment
3b. Trajectory inference

Gather Demonstrations

- Gathered expert demonstrations are suboptimal and inconsistent.

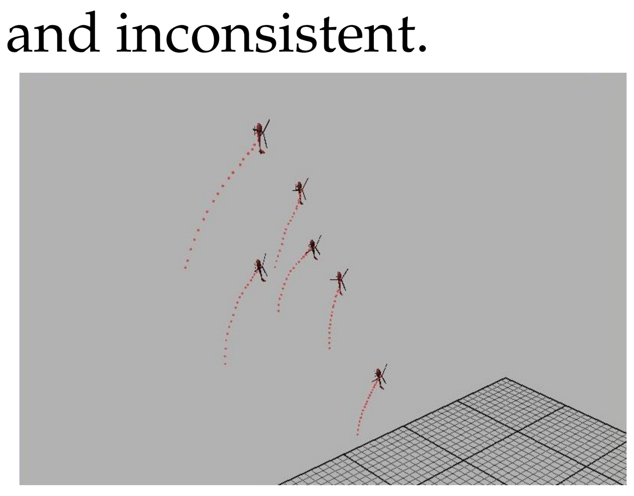

Time Alignment

- Align demonstrations so important structure is not smoothed away.
- Dynamic Time Warping [2].

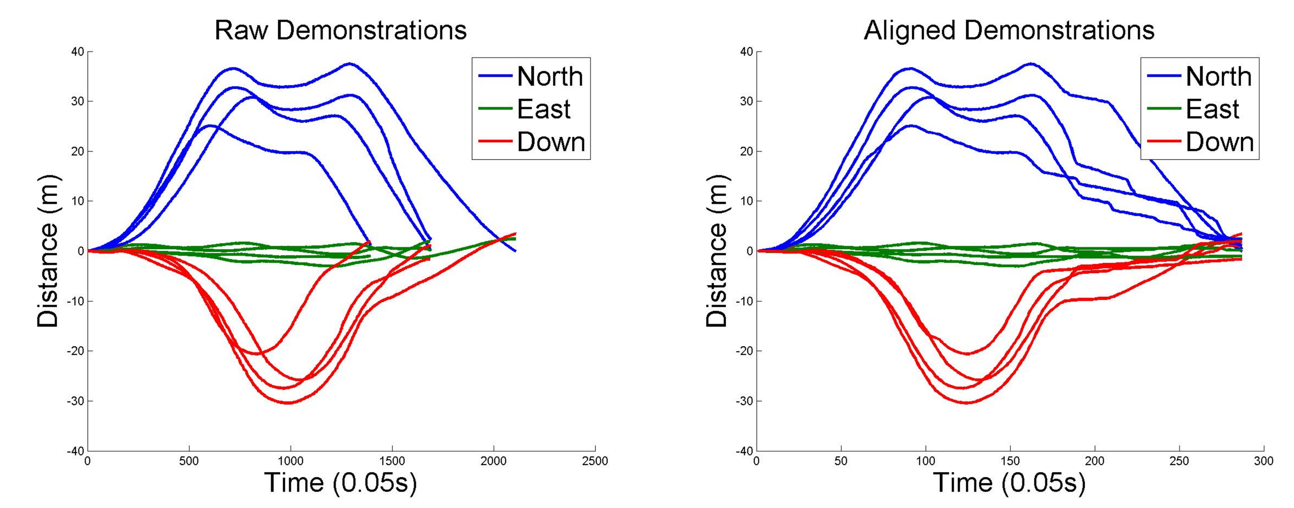

Trajectory Inference

- Infer target trajectory based on demonstrations
  - Structure of graphical model for a given alignment.

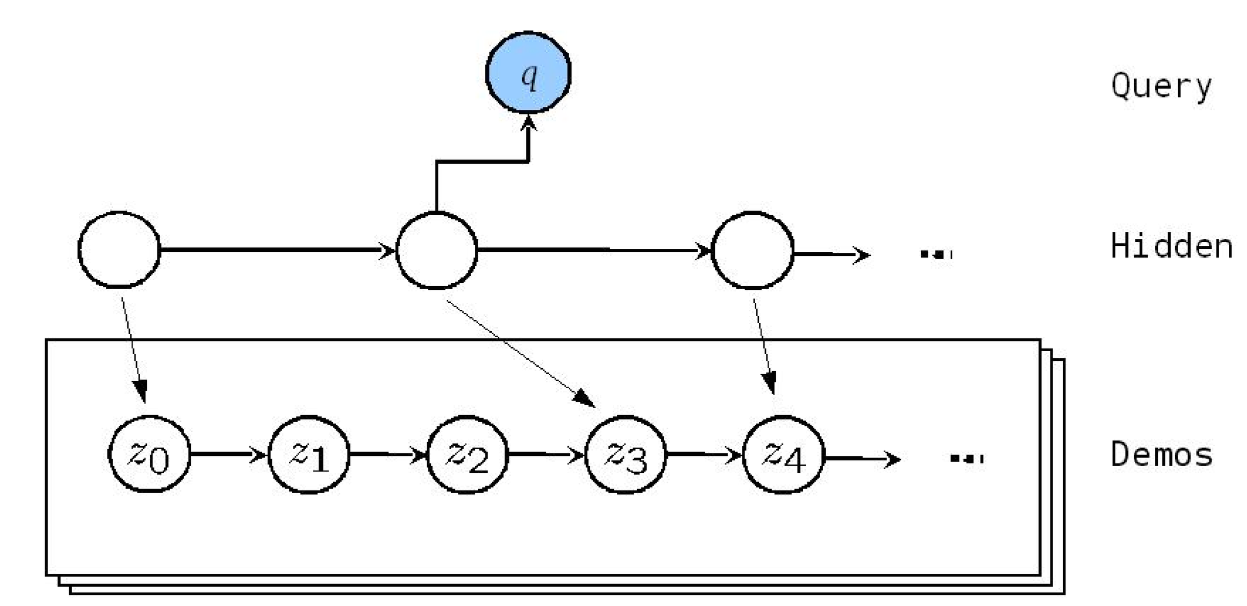

- Demonstrations are noisy "measurements."
- Fake "measurements" for query waypoints
- Run EKF and smoother to compute best posterior estimate for target trajectory.

Results

- Baseline: Interpolate using convex weights which generate the target waypoints. (Algorithm CX)
- Baseline: Run the result of CX through a Kalman filter. (Algorithm SM)
- Baseline: No DTW (Algorithm EM)

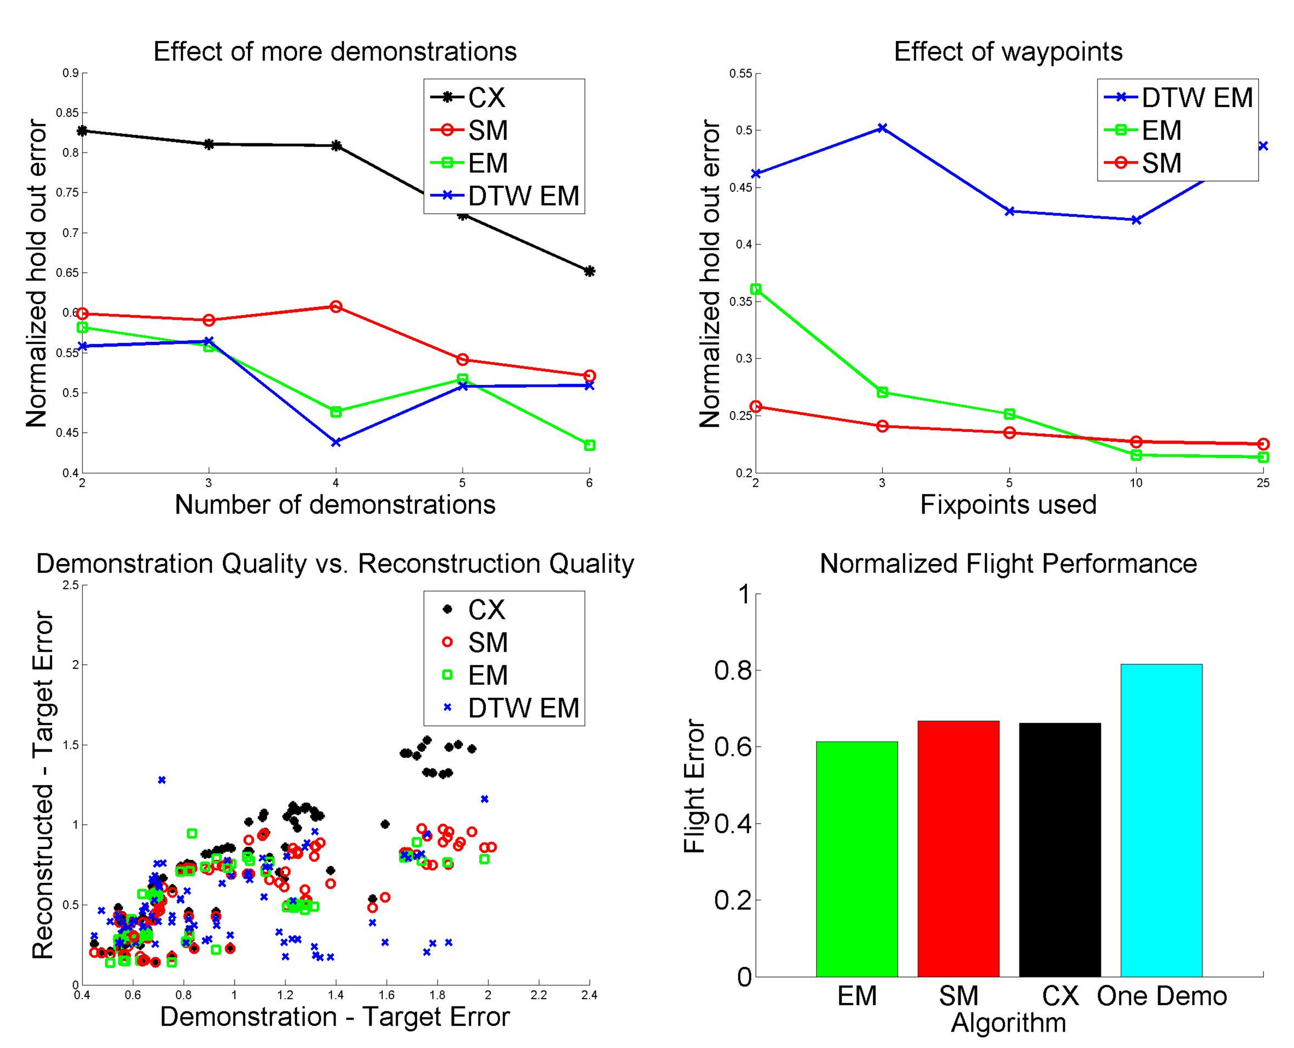

- Flew 3 aggressive maneuvers: stallturns, tictocs, loops.

Dynamics Modeling

- State: position, velocity, orientation, angular velocity.
- Controls: 4 inputs controlling (pitch rate,roll rate,yaw rate,vertical thrust)

$$
\begin{aligned}
\dot{u} &= v \times r - w \times q + g_u + C_u \times [u] \\
\dot{v} &= w \times p - u \times r + g_v + C_v \times [1; v] \\
\dot{w} &= u \times q - v \times p + g_w + C_w \times [1; w; u_4] \\
\dot{p} &= C_p \times [1; p; u_1] \\
\dot{q} &= C_q \times [1; q; u_2] \\
\dot{r} &= C_r \times [1; r; u_3]
\end{aligned}
$$

- Learn model biases for each trajectory by observing deviations during real demonstrations. These biases are remarkably consistent after alignment.

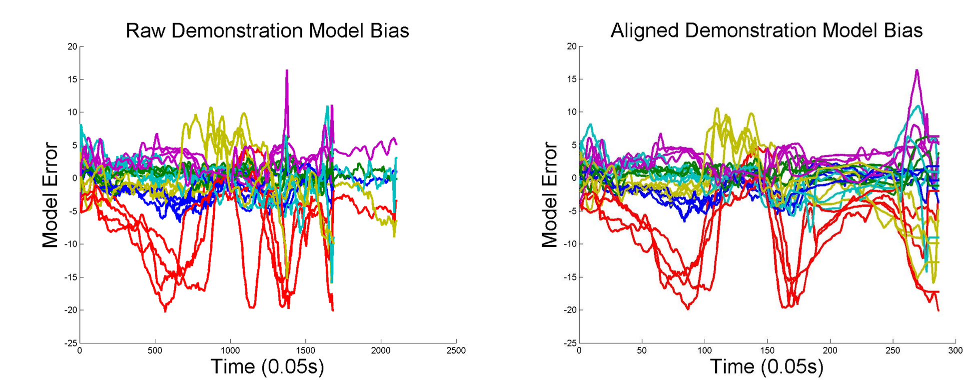

Helicopter Setup

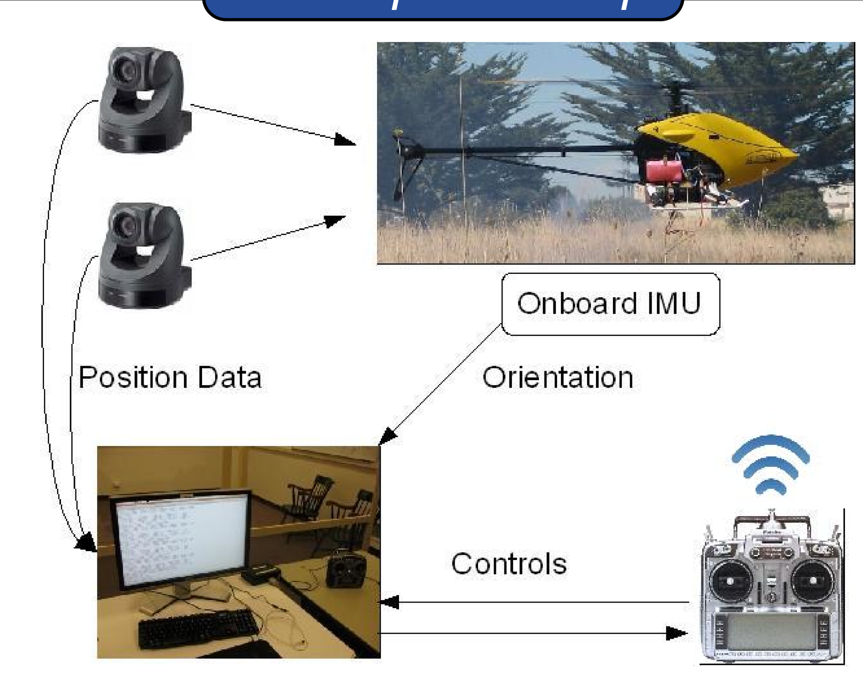

- EKF for state estimation
- LQR for trajectory following

$$\min_{u_0,u_1,\ldots,u_{H-1}} \sum_{t=0}^{H-1} (x_t - x_t^*)^\top Q (x_t - x_t^*) + (u_t - u_t^*)^\top R (u_t - u_t^*)$$

$$\text{s.t. } x_{t+1} - x_{t+1}^* \approx A_t (x_t - x_t^*) + B_t (u_t - u_t^*)$$

  - Linearize dynamics around target to get A_t, B_t
- Receding Horizon iLQR control problem solved online

Conclusions

- This technique allows us to generate parameterized, flyable trajectories for challenging robotic platforms.
- Learn representation for large class of similar trajectories.
- Learn locally tuned dynamics models for control.

References

[1] A. Coates, P. Abbeel, and A. Y. Ng. Learning for control from multiple demonstrations (Full version). *ICML*, pages 144–151, 2008. http://heli.stanford.edu/icml2008.

[2] H. Sakoe and S. Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 1978.