A Machine Learned Model of a Hybrid Aircraft

Kevin Jenkins, Brandon Jones
CS229: Machine Learning Class Project

Summary

Building an aircraft model from flight test data has been part of aircraft development programs for over 100 years [1]. Industry methods rely heavily on domain knowledge of aircraft dynamics for manual data generation and feature identification [2].

- This project aims to simplify the modeling process with a machine learning pipeline that generates a non-linear vehicle model for simulation, control design, and analysis with minimal reliance on domain knowledge.
- A **locally weighted linear regression** [3] approach was investigated to provide full flight envelope modeling of an unconventional Vertical Takeoff and Landing (VTOL) aircraft with lift motors and fixed-wing forward flight capability.

Data

- Data was generated from a high-fidelity aircraft flight simulation throughout the full airspeed envelope, hover to cruise, and sampled at 10 Hz.
- The vehicle flight software was modified to stochastically command actuators throughout the flight envelope to generate diversity in the data to observe response variables.
- Data included long periods of non-maneuvering flight and redundant states.

Vehicle States (x15)	Actuator States (x22)
Airspeed	Vertical and Forward Motor Thrust (x6)
Air Density	Vertical and Forward Motor Torque (x6)
Angle of Attack	Vertical and Forward Motor Speeds (x6)
Attitude (x3)	Aerodynamic Control Surface Deflection (x4)
Rotational Acceleration (x3)	
Body Acceleration (x3)	
Velocity (x3)	

Model

- The aircraft time series data set was can be formulated as the discrete Markov Decision Process,

$$x^{(t+1)} = f(x,u) \cdot x^{(t)} + g(x,u) \cdot u^{(t)}$$

with state $x^{(t)} \in \mathbf{R}^{15\times m}$, actuator inputs $u^{(t)} \in \mathbf{R}^{22\times m}$, and the one-step state response $x^{(t+1)}$. The objective is to learn the non-linear functions $f(x,u)$ and $g(x,u)$ to minimize the prediction error of $x^{(t+1)}$.

- The training examples were formulated as,

$$X \in \mathbf{R}^{37\times m} = \begin{bmatrix} x^{(t)} & \dots & x^{(m-1)} \\ u^{(t)} & \dots & u^{(m-1)} \end{bmatrix} \quad Y \in \mathbf{R}^{4\times m} = \begin{bmatrix} \alpha^{(t+1)} & \dots & \alpha^{(m)} \\ \theta^{(t+1)} & \dots & \theta^{(m)} \\ \dot{\theta}^{(t+1)} & \dots & \dot{\theta}^{(m)} \\ V_a^{(t+1)} & \dots & V_a^{(m)} \end{bmatrix}$$

with input features X and target variables Y consisting of aircraft longitudinal states only, angle of attack, α, pitch, θ, pitch rate, $\dot{\theta}$, and airspeed, V_a.

Algorithm

Selectively Weighted Linear Regression

$$\theta = \underset{\theta}{\arg\min} \left(\sum_i \exp\left(-\frac{\left\| X_w^{(i)} - X_w \right\|^2}{2\tau^2} \right) \left\| Y^{(i)} - \theta^T X_r^{(i)} \right\|^2 + \|\lambda\theta\|_2^2 \right)$$

where X_r and X_w contain selected elements from X,

$$X_w^{(i)} \in R^{n_w} = \left[X_{j_1}^{(i)} \; X_{j_2}^{(i)} \; \dots \; X_{j_{n_w}}^{(i)} \right]^T$$
$$X_r^{(i)} \in R^{n_r} = \left[X_{k_1}^{(i)} \; X_{k_2}^{(i)} \; \dots \; X_{k_{n_r}}^{(i)} \right]^T$$

with a L_2 regularization constant λ and bandwidth parameter τ. The optimization problem was solved using $\theta = (X_r W X_r^T + \lambda)^{-1}(X_r W Y^T)$, with training examples with small $w^{(i)}$ removed.

Results

Four variants are presented all with $\tau = 50$:

- Weighted, $\lambda = 0$
- Weighted, $\lambda = 10^{-3}$
- Weighted, $\lambda = 10^{-3}$, with PCA performed on X to select X_w based on largest principal coefficients.
- Weighted "smart removal", $\lambda = 10^{-3}$, where domain knowledge was used to select a classic subset of specific X_r features that should impact longitudinal motion.

	Training		Test	
	Error	**Samples**	**Error**	**Samples**
Weighted $\lambda = 0$	0.55	15120 (25 minutes)	4.13	5537 (9.23 minutes)
Weighted $\lambda = 10^{-3}$	0.55		4.13	
Weighted Ridge PCA	1.3		4.38	
Weighted "smart removal"	16.89		10.92	

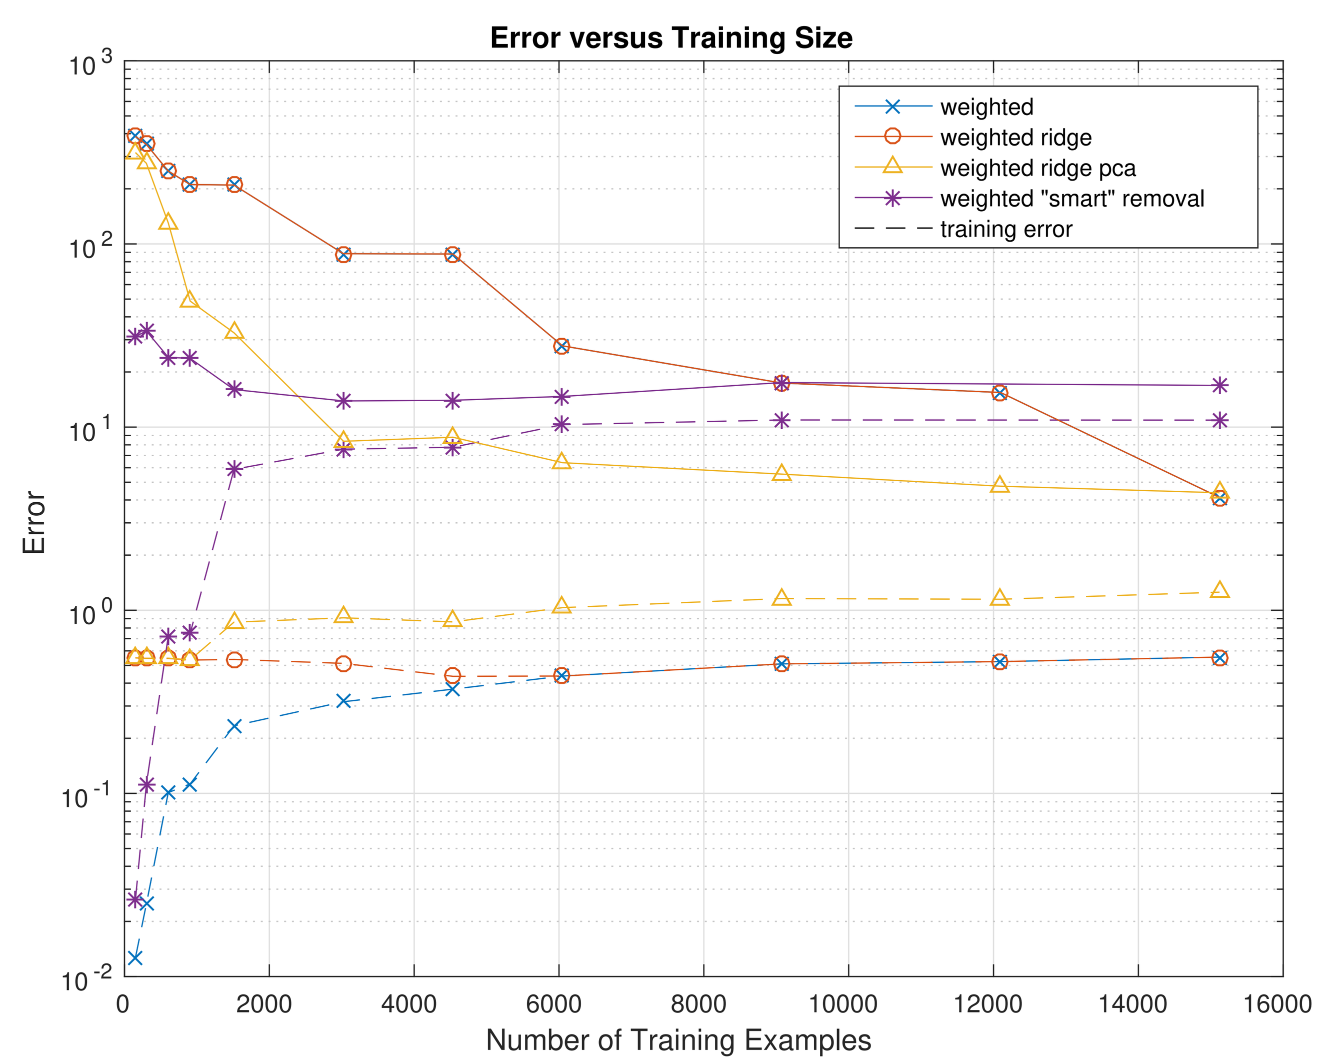

Discussion

- Inspection of time series results showed good time-step to time-step predictions.
- "Smart removal" of states resulted in a model that did not improve with more training examples, as it did not leverage the non-standard dynamics of the hybrid aircraft observed in a more general training set.
- The Weighted PCA model also resulted in less error with less training examples, however with larger training sets did not provide benefit.
- The matrix XWX^T was poorly conditioned due to multicollinearity of samples. L_2 regularization helped with the conditioning of the problem, but did not result in lower overall error.

Future

- Perform analysis on resulting $\theta(x,u)$ to generate classic aerodynamic parameter sets that could be used for physical insight and fast simulation.
- Reformulate Locally Weighted Least Squares with an alternative optimization algorithm to scale to very large datasets (thousands of hours).

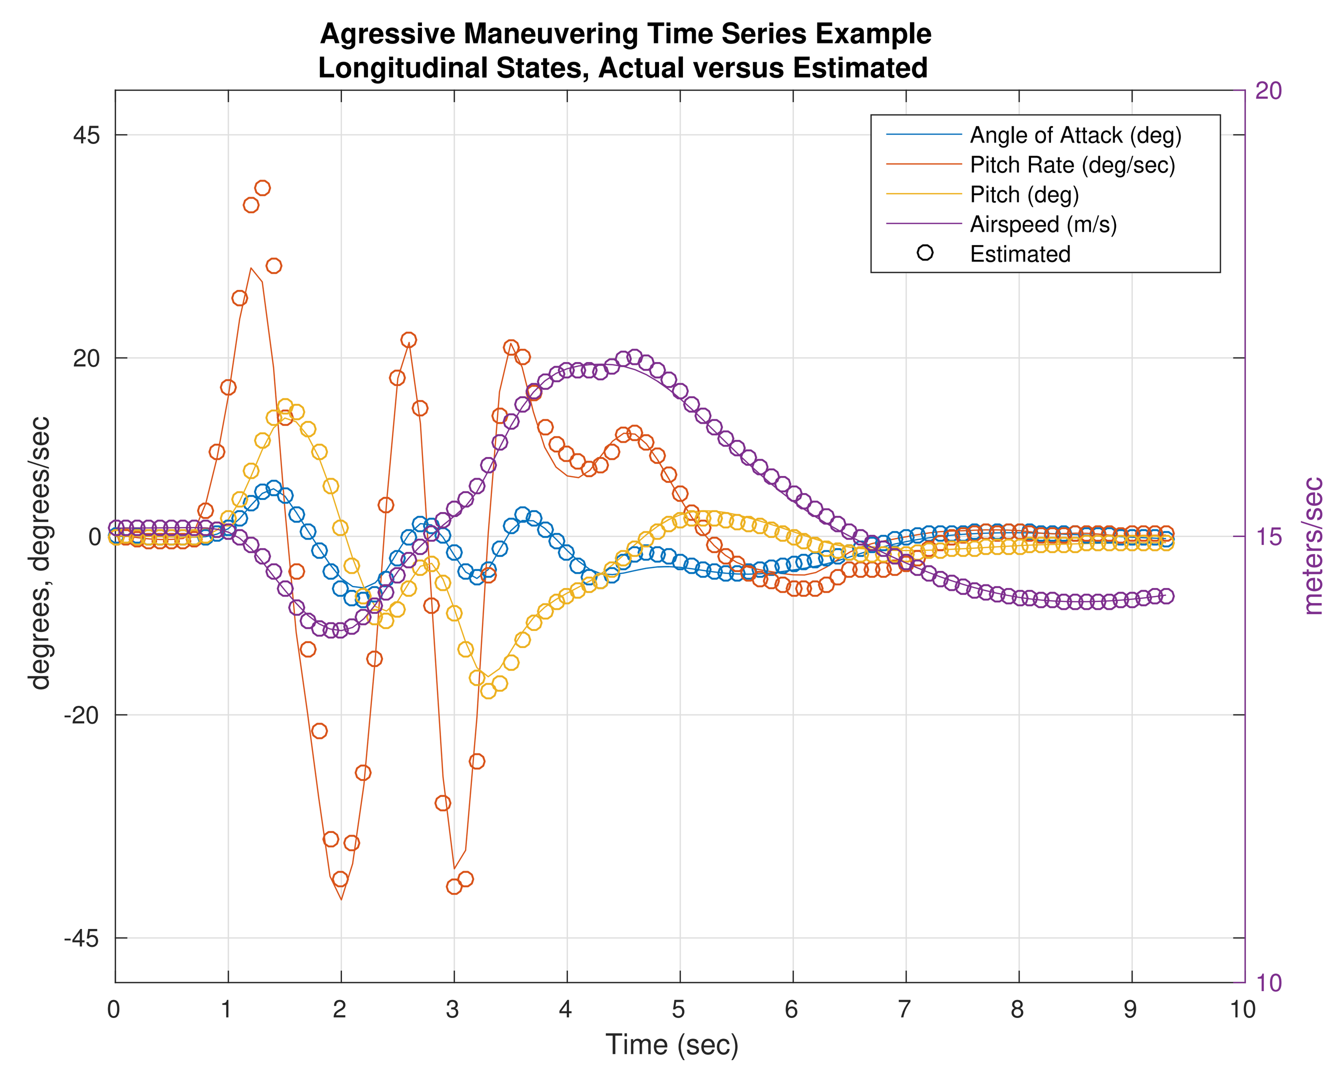

References

[1] E. A. Morelli and V. Klein, "Application of system identification to aircraft at nasa langley research center," *Journal of Aircraft*, vol. 42, no. 1, pp. 12–25, 2005.

[2] V. Klein and E. A. Morelli, *Aircraft system identification: theory and practice*. American Institute of Aeronautics and Astronautics Reston, Va, USA, 2006.

[3] W. S. Cleveland, "Robust locally weighted regression and smoothing scatterplots," *Journal of the American statistical association*, vol. 74, no. 368, pp. 829–836, 1979.