Multi-Vehicle and Assured Autonomous Control for Aerospace Applications
IEEE TCAC Workshop, 2021 CCTA,
August 8th 2021

Towards Trustworthy Autonomy:

How AI can help address fundamental learning and adaptation challenges?

Gokhan Inalhan
BAE Systems Chair
Professor of Autonomous Systems and Artificial Intelligence

www.cranfield.ac.uk

Cranfield Aerospace

Centre for Autonomous and Cyber-Physical Systems

Autonomy & AI Research Theme

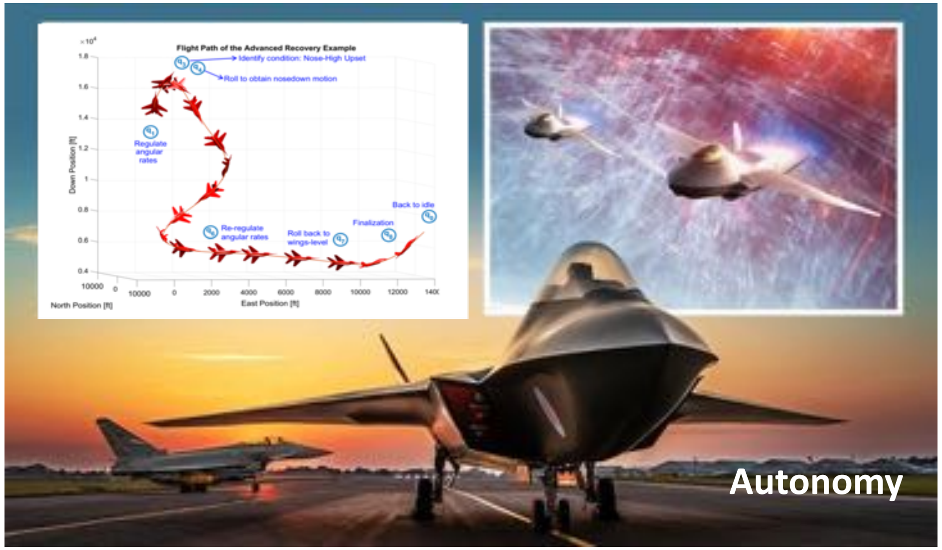

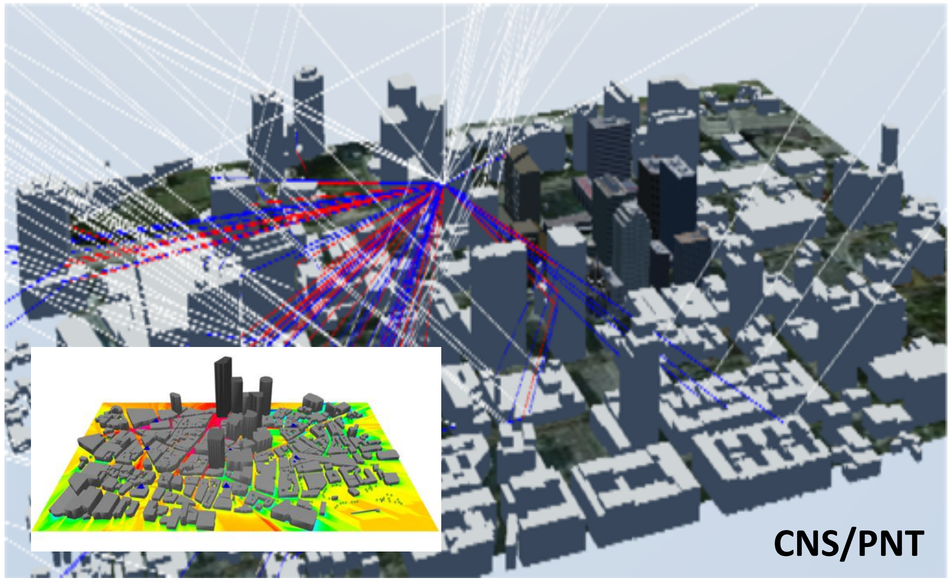

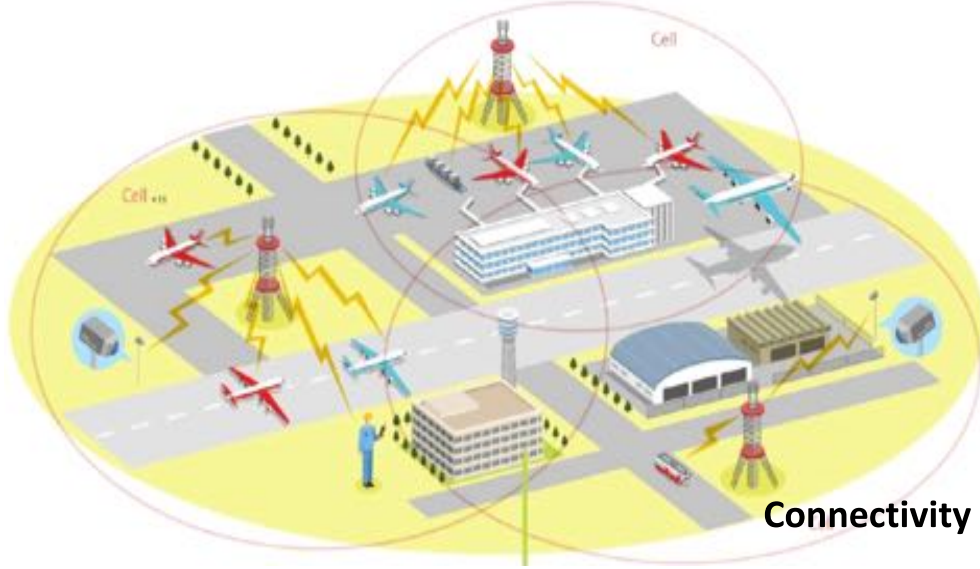

EPSRC Trustworthy Autonomous Systems Research Nodes

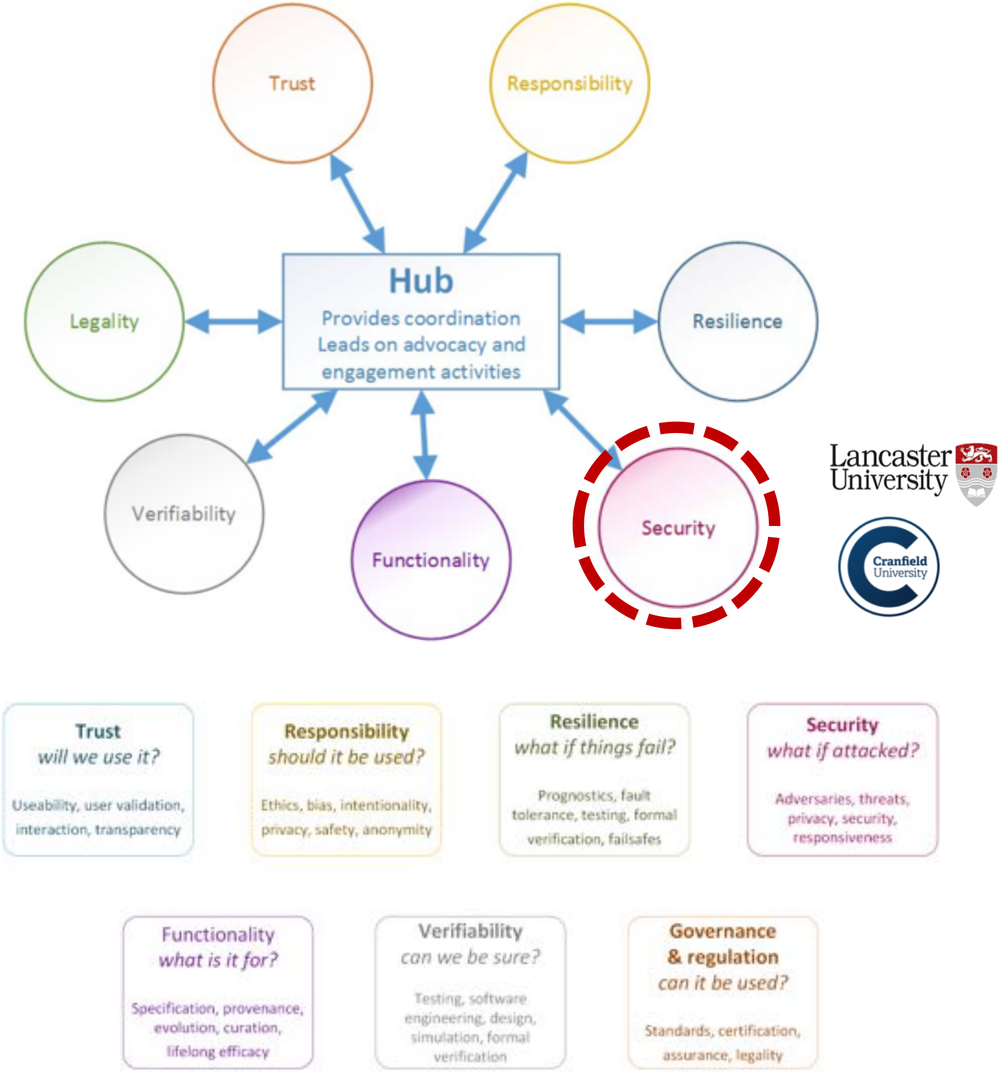

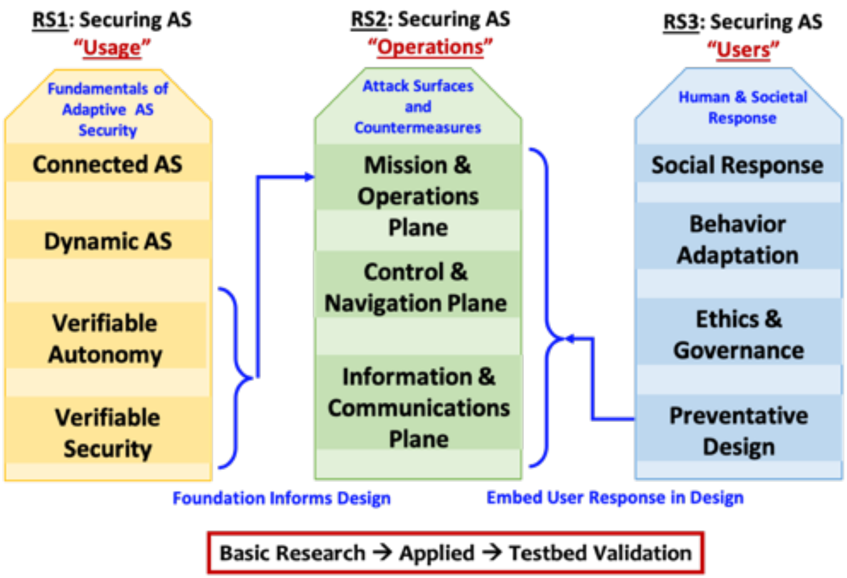

Trustworthy Autonomous Systems(TAS) Node on Security : The Control Challenge

- Autonomous Systems rely on the ability to conduct **run time adaptations of control decisions** over attacks or “perceived” attacks:
  - Adversaries
    - Physical
    - Information-plane
  - Information and dynamic environment uncertainties
  - Degraded performance
    - CNS and Infrastructure
    - Actuators
- How to do this in a “**trustworthy**” fashion in a “**learning-enabled context**”?
  - Safe
  - Secure
  - Reliable

Evolution of Attacks or “Perceived” attacks

- Sensing and COMM errors
- Loss of an actuator
- Environmental conditions
  - Wind
- Electronic Attacks
  - Jamming
  - Spoofing
- Electromagnetic deception
  - false/duplicate target generation

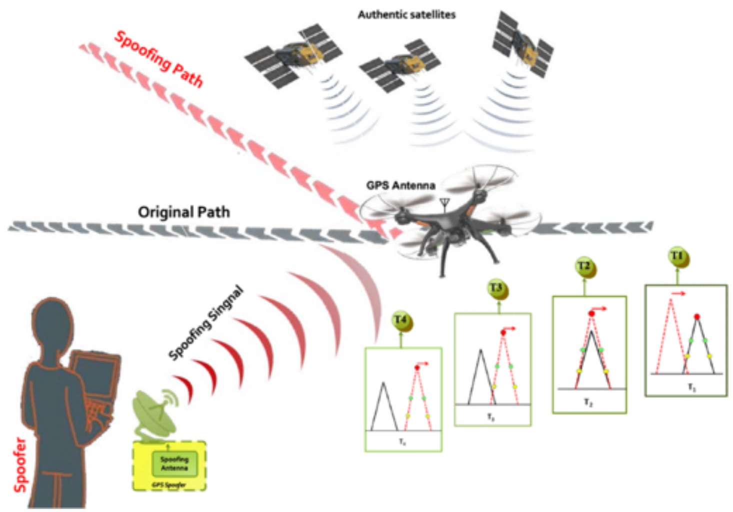

- Generative Adversarial Networks
  - DNN perception and classification
- Injecting false patterns into data

Key cornerstones in AI-Driven Design

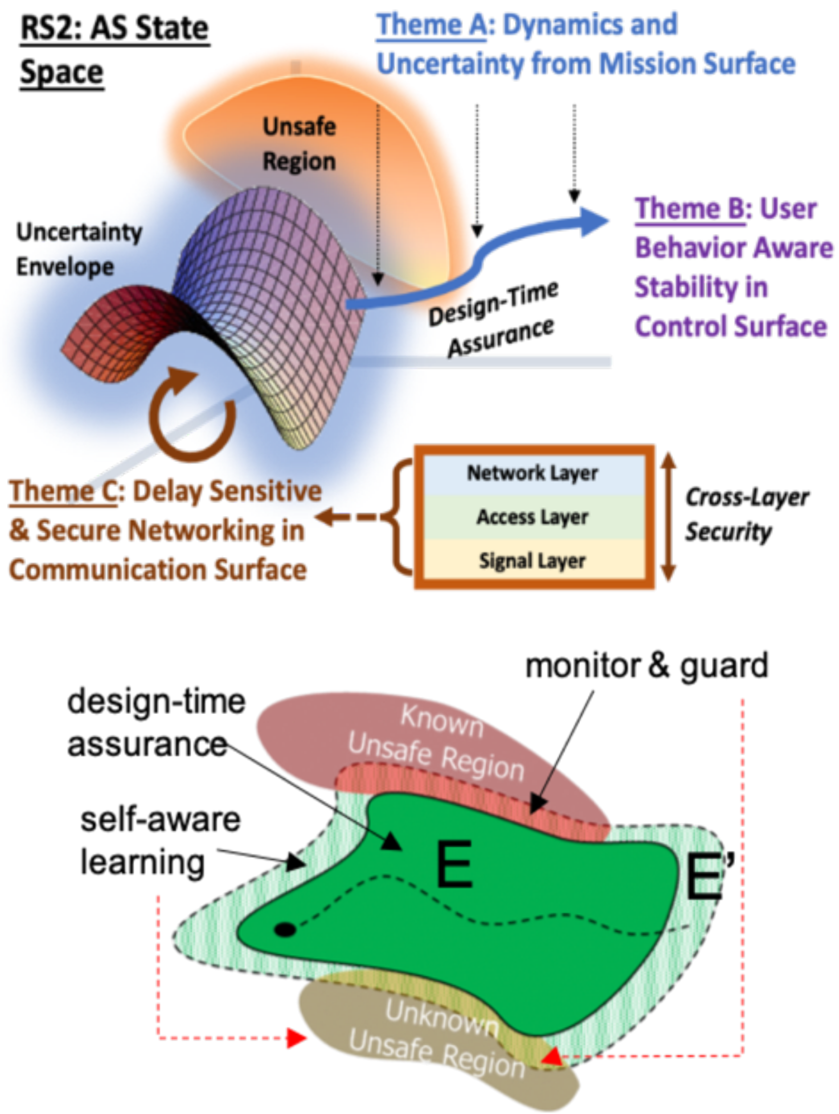

- Provide **quantifiable safety and feedback** to the mission surface when the limits of secure controllability are compromised within a time horizon under current policies and adversarial situations.
- Key Solution Cornerstones in Learning-Enabled Context
  - **Interpretability** => Explainable and Trustworthy AI
  - **Continual Assurance** => Dynamic Verification & Validation
  - **Adaptive Security Strategies**

Adaptive Security Strategies

Air Mobility Urban - Large Experimental Demonstrations (AMU-LED)

- Europe's main AAM demonstration project with CORUS XUAM (2021-2022)

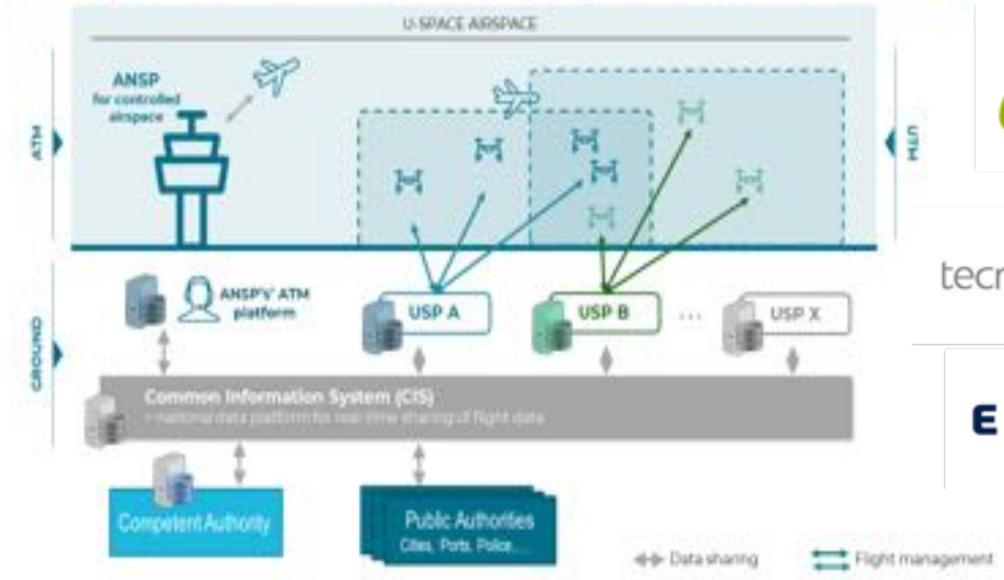

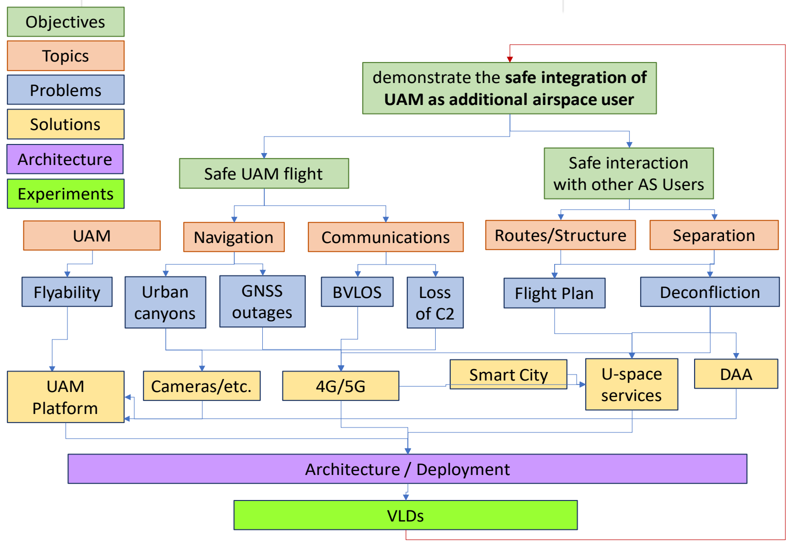

Adaptive Security Strategies

- Deep Reinforcement Learning Based Adaptive Controls
  - Learn adaptation strategy through observation between reference model and the reality

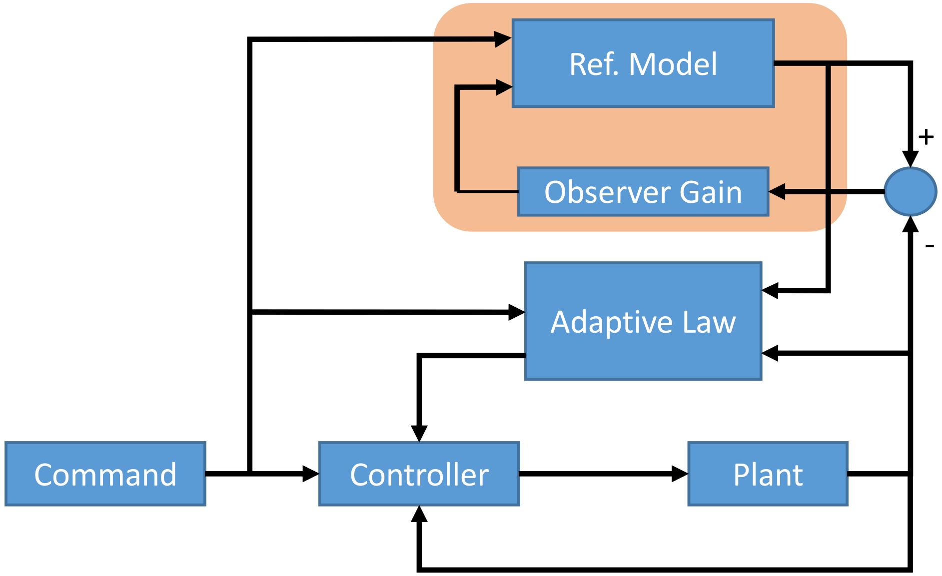

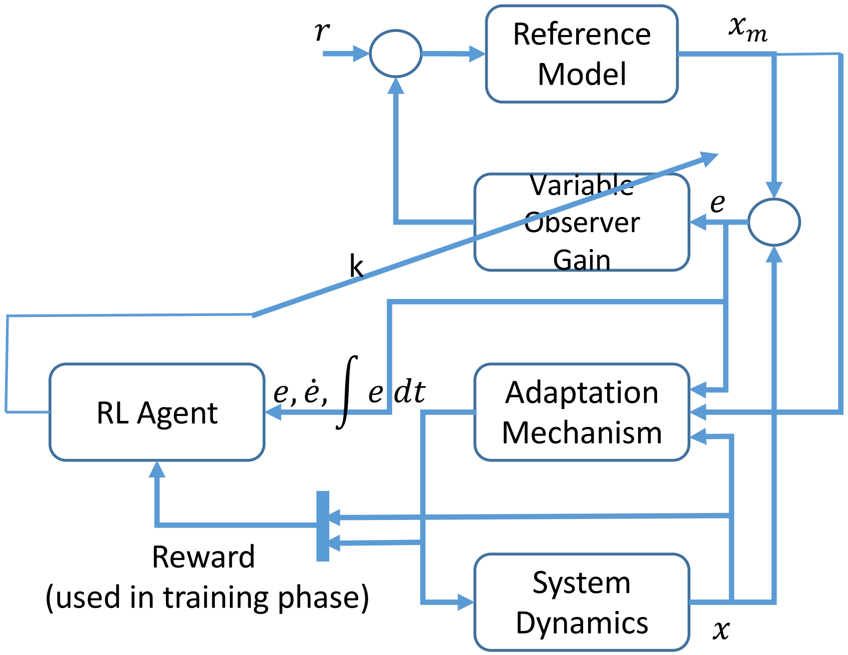

Yuksek B, Inalhan G. Reinforcement Learning Based Closed-loop Reference Model Adaptive Flight Control System Design. International Journal of Adaptive Control and Signal Processing. 2020;1–21.

State-of-Art Outlook

Model Reference Adaptive Control and Improvements

To provide robustness:
- Wise, Lavretsky*, Annaswamy**
  - mu Modification
  - Epsilon-modification
  - Deadzone adjustment
  - Projection algorithm
- Naira Hovakimyan***
  - L1 Adaptive Control

To improve transient performance:
- Lawrestky*, Annaswamy**, Gibson
  - *Combined/Composite MRAC (CMRAC)*
  - *Closed loop reference model (CRM) (observer-like reference model)*
  - *CRM + CMRAC*
- Naira Hovakimyan***
  - L1 Adaptive Control

- Trade-off in adaptive control systems between;
  - Improved transient performance vs decreased convergence speed of adaptation parameters.

*Lavretsky, E. and Wise, K. A., *Robust and Adaptive Control*, Springer, London, 2013.
**Narendra, K. S. and Annaswamy, A. M., *Stable Adaptive Systems,* Dover Publications, 2012.
***Hovakimyan, N. and Cao C., *$\mathcal{L}1$ Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation*, Society for Industrial and Applied Mathematics, 2010.

MRAC vs CRM

- Model Reference Adaptive Control (MRAC)
  - A universal observation in adaptive systems:
    - Convergent, yet oscillatory adaptation behavior in the presence of modeling errors.
    - Speed of adaptation can be increased by increasing the adaptation gain at the cost of increased oscillation frequency.
- MRAC with Closed-loop Reference Model (CRM)
  - Transient performance is improved.
  - Unlike the MRAC structure, Luenberger-like reference model is used in CRM adapitve systems*.

$$\dot{x}_{ref} = A_{ref}\, x_{ref} + \underbrace{\boxed{L_v \left(x - x_{ref}\right)}}_{\text{Error Feedback Term}} + B_{ref}\, y_{cmd}$$

*Eugene Lavretsky and Kevin A. Wise, *Robust and adaptive control* (pp. 317-353), Springer, London, 2013.

CRM Adaptive Control Systems Implementation

- General Scheme of the MRAC and CRM-Adaptive Systems

Model Reference Adaptive Control System (MRAC)

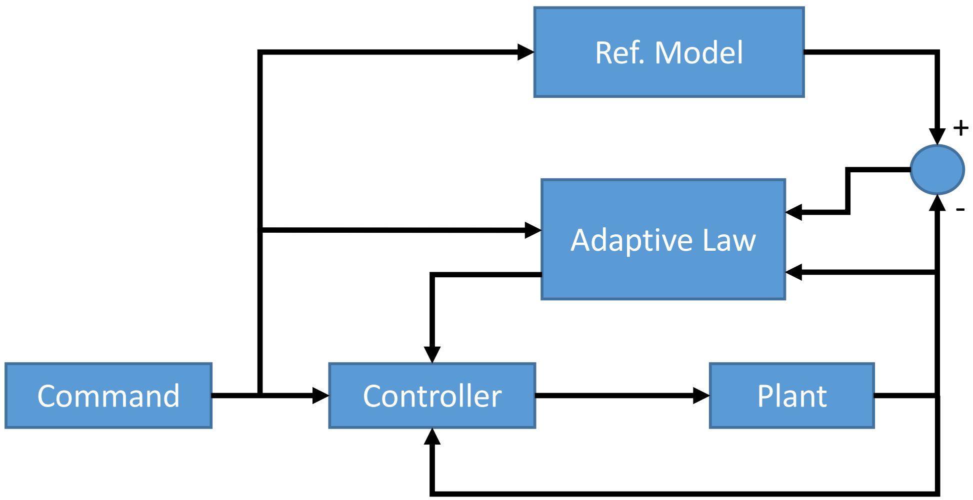

Closed-loop Reference Model (CRM) Adaptive System (Observer-like Reference Model)

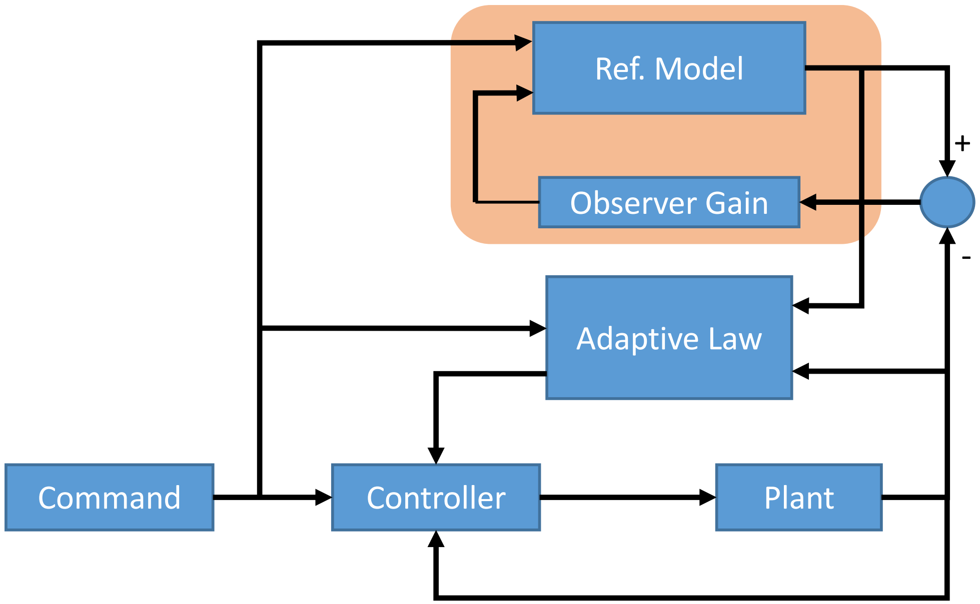

CRM Adaptive Control Systems Double Edge Sword

- Another important feature of the CRM-adaptive systems is <u>water-bed effects</u>
  - A badly chosen design parameters (learning rate and observer gain) can significantly worsen the adaptive system performance in terms of $\dot{u}(t)$

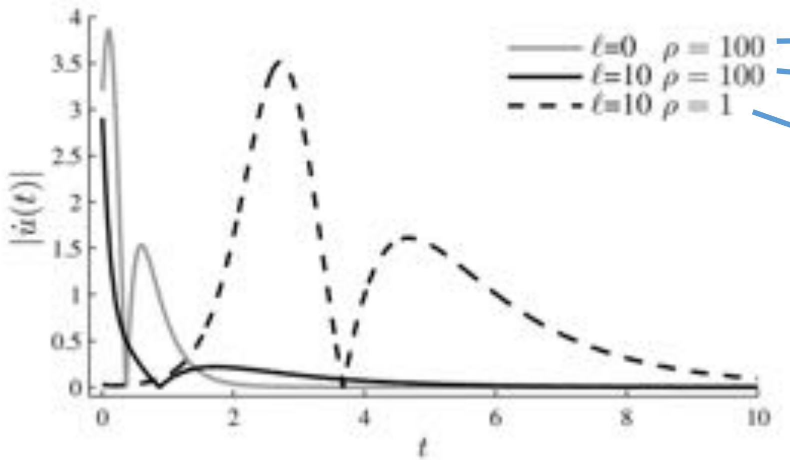

Open-loop ref. model

Closed-loop ref. model **with** parameter optimizaion

Closed-loop ref. model **without** parameter optimizaion

Travis E. Gibson, Anuradha M. Annaswamy, and Eugene Lavretsky. "Adaptive systems with closed-loop reference-models, part I: Transient performance." *2013 American Control Conference*. IEEE, 2013.

CRM Adaptive Control Systems

- CRM-Adaptive Systems with Fixed Observer Gain :
  - Small amplitude L_v => High frequency oscillation
  - Large amplitude L_v => Slow Dynamics
- Trade-off in CRM-adaptive systems between;
  - Improved transient performance vs decreased convergence speed of adaptation parameters.
- Why do not we use Variable Observer Gain ?
  - Large amplitude L_v is used in the initial phase of the adaptation process => to improve the transient dynamics
  - Small amplitude L_v is used after the adaptation process is completed => to speed up the system response
  - *Can we learn the adaptation policy of the observer gain magnitude by using Reinforcement Learning?*
    - *RL-CRM Adaptive Control Systems*

Yuksek B, Inalhan G. Reinforcement Learning Based Closed-loop Reference Model Adaptive Flight Control System Design. International Journal of Adaptive Control and Signal Processing. 2020;1–21.

Reinforcement Learning - CRM Adaptive Control System

Actor-Critic Structure Trained by utilizing DDPG Algorithm

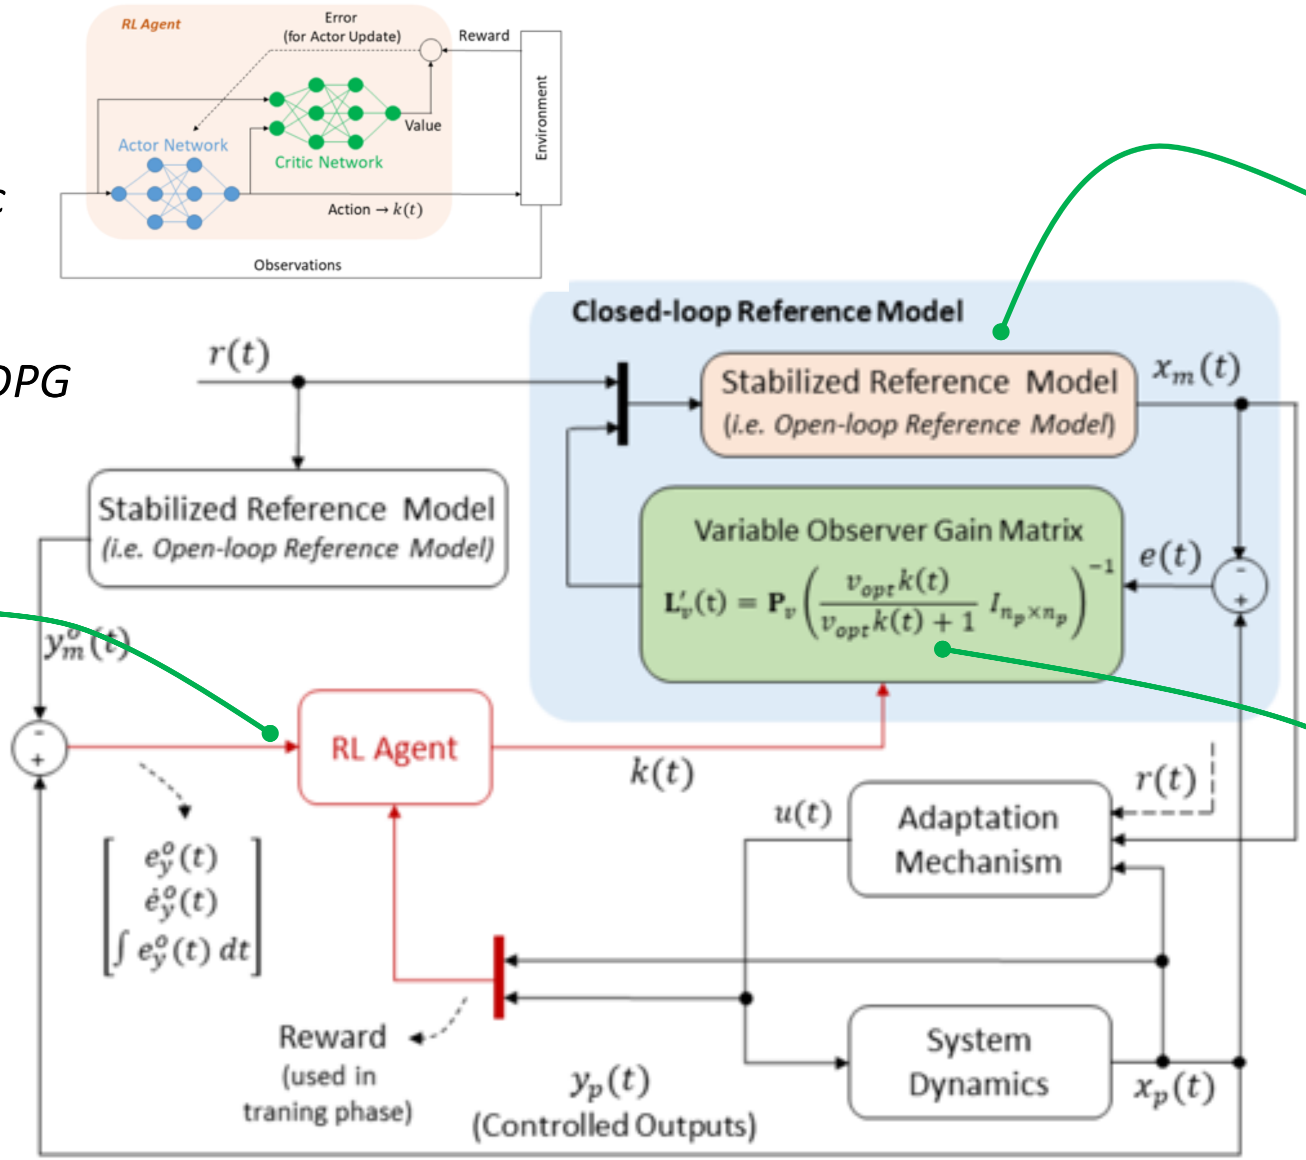

Stabilised model is required if the open-loop dynamics is unstable

Time-varying $k(t)$ provides scaling policy of the observer gain parameter v_{opt}

Learning of RL-CRM Adaptive Control Systems

- Learning algorithm is Deep Deterministic Policy Gradient (DDPG)
- Agent is based on an actor – critic neural network structure

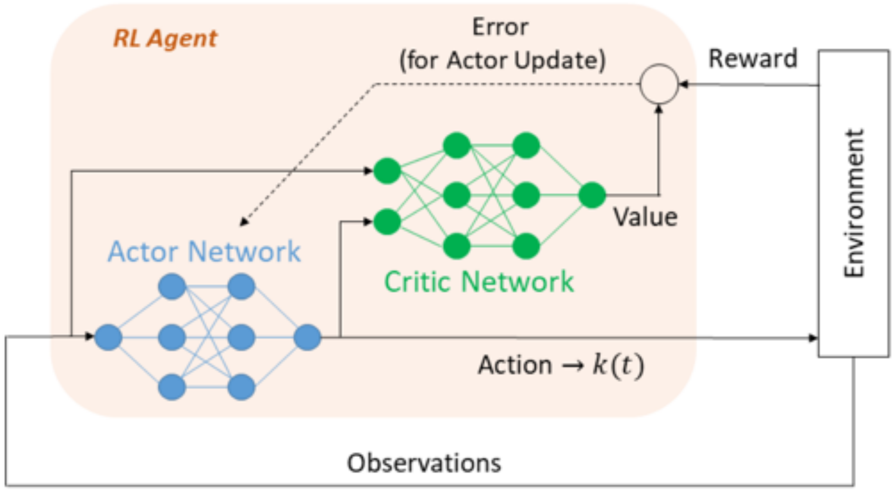

Additional questions about actor-critic agent:

- Can we use the trained agent on another platform which has similar mechanical structure but different dynamical parameters? Is <u>transfer learning method</u> a suitable solution to improve the performance of the trained RL agent on another platform?

NN and Reward Function Design for RL-CRM

- Neural Network Parameters

Network	Parameter	Value
Actor	Number of Hidden Layers	1
	Number of Nodes in Hidden Layers	10
	Activation Functions	Tanh
	Learning Rate	0.002
	Gradient Threshold	1
Critic	Number of Obs. Path Hidden Layers	2
	Number of Nodes in Obs. Path Hidden Layers	10
	Number of Action Path Hidden Layers	1
	Number of Nodes in Action Path Hidden Layers	10
	Activation Functions	Tanh
	Learning Rate	0.002
	Gradient Threshold	1

- Reward Function:

$$R(t) = w_1 R_p(t) + w_2 R_{e_y}(t) + w_3 R_u(t) + w_4 R_{e_{cmd}}(t) + w_5 R_o(t)$$

$$R_p(t) = \begin{cases} -1, & \text{if } \|y_p(t)\|_\infty \geq 0.105 \\ 0, & \text{otherwise} \end{cases}$$

$$R_{e_y}(t) = \begin{cases} 4, & \text{if } |e_y(t)| \leq 0.0005 \\ 0, & \text{otherwise} \end{cases}$$

$$R_u(t) = \begin{cases} 2, & \text{if } |\dot{u}(t)| \leq 0.02 \\ 0, & \text{otherwise} \end{cases}$$

$$R_{e_{cmd}}(t) = \begin{cases} 2, & \text{if } |e_{y_{cmd}}(t)| \leq 0.01 \text{ and } t \geq 0.3 \text{ sec} \\ 0, & \text{otherwise} \end{cases}$$

$$R_o(t) = \begin{cases} 1, & \text{if } |e_y^o(t)| \leq 0.02 \\ 0, & \text{otherwise} \end{cases}$$

$$w_i = 1 \ \ \forall i \in \{1, 2, 3, 4, 5\}$$

Ability to span the whole Pareto-optimal frontier across millions of different scenarios Including failures and variations.

RL-CRM Adaptive Control System Design on Scalar Pitch Dynamics of a Helicopter

- **Mathematical Model**

$$\dot{q} = M_q q + M_{\delta_e}\left(\delta_e + f(q)\right)$$

M_q: Vehicle pitch damping
M_{δ_e}: Elevator effectiveness
δ_e: Control input
$f(q)$: Inherent uncertainties in the helicopter dynamics

$$f(q) = -0.01 \tanh\left(\frac{360}{\pi} q\right) = \theta\, \Phi(q)$$

θ: Unknown constant
$\Phi(q)$: Known regressor vector

Pitch Dynamics Model of a Transport Helicopter in Hover Flight

(Lavretsky, 2013, p. 270)

*Eugene Lavretsky and Kevin A. Wise, *Robust and adaptive control*, Springer, London, 2013.

RL-CRM Adaptive Control System Design on Scalar Pitch Dynamics of a Helicopter

- Step Response Comparison of MRAC, CRM and RL-CRM

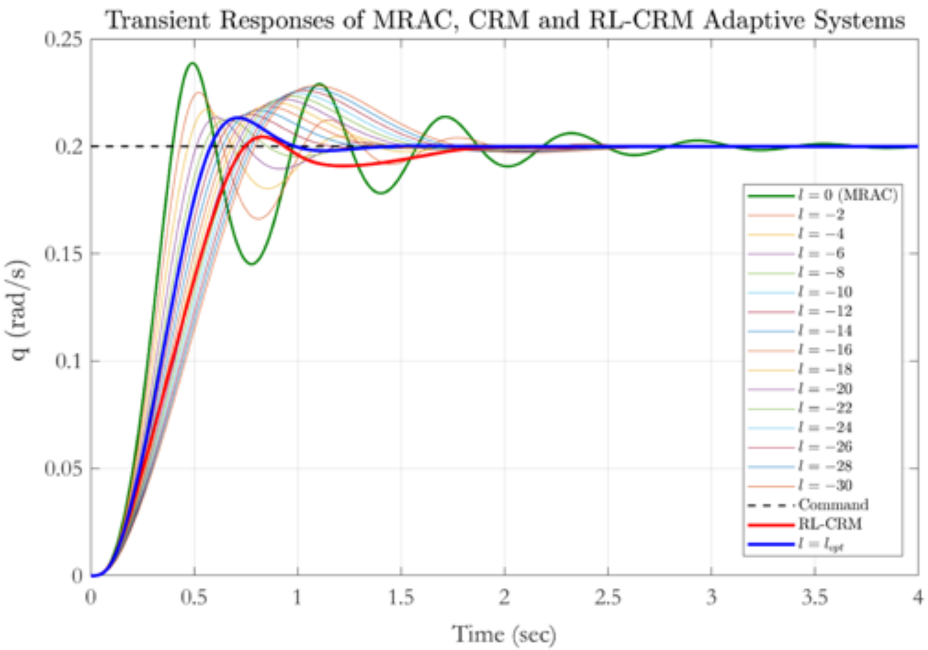

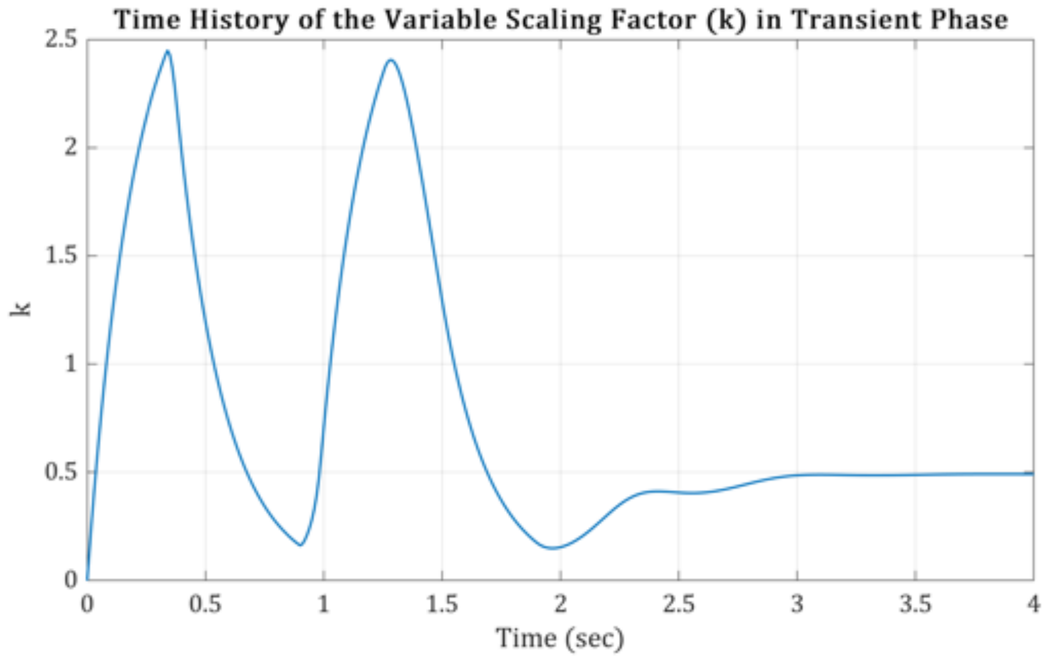

RL-CRM Adaptive Control System Design on Scalar Pitch Dynamics of a Helicopter

- Water-Bed Effect Comparison on MRAC, CRM and RL-CRM

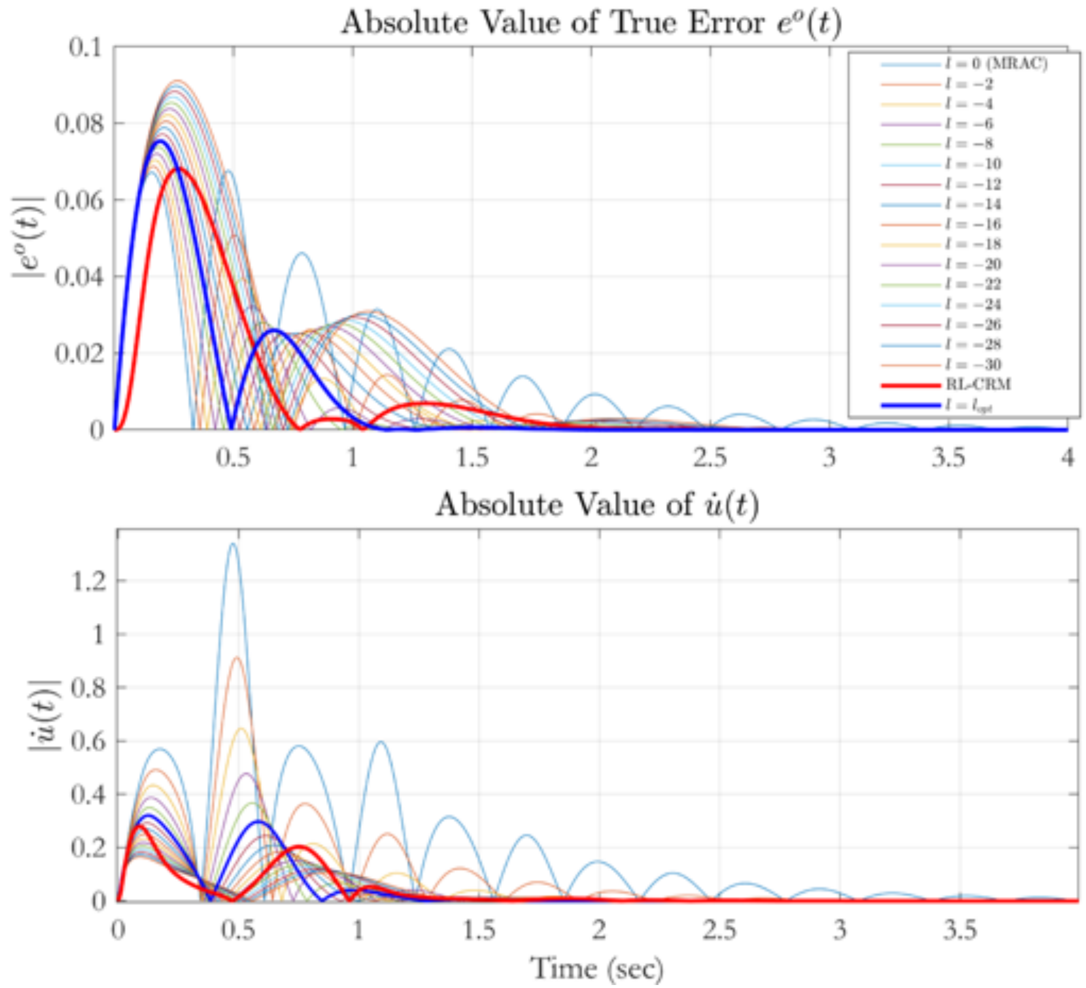

RL-CRM Adaptive Control System Design on Scalar Pitch Dynamics of a Helicopter

- 500-run Monte-Carlo Analysis for $\pm 35\%$ Parametric Uncertainty on M_q and M_{δ_e}

Performance Metrics	MRAC	CRM	Improvement (%)	RL-CRM	Improvement (%)
$\overline{\|\dot{\hat{K}}_x\|}$	15.2114	3.7341	75.4520	2.4489	83.9008
$\overline{\|\dot{\hat{K}}_r\|}$	18.4647	7.8298	57.5958	5.5146	70.1344
$\overline{\|\dot{\theta}\|}$	0.0888	0.0338	61.9369	0.0207	76.6892
$\overline{\|y_m\|_\infty}$	0.2	0.2064	-3.2	0.2	-
$\overline{\|e_y\|}$	0.4616	0.1957	57.6039	0.1379	70.1256
$\overline{\|e_y^o\|}$	0.4616	0.3928	14.9047	0.3886	15.8145
$\overline{\|\dot{u}\|}$	6.5704	2.0811	68.3262	1.4163	78.4290

RL-CRM Adaptive Control System Design on Scalar Pitch Dynamics of a Helicopter

- The Worst Case Analysis for -35% Parametric Uncertainty on M_q and M_{δ_e}

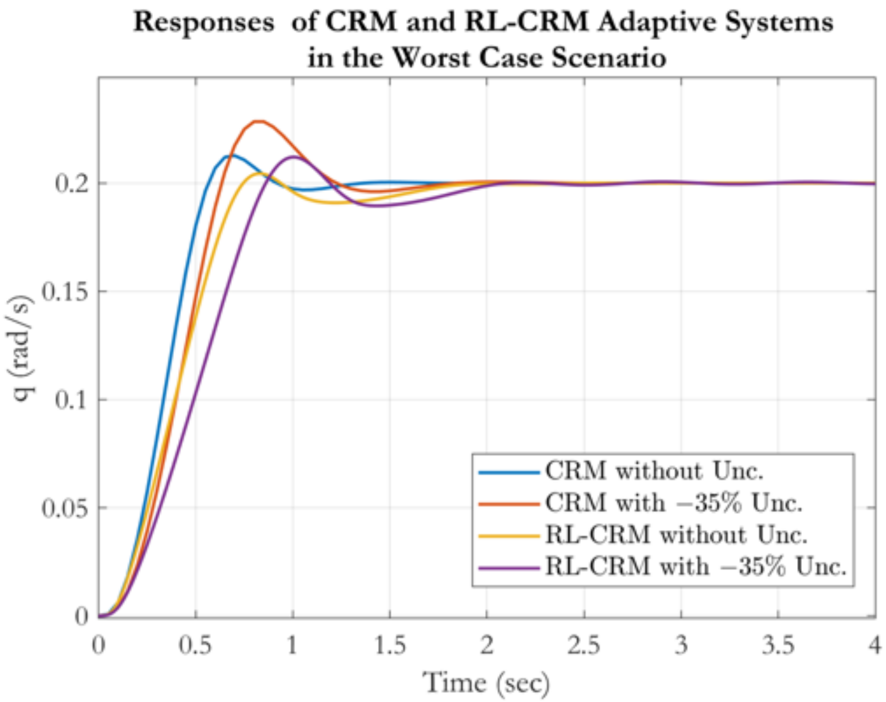

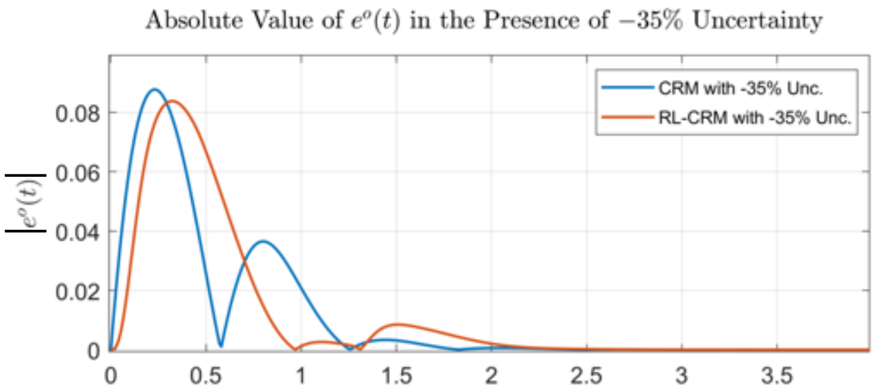

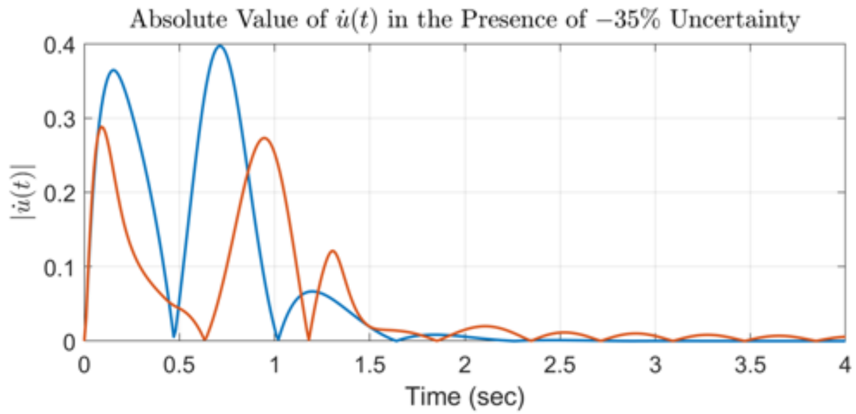

RL-CRM Adaptive Control System Design on Scalar Pitch Dynamics of a Helicopter

- The Worst Case Analysis for -35% Parametric Uncertainty on M_q and M_{δ_e}

Performance Metrics	MRAC	CRM	Improvement (%)	RL-CRM	Improvement (%)
$\|\dot{\hat{K}}_x\|$	19.7655	4.9225	75.0955	3.4801	82.3931
$\|\dot{\hat{K}}_r\|$	22.9284	9.4137	58.9431	6.4318	71.9483
$\|\dot{\hat{\theta}}\|$	0.1103	0.0407	63.1010	0.0246	77.6972
$\|y_m\|_\infty$	0.2	0.2171	-8.5500	0.2005	-0.2500
$\|e_y\|$	0.5732	0.2353	58.9498	0.1608	71.9470
$\|e_y^o\|$	0.5732	0.5101	11.0084	0.5214	9.0370
$\|\dot{u}\|$	8.5403	2.6274	69.2353	1.8001	78.9223

Major Challenges in Advanced Air Mobility Concept and Our Autonomy Focus

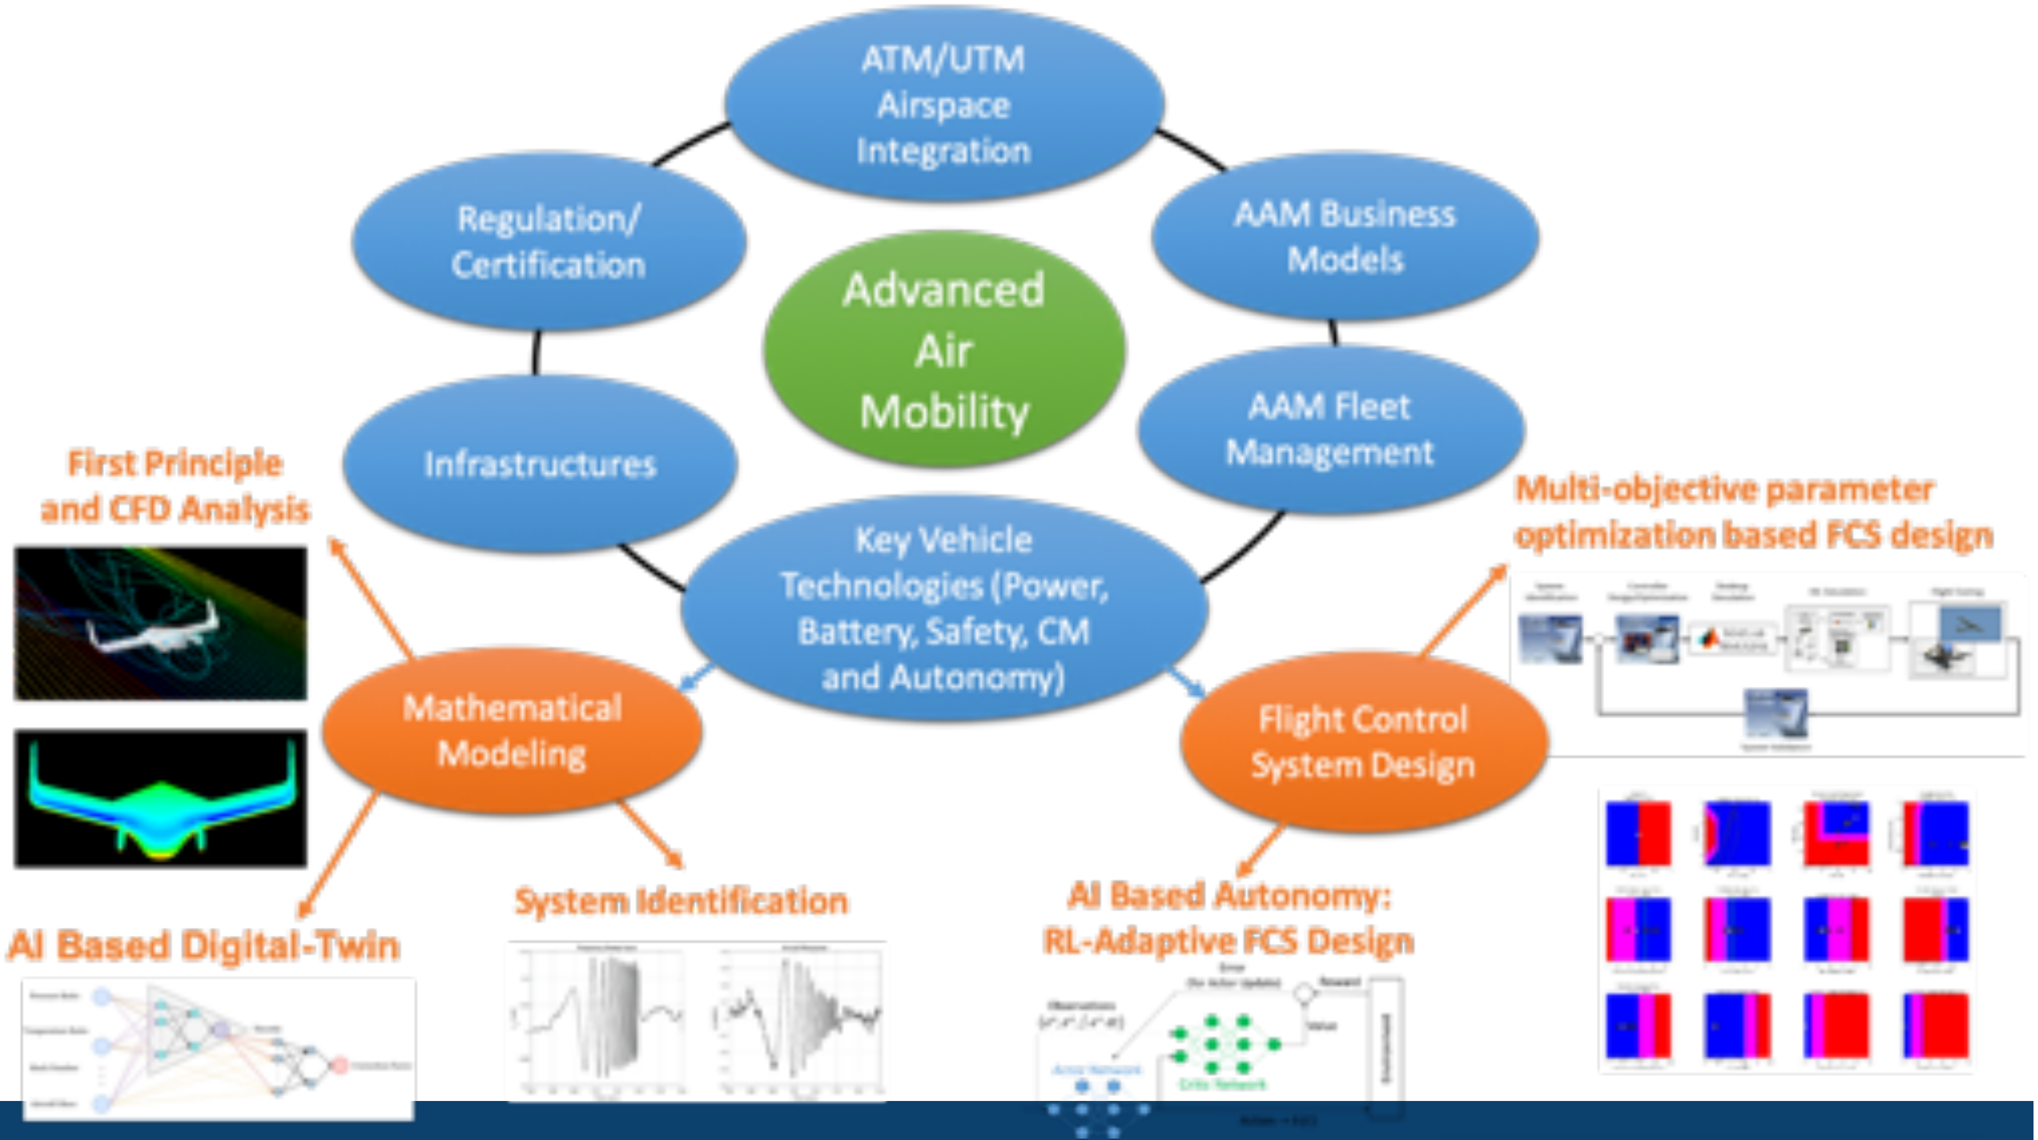

Desktop-to-Flight Design Workflow*

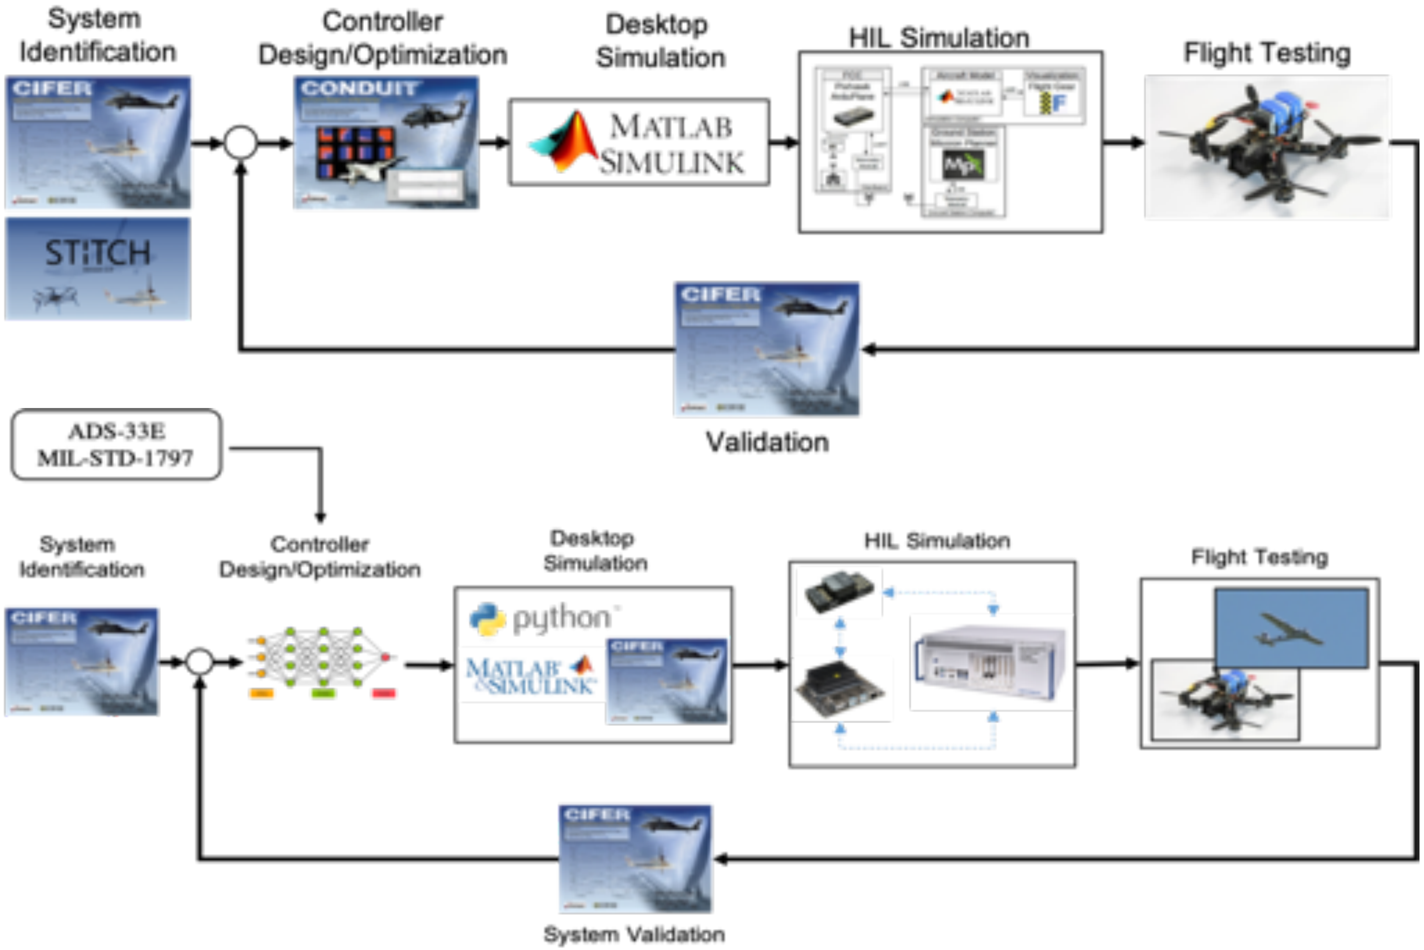

*Tischler, M. B., Berger, T., Ivler, C. M., Mansur, M. H., Cheung, K. K., and Soong, J. Y., “Practical Methods for Aircraft and Rotorcraft Flight Control Design: An Optimization-Based Approach,” AIAA education series, 2017.

Reliable performance under large variations

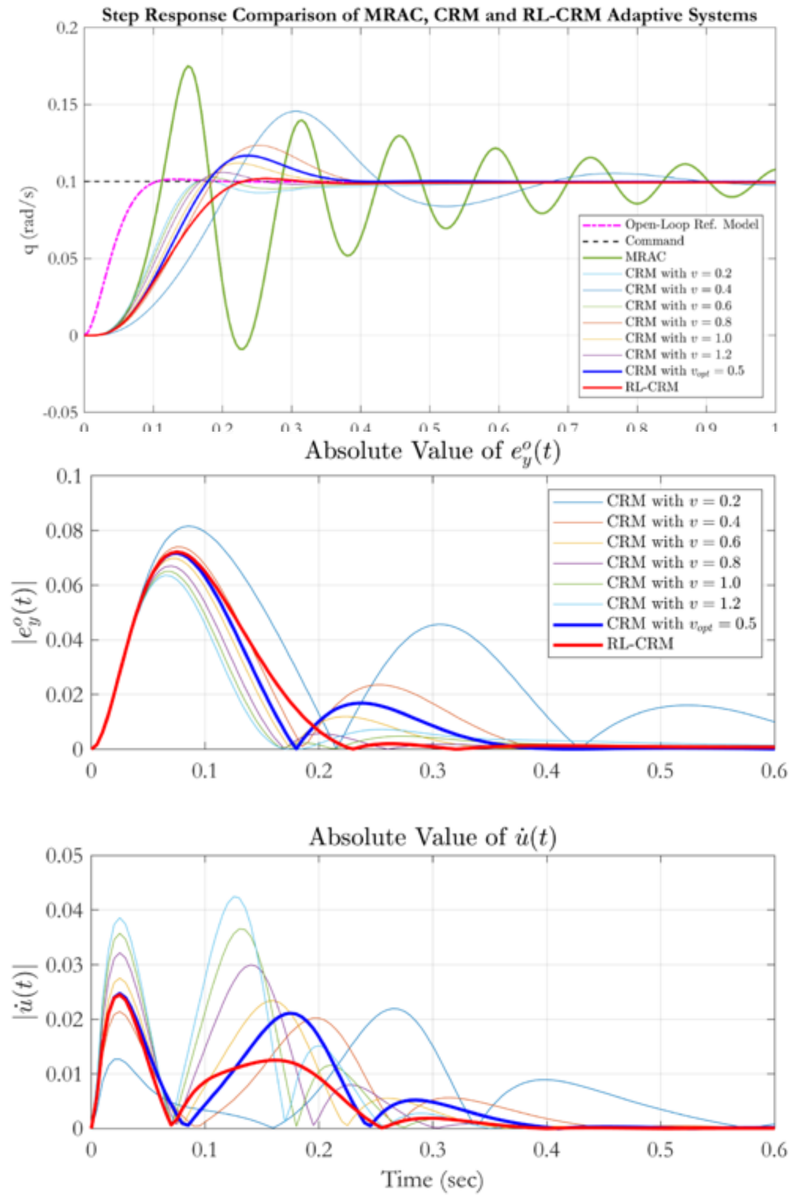

Performance Metrics	MRAC	CRM	Improvement (%)	RL-CRM	Improvement (%)
$\overline{\lVert \dot{\hat{K}}_x \rVert}$	15.2114	3.7341	75.4520	2.4489	83.9008
$\overline{\lVert \dot{\hat{K}}_r \rVert}$	18.4647	7.8298	57.5958	5.5146	70.1344
$\overline{\lVert \dot{\hat{\theta}} \rVert}$	0.0888	0.0338	61.9369	0.0207	76.6892
$\lVert y_m \rVert_\infty$	0.2	0.2064	-3.2	0.2	-
$\overline{\lVert e_y \rVert}$	0.4616	0.1957	57.6039	0.1379	70.1256
$\overline{\lVert e_y^o \rVert}$	0.4616	0.3928	14.9047	0.3886	15.8145
$\overline{\lVert \dot{u} \rVert}$	6.5704	2.0811	68.3262	1.4163	78.4290

Towards Certification of Hybrid (AI/Classical) Controllers

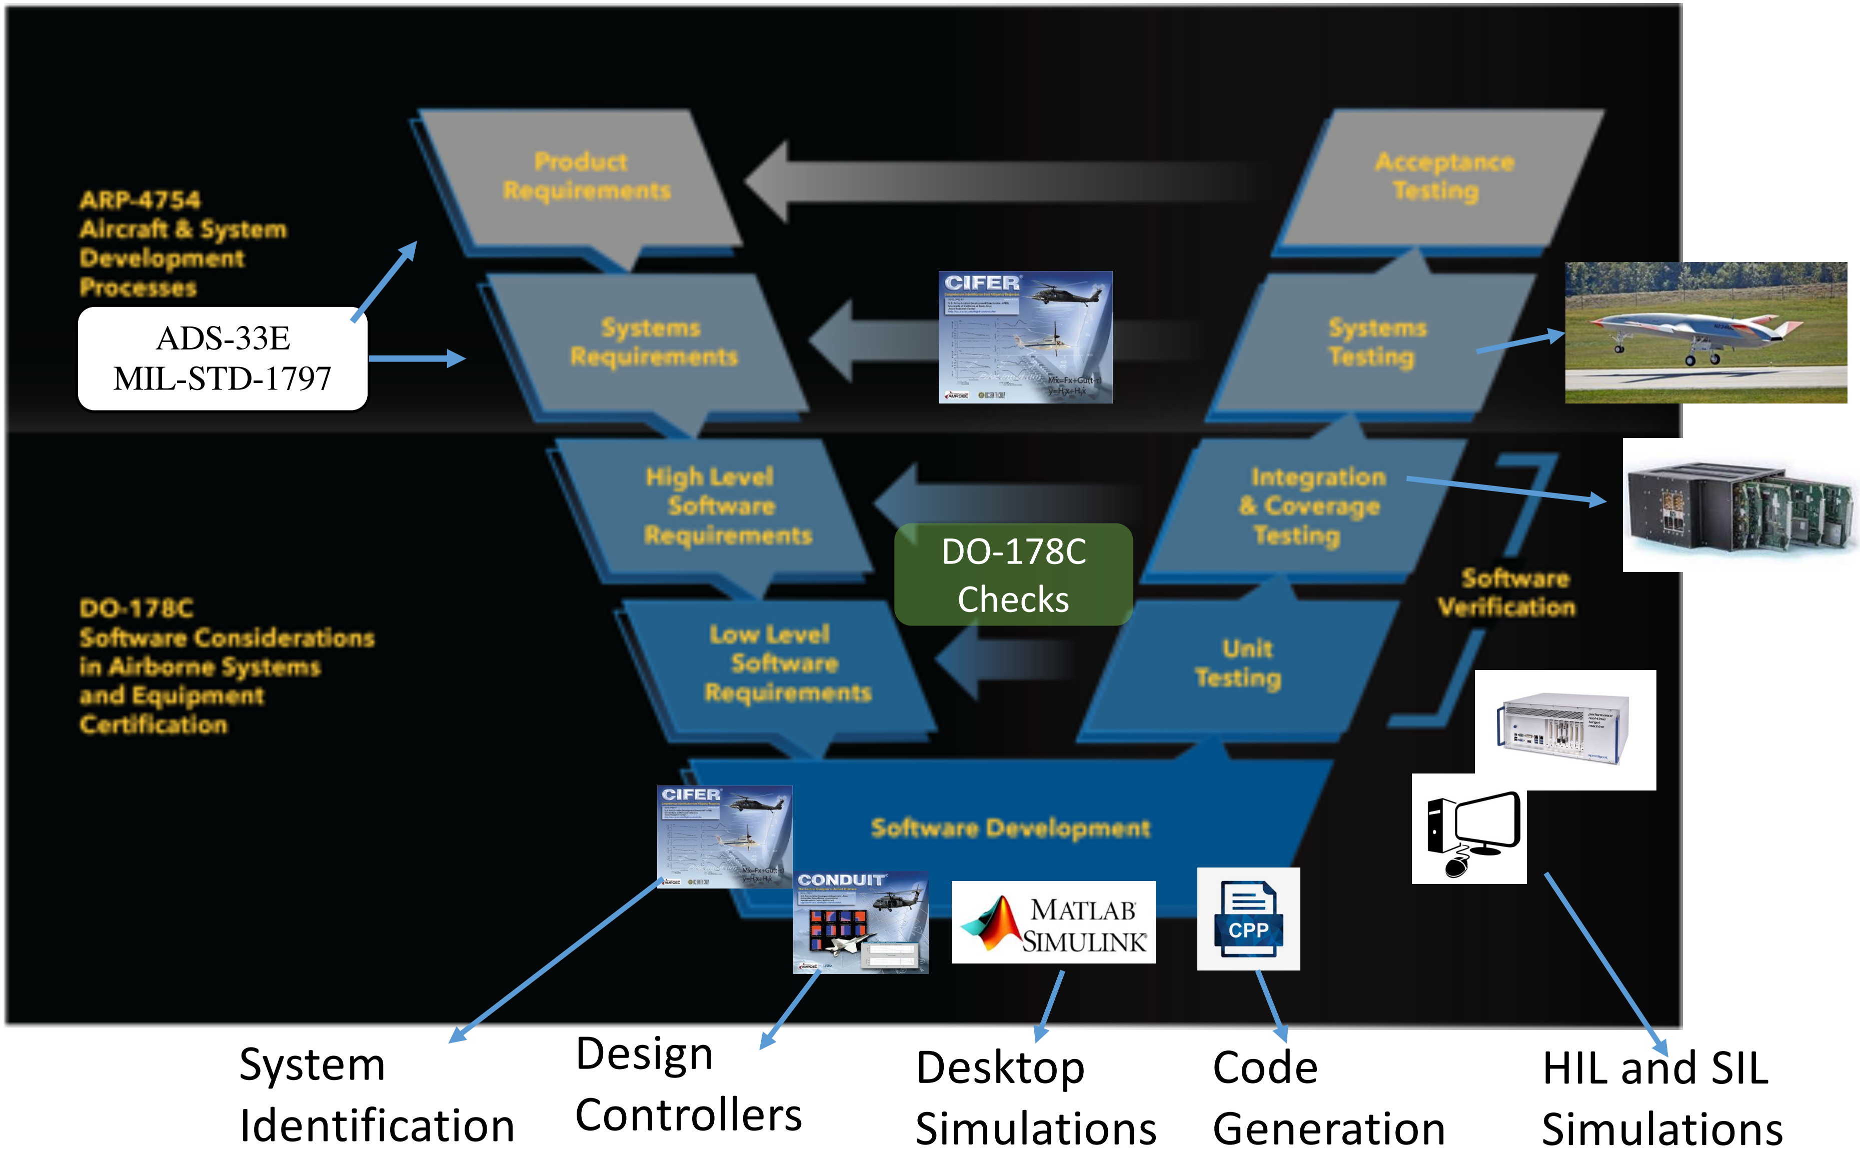

Next Steps….

- Design and VVQC for AI-Driven Safety Critical Systems
    - Extensive usage of synthetics and digital-twins
- Reinforcement Learning in Uncertain Environments with Decentralized Decision-makers
    - Fusion of Tree-based decision algorithms and RL with learned models
    - Survivability and Lethality
- Human-Machine Teaming
    - Hybrid-system models as descriptive for behaviour taxonomy

Next Steps…

- Explainable AI for Reinforcement Learning (XAI-RL)
  - Asynchronous Advantage Actor-Critic (A3C)
  - Explanation (Visualization) Methods
  - GradCam

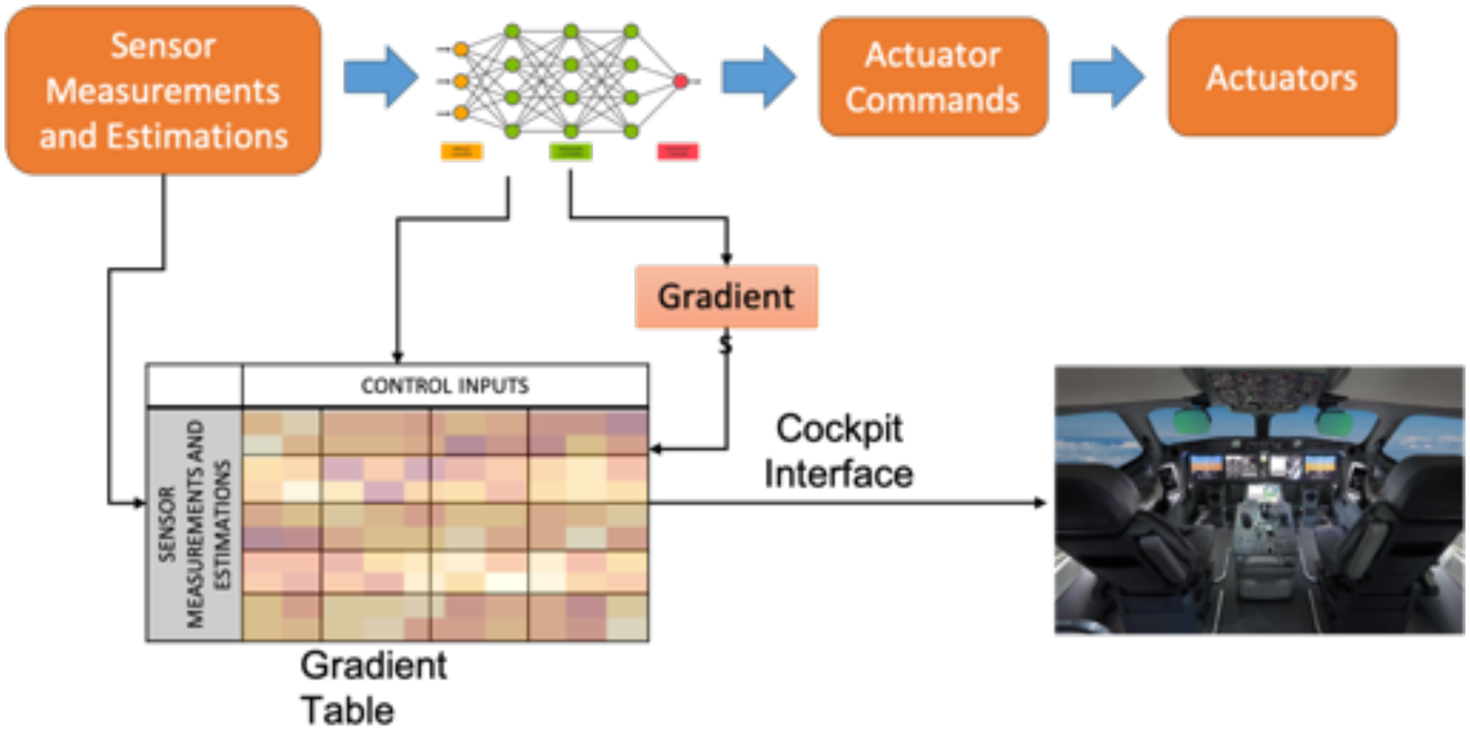

Key cornerstones in AI-Driven Design

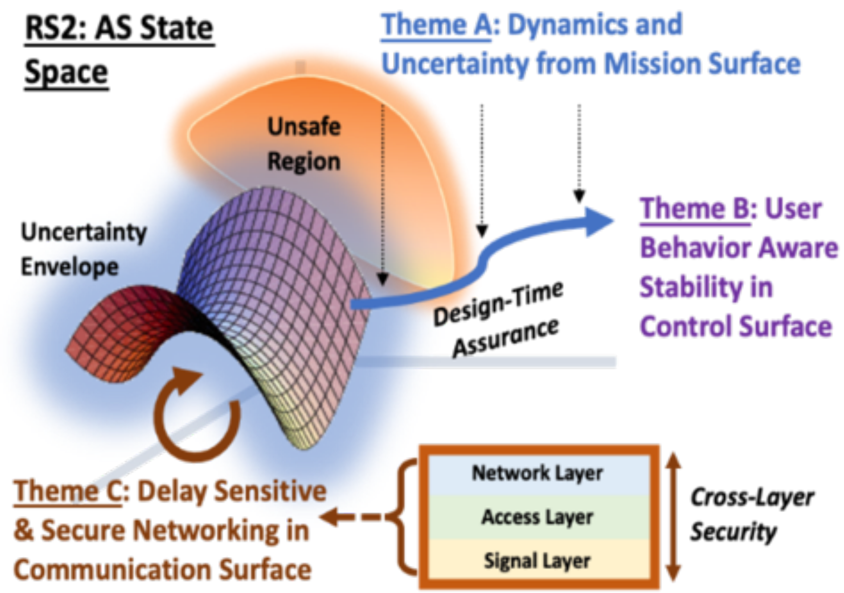

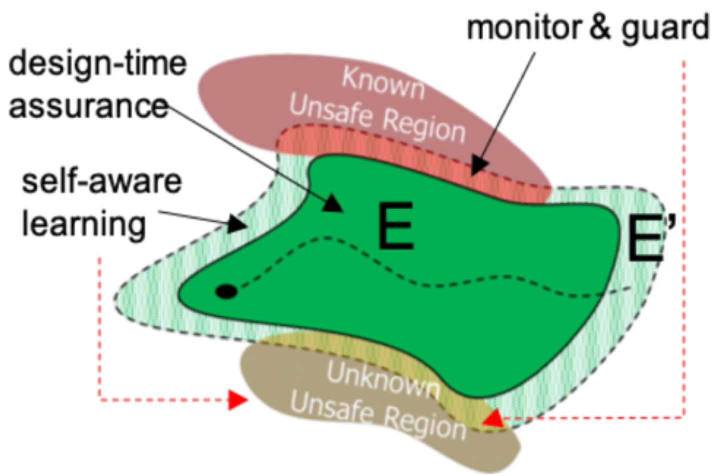

- Provide **quantifiable safety and feedback** to the mission surface when the limits of secure controllability are compromised within a time horizon under current policies and adversarial situations.
- Key Solution Cornerstones in Learning-Enabled Context
  - **Interpretability** => Explainable and Trustworthy AI
  - **Continual Assurance** => Dynamic Verification & Validation
  - **Adaptive Security Strategies**

Continual Assurance: Dynamic Verification and Validation

The major challenge of commercial flight planning

- Key factors (and uncertainty) in commercial flight planning
  - Wind
  - Tail-number specific fuel consumption
- Essentially ”the cost” boils down to fuel usage/cost

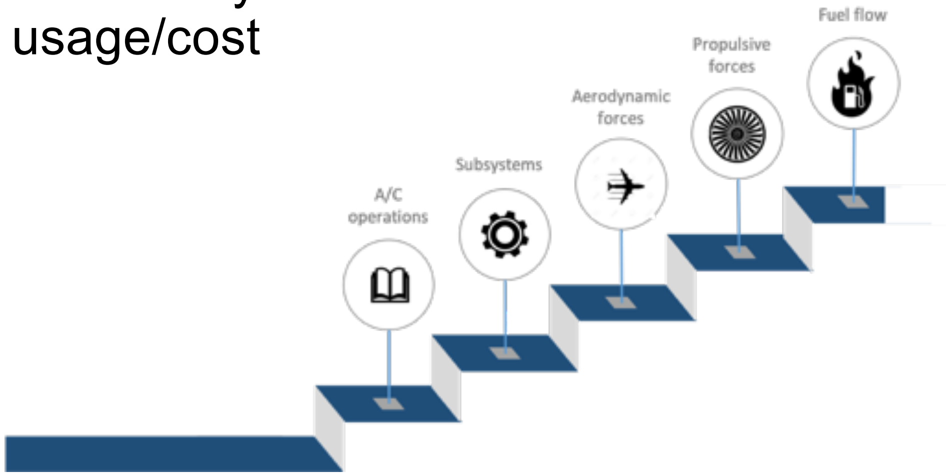

- Significant impact towards “sustainable aviation” concept
  - Cost
  - Emissions

JetPlanner Pro/ FlitePlan Core (Jepp/Boeing)

Aircraft Performance and Wind Calibration Scheme

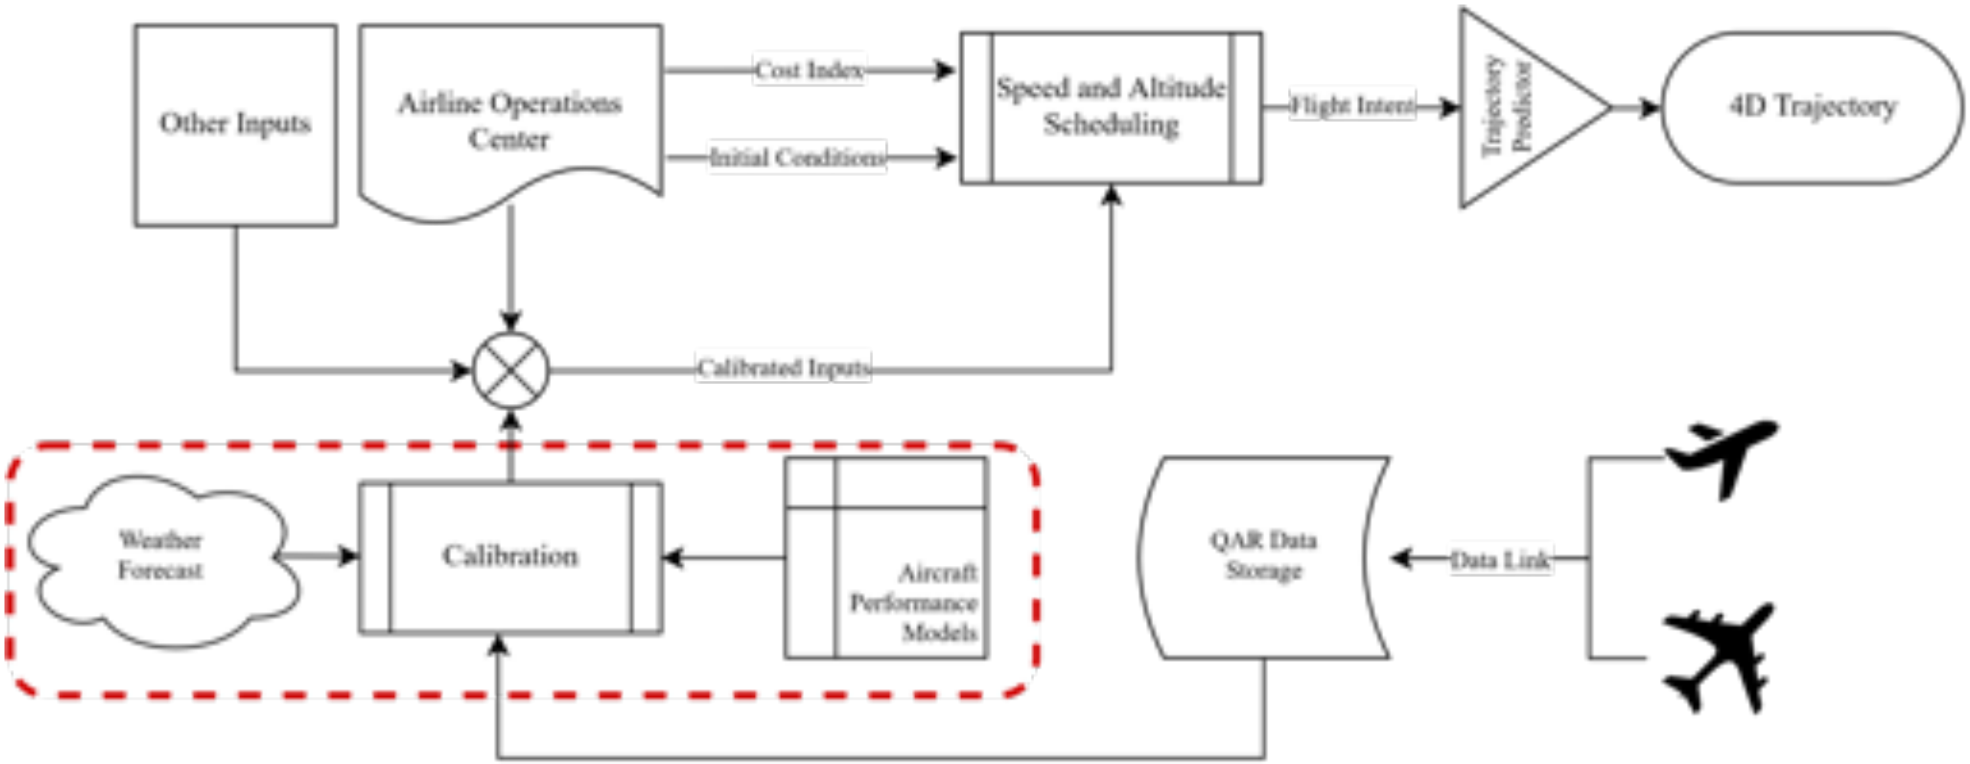

Developing Digital-Twin Performance Models

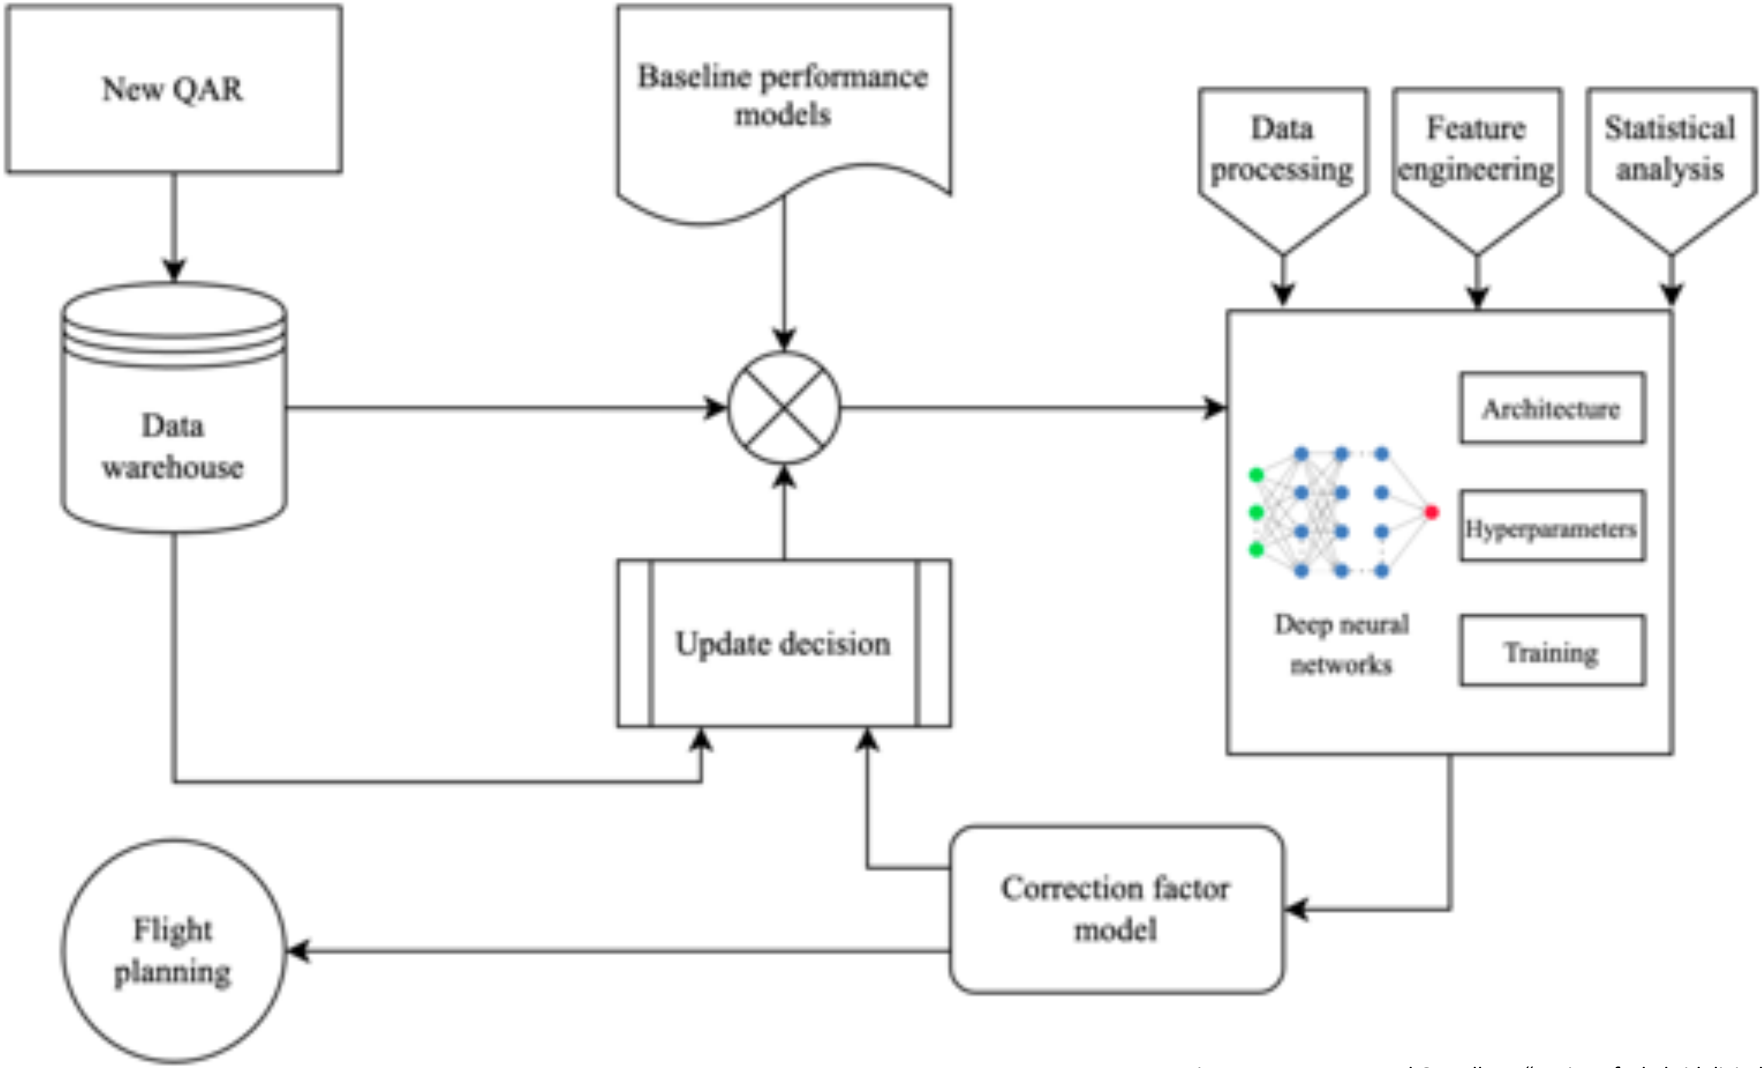

M. Uzun, M. U. Demirezen, E. Koyuncu, and G. Inalhan, "Design of a hybrid digital-twin flight performance model through machine learning," in *2019 IEEE Aerospace Conference*. IEEE, 2019, pp. 1–14.

Digital-Twin Aircraft Performance Model

- Accurate trip fuel calculation.
- Why high precision digital twins are important?
  - High fidelity performance model means correct estimation of take-off fuel weight.
    - Less take-off fuel stands for lest take-off weight, hence less total fuel consumption.
    - The ratio is approximately 3/1 (take-off gross weight / take-off fuel) for long haul and 6/1 for short haul flights.
      - Example B777-300ER: 322 tons / 99 tons / 11 h
      - Example B737-800: 66 tons / 11 tons / 3 h

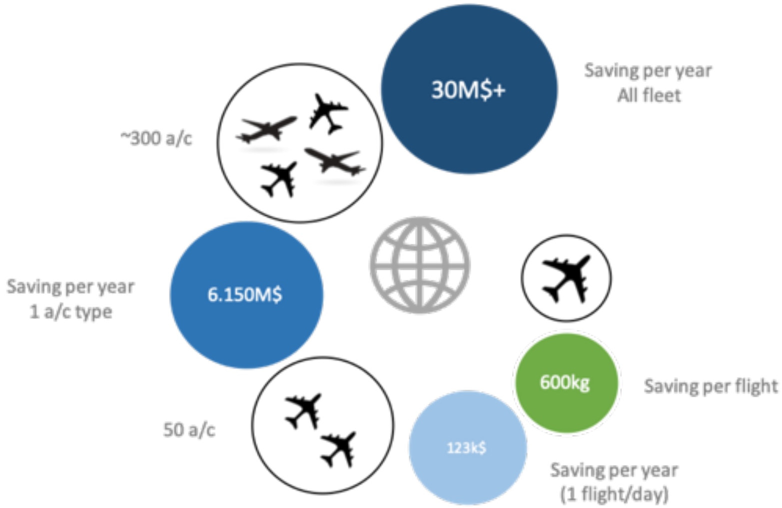

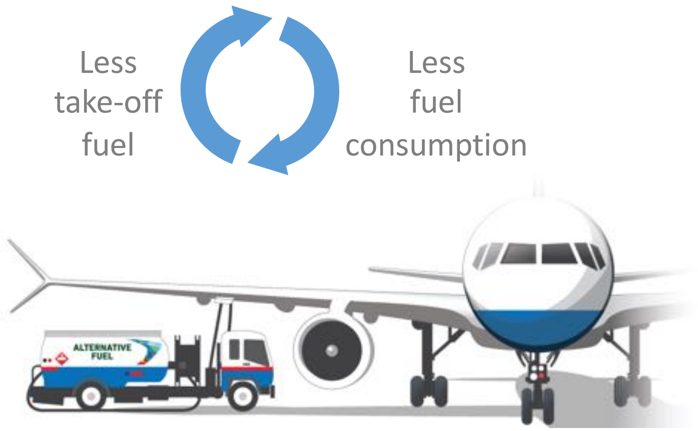

Fundamental behind our solution

Economy Cruise Cost Function [nm/kg]

Fixed

$$ECCF = \frac{\boxed{CI} + \boxed{F(h, M)}}{\boxed{v_{TAS}(h, M)} + \boxed{w(h)}}$$

Variable

Aircraft Performance Model

Wind Forecast

Uncertainties in

- Aircraft performance model
- Wind forecast

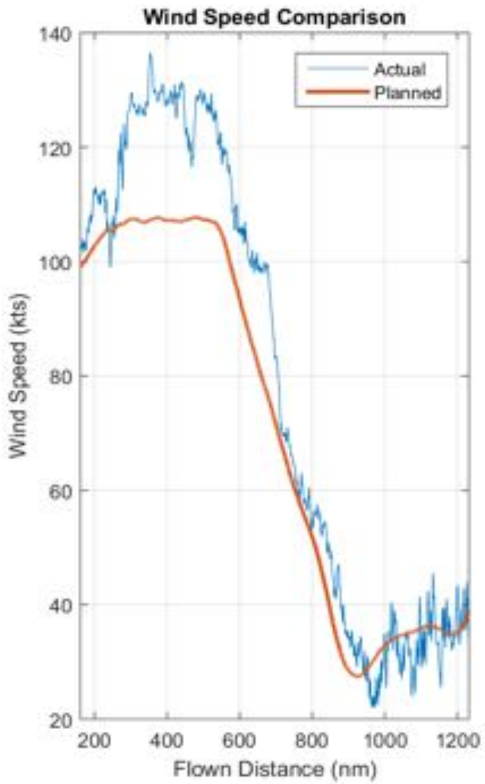

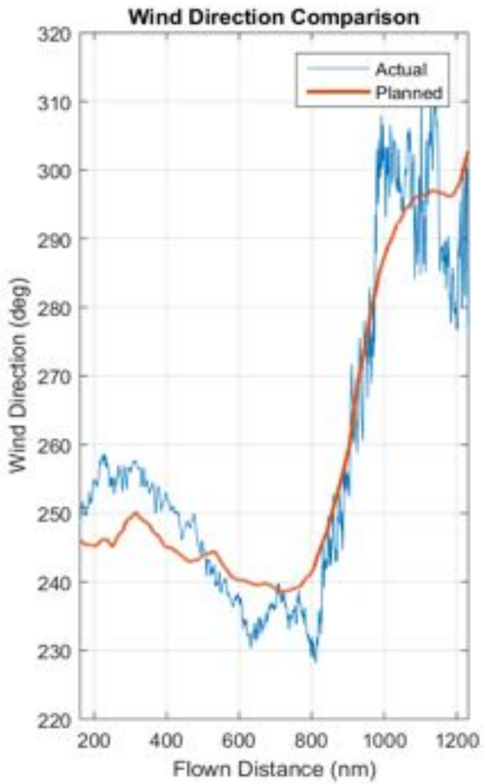

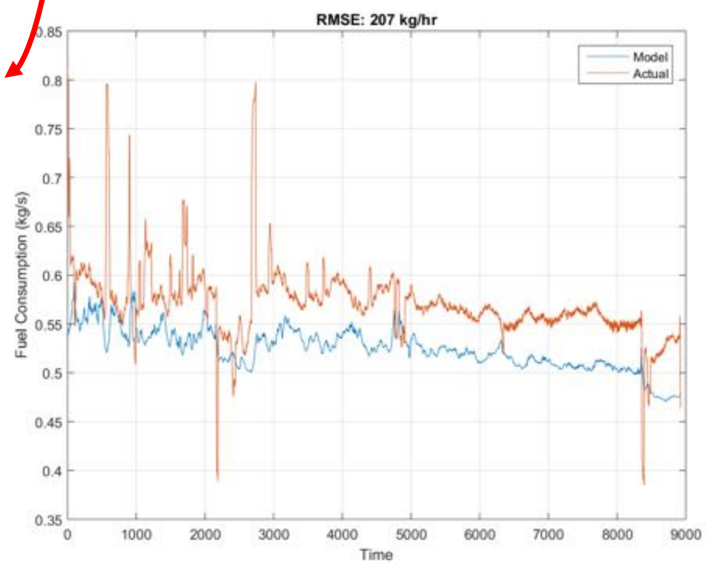

State-of-art in Performance Modeling

- Top aircraft performance models widely used in real world operations:
  - Aircraft manufacturer’s models (highest fidelity?):
    - Performance charts to be utilized in ground based planning tools.
    - Flight Management Computers.
    - Look-up tables.
  - Generic. Only customization is through performance factor which is calculated by aging of aircraft.
    - Boeing’s BPS (Boeing Performance Software) - INFLT (In flight)
    - Airbus’ PEP (Performance Engineer’s Program)
    - Eurocontrol’s BADA (Base of Aircraft Data) Family 3
    - BADA Family 4
- BPS and PEP are composed of look-up tables.
- BADA4 is a result of curve fitting to the synthetic data generated by BPS and PEP.
- BADA3 is based on empirical approaches.
- They are designed for ”zero” condition. However, aircraft tend to deviate from their original performances.
  - Operating at different regions, routes.
  - Maintenance.

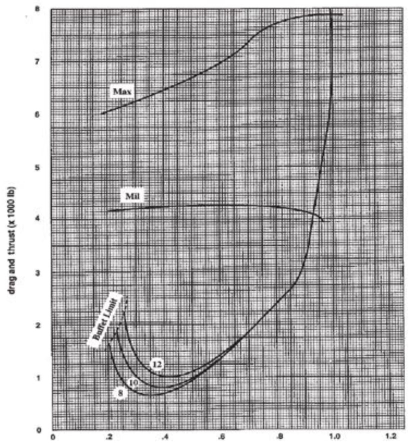

DISPATCH LOAD:			
		PAYLOAD:	46054
EZFW:	219623	MZFW:(S)	237682
ETOW:	302679	MTOW:(S)	351534
ELDW:	226481	MLDW:(S)	251290
REMF:	6858	**MIN DIV:**	**6101**

FMS INIT LOAD:	
KORD/LTBA	
LDG ELEV:0163FT	PRF FACTOR%:2.9
CI: 46	TROPO:31710

ALTN	DIST	TIME	FL	FUEL
LTBR (F)	106	00:24	130	3011

State-of-art in Performance Modeling

- We observe two types of discrepancies:
  - Operational
    - In BADA based trajectory predictions, a single type of thrust setting is assumed: Maximum climb for climb mode, Low-idle for descent mode.
    - Accelerations during cruise also cause differences.
  - Parametric
    - Projected as bias from the actual fuel flow.

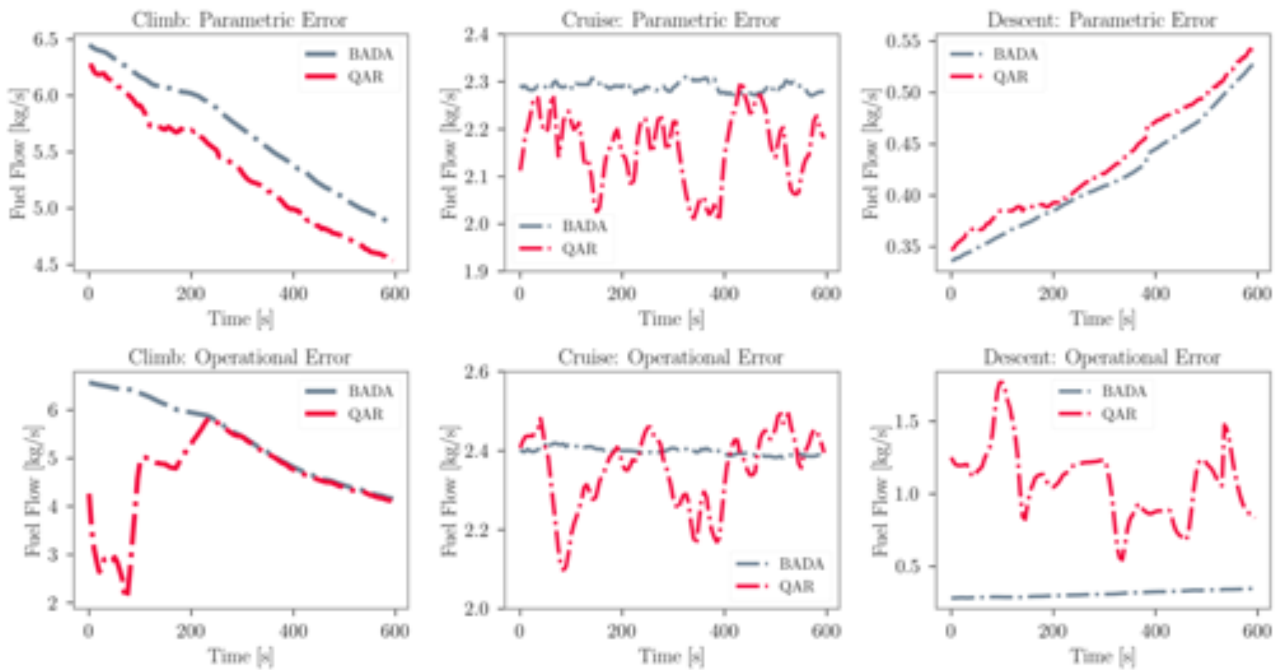

Developing Tail Number Specific Digital-Twin Performance Models

- Proposed network:

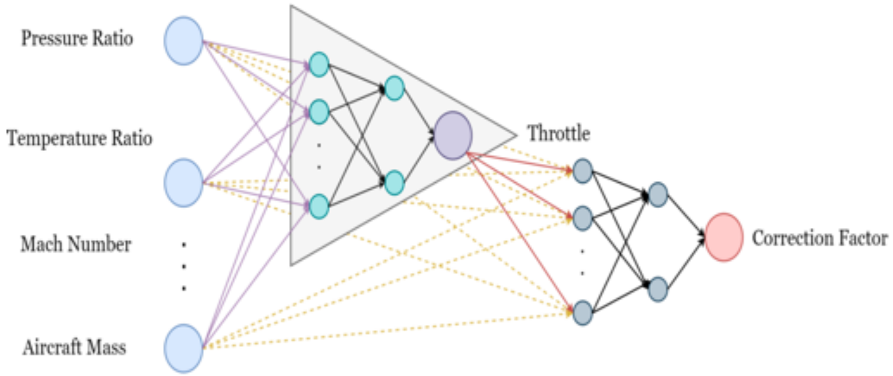

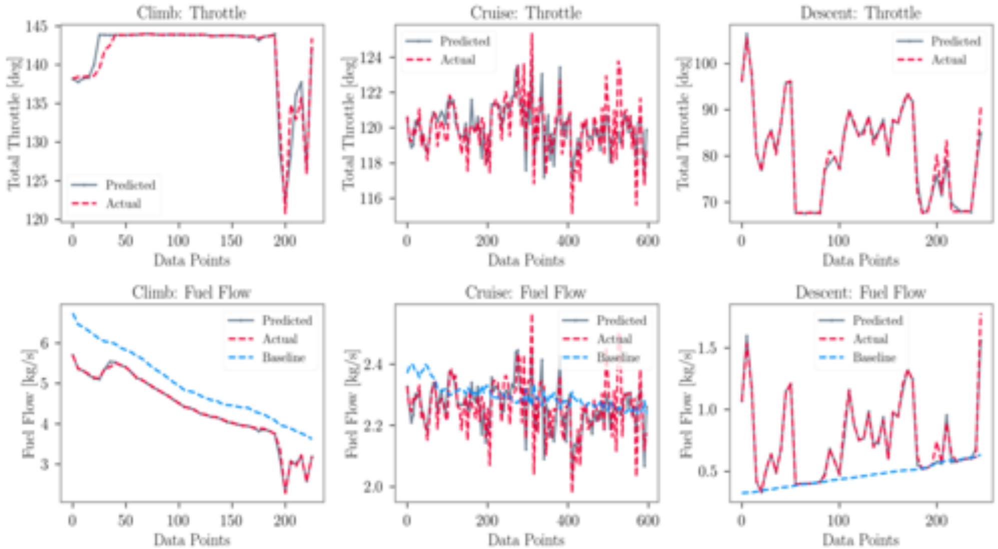

- Pressure ratio, temperature ratio, Mach number and aircraft mass are the baseline features that BADA, BPS and PEP use to calculate fuel flow.
- Deep learning techniques are utilized: Mini-batch, Yogi (another version of Adam optimization), L2 regularization.
- 98 tail-specific flights of a B777-300ER. 100k points for climb, 2M points for cruise, 150k points for descent.

AI Based Methodologies with Dynamic V&V Towards Fuel Efficiency

- **Aircraft:** B737, B777, B787
- **Data:** QAR data of 10,000+ flights.
- **Methodology:** Develop Deep Neural Networks to estimate fuel flow as a function of:
    - **Altitude**
    - **Mass**
    - **Temperature**
    - **ISA Deviation**
    - **Mach**
- **Evaluation:**
    - Compare the estimated fuel flow with the actual one, on unseen flights.
    - Benchmark with other aircraft performance models.

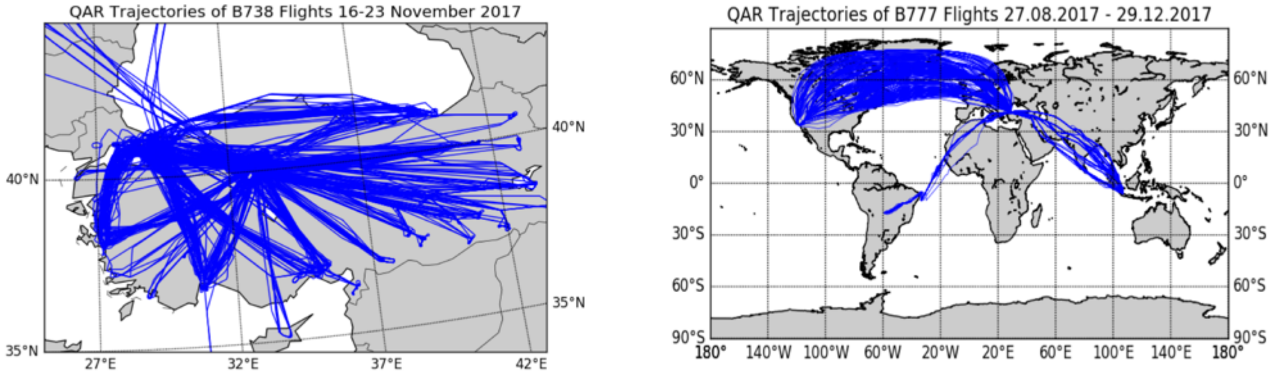

Short and long haul trajectories

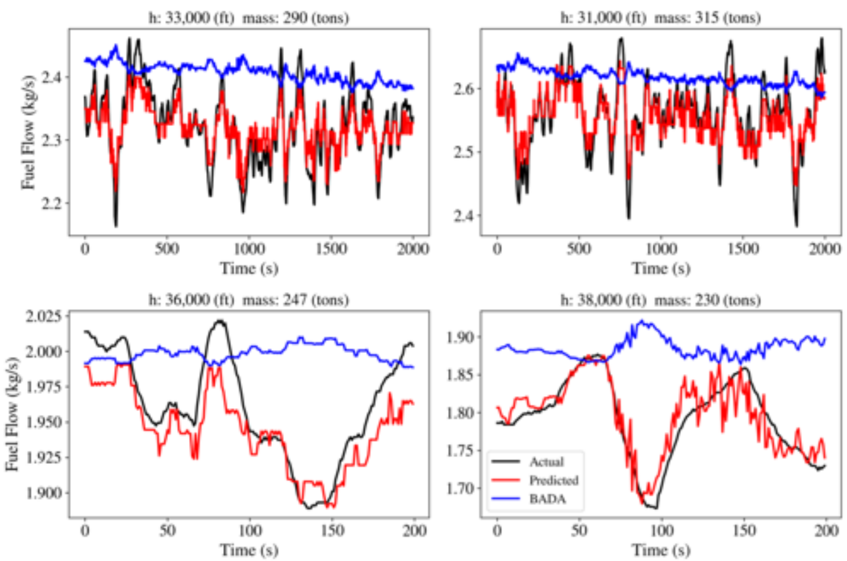

	B738W (3300 flights)		B773 (100 flights)	
	MAE (kg/h)	MAPE %	MAE (kg/h)	MAPE %
BADA	162. 78	6.99	289.11	3.75
INFLT	85.11	3.62	216.41	2.78
PF Update	56.79	2.45	222.95	2.95
AI	**52.46**	**2.27**	**137.15**	**1.46**

Application of Data-driven Models

- The updated baseline perfromance model is applied to the flight planning.
- Historical flight plans are re-generated using the update model as fuel burn estimator.

OPERATIONAL FLIGHT PLAN PAGE 4/11 RLSD 12APR18 0609.48Z

AWY MOCA	WPT FRQ NAME/FIR LAT/LONG	FL TRO SHR	MT TT VAR	WIND SAT TDV	TAS MN G/S	DIST REMD ACCD	TIME ACCT REMT	ETA ATA REV	RQRD ACCF FOB
UL602 040	**FIR EDVV** HANNOVER UIR FIR N51200E009076	320 347 02	299 302 03E	172/058 M53 M5	478 .827 511	24 3453 1076	3 0218 0719		62877 23919
UL602 038	**FIR EDGG** LANGEN FIR N51265E008505	320 329 05	299 301 02E	168/076 M54 M6	478 .829 514	13 3440 1089	2 0220 0717		62674 24122
UL602 040	HMM D115.65 HAMM N51514E007425	320 329 05	299 301 02E	168/075 M54 M6	478 .829 525	49 3391 1138	5 0225 0712		61910 24886
UL602 021	REBGU N52044E006582	320 328 01	295 296 01E	136/041 M54 M6	478 .829 516	30 3361 1168	4 0229 0708		61436 25360
UL602 020	**RELBI** EHAA AMSTERDAM FIR N52071E006488	320 328 01	293 295 02E	136/041 M54 M6	478 .829 515	6 3355 1174	0 0229 0708		61341 25455

	Count	Max difference [kg]	Mean difference [kg]
Under burn	163	3738	1039
Extra burn	75	3085	854

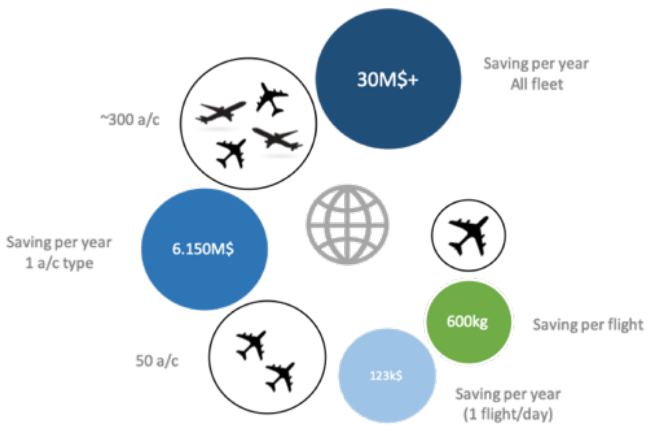

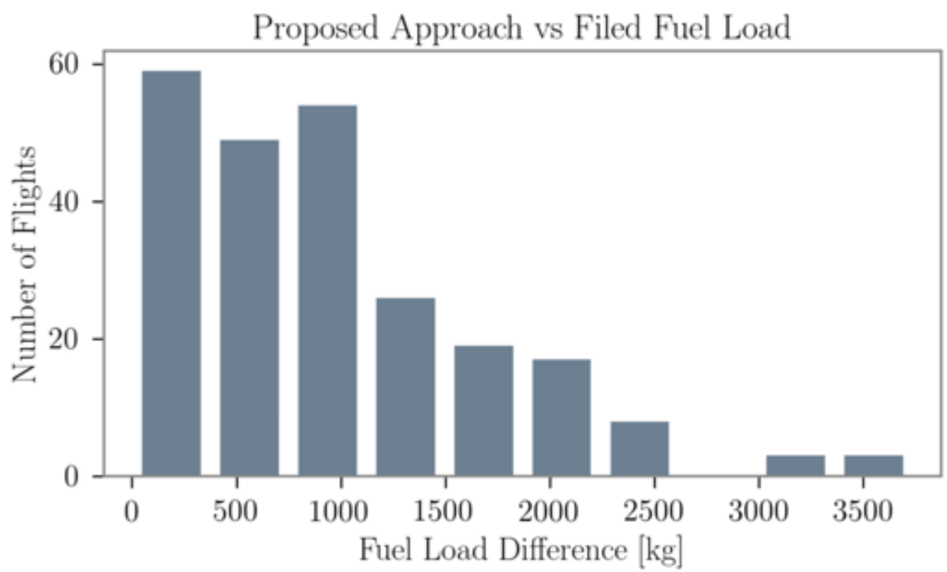

Aircraft Performance and Wind Calibration Scheme

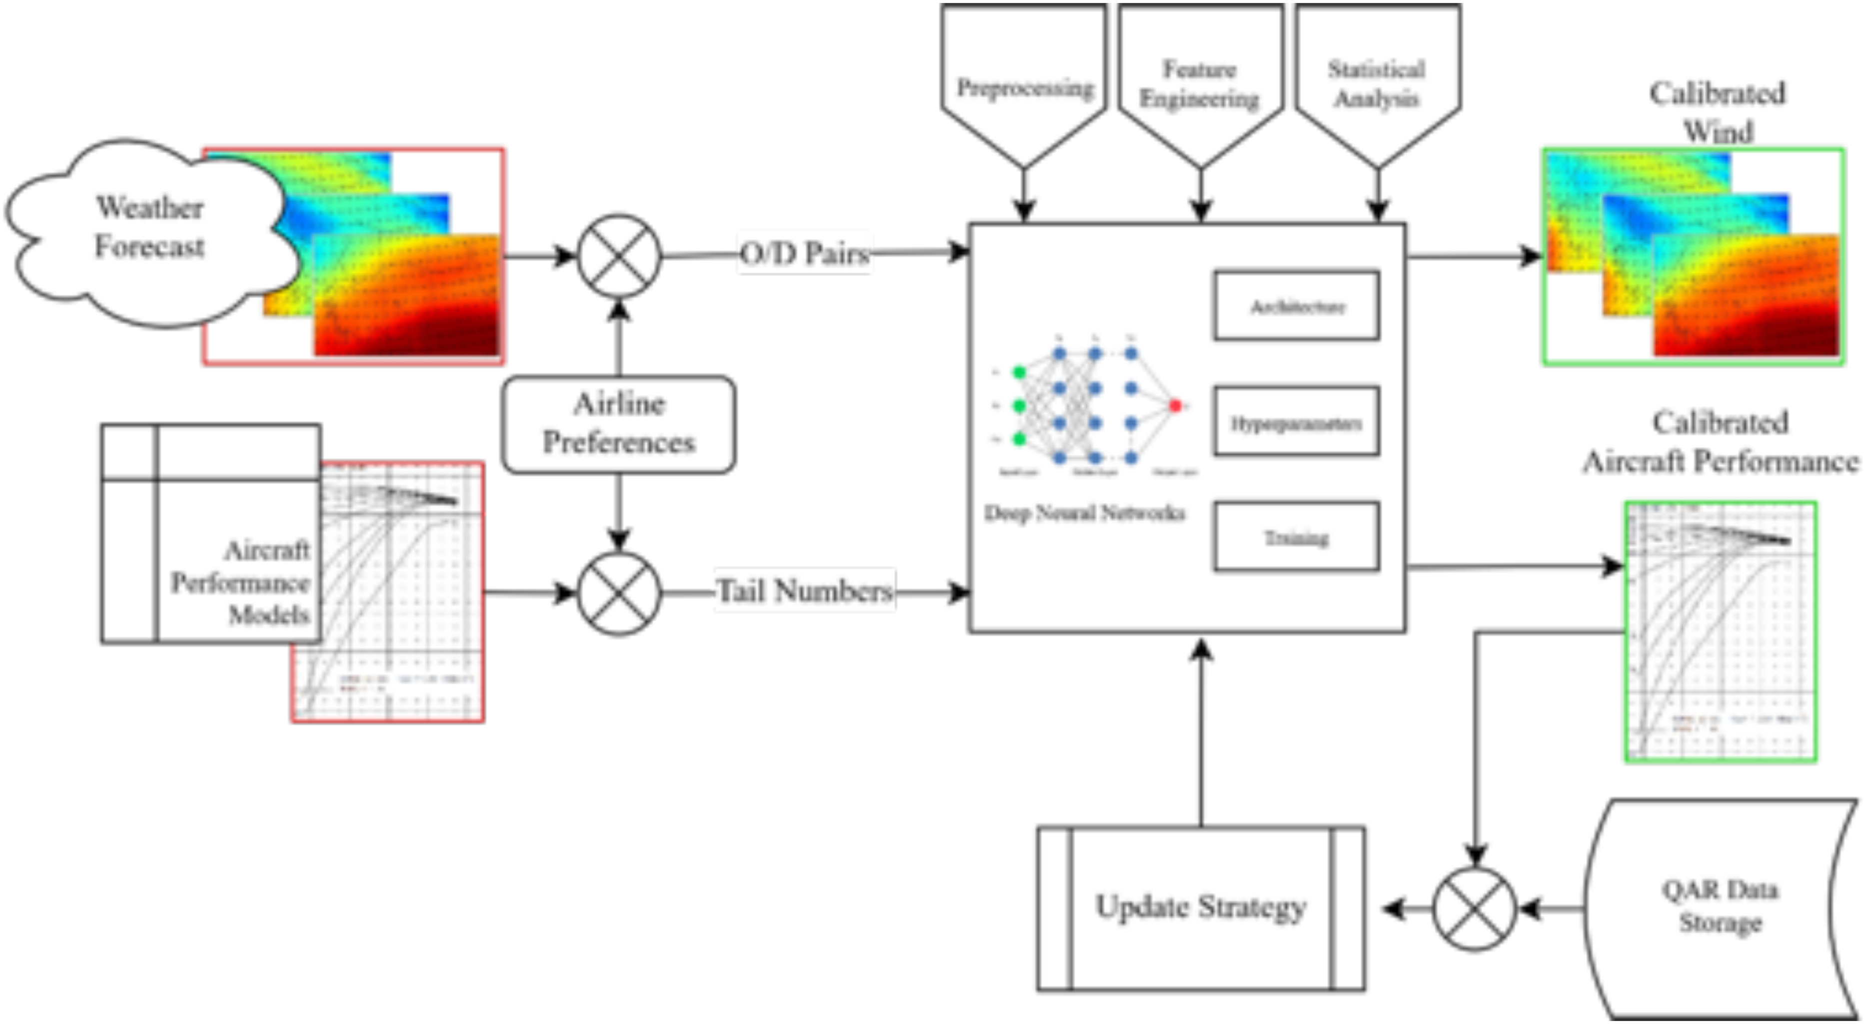

Pros and Cons

- What has been achieved:
  - State-of-the-art deep learning techniques are good at approximating non-linear mappings given a proper dataset.
  - Our fuel flow estimator represents the data quite well.
  - The models is applicable to flight planning.
- Drawbacks of ML:
  - The model «*naturally overfits*» to the data.
  - The model works fine at the seen flight regimes. What would be the fuel flow in flight conditions that are not in the data?
  - Having data from these regions would be ok, but it limits the applicability. How can we solve this without data?

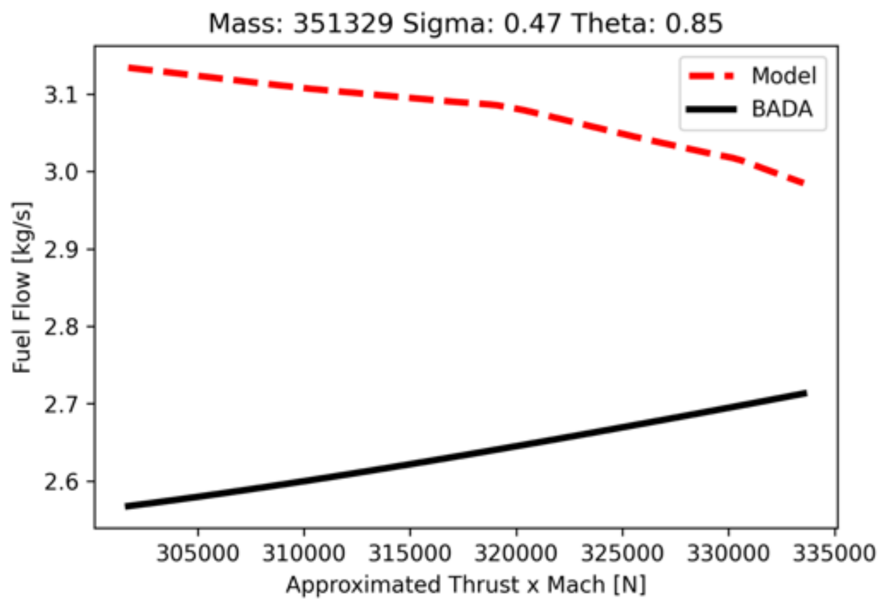

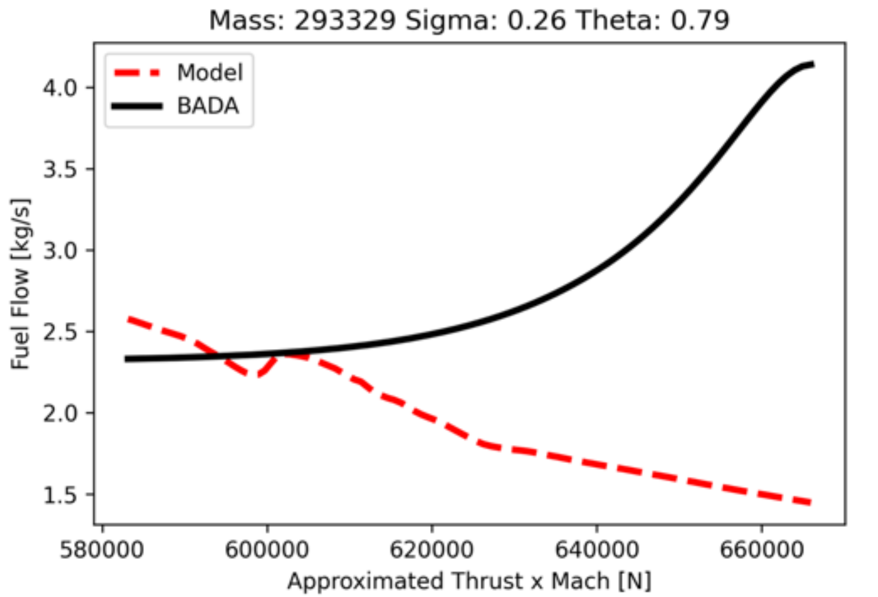

Physics-guided Neural Networks

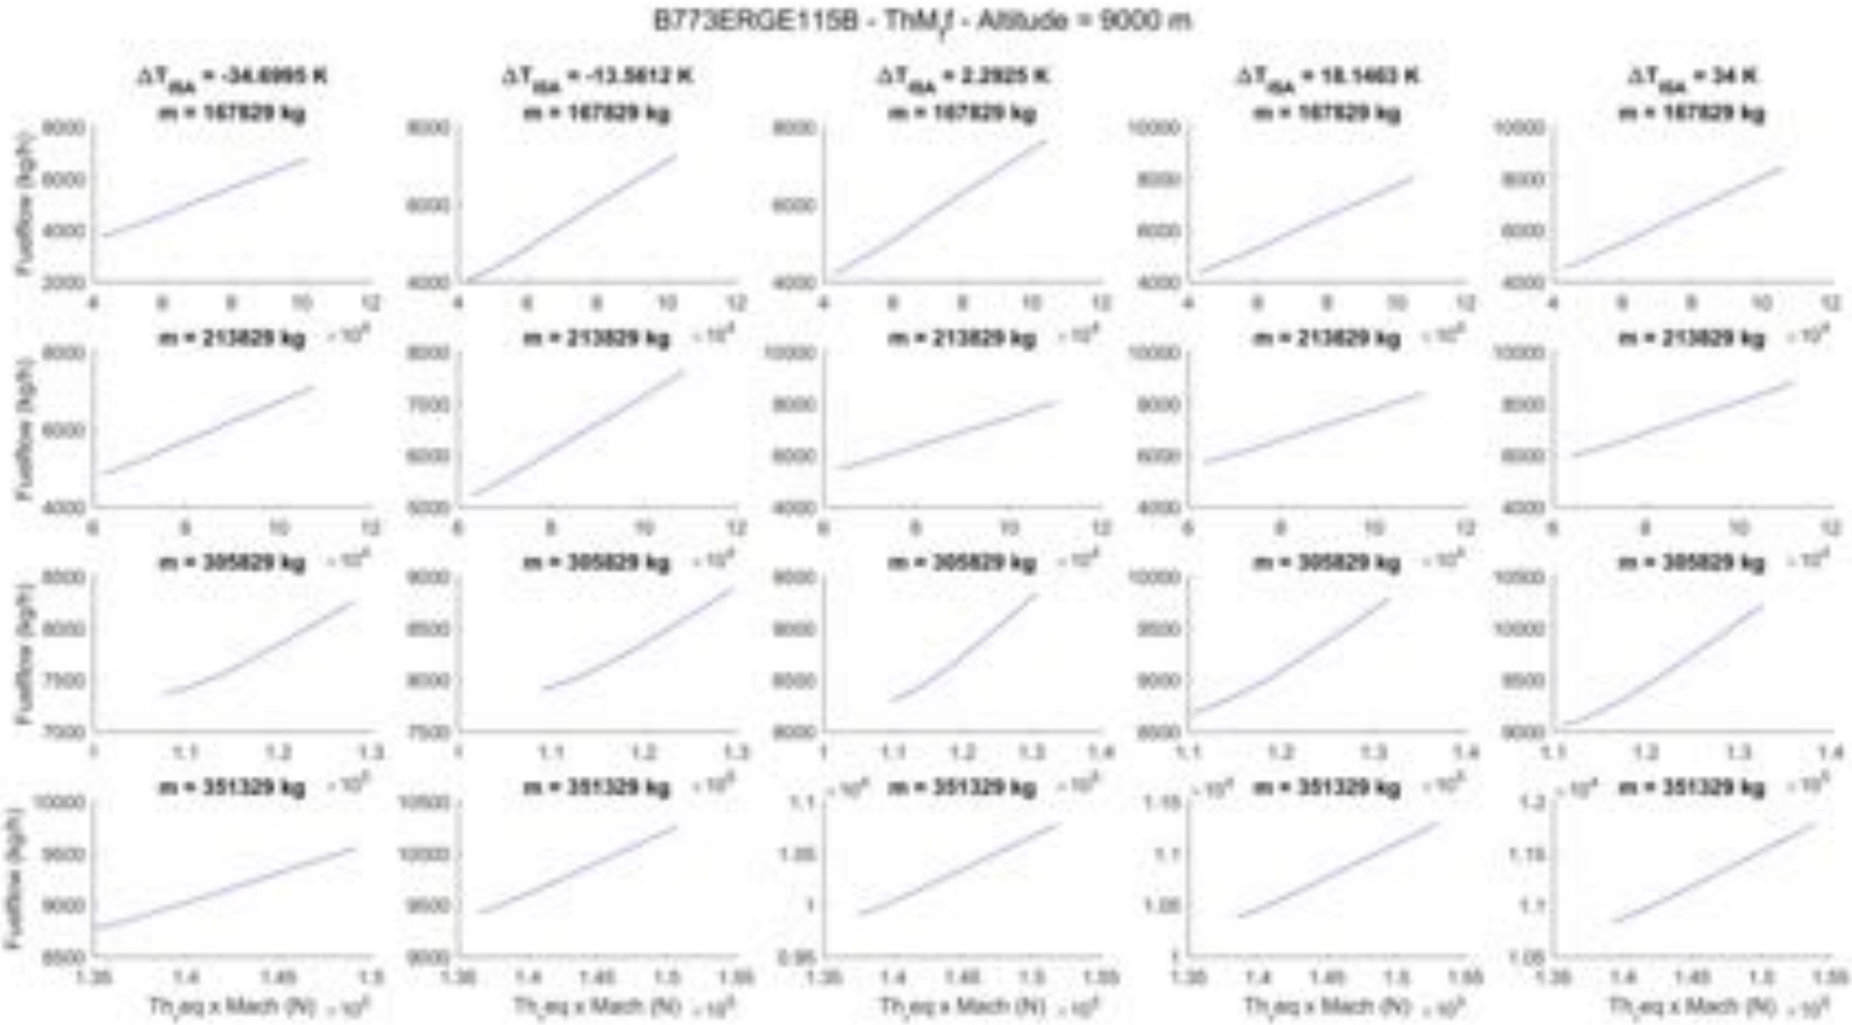

These plots are outpus of Boeing Performance Software for cruise flight

Physics-guided Neural Networks

- The labeled data do not cover the complete envelope.
- Include a physics based constraint to the optimization problem, so that the model also learns that physical intuition. It needs to be implementable to the loss function [1].
- In our case, the physical guidance for cruise flight is the following equation:

$$F \propto \frac{M}{\sqrt{\theta}}\left(a_1 M^2 + a_2 \frac{m^2}{M^2 \delta^2}\right)$$

- Which stands for that fuel flow is proportional to the thrust required multiplied by the Mach number. Thrust required is approximated through this equation.
- Any negative prediction of fuel flow is penalized.
- Final loss function is:

$$J = \lambda_1 MSE\left(y_{actual}, y_{pred}\right) + \lambda_2 J_{phy} + \lambda_3 J_{sign}$$

Uzun M, Demirezen MU, Inalhan G. Physics Guided Deep Learning for Data-Driven Aircraft Fuel Consumption Modeling. Aerospace. 2021; 8(2):44.

Physics-guided Neural Networks

- What difference does it make?

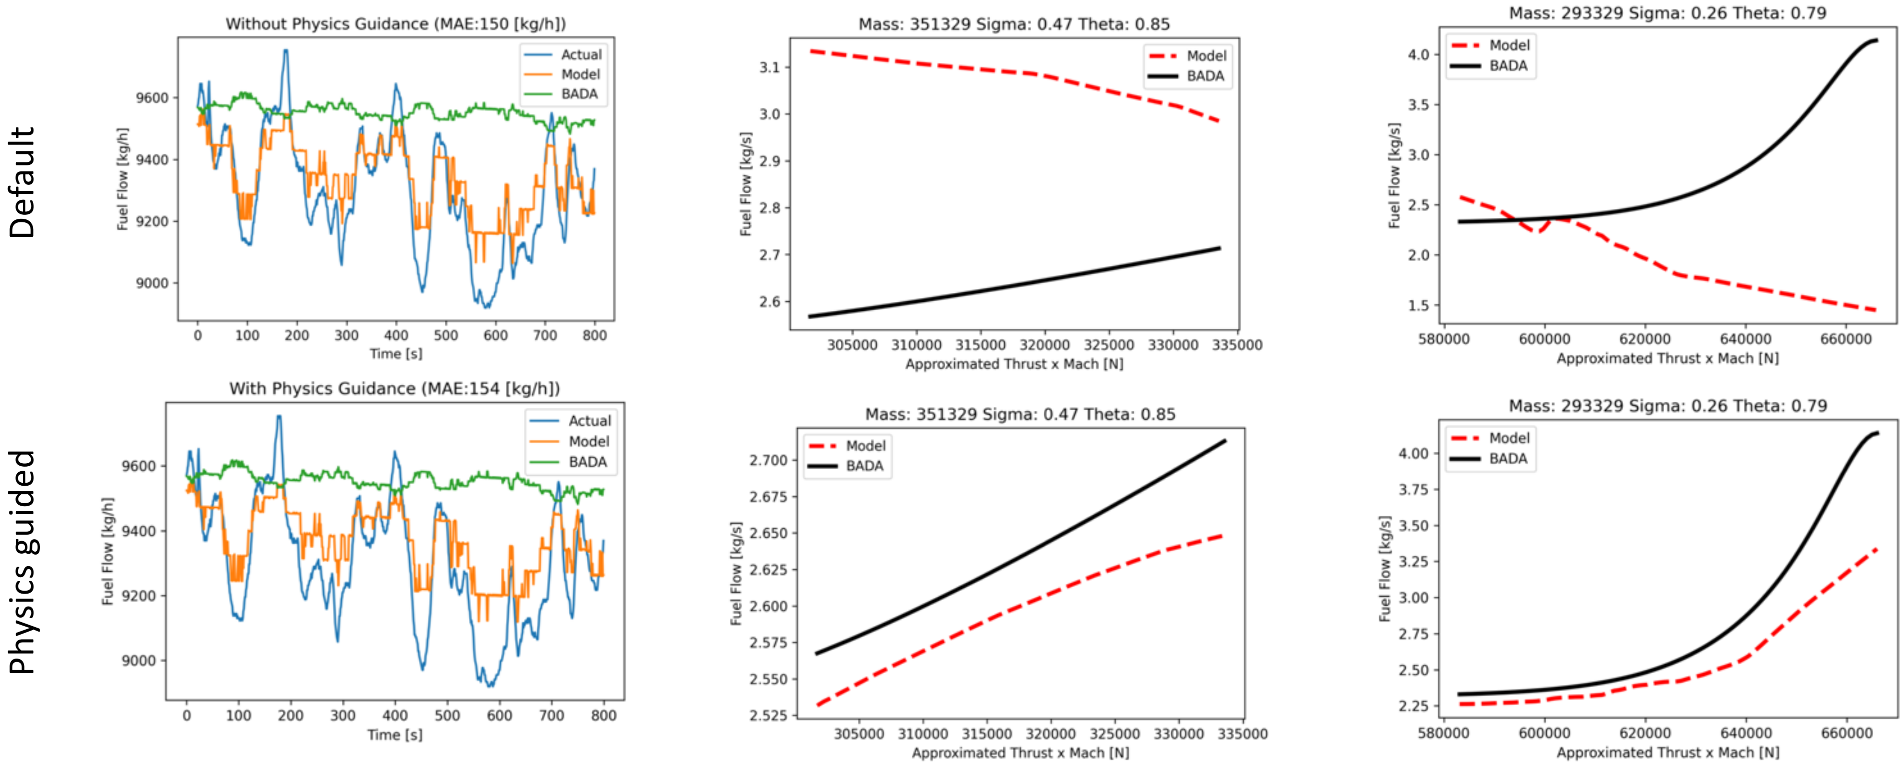

Next Steps….

- Aircraft performance calibration and events from surveillance data
  - Aircraft Health Monitoring
- Advanced flight planning
  - High precision integrated solution
  - Emission sensitive
  - Noise sensitive
- Advanced CCO/CDO
  - Noise
  - Fuel

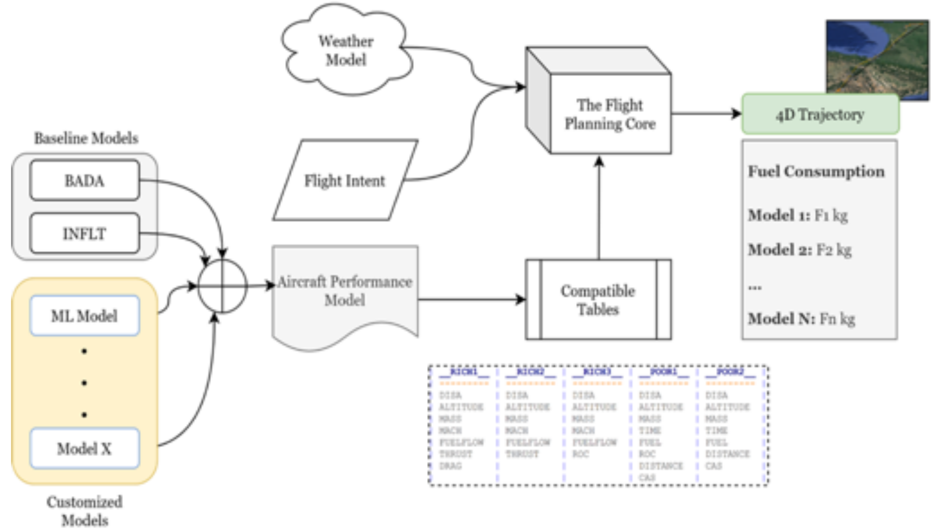

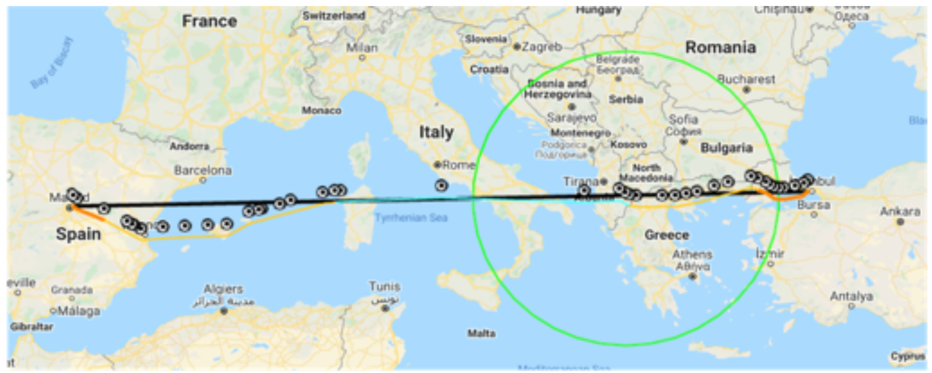

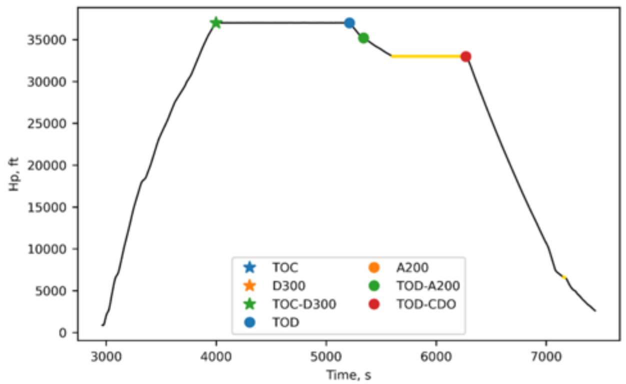

Further thanks to some key researchers @ Autonomy & AI Theme

- Dr. Burak Yüksek (TAS, GNC, AI)
- Dr. Mevlüt Uzun (AI, Future Air Mobility)
- Dr. Yan Xu (ATM/UTM)

Cranfield
Aerospace

www.cranfield.ac.uk