AI4GNC

Artificial intelligence techniques for GNC design, implementation, and verification

Final Review Meeting

December 12, 2022

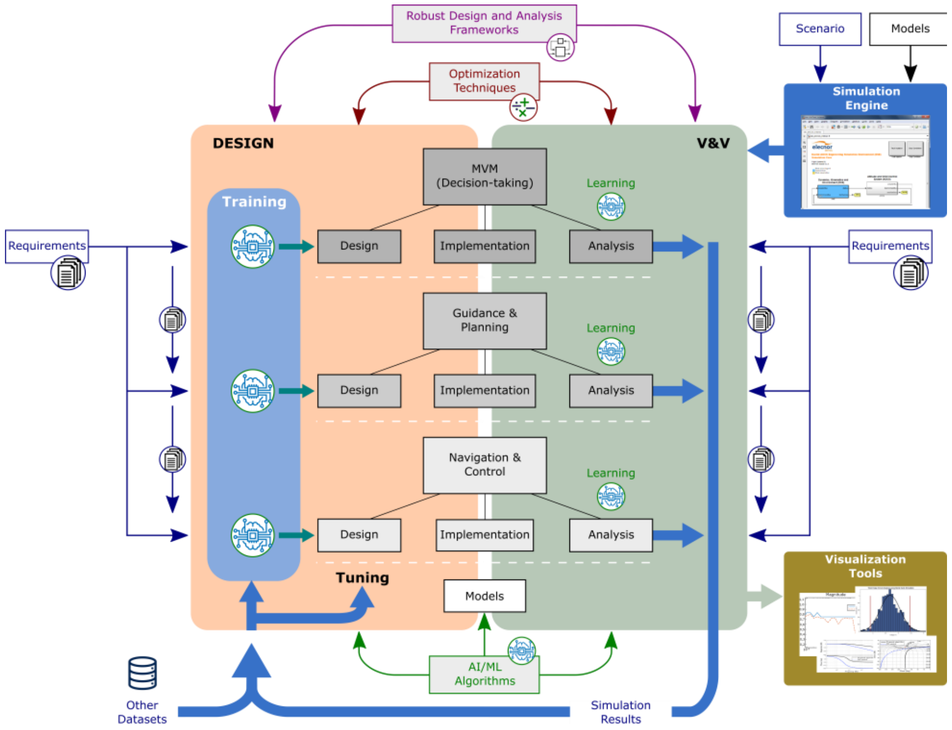

- Overview of the status of the project (10.00 – 10.15 CET)
- Assessment and Synthesis of the Results (10.15 – 10.30 CET)
- Way Forward – Maturation Plan & Roadmap (10.30 – 10.50 CET)
- Lessons Learnt and Recommendations (10.50 – 11.10 CET)
- Coffee break (11.10 – 11.30 CET)
- Discussions on RIDs (11.30 – 12.00 CET)
- Conclusions (12.00 – 12.15 CET)
- AoB (12.15 – 12.30 CET)

Scope and Objectives of the Activity

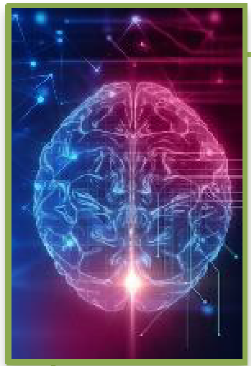

Goal 1: Implement ESA-iGNC, an AI-based GNC E2E design & analysis framework for layered architectures

- Cover the GNC system modeling, design and V&V process as per the SoW
- Supported by efficient optimization algorithms and formal mathematical techniques
- Ensuring robustness, performance, convergence, and explainable results

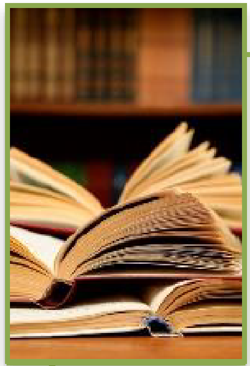

Goal 2: Exploit recent advances in control and AI

- Revisit the theory and techniques developed in the last two decades, including, but not limited to, fields such as IQCs, robust control, adaptive control, safe and robust reinforcement learning, and system identification
- Increase autonomy through onboard intelligence

Goal 3: Perform Trade-off analyses

- Different concepts to be considered, including full dedicated design architectures and augmentation strategies for already-existing control architectures
- Trade-off the offline design effort with the online real-time implementation requirements

Goal 4: Evaluate the proposed AI-based GNC design and V&V tool in a representative benchmark

- Define the criteria to sleect the benchmark
- Derive study cases and apply the tool to those
- Apply the tool to the benchmark

Overall of the Status of the Project

Planning
Schedule

Predecessors	Task Name
	AI4GNC - Artificial Intelligence Techniques for GNC Design,
52	**WP 1000 Literature Review**
	WP 1100 Safe, Robust, Adaptive Ctrl, and Reinf Learn
	WP 1200 Robust/ML Modeling and Analytical V&V Approac...
	WP 1300 ML Techniques for GNC Design
	WP 1400 ML Techniques for Embedded GNC Systems
	WP 1500 Form. Guarantees for Learning-based GNC
1	**WP 2000 Mission Definition and Req.s Capture**
	WP 2100 Development of the Design Framework
	WP 2200 Development of the Simulation Framework
	WP 2300 Mission Definition and Unitary Studies
8SS+10d	WP 2400 Requirements Definition and DDV Plan
53	**WP 3000 Comparative & Trade-off Analyses**
	WP 3100 Data-Driven vs. Model-Based Approaches
	WP 3200 Trade-off of Optimization Approaches
	WP 3300 Learning Approaches Trade-off
	WP 3400 Trade-off on Modeling & Learning Approaches
55	**WP 4000 Preliminary Design**
	WP 4100 MVM and Guidance
	WP 4200 AI-based Navigation and SysID
	WP 4300 AI-based Control Design
	WP 4400 Robust/ML Modeling and V&V Techs
	WP 4500 Support to Control Design
	WP 4600 Req.s, Models, GNC Des., V&V Fram.
54	**WP 5000 Benchmark and Study Cases Definition**
	WP 5100 Benchmark and Study Cases Definition
27	WP 5200 GNC Design for the Benchmark: Iteration 1
25	WP 5300 Study Cases Consolidation and Test Plan
26SS	WP 5400 Support to GNC Design
26SS	WP 5500 Support to GNC Design
56,57FS-1...	**WP 6000 Detailed Design**
	WP 6100 ML-based GNC Design & Verif. Fwork Imp.
	WP 6200 ML Detailed Design and Implementation
	WP 6300 Simulation Framework Implementation
	WP 6400 Robust/ML Modeling and V&V Techs
	WP 6500 Support to GNC Design and Verification
	WP 6600 Support to GNC Design and Verification
58	**WP 7000 Verification & Validation**
	WP 7100 GNC Design for the Benchmark: Iteration 2
38SS+30d	WP 7200 V&V Applied to the Benchmark
39SS	WP 7300 Support to V&V
59	**WP 8000 Study Synthesis**
	WP 8100 Assessment and Synthesis of the Results
42	WP 8200 Roadmap for Future R&D & App.
	WP 8300 Support to Study Synthesis - TASC
	WP 8400 Support to Study Synthesis - LUND
52	**WP 9000 Management**
	WP 9100 Management - DME
	WP 9200 Management - TASC
	WP 9300 Management - LUND
	WP 9400 Management - INESC-ID
	WP 9500 Management - DES
	KO
1,7	PRR Meeting
53	PRR Closeout
12	CRR
17,24	PDR
56	PDR Closeout
30	DDR
37	VR
41	FR

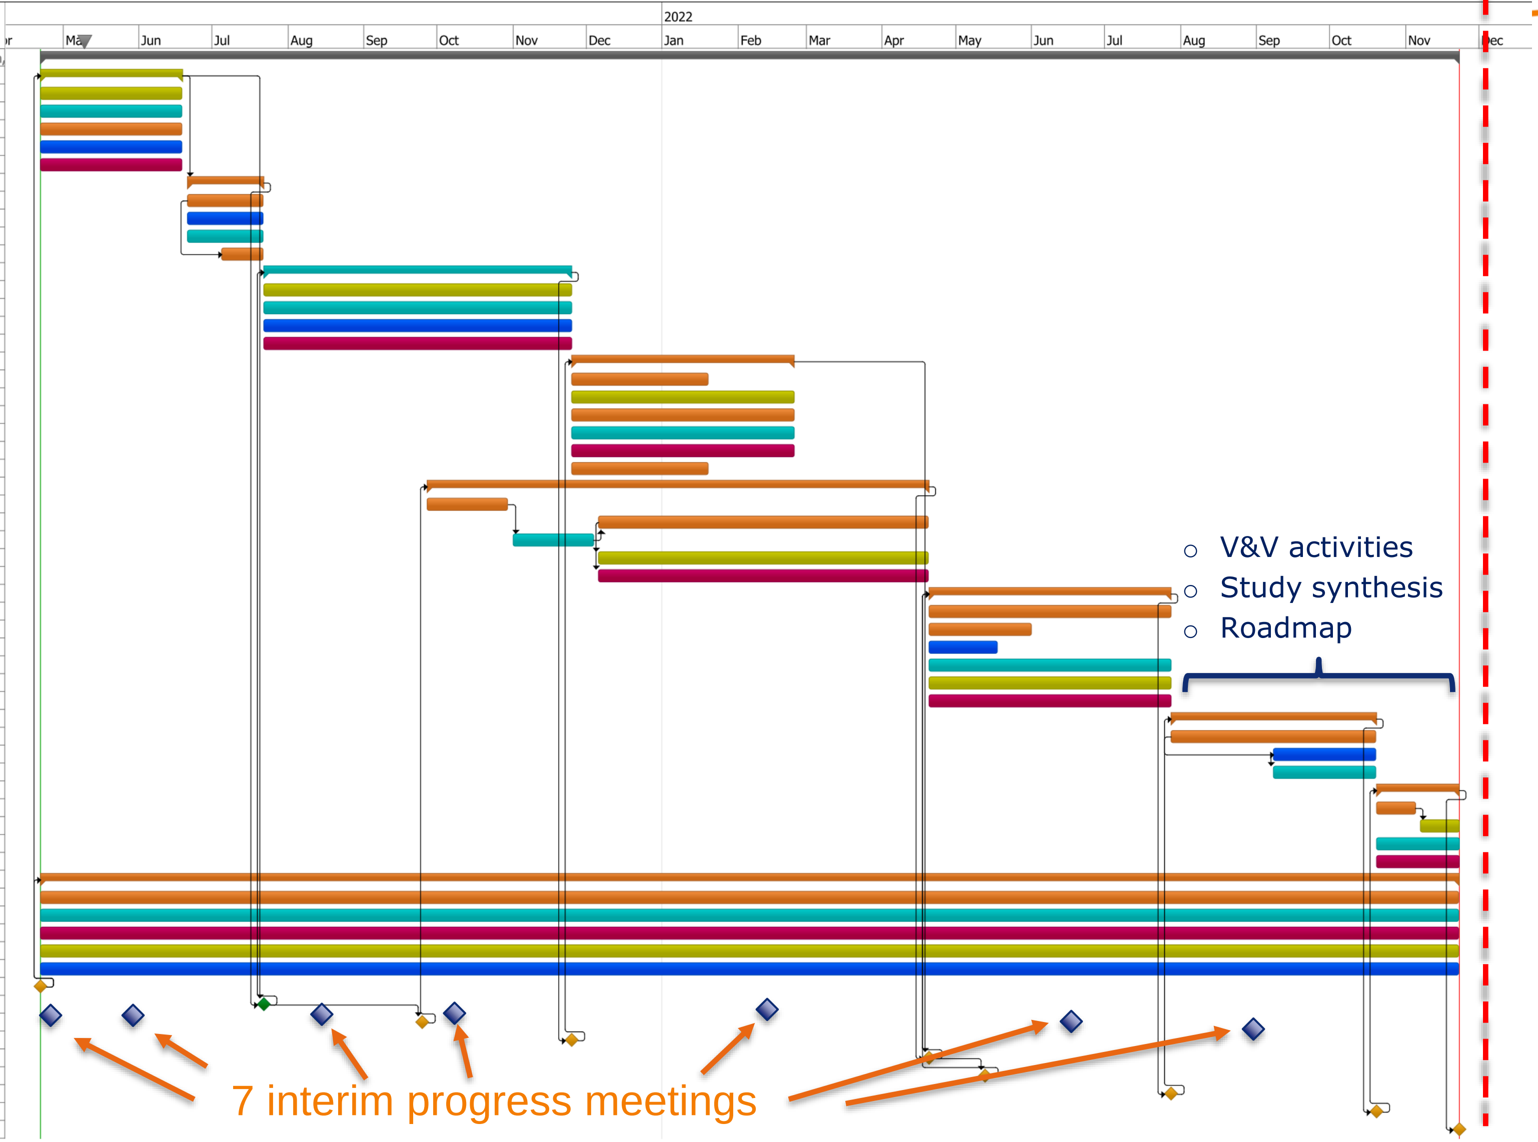

Project Status

Recalling the Rationale Behind the Work Logic

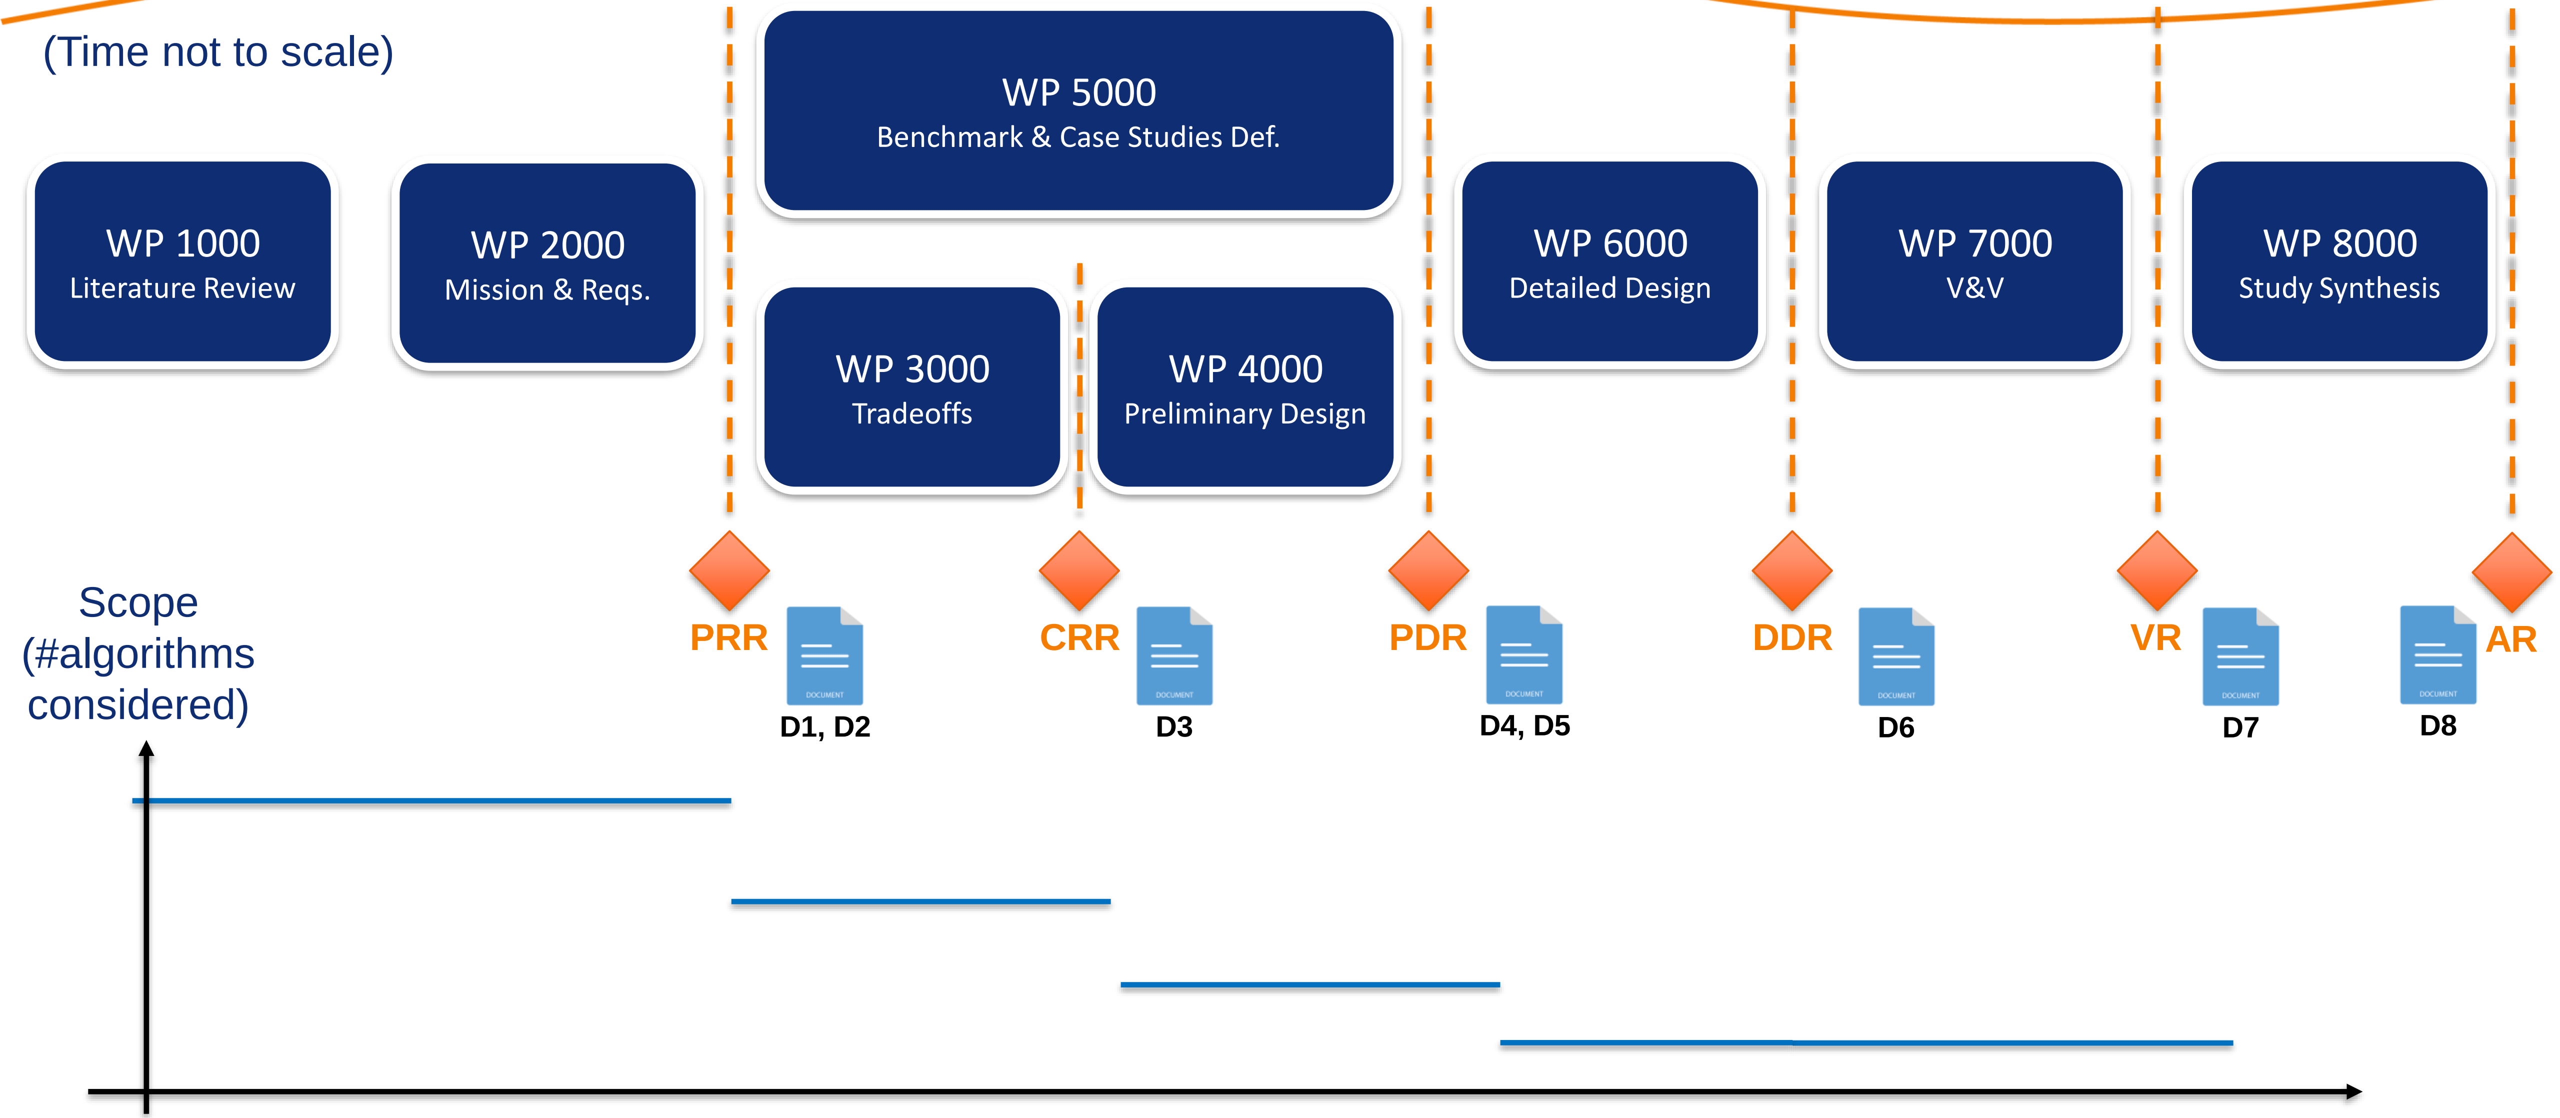

Project Status

Recalling the Rationale Behind the Work Logic (cont.'d)

WP 1000
Literature Review

WP 2000
Mission & Reqs.

WP 5000
Benchmark & Case Studies Def.

WP 3000
Tradeoffs

WP 4000
Preliminary Design

WP 6000
Detailed Design

WP 7000
V&V

WP 8000
Study Synthesis

PRR | CRR | PDR | DDR | VR | AR

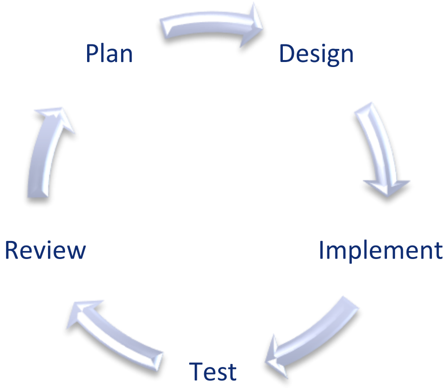

Preliminary design:
- Case studies based on the benchmark
- Design and evaluation performed for nominal conditions and simplified scenarios

1st detailed design loop:
- Design and evaluation performed also for off-nominal conditions
- Preliminary V&V
- All running with ESA-i4GNC

2nd detailed design loop and V&V:
- Design consolidation
- V&V
- All running with ESA-i4GNC

Project Status
Work Packages

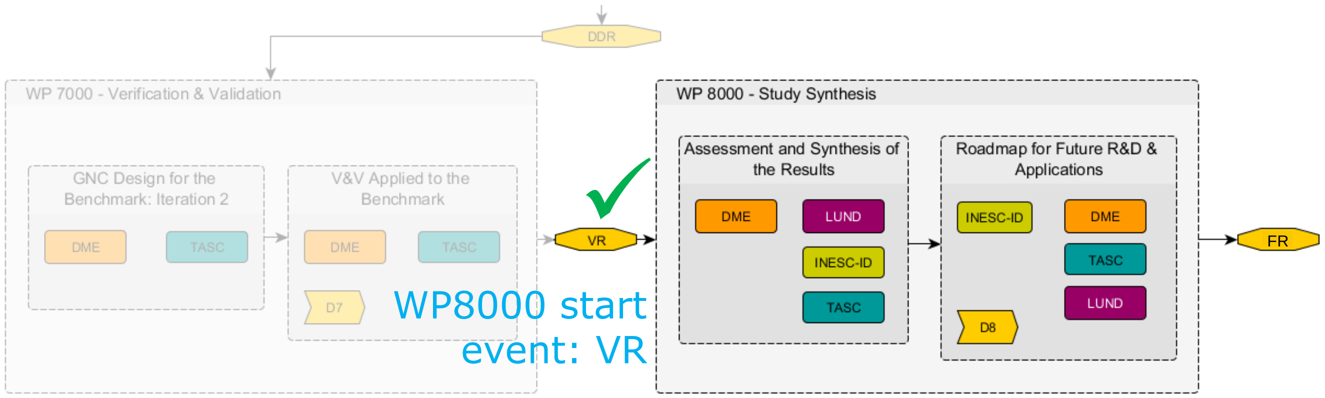

WP8000 end event: FR

Outputs:
D8 deliverable

- **WP8000: Study Synthesis**
  - Assessment and Synthesis of the Results
  - Roadmap for Future R&D & Applications

- **Items delivered at FR**
  - **D8**: Study Synthesis and Way Forward, Maturation Plan
  - **Technical Data Package**
  - **Brochure**
  - **Abstract**
  - **Technology Achievement Summary**
  - **Summary Report**
  - **Executive Summary Report**
  - **Final Report**

❑ Items delivered at FR

- **D8**: Study Synthesis and Way Forward, Maturation Plan
- Assessment and Synthesis of the Results
- Way Forward
- Roadmap
- Lessons Learnt and Recommendations
- Conclusions

- Trade-offs between techniques
- ESA-i4GNC overview
- Benchmark selection
- Case studies drivers and definition
- Techniques and Results

- Points to improve, how to improve, application scenarios and potential indicator per technique

- Roadmap summary

- Lessons Learnt throughout the project
- Recommendations

- Summary of the work done

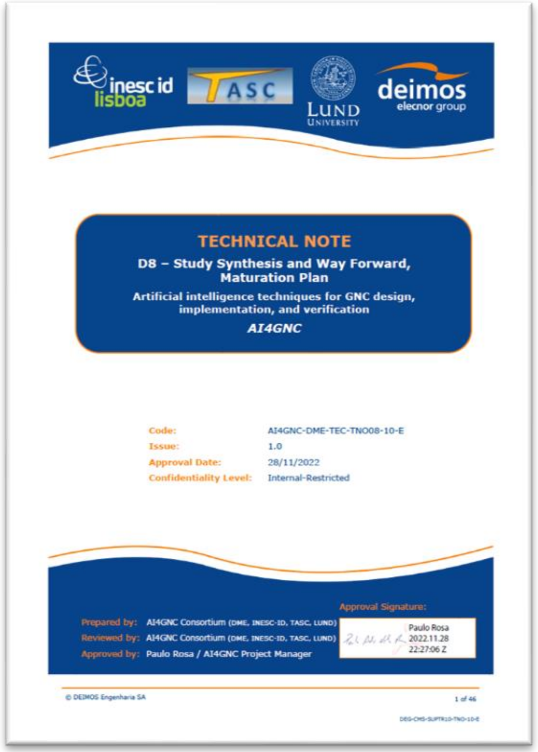
TECHNICAL NOTE

D8 – Study Synthesis and Way Forward, Maturation Plan

Artificial intelligence techniques for GNC design, implementation, and verification

AI4GNC

Code:	AI4GNC-DME-TEC-TN008-10-E
Issue:	1.0
Approval Date:	28/11/2022
Confidentiality Level:	Internal-Restricted

Approval Signature:

Prepared by: AI4GNC Consortium (DME, INESC-ID, TASC, LUND)
Reviewed by: AI4GNC Consortium (DME, INESC-ID, TASC, LUND)
Approved by: Paulo Rosa / AI4GNC Project Manager

Paulo Rosa
2022.11.28
22:27:06 Z

© DEIMOS Engenharia SA
1 of 46

Project Status
FR Deliverable Items

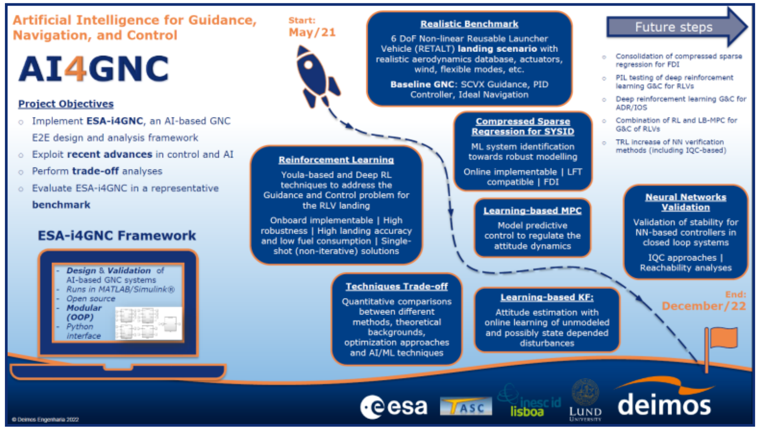

Actions from the last meetings

Action #	Action	Actionee	Due
VR_001	DEIMOS to send a Doodle poll for the selection of the date to discuss the potential special session at a conference	DEIMOS	Closed
VR_JB-01	Add ToF penalization in the reward function	DEIMOS	Closed
VR_JB-02	Assess the impact of penalizing negative thrust derivative profiles	DEIMOS	Closed
VR_JB-03	Summarize the pros and cons of the technique and indicate potential improvements as part of the roadmap definition in WP8000	TASC	Closed
VR_JB-08	Run MC for the linear case and compare it with the one with GP, and to increase the dispersion in the initial condition	LUND	Closed
VR_JB-09	Discuss the pros and cons of AI-based methods when benchmarked against heritage techniques as part of the outcomes of WP8000	LUND	Closed
VR_JB-12	Increment the verbosity of comments in the code	LUND	Closed
VR_JB-13	Add a paragraph in D8 on the trade-off of constraint in accuracy and the approach of Fazlyab; add discussion in D8 on the uncertainties and on why this is a strength of the IQC framework	LUND	Closed
VR_JB-14	Include a discussion in D8 on the use of IQCs for possibly other AI/GNC applications	DEIMOS/LUND	Closed

Project Status
Building Blocks Statuses

Benchmark selected and consolidated

Compressed sensing for LFTs

Successive convexification and ability to use many optimization toolboxes

RLV simulator implemented

Requirements available

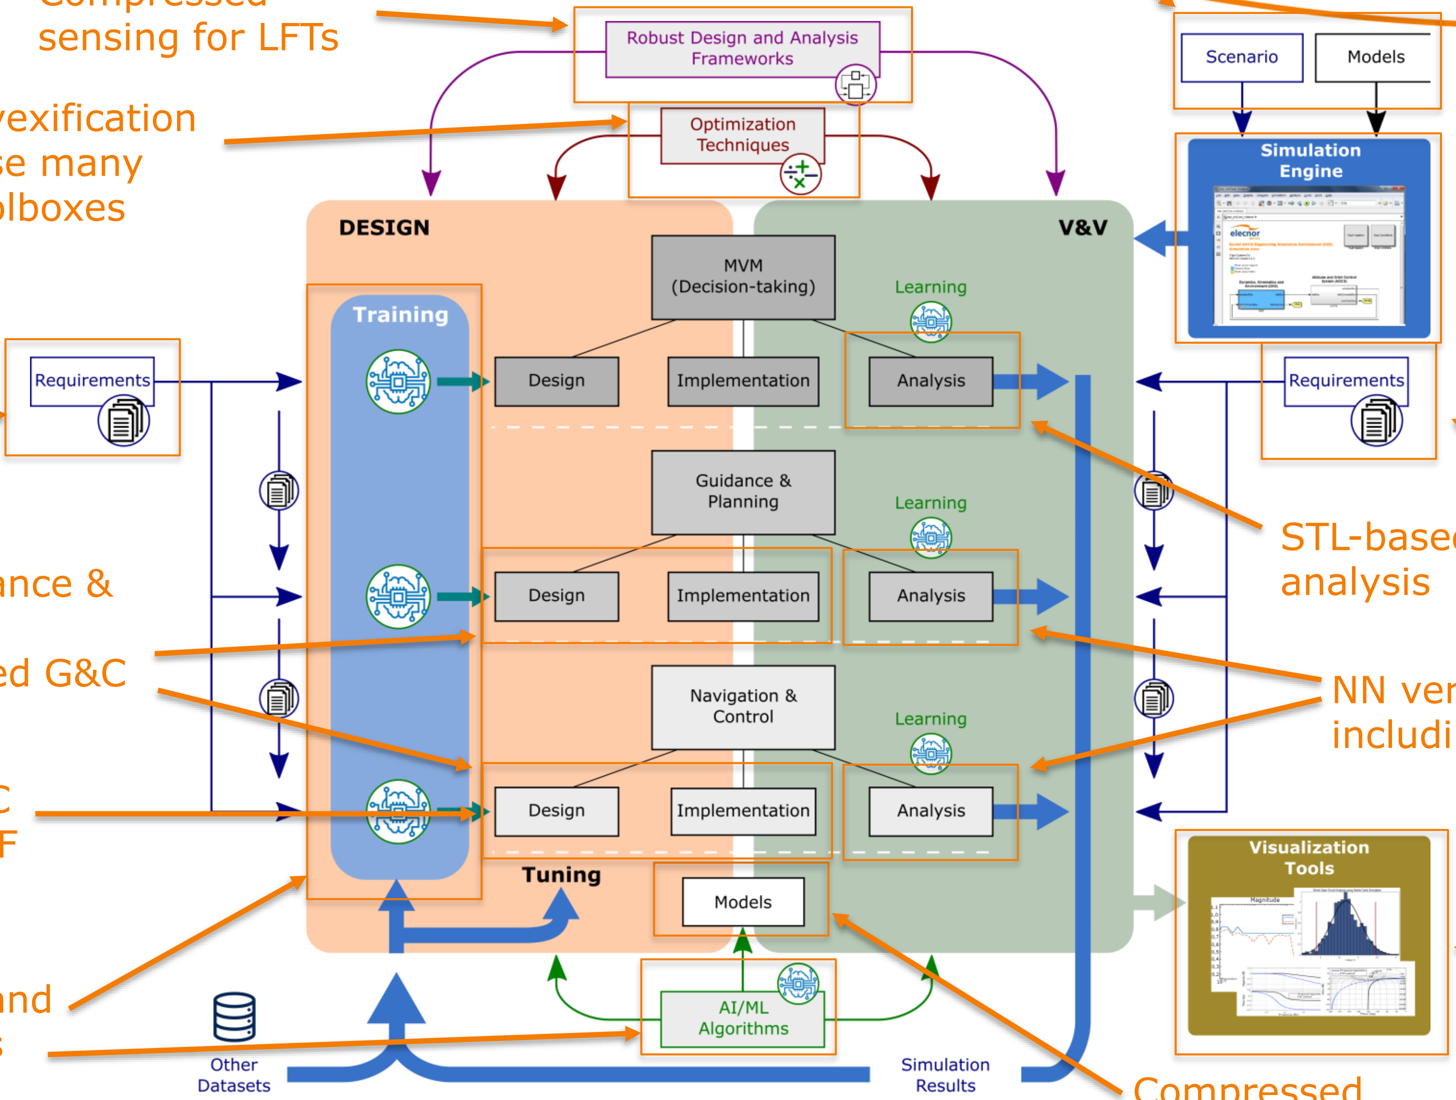

Requirements available

STL-based analysis

RL Guidance & deep-RL integrated G&C

NN verification implemented, including IQC-based

LBMPC & LBKF

Training process and AI/ML/RL methods implemented

Visualization tools for the simulator and the AI/ML/RL implemented

Compressed sensing for LFTs

Assessment and Synthesis of the Results

- Extensive state of the art literature review
  - Safe, robust, adaptive control and RL; Robust ML modelling and analytical V&V approaches; ML for GNC design and embedded GNC systems; formal guarantees for LB GNC.
- Qualitative trade-off between *data-driven* and *model-based* approaches.
- Quantitative trade-off between the most relevant *ML techniques for GNC applications*
  - Explainability, on-board implementability, training dataset need, performance, generalisability.
- Quantitative trade-off between the *ML toolboxes*
- Quantitative trade-off between *ML libraries*
  - Algorithm availability, documentation, code readability, support for embedded devices, supported environments, integration with MATLAB.

Assessment and Synthesis of the Results

ESA-i4GNC overview

- Tool implemented using the Object-Oriented Programming (OOP) approach
- Besides the tool design, other libraries and functionalities were implemented such as:
    - CVX
    - MPT 3.0
    - S-TaLiRo Runtime and Falsification
    - M2HTML
    - esai4gnc_install
    - esai4gnc_clean
    - autoDoc_ESAi4GNC
    - Profiler

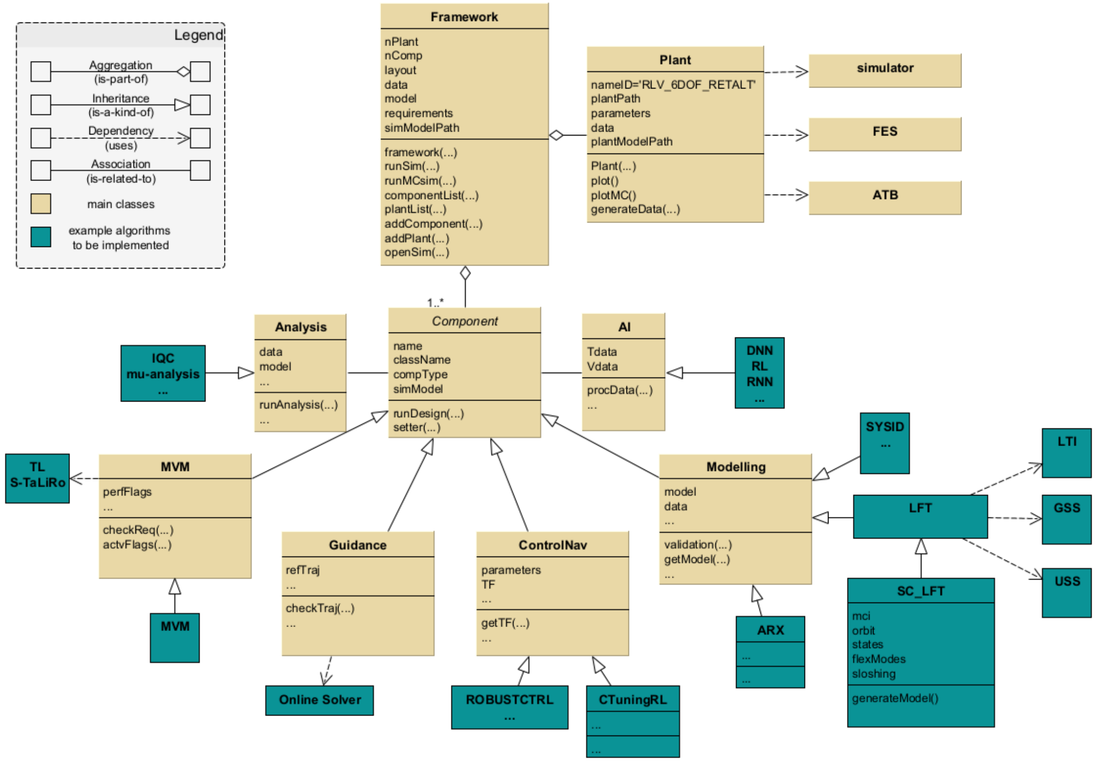

- 6 DOF **Landing Burn Scenario** of a Reusable Launcher Vehicle (**RETALT**)
  - Realistic Aerodynamics DB
  - Actuator (TVC) model
  - Wind model
  - Flexible modes

- **<u>Baseline GNC</u>:**
  - SCVX guidance
  - Ideal navigation
  - PID controller

Initial Conditions

Mass [kg]	Position [m]	Velocity [m/s]	Attitude [q]	Angular vel. [rad/s]
80334	[2874,-1288, -82.2]	[-189.9, 151.3, 9.6]	[0.943, 0.006, 0.018, -0.329]	[0, 0, 0]

1st Stage features	Value
Height [m]	71.2
Diameter [m]	6
Dry mass [kg]	59300
Propellant mass (incl. descent propellant) [kg]	621500
Specific Impulse SL [s]	401.6
Thrust SL [kN]	11453

RLVs Challenges

To Support the Case Studies Definition

ID	Challenge	Why is it a challenge?
1	**Trajectory generation for the landing phase**	Large uncertainty in the initial condition; wind disturbances; potentially large reference tracking errors
2	**Attitude control**	Large disturbances (wind, flexible modes, fuel sloshing); uncertain CoG (see Challenges #3 and #10)
3	**Center of gravity estimation**	Can be affected by uncertainty, but also by the payload; this means that the torque generated by actuators (and thruster misalignment) will be uncertain
4	**RCS failure detection**	Needs to be decoupled from other failures and from disturbances
5	**TVC failure detection**	Needs to be decoupled from other failures and from disturbances
6	**Aerodynamic fin failure detection**	Needs to be decoupled from other failures and from disturbances
7	**RCS failure tolerance**	Roll compensation, for instance, may become impossible; if not compensated for sufficiently fast, can jeopardize the mission
8	**TVC failure tolerance**	If not compensated for sufficiently fast, can jeopardize the mission; can generate very large spurious torques; may require non-nominal trajectory
9	**Aerodynamic fin failure tolerance**	Can generate constant torques; may render the system unstable during the unpowered phases
10	**MCI estimation**	More generic than Challenge #3; can be affected by uncertainty, but also by the payload, and by fuel sloshing
11	**Thruster misalignment**	Can impact control authority
12	**Main engine re-ignition**	When the vehicle is returning the fuel will be displaced in the tanks toward the top of the booster; a dedicated maneuver (with the thrusters) could be necessary during the return phase before the reignition of the main engine(s) to relocate the propellant inside the tanks
13	**Trajectory tracking for the aerodynamic phase**	Large uncertainty in the initial condition; wind disturbances; potentially large reference tracking errors
14	**Inertia estimation during the flip over before the boostback burn**	A quick and accurate flip over maneuver is necessary to limit the propellant consumption during the boostback burn, the inertia of the vehicle is affected by the position of the fuel in the tanks, and it can impact significantly the control performance during the flip over maneuver
15	**Fuel sloshing**	Fuel sloshing may induce disturbances that are hard to model in the linear realm
16	**RLV navigation**	The use of hybrid navigation schemes can reduce costs (as lower accuracy IMUs can be used, for instance, if combined with GNSS measurements), although they may also pose assumptions on the vehicle or the sensors’ characteristics that may not be satisfied in practice.

Case Studies
Definition of Baseline Case Studies

Case Study #	Description	Challenge(s) Addressed
1	RL-based adaptive control to regulate the attitude in response to disturbances	2
2	RL-based adaptive control to regulate the trajectory in the aerodynamic phase with respect to the reference trajectory	13
3	NN approximation of the QUEST algorithm for three axis attitude estimation from vector observation data	16
4	Sparse regression, compressed sensing, compressed learning and potential connections with LFT modelling	3, 4, 5, 6, 10, 11, 14
5	Learning Based Model Predictive Controller (LB-MPC) for attitude control	2
6	IQC formalism for NN-based attitude control verification	2
7	Learning-based Kalman filtering with Gaussian process for attitude estimation	16
8	Deep RL for trajectory tracking (integrated G&C)	13

Proposed GNC Architecture

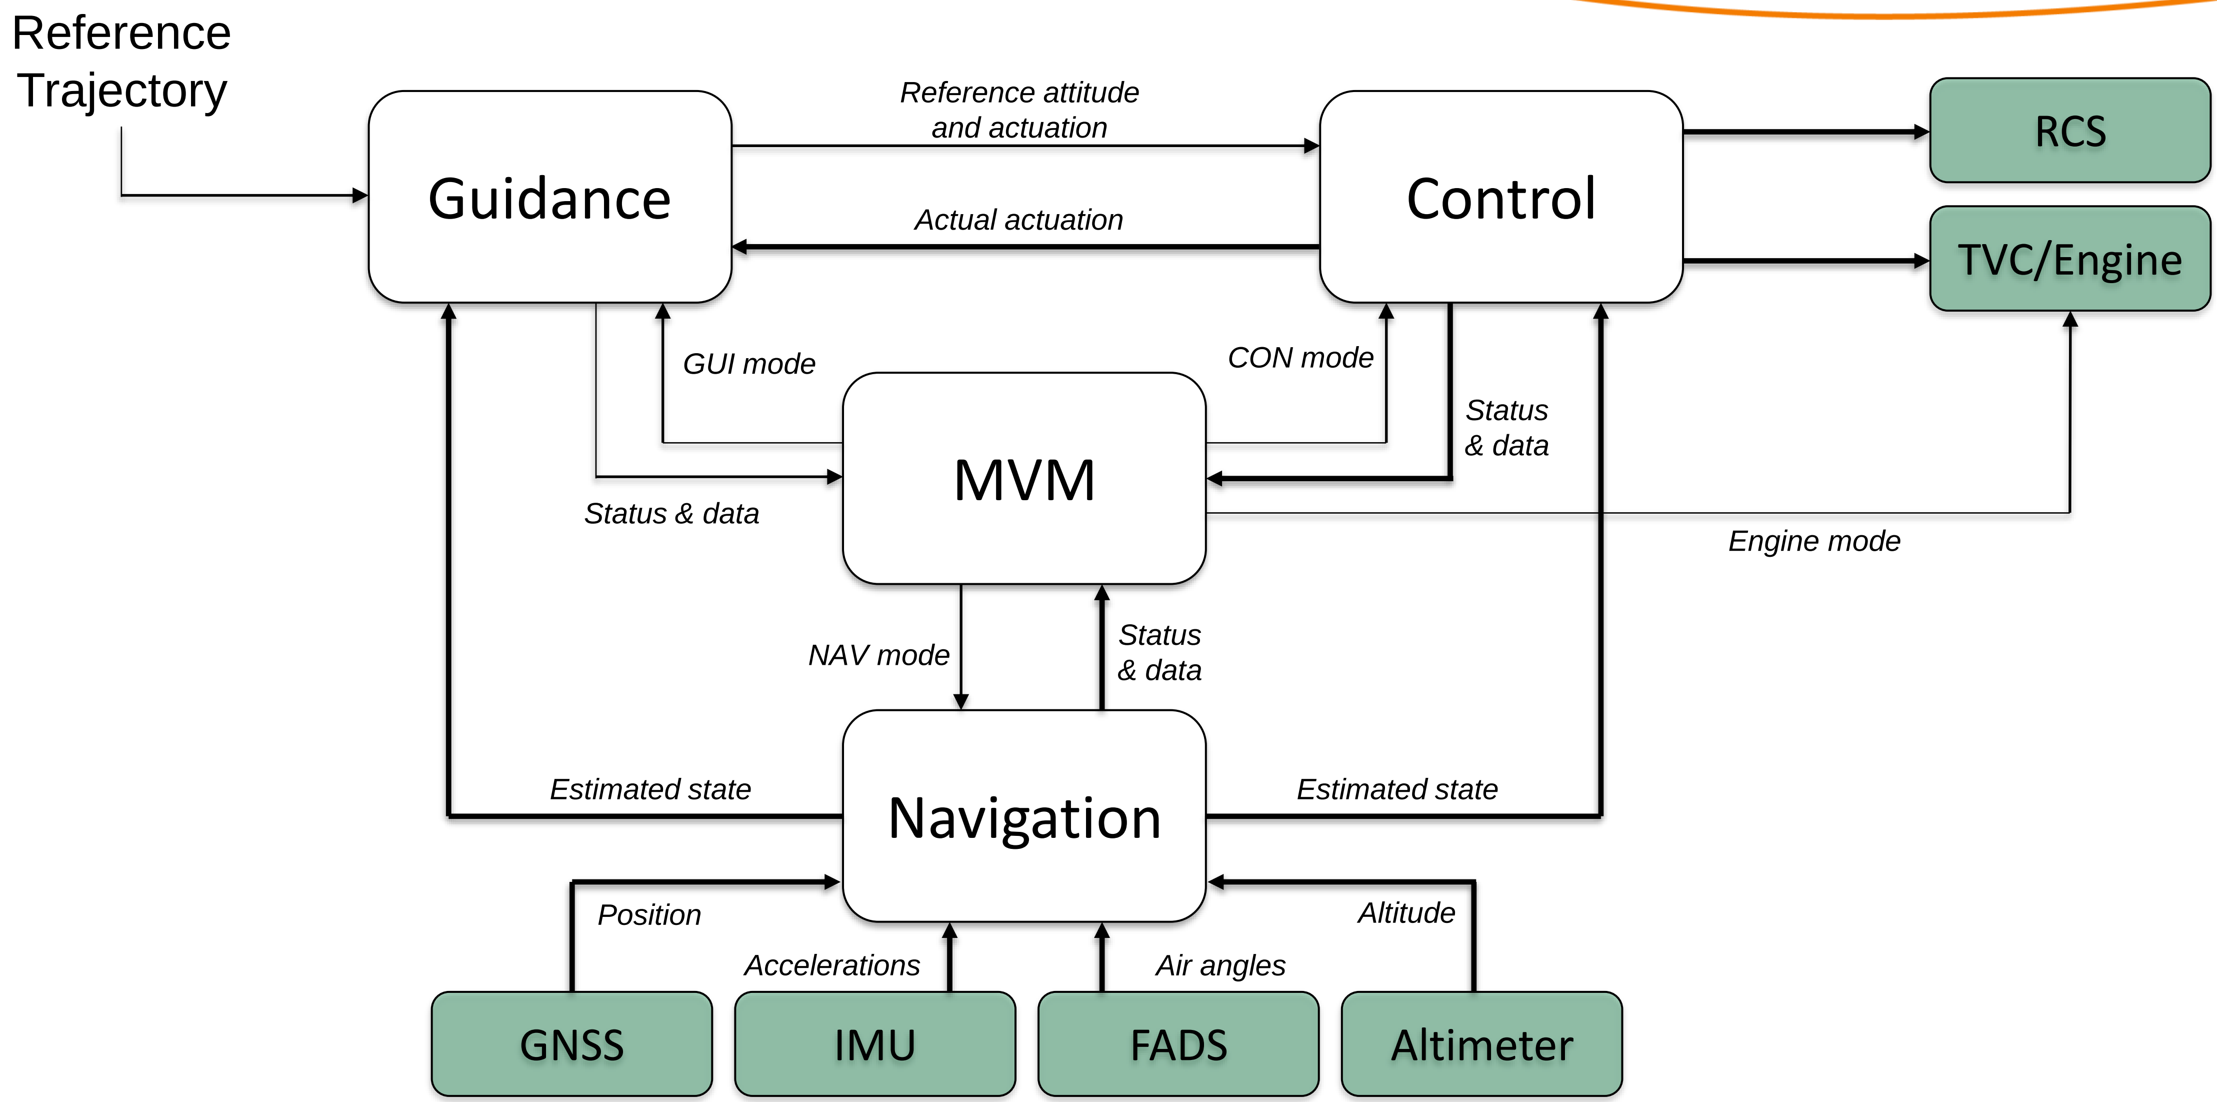

Assessment and Synthesis of the Results
Techniques and Results

- **Case study #1 and #2** are concerned with **Reinforcement Learning adaptive control** algorithms developed on the attitude and position tracking problems

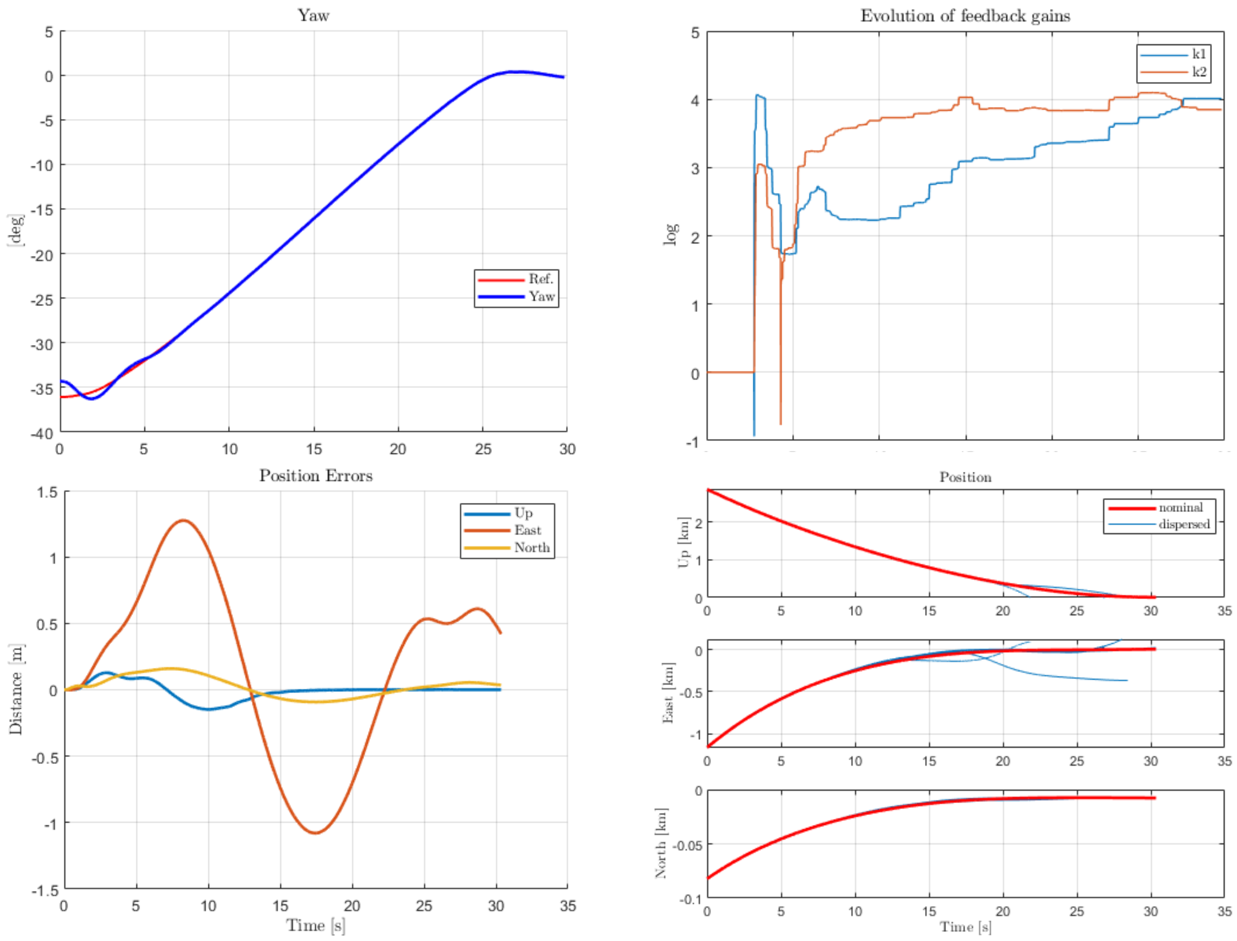

- **Case study #3** deals with **attitude estimation** (Wahba's problem) from vector sensor observations by using a Deep NN
- Mainly a tutorial case study, that reduces the computational power on-board, while being less sensitive to noise than traditional algorithms

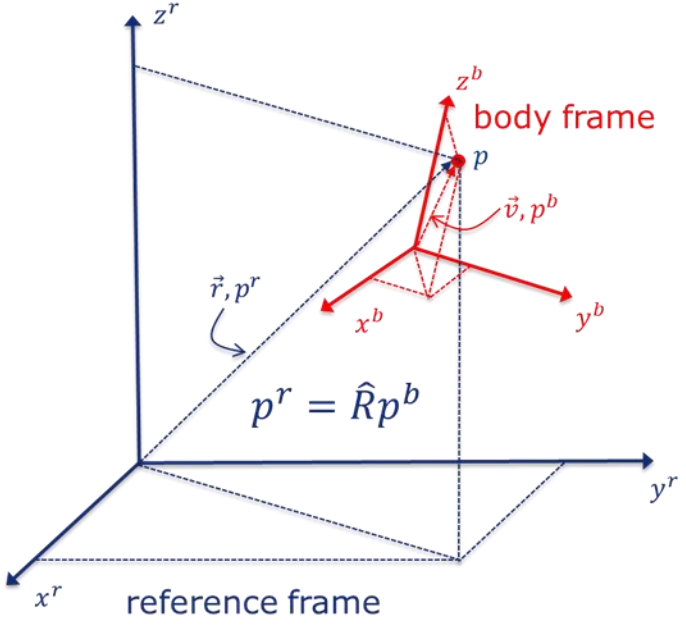

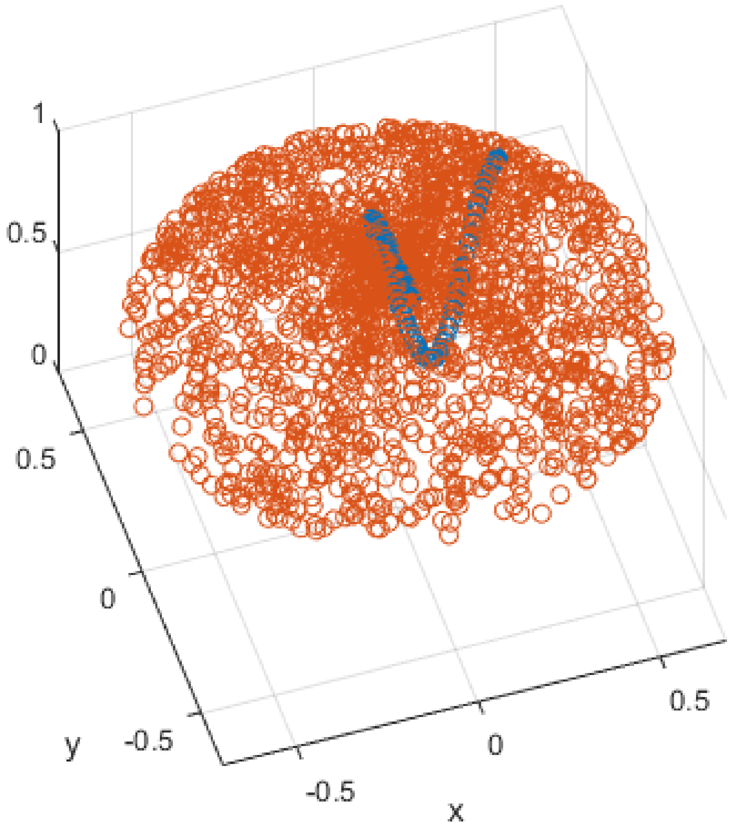

- **Case study #4** aims to explore machine learning system identification techniques towards robust modelling using a combination of data-driven techniques and consolidated model-based techniques such as LFT theory.

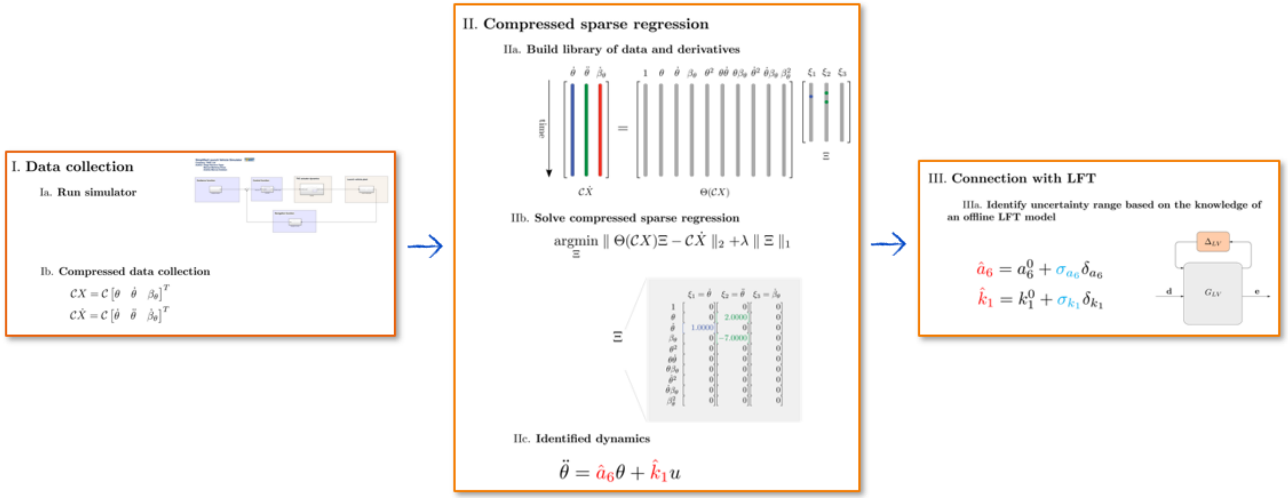

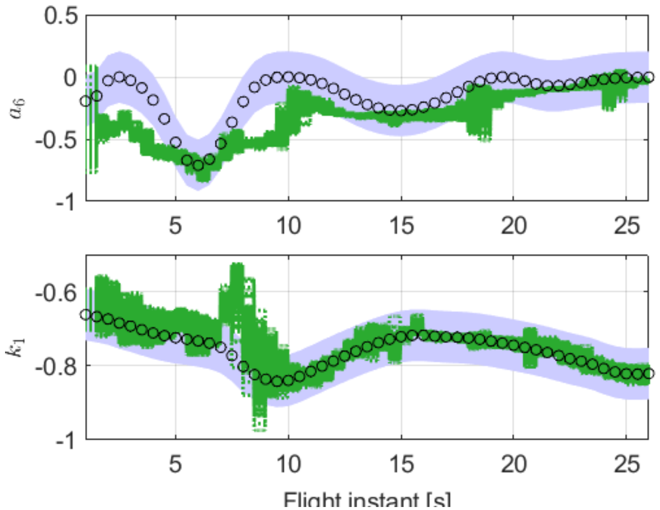

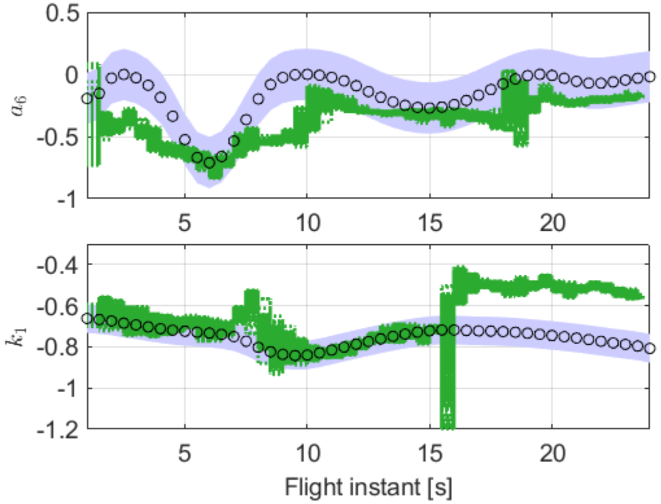

- **Case study #5** involves controlling the attitude dynamics of the RETALT vehicle during its landing phase using **learning-based MPC**
- The LBMPC algorithm design approach considers the modeling approach of case study #4 to obtain a sequence of linear models at different points of the nominal reference trajectory

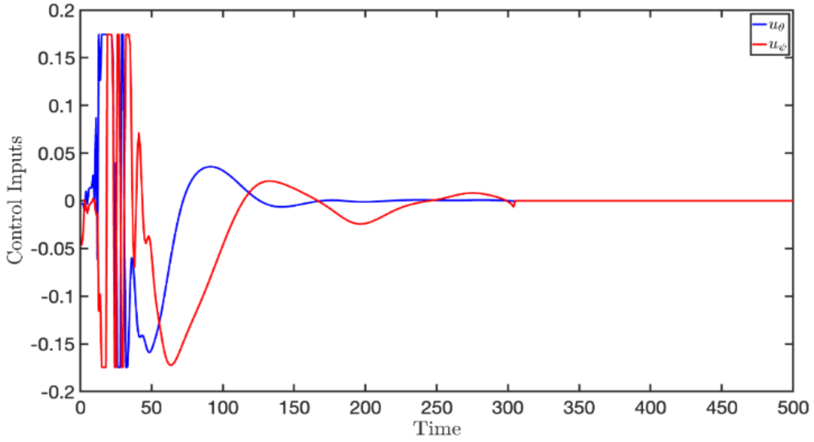

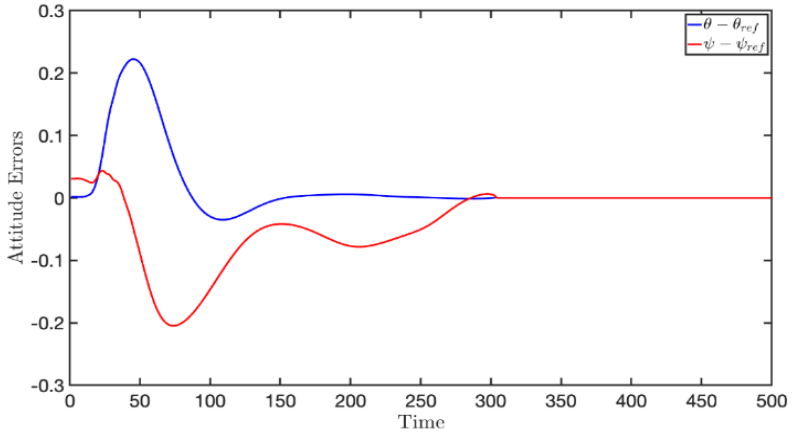

- **Case study #6** uses **Integral Quadratic Constraints** (IQCs) to validate the closed-loop behavior of a reusable launch vehicle, during the landing phase, controlled by a neural network (for the 1D scenario)

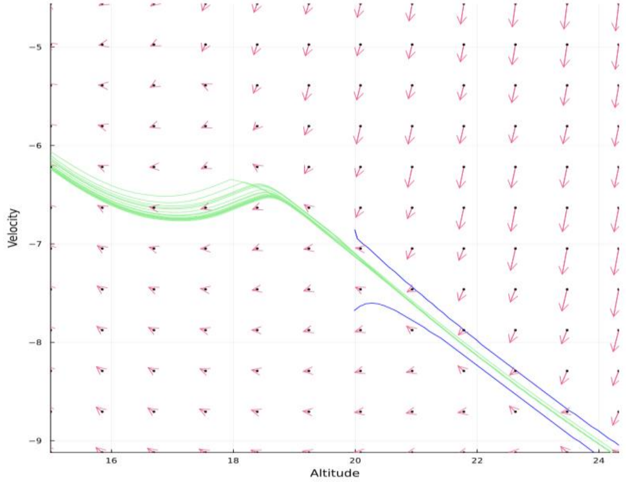

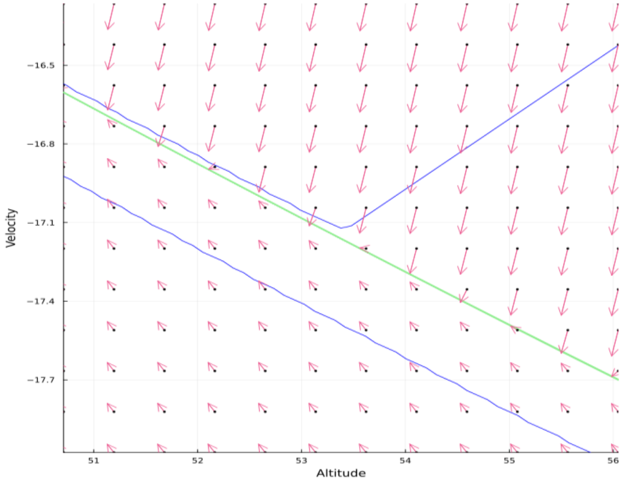

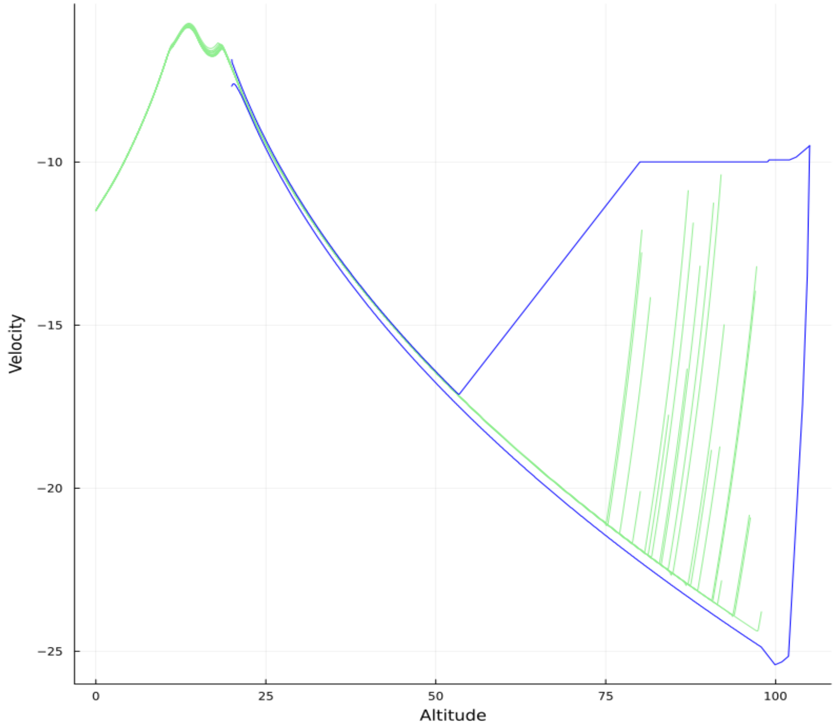

- **Case study #7** demonstrates the usage of **Learning-based Kalman filtering** (LBKF) for attitude estimation

Filter type	Angle MSE	Angular rate MSE
Standard Kalman filter	$1.09 \cdot 10^{-2}$	$4.18 \cdot 10^{-2}$
Learning-based Kalman filter with linear model	$6.92 \cdot 10^{-3}$	$1.77 \cdot 10^{-2}$
Learning-based Kalman filter with Gaussian process	$7.68 \cdot 10^{-3}$	$2.31 \cdot 10^{-2}$

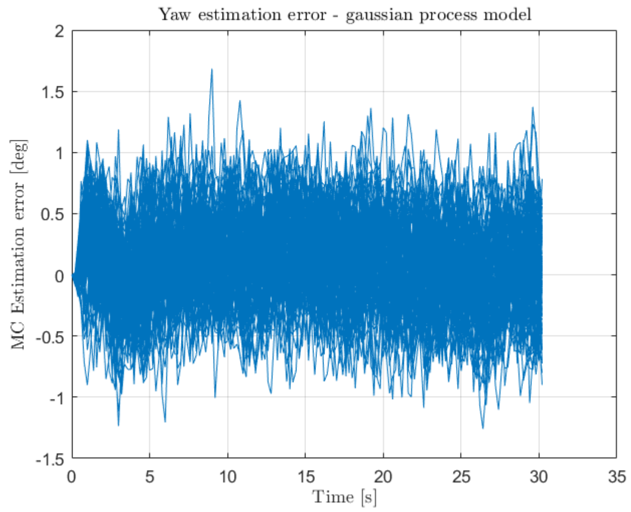

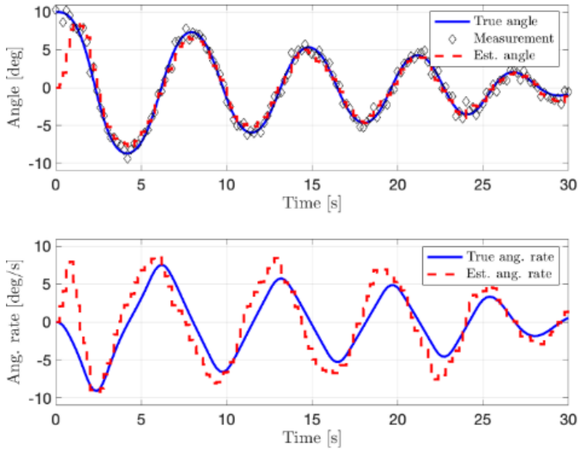

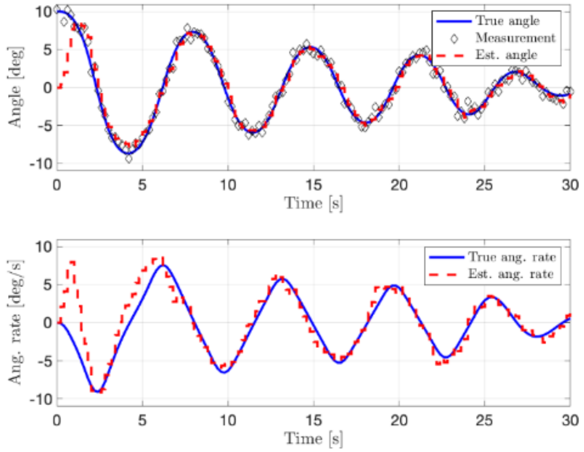

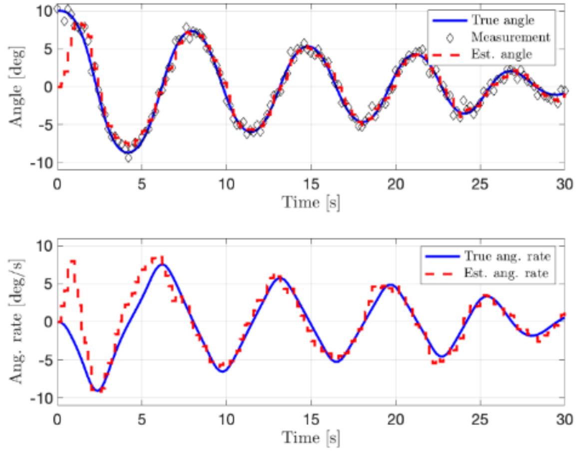

- **Case study #8** address the Guidance and Control (G&C) problem for the landing phase of a Reusable Launcher Vehicle (RLV) using **Deep Reinforcement Learning**

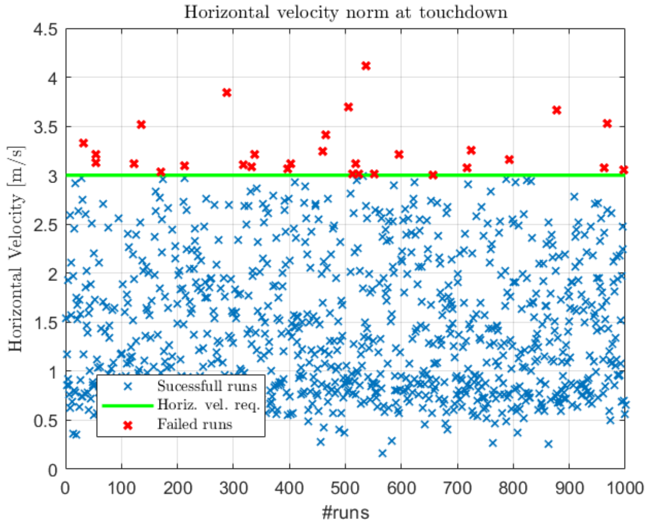

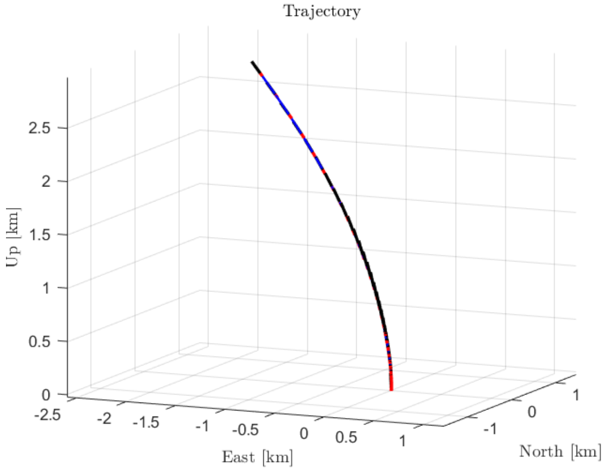

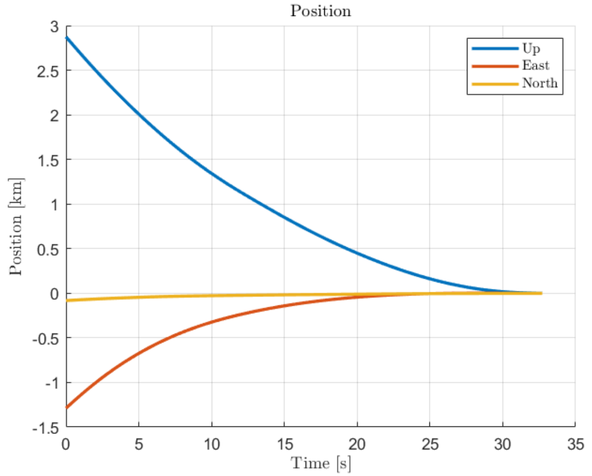

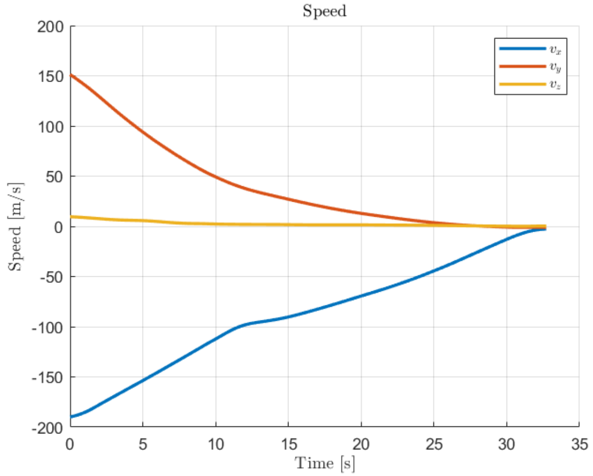

- The validation of the closed-loop behavior for the 1D scenario is also performed through reachability analysis

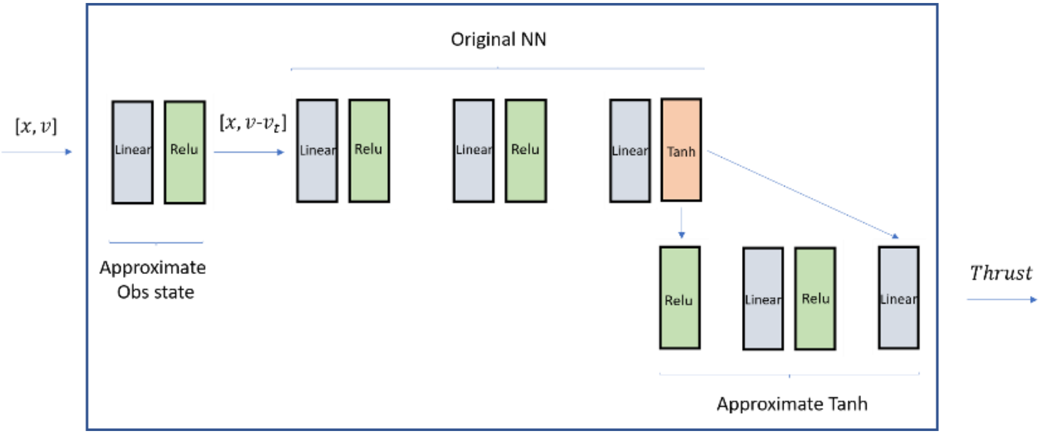

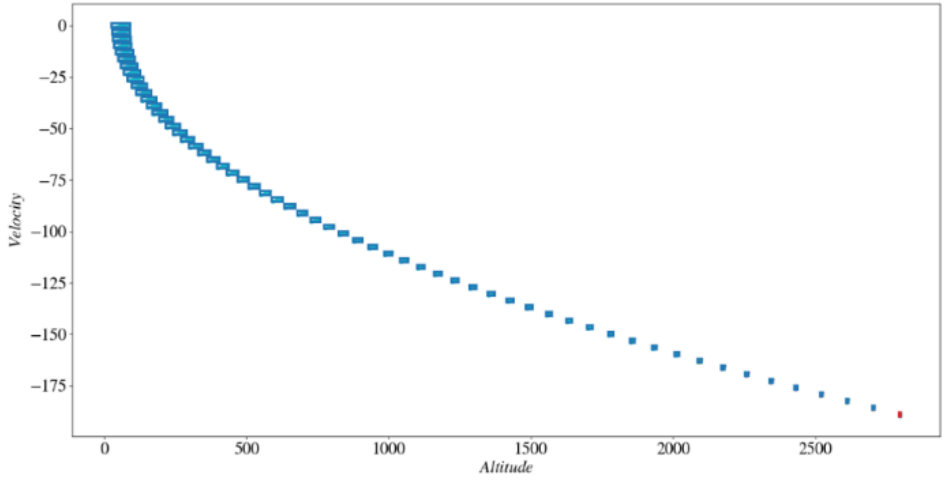

Way Forward – Maturation Plan & Roadmap

Way Forward - Maturation Plan & Roadmap
Compressed sparse regression approach

❑ Improvement areas and how to improve

- **Quality of the sensors**

  Estimation results can be improved by using higher quality sensors

- **Uncertainty identification**

  Uncertainty level identification was not possible due to low dispersions. Larger uncertainty ranges will allow to distinguish between uncertainty levels.

- **Linearized vs nonlinear estimates of a_6**

  To determine the quality of $\boldsymbol{a_6}$ estimation:

  i) Further analyses on the aerodynamic database of RETALT
  ii) Compare with other estimation techniques
  iii) Use the estimates to feed an ideal controller and evaluate in simulation which one is better

- **Thrust effect on the estimates of k_1**

  The estimates of $\boldsymbol{k_1}$ can be improved as follows:

  i) Using the knowledge of thrust changes to modify or temporarily stop the estimation flagging
  ii) Use of variable window length

- **Variable estimation window length**

  Sensitivity analysis of different window lengths was performed but it is recommended the use of variable window lengths depending on flight conditions

- **Benchmarking against classical system ID approaches**

  Several approaches can be used: i) least-squares methods; ii) Kalman filtering; iii) MATLAB's system ID toolbox (ls search or subspace methods)

- **Selection of candidate functions library**

  The effect of the selection of candidate functions library can be further investigated

Way Forward - Maturation Plan & Roadmap
ML-based failure detection approaches

Improvement areas and how to improve

- **Benchmarking against classical and other advanced FDI approaches**

  The compressed sparse regression approach could be compared to other well-established FDI methods:

  i) Model-based ($\mathcal{H}_\infty$-based methods)

  ii) Potentially with other data-driven, ML-based FDI techniques (NNs or SVM)

- **Fault assessment strategy**

  Connect the compressed sparse regression technique with well-established FDI functions to improve the fault detection assessment

Way Forward - Maturation Plan & Roadmap
Deep Reinforcement Learning G&C

- **Description**
  - Investigate possible advantages and/or complementarities of this AI method with respect to the classical G&C approaches
  - Generate a policy to map the sensor measurements directly to the action commands

- **Points to improve**
  - Robustness to wind
  - Robustness to initial condition
  - Fuel consumption
  - RL hyperparameters fine tuning
  - Extension of the convergence analysis to the 3D scenario

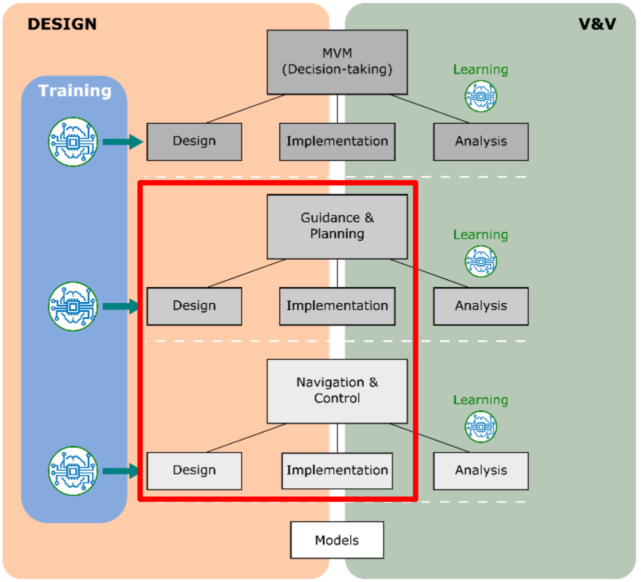

- **How to improve**
  - Robustness to wind: include a sensor to estimate wind to provide measurements to the NN; use a NN to estimate the wind itself, leading to a G&C composed by a concatenation of NNs
  - Robustness to initial conditions: increase the dispersion considered during RL training (and possibly the number of training steps)
  - Fuel consumption VS Landing accuracy: modify the reward used by RL algorithm to look for solutions with less consumption maintaining, at the same time, a good landing accuracy
  - RL hyperparameters fine tuning: manual fine tuning; define an optimization problem to optimize the hyperparameters
  - Extension to the 3D scenario: adapt the implementation of the robustness tool to handle the 3D scenario, in particular the model needs to be adapted to include the attitude dynamics.

Way Forward - Maturation Plan & Roadmap
Deep Reinforcement Learning G&C

- **Application scenarios**
  - In-Orbit Servicing (IOS)
  - Active Debris Removal (ADR)
  - Entry, Descent and Precision Landing (EDL)
    - Reusable Launch Vehicle (RLV); Re-entry vehicles with Inflatable Heat-Shields (IHS)
- **Potential indicator**
  - Identified to be **High, since:**
    - Remarkable results of the extensive Monte-Carlo campaign: the NN obtained passed all of the V&V tests with high level of confidence, and its results are comparable to SCVX.
    - The NN training can be done and repeated whenever the dynamics change
    - Non-iterative algorithm with guaranteed computational time
    - NN validation approaches exist

Way Forward - Maturation Plan & Roadmap
ML-based Guidance Optimization Surrogates

- **Description**
  - Train a NN with demonstrations of the optimal guidance, using a Supervised Learning approach

- **Points to improve**
  - Training performance
  - Expert guidance optimizer
  - Assess performance in simulation
  - Validation

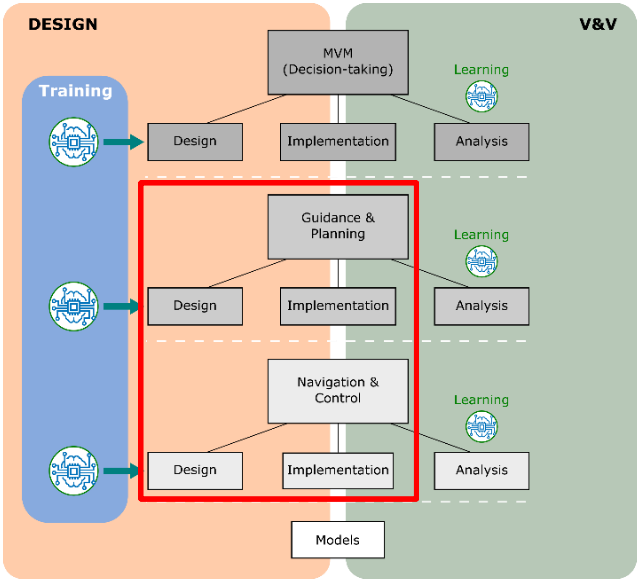

- **How to improve**
  - Training performance: test other optimizers in Keras; test other open-source libraries and tools
  - Expert guidance optimizer: consider a different sub-problem solver, such as ECOS, instead of CVX; consider other external and open-source tools (SCP by Danylo Malyuta et al.)
  - Assess performance in simulation: after the training, the NN should be tested in the high-fidelity simulator; iterative design process may be necessary for tuning the expert guidance
  - Validation: the validation tools used in other case studies may be used to validate the resulting NN

Way Forward - Maturation Plan & Roadmap
ML-based Guidance Optimization Surrogates

- **Application scenarios**
  - In-Orbit Servicing (IOS)
  - Active Debris Removal (ADR)
  - Entry, Descent and Precision Landing (EDL)
    - Reusable Launch Vehicle (RLV); Re-entry vehicles with Inflatable Heat-Shields (IHS)
- **Potential indicator**
  - Identified to be **High, since:**
    - Good results in the approximation of an online optimization algorithm.
    - The training process is typically easier than the Deep RL method, although it requires a very high number of expert demonstration
    - NN validation approaches exist

Way Forward - Maturation Plan & Roadmap

Reinforcement Learning-based Adaptive Control for Attitude Regulation

- Aims at developing a **model-free adaptive control** law for attitude regulation.
- Learns state feedback gains to solve a linear quadratic problem, without the assumption of knowing the plant model, by using **Q-learning, and policy RL**.
- A **state made of derivatives** of the plant output avoids the need for state observers
- **Combination of RL with LB-MPC** is a promising issue.
- Include a supervisor that **monitors gain learning**
- Develop a **Validation procedure** that is not just based on Monte Carlo, by getting inspiration from validation for Deep Reinforcement Learning and adaptive control.

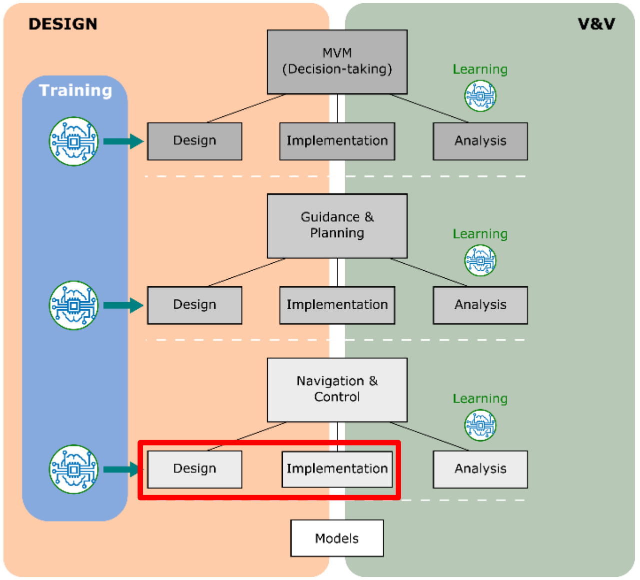

Way Forward - Maturation Plan & Roadmap
Learning-based Model Predictive Control for Attitude Control

- The results obtained using the integration of the Learning-based Model Predictive Control (LBMPC) into the ESAi4GNC tool had its own merits and demerits.
- Though attitude regulation of RETALT vehicle was achieved by the LBMPC algorithm, it came at the expense of high landing errors (in the range of 150m).
- In principle, such shortcomings can be rectified by proper choices of constraints and problem parameters although tuning and selection of parameters and constraints value was not straightforward to do for this application.
- Though such practical difficulties resulted in mediocre Monte-carlo simulations, the potential to solve this using informed choices of constraints and using the well established and rigorously proven LBMPC strategy gives us great hope for future work.
- Incorporate additive disturbance into the design of RETALT vehicle simulation using the ESAi4GNC tool.

Way Forward - Maturation Plan & Roadmap

Learning-based Kalman Filter

- As compared to a standard Kalman filter, the LBKF can be used to **slightly improve** the navigation **performance** but at the expense of a **larger modelling effort** and rather **expensive online calculations** and a **lack of guarantees**. These are serious drawbacks considering applications in GNC applications for spacecraft.

- **More research** is required to make the method useful and worthwhile for space applications as compared to heritage techniques.

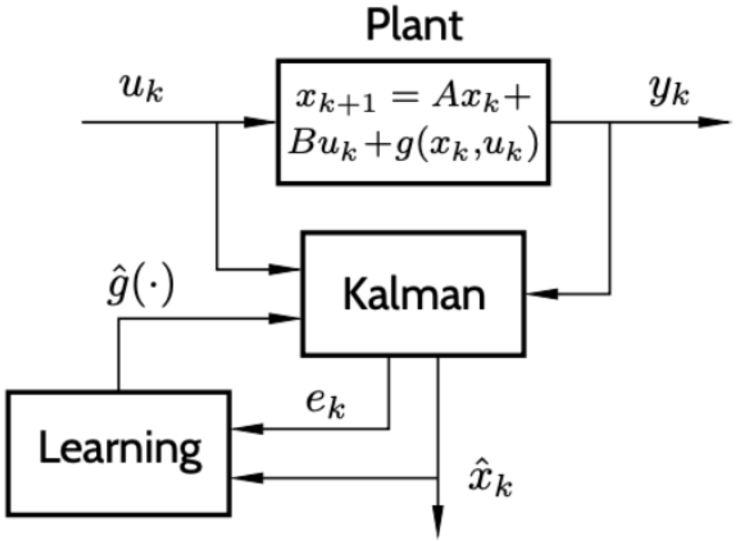

Way Forward - Maturation Plan & Roadmap
Integral Quadratic Constraints

- **Main Obstacle: Dimensionality**
  - We validated the ”1D”-model without attitude dynamics
  - Work needed towards ”3D”-model
  - Partitioning of state space has scaling issues, unclear to what extent.
  - Try to partition a narrow tube around trajectories, which might scale well enough
- **IQC Benefit: Uncertainties do not need to be enumerated**
  - Useful when Monte-Carlo has issues testing all possible values of uncertainties.
  - Our focus: Nonlinearities in neural networks.
  - Look for other applications witin the GNC problem space.
- **Roadmap:**
  - Work towards higher-dimensional models.
  - Consider other validation problems with high-dimensional uncertainty sets.

Way Forward - Maturation Plan & Roadmap
Roadmap

#	Technique name	GNC level	Typical application	Potential indicator
1	Compressed sparse regression approach	i-L-DI	1. Consolidation for estimation 2. Further improvements for FDI 3. Use for FTC applications	H
2	Deep Reinforcement Learning G&C	i-GC-DI	1. Entry, Descent and precision landing (EDL) both on Earth and other celestial bodies 2. In Orbit Servicing (IOS) 3. Active Debris Removal (ADR)	H
3	ML-based Guidance Optimization Surrogates	i-GC-DI	1. Entry, Descent and precision landing (EDL) Guidance on Earth and other celestial bodies 2. In Orbit Servicing (IOS) Guidance 3. Active Debris Removal (ADR) Guidance	H
4	Reinforcement Learning based adaptive control for attitude regulation	i-C-DI	Attitude regulation around a reference provided by the guidance system	M
5	LBMPC for attitude tracking of RETALT Vehicle	i-C-DI	Attitude tracking to enable soft landing during the re-entry phase	M
6	Learning-based Kalman filter	i-N-DI	Adapt to component imperfections	L
7	IQC	i-C-A	Validate full system with NN controller	L

Coffee break?

Lessons Learnt and Recommendations

Lessons Learnt
Reinforcement Learning

- RL provides a way to develop **adaptive regulation** controllers that approximate the LQR for attitude angle regulation.
- Adjusting the **dither level** is a very important tool.
- The use of RL in an architecture based on the Youla parameterization does not yield a significant advantage.
- The **short time period** of operation poses a difficulty in learning the controller gains.
- The gimbal rate constraint poses great difficulties
- When using a cascade control structure, it is very important to have a **clear separation in the time scales** between the inner loop and the outer loop
- It was not possible to simultaneously use 2 RL based adaptive regulators.
- The use of a RL based adaptive regulator in the attitude controller allows to tackle some **actuator faults**. Strong offsets may cause the learning algorithm to diverge.
- **Deep Learning Techniques** for attitude determination during landing is a more robust alternative compared to algorithmic based solutions.

Lessons learnt

- Compressed sparse regression approach demonstrated the feasibility of online system identification and failure detection
- The proposed approach successfully captured the variability of the main rotational parameters of a launch vehicle
- The implementation is compatible with Simulink code generation
- The results from an intense Monte Carlo campaign highlighted the FDI capabilities of this approach of engine thrust failures

Recommendations

- Use of compressed sparse regression approach for other phases of a launch vehicle flight
- Benchmark against other classical system identification approaches
- Consider the use of flight-dependent (variable) window length
- The estimation results can be improved by using higher quality sensors

Learning-based Model Predictive Control

- Attitude Control of RETALT vehicle can be effectively performed using LBMPC if the nominal model and the bounds on the unmodeled dynamics are exactly known.
- LBMPC gives us the opportunity to learn the unmodeled dynamics apart from its known bounds, and based on the usage of different oracles, even complex unmodeled dynamics can be learned up to a desired accuracy.
- LBMPC brings the best of both worlds with enhanced performance due to the learning process and ensures constraint satisfaction using the nominal model.
- The hinderances of implementing the LBMPC can be rectified by using informed choices of problem parameters and system constraints and by dedicated tuning efforts.
- Combination of LBMPC with reinforcement learning is an interesting future aspect to investigate to get both safe learning and control of RETALT vehicle.

- The LBKF technique may be useful in cases where some type of **nonlinear dynamics can be anticipated, but not known beforehand**. In the RLV case study, we adapted the Kalman Filter to actuator degradations in the form of magnitude and rate limited control signals, which were not known at design time.

- Quite severe nonlinearities must be present for the technique to deliver any noticeable performance improvements. **A well-tuned standard Kalman filter should already be robust to modest perturbations**.

- Learning or adaptation always includes the risk of overfitting to noisy or temporary problems, which in the end may lead to worse performance. We saw that the LBKF could introduce additional artificial disturbances when it tried to learn nonlinear dynamics that were not actually present.

- **Recommendation:** The small potential estimation performance improvement probably does not **outweigh the risk and runtime cost of adding a learning component to the Kalman Filter**.

- In a **Deep-RL framework**, the reward should be defined such that the final step has a positive reward if the landing is successful. The results were observed to be more benign if this positive reward at the end is constant, as long as a satisfactory region (for the state) is attained.
- Using a target velocity helps accelerate the training process.
- Selecting the initial state as random points around the reference trajectory significantly improves learning.
- The longer the training, the better the final results in terms of robustness and accuracy.

Recommendations

- Within **reinforcement learning** scenarios it is recommended to perform an intensive simulation study to define the dither noise variance and the weight on the control variable.
- If multiple control loops based on adaptive control are to be used, it is recommended to tune each one at a time, while the controller gains of the other are forced to be constant.
- The estimation results can be improved by using higher quality sensors with lower noise levels and higher sampling rates.
- The use of **compressed sparse regression approach is recommended** for other phases of a launch vehicle flight, and to benchmark it against other (classical) identification methods, as well as to assess the performance of this estimation approach using variable window lengths.
- For the **LBMPC approach**, consider the additive disturbance in the integration.
- It is recommended to study the attitude dynamics carefully to infer proper bounds for the unmodeled dynamics to be used for constraint satisfaction (tightening).
- It is recommended to relax the unmodeled dynamics from linear to nonlinear models and learn them from input-output data using neural networks (nonlinear oracle modules) in the future.
- It is recommended to learn how to switch between models and try using the respective learning-based MPC controller.
- It is recommended to study the regret incurred by the learning-based MPC for operating without knowing the exact unmodeled dynamics

Recommendations

- The small potential estimation performance improvement probably does not outweigh the risk and runtime cost of adding a learning component to the **Kalman Filter**, except in very special cases.
- In the case of possible actuator degradations (as studied here), a safer, alternative approach could be to add additional sensors to the vehicle. Another approach is to design a standard Kalman Filter to be robust towards more plant variability.
- In **Deep-RL**, the exploration noise should be adjusted to the magnitude and type of actions. In particular, improved results were obtained by reducing the exploration noise for the gimbal action.
- The target velocity profile is the main driver for the trajectory to be followed. To have a different trajectory, e.g., one with a better approximation to a bang-bang solution of the thrust, the target velocity should be different.
- The RL hyperparameters should be adjusted every time there is a significant modification in the overall setting. For instance, a different reward shape or a modification in the environment dynamics (adding or removing actuation, etc).

Review of RIDs

Next Steps

Next Steps
Planning

- Implement actions from RIDs
- Update D8, and other items accordingly

AoB
Dissemination and Exploitation of the Results

- We propose to disseminate the results of AI4GNC in one of the following conferences (e.g. in a special session):
  - ESA GNC Conference 2023
  - EUCASS/CEAS Aerospace Europe Conference 2023

- The consortium is internally discussing the possibility of making ESA-i4GNC (and possibly some of the case studies) publicly available

The inputs from ESA are very welcome on the two topics!

AoB

- ❑ Any other points to be discussed?

Thank you!

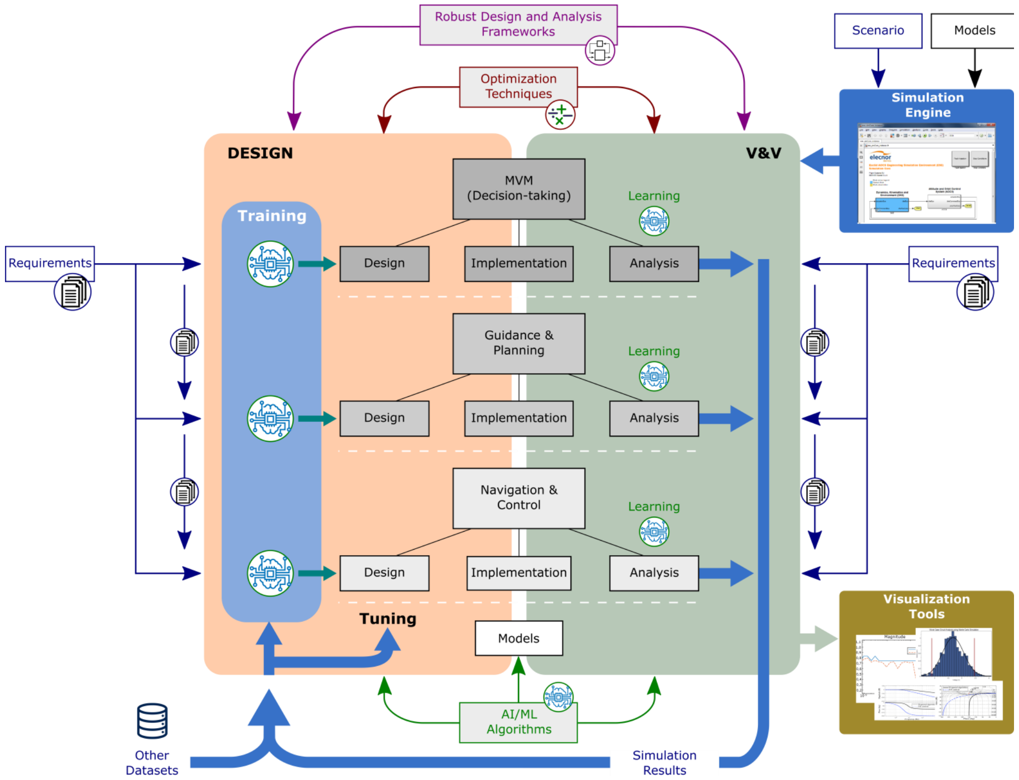

Robust Design and Analysis Frameworks
Optimization Techniques
Scenario
Models
Simulation Engine
DESIGN
V&V
MVM (Decision-taking)
Training
Learning
Requirements
Design
Implementation
Analysis
Requirements
Guidance & Planning
Learning
Design
Implementation
Analysis
Navigation & Control
Learning
Design
Implementation
Analysis
Tuning
Models
Visualization Tools
AI/ML Algorithms
Other Datasets
Simulation Results

BACKUP SLIDES

- Two simulations were performed following the suggestion and the thrust profile was still not the desired one;
- The **target velocity** profile is the **main driver** for the thrust profile
- Therefore, the target velocity should be **modified** to allow for a **bang-bang thrust profile**.
- Modifying the target velocity may lead to a complete new reward shaping procedure, which could entail a large effort.

Answer to RIDs VR_JB-01 & VR_JB-02
Deep Reinforcement Learning G&C

- Penalization added for the Tof in the reward function
- Improvement in the overall Tof (less than 1 sec) but still not bang-bang solution