AE6580-Spring 2025 Course Project Adaptive Control For Catastrophic Loss of Effectiveness

1st Keith Gibson
George W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
kgibson43@gatech.edu

1st Connor Bossard
George W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
cbossard3@gatech.edu

I. Introduction

We looked at different methods of adaptive control (Model Reference Adaptive Control (MRAC), Neural Network based Adaptive Control, Sliding Mode Control) for nonlinear flight systems - specifically an F-16 - and how it is able to respond to a catastrophic failure (loss of effectiveness of the elevator).

Aircraft dynamics are highly non-linear [1]. For example, Figure 1 shows the forces that act on an aircraft.

The force equations can thus be written as

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} + mg_0 \begin{pmatrix} -\sin\Theta \\ \cos\Theta\sin\Phi \\ \cos\Theta\cos\Phi \end{pmatrix} = m \begin{pmatrix} \dot{u} + qw - rv \\ \dot{v} + ru - pw \\ \dot{w} + pv - qu \end{pmatrix} \tag{4.12}$$

Fig. 1. Forces acting on an aircraft

We used the F-16 model cited below [2], with the initial conditions as follows:

TABLE I
Initial Conditions (IC) State Variables

h	10,000 ft
θ	0 degrees
v	310 ft/sec
α	0 degrees
q	0 deg/sec
δ_t	0
δ_e	0

We also used the following state for the equilibrium and trim state:

TABLE II
Trim / Equilibrium State Variables

h	10,000 ft
θ	0 degrees
v	300 ft/sec
α	0 degrees
q	0 deg/sec
δ_t	0
δ_e	0

A. Related Work

Adaptive controllers in aircraft are not a new phenomenon, with an early example being an experimental controller on NASA's X-15 in 1967 [3]. Generally, the advantage of these types of controllers is that they are able to adapt to uncertainty in the model, at the expense of extra computation power. There are a wide range of adaptive controllers, ranging in complexity from simply choosing gain and switching between them, to fuzzy logic based neural networks (NN) [4]. Since 1967, advancements have been made in the stability of adaptive control systems, as well as in the development and stability analysis of NN based adaptive control systems [5] [6]. Due to the wide breadth of available controllers, and the developments of stability in many of these systems, recent literature has looked into using these control systems for unexpected damage to aircraft [7], with this paper also providing one of the possible failure modes.

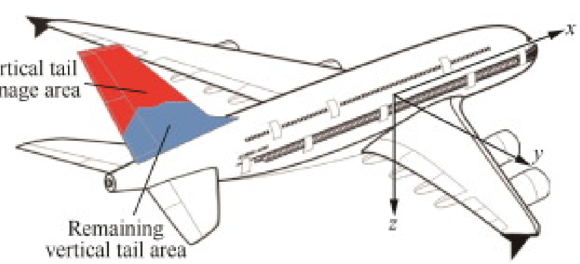

Fig. 2. Airplane vertical tail failure [7]

An interesting topic also evolves from this discussion of adaptive controls, namely whether to use model-free or model-based adaptive controls [8]. Both methods will be explored, however the focus will mainly be on the model-based architectures, mainly due to the case of interest being an aircraft which have relatively known nominal models. While recent literature has explored learning the dynamics online with predictive cost methods, the learning of the damaged aircraft dynamics will not be a main focus of this project [9]. Another interesting extension of adaptive control is robustness to malicious attacks on the controller themselves, with this presenting a possible

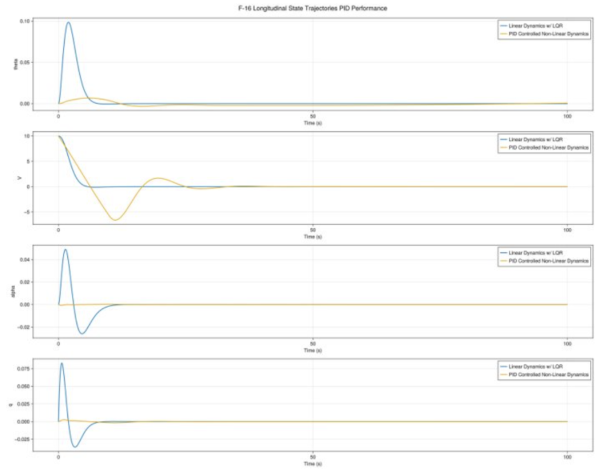

Fig. 3. PID Performance Without Disturbance

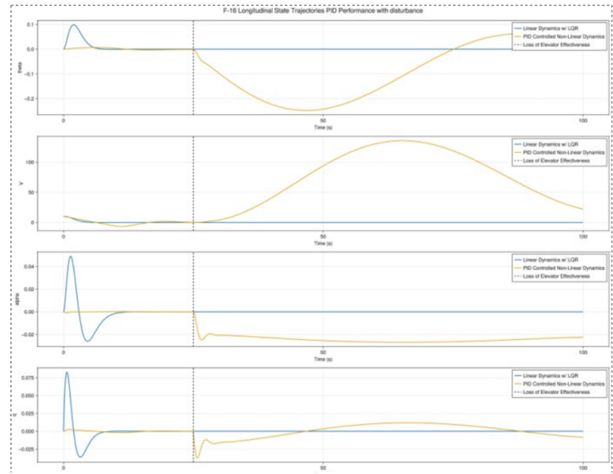

Fig. 4. PID Performance With loss of elevator effectiveness

catastrophic failure [10]. This method of failure may be considered in a very basic sense, such as unknown input changes, but will not be a main consideration for much of the work.

The authors would like to compare modern adaptive control methods over multiple levels of severity of changed dynamics, represented by these failure modes and investigate the strengths and weaknesses of each control method.

B. Controllers

1) PID: First, we looked at PID and LQR controllers as non-adaptive control algorithms as a baseline. We were able to tune the gains on the PID controller to get a good response to the non-linear dynamics without the loss of effectiveness. The LQR response on the linearized system was added to the figures for the different controller responses to act as a baseline and ideal response across, allowing easy comparison across controllers.

However, when we added the loss of effectiveness of the elevator at time t = 25 seconds ($\delta_e = \delta_e/2$), the controller was not able to respond to the loss of effectiveness and led to large oscillations and offsets, as shown in Figure 4.

2) LQR: Next, we calculated the LQR gains from the linearized model around the trim condition using the equations below:

$$J(u) = \int_0^\infty (x^T Q x + u^T R u)\, dt \quad (1)$$

$$\text{such that } \dot{x} = Ax + Bu \quad (2)$$

$$A^T P + PA - PBR^{-1}B^T P + Q = 0 \quad (3)$$

$$K = R^{-1}(B^T P + N^T) \quad (4)$$

When applying the control input u=Kx on the non-linear system, we see an offset in both the theta and velocity. Which is expected being that the gains were determined for a simplified dynamics model.

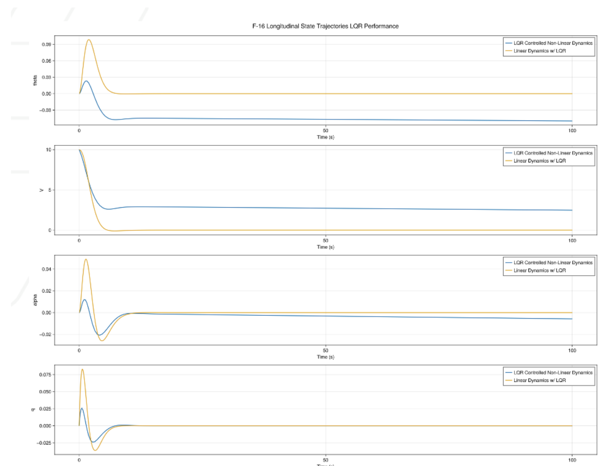

Fig. 5. LQR Performance Without Disturbance

We also saw offsets in both theta and velocity, however they were even larger after the loss of effectiveness of the elevator.

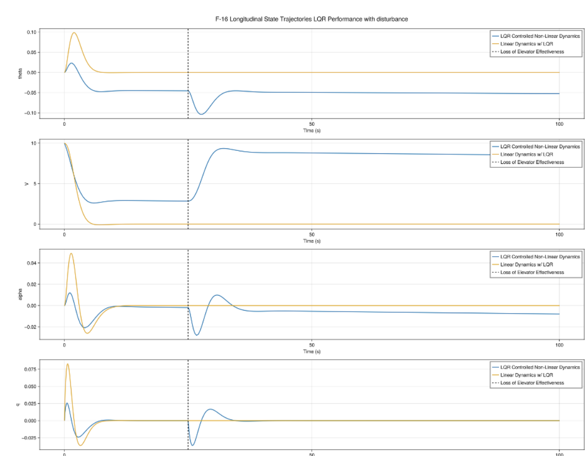

Fig. 6. LQR Performance With loss of elevator effectiveness

3) LQR-I: To handle the offsets above, we implemented LQR-I, adding an integral term to try and mitigate these offsets. It was implemented using the following equations:

$$A_I = \begin{bmatrix} A & 0_{n\times(n-p)} \\ I_{p\times p} & 0_{p\times(n-p)} \end{bmatrix} \qquad B_I = \begin{bmatrix} B \\ 0_{p\times m} \end{bmatrix}$$

$$x_I = \begin{bmatrix} x \\ \int x\,dt \end{bmatrix}$$

$$\text{minimize}$$
$$J(u) = \int_0^\infty x_I^T Q_I x_I + u^T R u\,dt$$
$$\textit{such that:}$$
$$\dot{x}_I = A_I x + B_I u$$
$$A^T P_I + P A_I - P B_I R^{-1} B_I^T P + Q_I = 0$$
$$K = R^{-1}(B_I^T P)$$

Fig. 7. LQR-I Equations

This helped with the offset in the pitch angle, but was still prevelant in the velocity state. ADD PERFORMANCE OF LQR-I IF WANTED / SPACE

4) MRAC: As the first adaptive control method, we implemented MRAC (Model Reference Adaptive Control). The block diagram of the algorithm is included below.

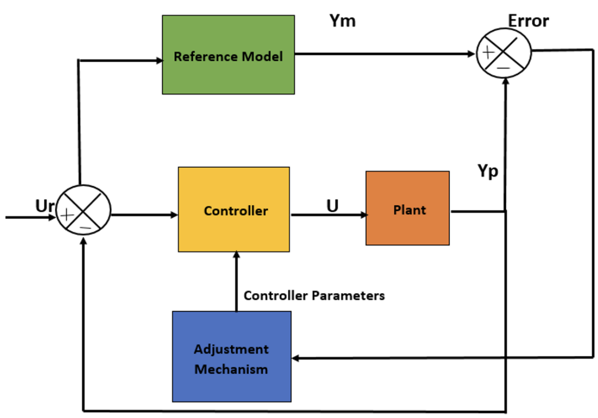

Fig. 8. MRAC Block Diagram [11]

For our implementation, we used the linearized model around the trim state as the reference model, and the control inputs for the LQR linearized states as the reference. The plant was the Non-linear model with loss of elevator effectiveness at 25 seconds. The goal of MRAC is to find a mapping of X and r (the LQR control input) from the reference model to a U that has a similar response for the non-linear, loss of effectiveness dynamics model. This is done by learning weights for the gains of specific states, control inputs and basis functions using the equations below, where Γ is the learning rate for each gain. This implementation was adapted from the Robust and Adaptive Control with Aerospace Applications textbook [12].

$$u = K_x^T x + K_r^T r - \Theta^T \Phi(x) \tag{5}$$
$$\dot{K}_x = -\Gamma_x x e^T P B \tag{6}$$
$$\dot{K}_r = -\Gamma_r r(t) e^T P B \tag{7}$$
$$\dot{\Theta} = \Gamma_\Theta \Phi(x) e^T P B \tag{8}$$

Below is the response of the MRAC controller on the non-linear dynamics with the elevator loss of effectiveness.

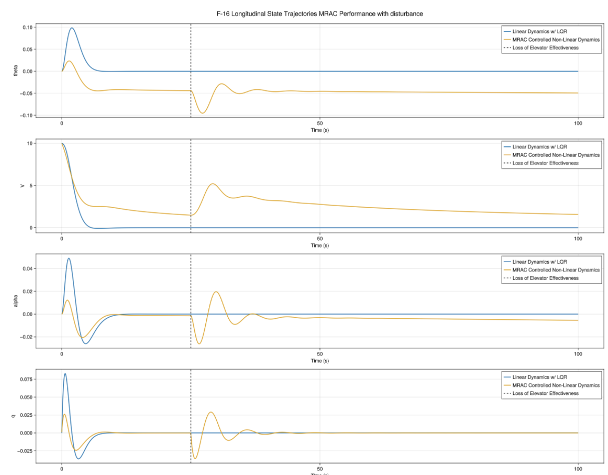

Fig. 9. MRAC Performance With loss of elevator effectiveness

We see that MRAC has a much better response to the disturbance than the non-adaptive methods implemented above. Specifically, we can see that the velocity state is not settling to an offset and rather driving towards the equilibrium state.

5) NN Based Adaptive Control: The next adaptive control algorithm we implemented was a Neural Network based adaptive control method. This is very similar to the MRAC implementation, but instead of learning weights for predefined basis functions, it is learning the weights of a shallow neural network. The architecture we used is shown below.

for this diagram, we are learning the weights $\hat{V}$ and $\hat{W}$. The output V is then being multiplied by the state to get the control input. Due to the instability of the F-16 and sensitivity to perturbations to the control sequence, we adapted the algorithm to have the output from the network to be small and added these to the control inputs using the linearized LQR-I gains. These updated equations are shown in Figure 11.

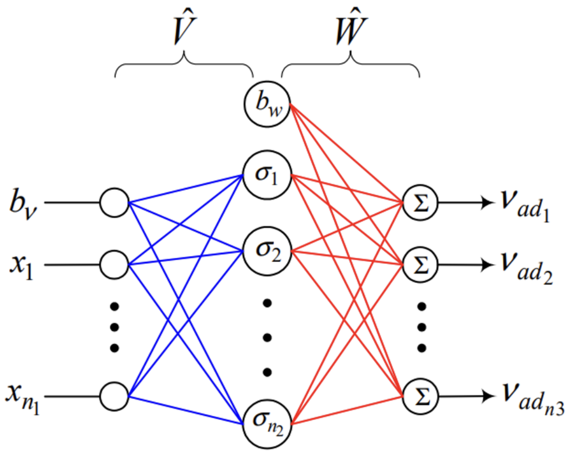

Fig. 10. Shallow NN Diagram [13]

Adaptive Control

$$z = Vx + b_v$$
$$\sigma(z) = \tanh(z)$$
$$u_{ad} = W\sigma(z) + b_w$$
$$u = u_{LQR-I} + u_{ad}$$

Parameter Update

$$\dot{x}_r = A_r x_r + B_r u_r$$
$$e = x - x_r$$
$$\dot{\hat{W}} = -\Gamma_w (B^T P e) z^T$$
$$\dot{b}_w = -\Gamma_{b_w} (B^T P e)$$
$$\dot{\hat{V}} = -\Gamma_v \left(W^T (B^T P e) \frac{d\sigma}{dx} \right) x^T$$
$$\dot{b}_v = -\Gamma_{b_v} \left(W^T (B^T P e) \frac{d\sigma}{dx} \right)$$

Fig. 11. NN Adaptive Control Equations [13]

The response of the NN Adaptive control on the non-linear dynamics without a disturbance shows that the performance exceeds that of LQR, and LQR-I without the adaptive portion from the Neural Network.

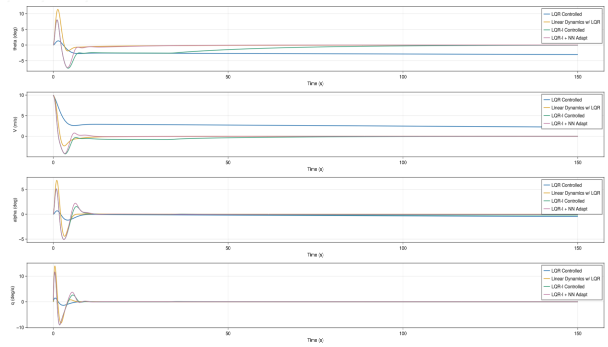

Fig. 12. NN Adaptive Control Response Without Disturbance

Furthermore, we see that the NN controller has improved performance compared to the LQR and LQR-I controllers, when there is a loss of effectiveness of the elevator. In particular, it is able to return to the equilibrium state quicker.

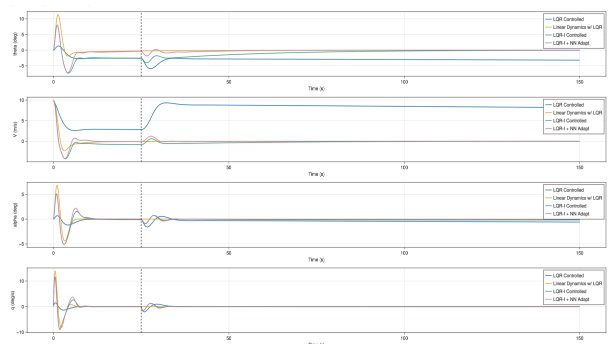

Fig. 13. NN Adaptive Control Response with loss of elevator effectiveness

6) Sliding Mode Control: Sliding Mode Control (SMC) is a nonlinear controller that drives the system to a pre-defined sliding surface, and then along that surface, to the equilibrium. This can be seen in Figure 14, which shows the phase portrait of the sliding mode controller. The flow of the states will be towards the pink line, and then along the pink line towards the equilibrium. This is similar to how one may expect a bang-bang controller to respond, and staying with that parallel, the sliding mode controller often suffers from the same chatter seen in bang-bang. Ultimately, sliding mode control is used due to its simplicity, with the two overarching governing equations shown below. SMC has been shown previously to be robust for highly unstable aircraft [15].

$$\dot{s} = -k|s|^a * sgn(s) = h(s(x)) \tag{9}$$
$$u = (C^T g(x))^\dagger (-C^T f(x) + h(s(x))) \tag{10}$$

This controller method is guaranteed to be robust, and the reaching law shown above is only one of many possible. ”C”, ”k”, and ”a” are tunable parameters used to control the response of the controller, however for ”C”, the LQR gain matrix ”K” was used for simplicity. This controller guarantees that the controller will converge to the sliding mode, and as such relies on the sliding mode being stable.

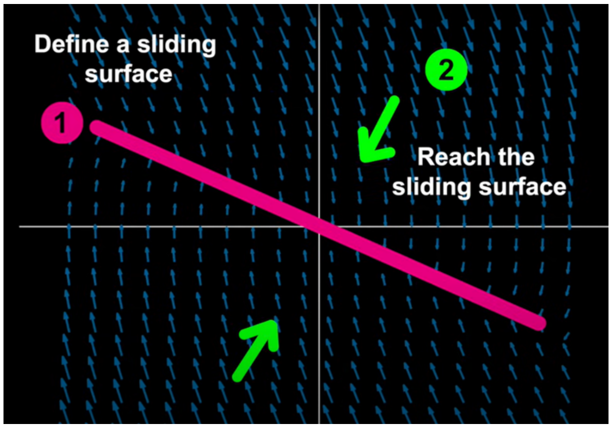

Fig. 14. Sliding Mode Control Overview [14]

C. Performance Comparison

To accurately compare the performance of the systems, as well as to try to drive the system to stay within the simulation bounds, all of the adaptive components are applied on top of LQR or LQR-I. The ”Linear Dynamics w/ LQR” represents the best possible response (with respect to our LQR weighting matrices), a forward simulation of the LQR controller on the linear dynamics with no loss of effectiveness. This helps give a good baseline of what perturbations from 0 error are to be expect at the beginning of the run.

The figure 15 shows the comparison of the adaptive elements when using a baseline LQR controller.

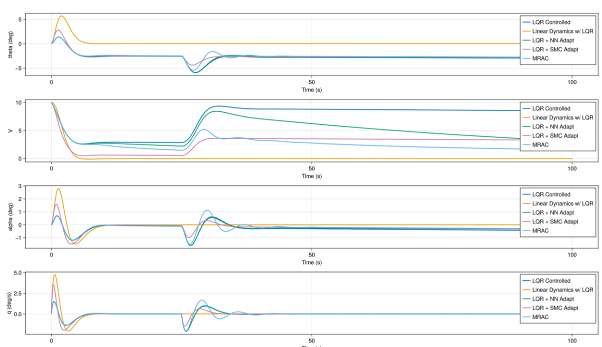

Fig. 15. Adaptive Controller Comparison with 50% Loss of Elevator Effectiveness and LQR Base Controller.

As expected, most of the controllers have a constant steady state error, which is highly undesirable. This error also grows when the 50% loss of effectiveness occurs for all but the NN and MRAC based controller, which slowly adapt the steady state velocity. To combat this, adding integral terms to the LQR and SMC controllers allows for minimization of steady state error.

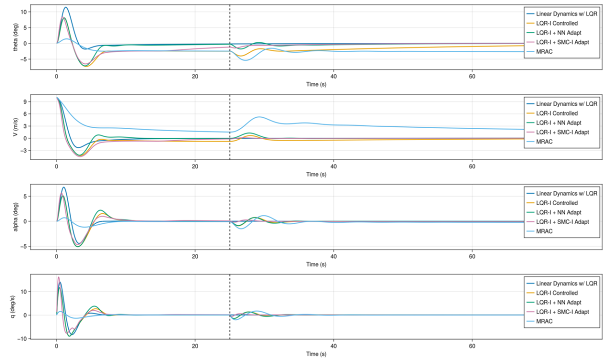

Fig. 16. Adaptive Controller Comparison with 50% Loss of Elevator Effectiveness and LQR-I Base Controller.

As seen in the above figure, the steady state errors of both LQR and SMC are virtually eliminated. It can also be seen that while the NN controller seems to converge the quickest, the sliding mode with integral seems to be the most resistant to the instantaneous loss of effectiveness. MRAC still has an offset in all but the velocity because it has no integral term to drive the smaller errors to 0.

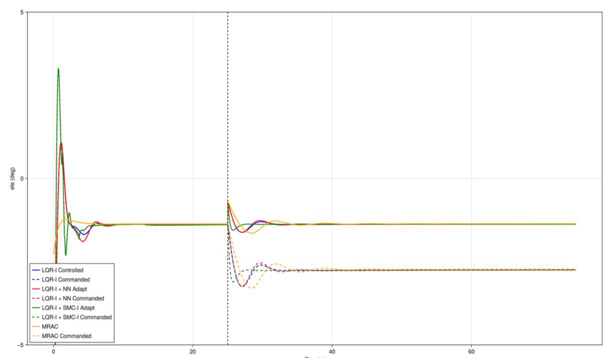

Fig. 17. Elevator Adaptive Controller Responses with 50% Loss of Elevator Effectiveness. The dashed lines represent the commanded deflection.

Looking at the controller response, it is clear that each of the adaptive elements are able to adapt to the loss of effectiveness, and change their commanded deflections to negate the effectiveness loss. All of the controllers are able to get the true deflection close to their pre-fault steady state values. The SMC-I seems to respond the quickest, which matches the earlier assertion that it is the most resilient to the loss of effectiveness.

Looking into a larger loss of effectiveness, the same test is ran with a 90% loss of effectiveness, a loss effectiveness that could often be catastrophic if not accounted for.

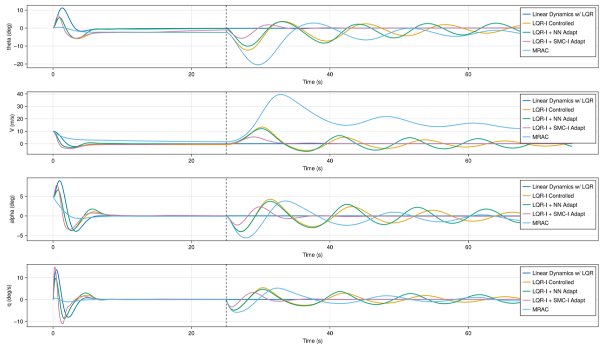

Fig. 18. Adaptive Controller Comparison with 90% Loss of Elevator Effectiveness and LQR-I Base Controller.

As seen in figure 18, the large loss of effectiveness causes oscillations in all of the controllers except LQR-I a + SMC-I. This may be due to the learning rates of MRAC and NN based controllers causing the weights to overshoot the baseline, leading to oscillations, as well as the integral anti-windup term preventing stabilization of the LQR-I controller. The anti-windup is required to keep errors from growing unreasonably and causing commands that drive the F-16 into unstable or unrecoverable regimes.

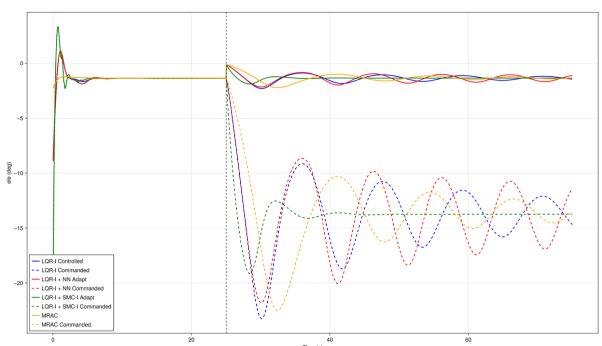

Fig. 19. Elevator Adaptive Controller Responses with 90% Loss of Elevator Effectiveness. The dashed lines represent the commanded deflection.

Interestingly, it seems that all of the controllers tend to respond in a second order fashion, however, the damping is directly related to the loss of effectiveness seen. Again, as in the 90% case, all of the controllers respond by immediately commanding larger deflections than the true response, but the SMC-I seems to have the best damping in both cases. This is likely due to the exponential reaching law used, which was chosen to reduce chatter and smoothly drive towards the sliding manifold. The other controllers can also be shown to be damped, with each subsequent peak being smaller than the last, but they do not stabilize as fast. This may be undesirable, as jets such as the F-16 are designed to be able to be highly maneuverable, so these oscillations may prevent that. Thus, while all of the controllers begin to damped out, it may be desirable for a fighter to take the controller that adapts the fastest.

Finally, it important to note that while all of the controllers do adapt to disturbance, it seems that the NN controller actually dampens the slowest of the bunch. This is likely due to the extremely small learning rates that had to be applied to prevent the controller from going unstable while learning. This was a simulation restriction, as there was limited aerodynamic data. However, when comparing the MRAC to the NN, the big difference is that the NN does not require tailored nonlinear terms. For an aircraft the governing nonlinear dynamics are relatively well known, with most of the uncertainty coming from external forces (such as wind) or unmodeled aerodynamics (such as turbulence). Thus the generalized nature of the NN adaptive control may not be as desirable, if this can be traded for more robust guarantees.

D. *Challenges*

One major challenges we faced was the fact that the F-16 is inherently unstable, leading to the divergence of dynamics for small changes to the controls. This made tuning the controllers (particularly the learning rates of the MRAC and NN based controllers) very difficult, as small changes would lead to instability. Another challenge we faced is that the aero table used were bounded, leading us to getting out of interpolation ranges for angle of attack values causing the simulation to crash. Finally, a more generalized challenge is that a loss of effectiveness may often lead to a constant steady state error, that is undesirable in the system. To combat this, generally some sort of integral term may be needed to drive the steady state error to 0. However, with this comes two challenges. The first being a way to prevent integral windup, which could drive us unstable. This provides an extra set of tuning parameters. The second is that the integral terms generally expand the state spaces, which caused problems when attempting to expand the MRAC and NN based controllers for the integral term. While we were able to add the integral term in LQR and SMC, the MRAC and NN based adaptive controllers need more careful consideration due to having to solve the Lyapunov eqauation.

E. *Next Steps*

Our next steps include adding the integral term to LQR for the lateral dynamics, as we saw a similar response with an offset in states in the lateral dynamics that were fixed in the longitudinal dynamics by implementing the integral term for LQR gains. We also want to incorporate the integral terms to the neural network and MRAC controllers, as we had good results when adding the integral term to the Sliding Mode Controller. Finally, we want to put our controller in a higher fidelity simulation (like FlightGear) and see if the results carry over from our dynamics model to a real simulation, with uncertainty from dynamic noise (such as wind gusts) and sensor uncertainty. While the focus of this was stabilization around a trim state, switching between trims (such as climb into steady level flight) could also change the conclusions as to which adaptive element may be best. Adding knowledge of control bounds and state bounds may also allow for safer adaptation, which is the ultimate prioirty in these systems.

F. Conclusion

Overall, throughout this semester, we have implemented 7 controller algorithms and compared their performances on a highly non-linear F-16 Model that underwent significant disturbances to it's elevator control. We compared their performances and concluded that the Sliding Mode Controller with Integral Action performed best, but we believe that MRAC and the Neural Network Adaptive Controllers could perform aswell if not better, when we integrate the integral action into those controllers. Generally, all of the controllers were able to adapt to the loss of effectiveness, even with 90% loss by commanding larger deflections, without any knowledge of the loss. This shows promise for keeping the F-16 flying despite critical loss of effectiveness that may otherwise be catastrophic, but much more testing would be needed for real-world implementation.

References

[1] Dynamical Equations for Flight Vehicles. (n.d.). https://courses.cit.cornell.edu/mae5070/DynamicEquations.pdf

[2] Intelligent Systems Research Lab, Julia F-16 Model. https://github.com/isrlab/F16Model.jl

[3] Z. T. Dydek, A. M. Annaswamy and E. Lavretsky, "Adaptive Control and the NASA X-15-3 Flight Revisited," in IEEE Control Systems Magazine, vol. 30, no. 3, pp. 32-48, June 2010, doi: 10.1109/MCS.2010.936292.

[4] Riccardo Marino, Adaptive control of nonlinear systems: Basic results and applications, Annual Reviews in Control, 1997, https://doi.org/10.1016/S1367-5788(97)00033-3.

[5] Anderson, Brian D. O. and Bitmead, Robert R. and Johnson, C. Richard and Kokotovic, Petar V. and Kosut, Robert L. and Mareels, Iven M.Y. and Praly, Laurent and Riedle, Bradley D., Stability of adaptive systems: passivity and averaging analysis, 1986, MIT Press, Cambridge, MA, USA

[6] Anran Li, John P. Swensen, Mehdi Hosseinzadeh, Provably-stable neural network based control of nonlinear systems, Engineering Applications of Artificial Intelligence, 2024, https://doi.org/10.1016/j.engappai.2024.109252.

[7] Jun Wang, Shaoping Wang, Xingjian Wang, Cun Shi, Mileta M. Tomovic, Active fault tolerant control for vertical tail damaged aircraft with dissimilar redundant actuation system, Chinese Journal of Aeronautics, 2016, https://doi.org/10.1016/j.cja.2016.08.009.

[8] Zhao M, Wang D, Li M, Gao N, Qiao J. A new Q-function structure for model-free adaptive optimal tracking control with asymmetric constrained inputs. Int J Adapt Control Signal Process. 2024; 38(5): 1561-1578. doi: 10.1002/acs.3761

[9] Riley Richards, Juan Paredes and Dennis Bernstein. "Predictive Cost Adaptive Control of Fixed-Wing Aircraft Without Prior Modeling," AIAA 2025-2081. AIAA SCITECH 2025 Forum. January 2025

[10] Z. Han, Y. Zhou, Y. Jiang, K. Bao and W. Yue, "Distributed Adaptive Consensus Control for Nonlinear Network Systems With Event-Based Switching Mechanism Against Malicious Attacks," in IEEE Transactions on Automation Science and Engineering, doi: 10.1109/TASE.2024.3484512.

[11] General Block Diagram of MRAC. — Download Scientific Diagram, www.researchgate.net/figure/General-block-diagram-of-MRACfig3331459937. Accessed 22 Apr. 2025.

[12] Robust and Adaptive Control with Aerospace Applications. Springer London, 2013.

[13] Shin, Yoonghyun. Neural Network Based Adaptive Control for Nonlinear Dynamic Regimes. 2005.

[14] MATLAB. "What Is Sliding Mode Control?" YouTube, 26 Sept. 2024, www.youtube.com/watch?v=RD-2oiwEbDo. Accessed 11 Dec. 2024.

[15] Vetter, Travis. Sliding-Mode Control Applied for Robust Control of a Highly Unstable Aircraft BY.