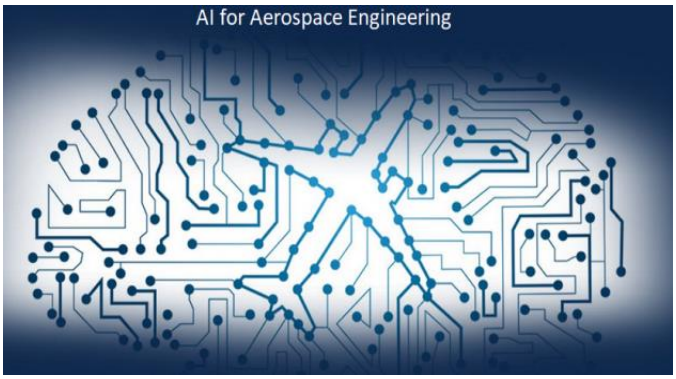


Abstracts of

“Advances in Artificial Intelligence for Aerospace Engineering” 5th workshop

May 19th to 21th 2025, ONERA Toulouse, Toulouse, France

Foreword

In the scope of an agreement for cooperation in Artificial Intelligence (AI), signed in 2019, DLR and ONERA jointly consider AI research challenges in a number of application fields of particular importance in the aeronautics and space domains such as robotics and autonomy in aerospace systems, hybrid simulation, digital transformation of aerospace industry, processing of geospatial information for earth observation.

In 2020, the decision to hold annual workshops launched the cooperation. The first and second editions were virtual conferences. Presentations of scientific contributions and cooperation proposals highlighted the interest of joint action of the two research institutions in the field of AI. Permanent researchers, post-docs and doctoral students discussed approaches and ideas forming the framework of the cooperation project and identified areas of common interest for using AI such as: the necessary robustness of AI, its interaction with humans, the use of AI in space robotics, for Earth observation, for improving the efficiency of computational fluid dynamics and for industry. In 2022, AI-related presentations were encouraged for ODAS. This led to the receipt of 18 presentation proposals on AI. Among these proposals, seven dealt with the collaboration in AI. During the third edition held in presence in Paris in 2023 a day before ODAS, 20 scientists presented their contributions. Twenty-five contributions were presented during the fourth AI4Aerospace workshop, held in Braunschweig in 2024 two days before ODAS. Finally, the fifth edition of the workshop, held in Toulouse in 2025 before and together with ODAS, featured 27 presentations.

The purpose of this document is to provide an overview of the fifth workshop held in May 2025 in Toulouse. Bellow the general objective of strengthening our collaborative work, the workshop has two sub objectives: networking and presentation of research. It addresses the following topics:

- Machine Learning approaches for physics and partial derivative equations,
- Certification of AI in aviation,
- AI for control and monitoring,
- AI for engineering and operation,
- AI for sensing and image processing,
- AI engineering methods and approaches and
- Other AI related contributions.

The structure of this document highlights those topics.

The organizing committee
Günther Waxenegger-Wilfing: Guenther.Waxenegger@dlr.de
Umut Durak: Umut.Durak@dlr.de
Markus Kintscher: Markus.Kintscher@dlr.de
Philippe Bidaud: Philippe.Bidaud@onera.fr
Stéphane Herbin: Stephane.Herbin@onera.fr
Jean-Loup.Farges: Jean-Loup.Farges@onera.fr

Machine Learning Approaches for Physics and Partial Derivative Equations

Contents

- Preconditioned FEM-based Neural Networks for Solving Incompressible Fluid Flows and Related Inverse Problems *Fabrice Von der Lehr, Katharina Rauthmann, Philipp Knechtges*.
- RANS-CNN: A Physics-informed Convolutional Neural Network for Solving RANS Equations in Duct Flows *Gaurav Bokil, Sebastian Merbold, Stefanie De Graaf*.

Preconditioned FEM-based Neural Networks for Solving Incompressible Fluid Flows and Related Inverse Problems

Fabrice von der Lehr (Fabrice.Lehr@dlr.de), Katharina Rauthmann, Philipp Knechtges (all DLR SC-HPC, Cologne, Germany)

Presentation type
Result of research work

Keywords
FEM-based Neural Network, Machine Learning, Preconditioning, Finite Elements Method, Parametric PDEs, Stokes, Navier–Stokes

Introduction

The numerical simulation and optimization of technical systems described by partial differential equations (PDEs) is expensive, especially in multi-query scenarios in which the underlying equations have to be solved for different parameters. A comparatively new approach in this context is to combine the good approximation properties of neural networks (NNs; for parameter dependence) with the classical finite element method (FEM; for discretization).

Methods and Results

Instead of considering the solution mapping of the PDE from the parameter space into the FEM-discretized solution space as a purely data-driven regression problem, so-called physically informed regression problems have proven to be useful. In these, the equation residual is minimized during the training of the neural network, i.e. the neural network “learns” the physics underlying the problem (1).

In this work (2), we extend the approach to saddle-point and non-linear fluid dynamics problems, respectively, namely stationary Stokes and stationary Navier–Stokes equations. In particular, we propose a modification of the existing approach: Instead of minimizing the plain vanilla equation residual during training, we minimize the equation residual modified by a preconditioner (Fig. 1). By analogy with the linear case, this also improves the condition in the present non-linear case. Our numerical examples demonstrate that this approach significantly reduces the training effort and greatly increases accuracy and generalizability (Fig. 2). Finally, we show the application of the resulting parameterized model to a related inverse problem (Fig. 3).

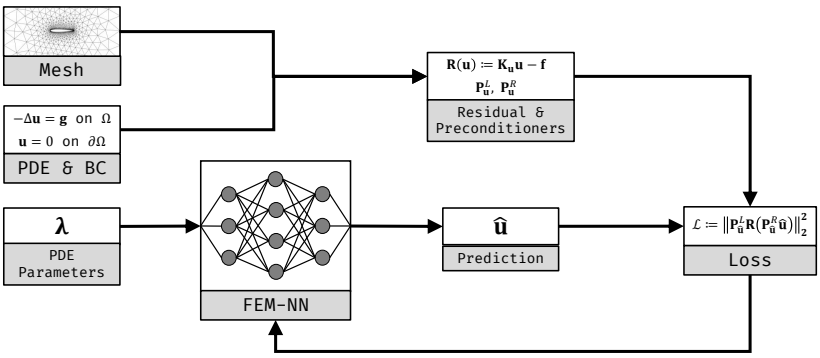
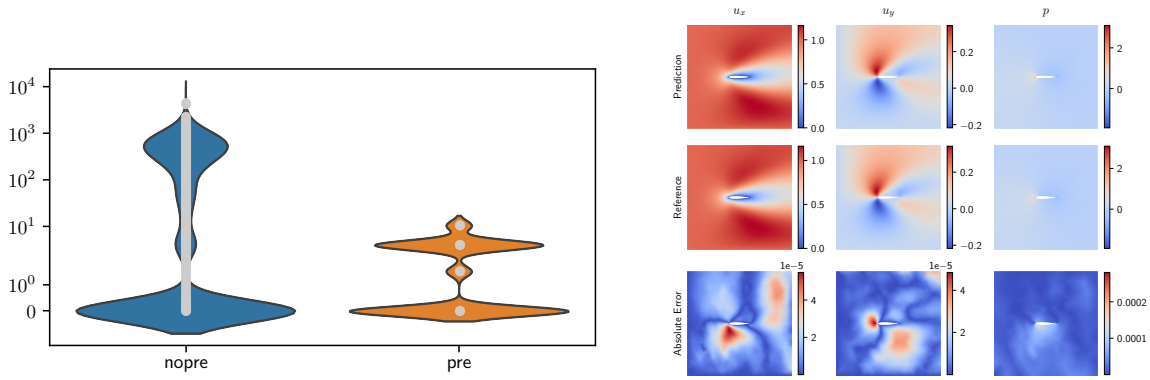
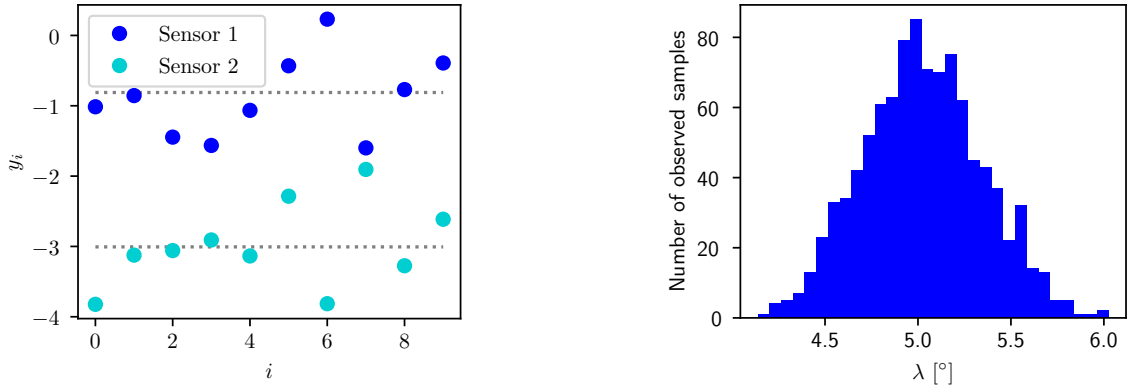


Figure 1: Schematic of a preconditioned FEM-NN. Given a PDE, boundary conditions (BC) and a mesh, the FEM module provides the residual $\mathbf{R}$ and preconditioners $\mathbf{P}_\mathbf{u}^L$, $\mathbf{P}_\mathbf{u}^R$. The FEM-NN, which takes the PDE parameters $\boldsymbol{\lambda}$ as input and predicts a solution $\hat{\mathbf{u}}$, is then trained by minimizing the preconditioned residual under the L2-norm.



(a) (b)

Figure 2: *Left:* Eigenvalues of the Hessian of the loss function, evaluated after training the FEM-NN with preconditioning (pre) and without preconditioning (nopre) for Stokes equations. *Right:* Results for Navier-Stokes equations. The top row shows the prediction by the preconditioned FEM-NN, the second row the reference solution obtained with FEM, and the third row the absolute error of the prediction.



(a) (b)

Figure 3: *Left:* Noisy pressure measurements y_i at two sensor locations of the airfoil. *Right:* Computed posterior distribution of the angle of attack λ.

Conclusion

In this work, we have provided a proof-of-concept how preconditioning enables the application of FEM-NNs to Stokes and Navier-Stokes problems, both forward and inverse. In an upcoming project named "Trustworthy Physics-informed AI for Aerospace and Transportation (TIARA)", we will investigate further stabilization methods and low-cost preconditioners.

References:

(1) Meethal, R, Kodakkal, A, Khalil, M, Ghantasala, A, Obst, B, Bletzinger, KU, Wuchner, R. "Finite element method-enhanced neural network for forward and inverse problems". Advanced modeling and simulation in engineering sciences 2023; 10(1):6.

(2) Franziska Griese, Fabian Hoppe, Alexander Rüttgers, Philipp Knechtges. "Preconditioned FEM-based neural networks for solving incompressible fluid flows and related inverse problems". Journal of Computational and Applied Mathematics 2025:116663.

RANS-CNN: A Physics-informed Convolutional Neural Network for Solving RANS Equations in Duct Flows

Gaurav Bokil (DLR Institute of Electrified Aircraft Propulsion), Sebastian Merbold (DLR Institute of Electrified Aircraft Propulsion), Stefanie deGraaf (DLR Institute of Electrified Aircraft Propulsion)

Presentation type
Result of research work

Keywords
Physics-informed machine learning, convolutional neural network, turbulent flow, surrogate modelling

Introduction

Numerical simulations have been the backbone of engineering design since its inception in the 1970s. Discretisation methods such as the Finite Difference Methods (FDM), Finite Element Methods (FEM) and Finite Volume Methods (FVM) are on the forefront in this regard. For aerospace applications, accurately simulating the fluid flow at high velocities is crucial. As a result, Computational Fluid Dynamics (CFD) simulations required to design the aircraft components become computationally expensive. Currently, electrification of aircraft propulsion systems is being examined as a sustainable alternative to combustion engines. Propulsion topologies primarily powered by hydrogen fuel-cells and batteries are sensitive with regards to operational temperatures and temperature gradients. With fuel cells operating at approximately 50 % efficiency, managing heat rejection poses a significant challenge for the Thermal Management System (TMS). Thus, the design of the cooling ducts and heat exchangers to efficiently dissipate the heat is of critical importance. To design cooling ducts and heat exchanger configurations suitable for specific applications, sensitivity analyses and geometry optimisation needs to be carried out across a large design space. However, the computational cost of conducting CFD simulations acts as a bottleneck. To overcome this, a surrogate modelling approach is chosen as an alternative. By training a surrogate model on simulations and experimental data, the design space can be interpolated and more configurations can be evaluated. In recent trends, various Machine Learning (ML) methods to construct surrogate models based on Artificial Neural Networks (ANN), Random Forest regression (RF), and Gaussian Processes (GP) are being put forth. Such techniques require tremendous datasets to accurately capture the design space. In most cases, where only sparse data or no data is available, these approaches cannot be applied. To address this challenge, physics-based machine learning methods are chosen as an alternative to develop surrogate models for cooling ducts and heat exchangers.

Methodology

A CNN framework is developed which takes geometry representation as input and predicts the velocity and pressure as an output. The CNN model is built using two independent U-nets, one each to output velocity components and pressure. Binary masks and distance function fields of the duct shape represent the geometry. Additionally, supplementary inputs of boundary masks representing the boundaries of the duct domain are employed to “strongly” enforce Dirichlet boundary conditions (BCs) on the network architecture. As a result, the CNN predictions inherently satisfy the Dirichlet BCs for velocity and pressure. To compute the spatial gradients in the flow accurately, sixth-order central difference scheme is implemented. Using these gradients, the loss function comprised of mass conservation, momentum conservation, turbulence model and Neumann BCs is formulated. Numerical diffusion based on Peclet number is added in the governing equations to ensure stability of the RANS equations. Followed by appropriate weighting of the loss terms, the physics-informed loss is constructed as a mean squared residuals. Finally, the Adam optimiser trains the RANS-CNN in an unsupervised manner. The entire outline is described in Fig. 1. To evaluate the performance of this method across various duct geometries, it is implemented on three case studies. The first study considers a typical S-shaped duct with experimental and simulation data from literature as reference. S-shaped ducts are prevalent in the aerospace industry primarily as air intakes for engines and as cooling ducts. The second case study deals with

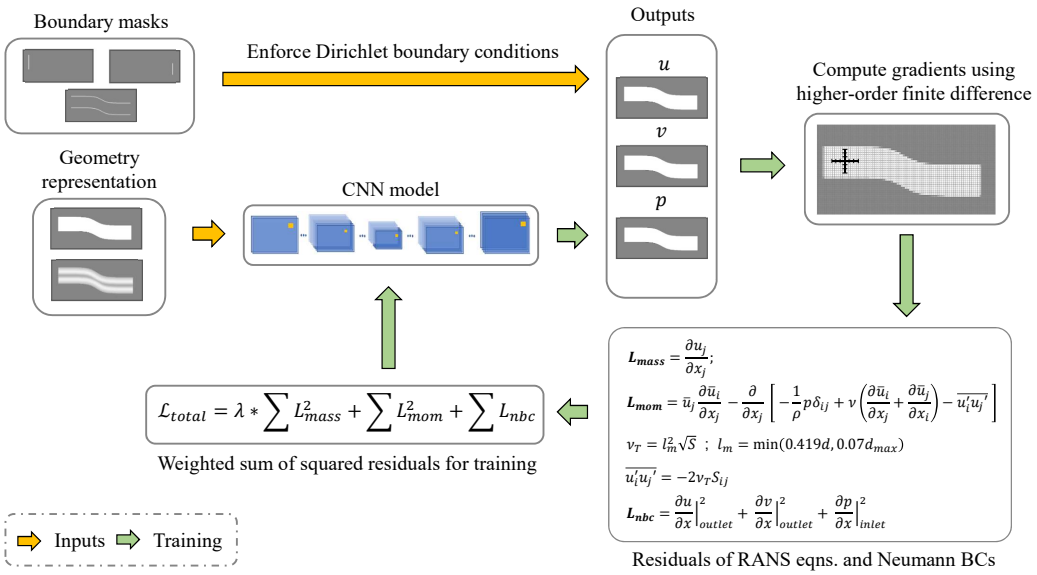


Figure 1: Schematic description of the RANS-CNN methodology

flow through a ducted heat exchanger modelled as a body force. The results of this case are compared with reference experiment data. Finally, the third case considers a forward facing step problem which is a standard benchmark problem for CFD codes.

Conclusion

The overall results observed demonstrate that the proposed approach of utilising:

- higher-order finite difference scheme for computing gradients,
- numerical stabilisation, and
- strong boundary enforcement

enabled the CNN model to simulate the RANS equations. Since the binary masks and distance function fields are created using the duct wall coordinates, the procedure is easily reproducible for a different duct geometry such as a convergent-divergent nozzle or a U-bend. The work aimed to develop a framework that produces accurate results which can be further improved via provided labelled data.
This work is carried out towards the development of cooling ducts for electric aircraft thermal management system. The aim of this study was to develop a physics-informed CNN which can simulate fluid flow through ducts at high Reynolds numbers. The results obtained from the case studies confirm the capability of the RANS-CNN method in achieving this goal. To this end, the next step is the development of a RANS-based geometry-adaptive surrogate model for ducted heat exchangers in electric aircraft. Such a surrogate model will then aid in designing optimal cooling ducts for electric aircraft.

Certification of AI in aviation

Contents

- A Case Study for Embedded and Certified Aeronautical Systems using AIDGE, an Open and Sovereign ML Tool *Filipo Perotto, Youcef Bouchebaba, Jean-Loup Farges, Anthony Pires, Benjamin Lesage, Claire Pagetti, Mohammed Belcaid* .

A Case Study for Embedded and Certified Aeronautical Systems using AIDGE, an Open and Sovereign ML Tool

Filipo Studzinski Perotto, Anthony Fernandes Pires, Jean-Loup Farges, Youcef Bouchebaba, Mohammed Belcaïd, Benjamin Lesage, Claire Pagetti (ONERA-DTIS)

Presentation type
Result of research work

Keywords
Machine Learning, Embedded Systems, Certification, ACAS-Xu.

Introduction

AI was developed in recent years based on giant neural networks and massive data processing, but today the challenge is to transpose these methods into small components as close as possible to industrial solutions. Two important issues are then highlighted: frugality and certificability, both necessary to allow the use of AI in embedded critical systems and infrastructures.

The DeepGreen project, funded by the French Research Council, is developing a new open-source software platform for embedded AI called AIDGE[1], designed to learn *Deep Neural Networks* (DNN), transform, and generate optimized code for target hardware architectures, in a completely transparent and traceable way, helping to avoid dependence on opaque and non-sovereign software, ensuring good performance, and favoring the cerficability of systems conceived with *Machine Learning* (ML).

In this context, the capabilities of AIDGE to produce an embedded and certifiable version of a Neural Network need to be evaluated (2). Our contribution to the project is defining how the tool can generate the code (in C) implementing the ML model for a particular use case, the ACAS-Xu system, ensuring to support certification following the latest guidelines provided by EASA. As a preliminary result, we present a comparison between the performance of ACAS-Xu from code generated with different ML platforms.

ACAS-Xu

ACAS-Xu is the new generation of Airborne Collision Avoidance System for fixed-wing unmanned aircraft (1). The function of that system is to generate resolution advisories (evasive actions) in case of imminent risk of collision between two airplanes. The system interrogates the transponders of nearby aircraft to determine the distance, altitude, and bearing of surrounding traffic. The resolution policy has been optimized through dynamic programming, represented within a very large lookup table: 600 million parameters (q-costs) for the single horizontal resolution case. In the literature, a strategy to reduce its memory footprint is using machine learning to learn a q-function approximation from the table data using a neural network.

However, an important aspect to make possible the use of a DNN version of ACAS-Xu in real embedded systems is the certification of the neural networks (4), (3). In this presentation, we show the preliminary results of performance analyses for different implementations of the ACAS-Xu DNN with the ones obtained with AIDGE. We also introduce the roadmap developed by ONERA and CS-Group within the DeepGreen project in order to use the explicit intermediary representation models that can be extracted anytime from AIDGE at each step of the optimization and transformation of a given neural network to the target code to support certification following the development objectives stated by the EASA concept paper (5).

Certification Qualities

The EASA defined a number of guidelines for the certification of ML-based systems in safety-critical systems. While those are not yet regulatory requirements, they highlight a number of desirable qualities from ML-related tooling. Based on the examination of the certification objectives in (5), we identified some basic characteristics necessary for any ML tool aiming to produce certified software:

[1] https://projects.eclipse.org/projects/technology.aidge

1. **determinism:** all the operations executed with the help of the platform (dataset split, learning, optimization, quantization, translation into intermediate model, decomposition into items, translation into source code, compilation, etc.) must be deterministic, i.e. all the necessary parameters and random seeds must be controllable and explicit, ensuring that repeating the operations will always lead to exactly the same results.
2. **traceability:** all the subsequent models, generated by chained transformations, must have their parts identified in order to trace their origin. For example, each neural network layer should be commented, annotated, and identified identically from the ML model to the generated source code.
3. **reproducibility:** the complete process leading from the data to the implemented model must be reported, identifying the order, inputs, and parameters used on each operation, allowing to reproduce all the activities of the development cycle obtaining the same results.
4. **formalization:** to each manipulated model, object, and operation, a complete, explicit, non-ambiguous formal description should be associated.

Preliminary Performance Evaluation

In preliminary tests, starting from the same neural network, with 6 fully-connected layers of 50 neurons and ReLU activation, represented as an ONNX file, we ran 30000 inferences using different instances, some compiled from generated code from different platforms, others directly using runtime tools, some for CPU, others using GPU (Nvidia T4). Those initial tests produced interesting results from the version compiled from Aidge generated C++ code.

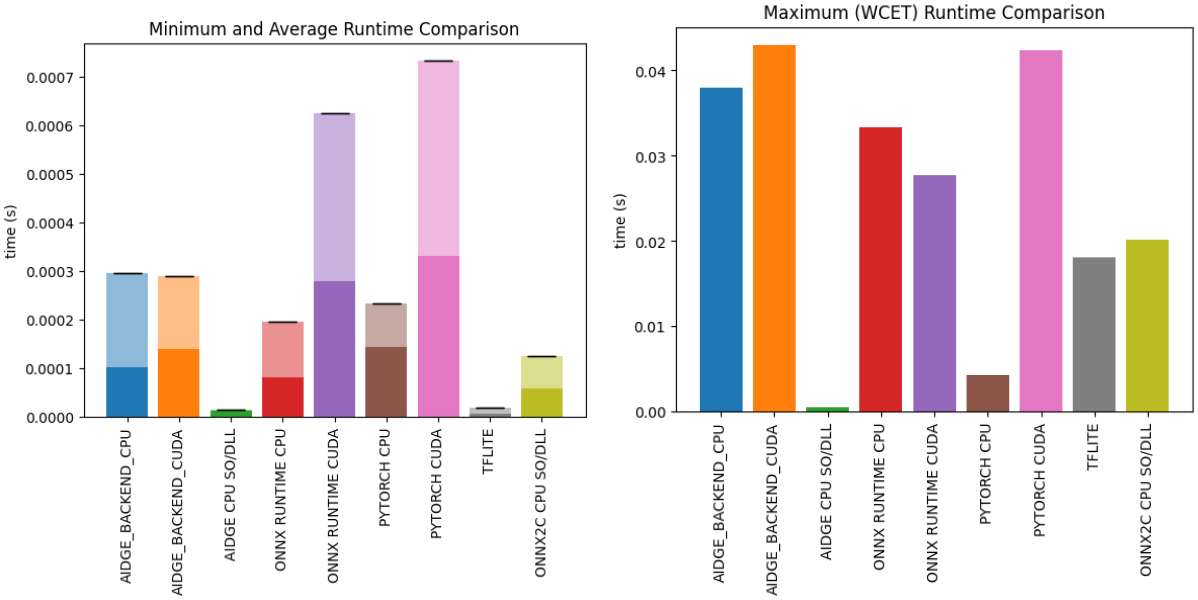


Figure 1: At left, average and best inference times using different code generation or runtime tools, and at right, the worst-case elapsed times.

References:

(1) Perotto. 2024. Using Deep RL to Improve the ACAS-Xu Policy: Concept Paper. In Proc. of the 2nd Int. Conf. on Cognitive Aircraft Systems: ICCAS. INSTICC, SciTePress, 110–117.

(2) Perotto et al. 2024. Thinking the certification process of embedded ML-based aeronautical components using AIDGE, a French open and sovereign AI platform. In Proc. of the 2nd Int. Conf. on Cognitive Aircraft Systems: ICCAS. INSTICC, SciTePress, 64–71.

(3) Gabreau et al. 2022. Toward the certification of safety-related systems using ML techniques: the ACAS-Xu experience. In Proc. of 11th European Congress on Embedded Real Time Software and Systems, ERTS.

(4) Damour et al. (2021). Towards Certification of a Reduced Footprint ACAS-Xu System: a Hybrid ML-based Solution. In Computer Safety, Reliability, and Security (40th SAFECOMP), pages 34–48. Springer.

(5) EASA (2024). Concept Paper: First Usable Guidance for Level 1 and 2 Machine Learning Applications, n.2. Eu- ropean Aviation Safety Agency (EASA), Cologne.

AI for Control and Monitoring

Contents

- A VAE-based approach for anomaly detection in experimental turbomachinery sensor data *Sven Durchholz, Marius Schwämmle, Baptiste Lambert*.
- Optimal Drone Control with Liquid Neural Networks *Tudor-Mihai Avarvarei, Mario Cassaro, Nicolas Drougard*.
- Improving Fault Diagnosis in Complex Systems: Investigating xAI Techniques for Abnormal Sensor Isolation *Baptiste Lambert*.
- Recent works at ONERA on the runway detection with Neural Network models for autonomous vision-based landing of an airliner *Sofiane Kraïem, Augustin Fuchs, Mario Cassaro, Tudor Avarvarei, Gustav Oman Lundin*.
- Ensuring the Safe Deployment of Neural Networks in Rocket Engine Control Systems *Jonas Dauer, Günther Waxenegger-Wilfing, Kai Dresia*.
- Investigating the limits of ad-hoc extension of single-intruder collision avoidance policies to multi-intruder scenarios *Thibault Roux, Filipo Perotto, Gauthier Picard*.

A VAE-based approach for anomaly detection in experimental turbomachinery sensor data

Sven Durchholz (DLR), Marius Schwämmle (DLR), Baptiste Lambert (DLR)

Presentation type
On-going collaborative work

Keywords
Machine Learning, Anomaly Detection, Data Reconstruction, Variational Autoencoder, Turbomachine

Introduction
Experimental turbomachinery operations are essential for aerospace research but are often constrained by high costs, complex setups, and limited testing time. To ensure efficient use of these valuable resources, rapid and reliable identification of anomalies in measurement data is critical. However, the diversity of sensor configurations across test facilities and the lack of standardized validation methods pose a significant challenge. To address this, we explore an unsupervised, data-driven anomaly detection approach, such an approach has been shown to be effective in similar use cases (1) (2). This work presents our progress in applying a Multi-Convolution Heads Variational Autoencoder (MCHVAE) for system health monitoring, using data from the Transonic Cascade Wind Tunnel (TGK) in Cologne.

Dataset
The dataset used for this project was collected by the TGK. The TGK is used to study the aerodynamic behavior of turbomachinery cascades and allows experiments under transonic conditions to investigate airflow phenomena (3). The dataset contains over four hours of data, consisting of low-frequency (10 Hz) and high-frequency (200 kHz) components. Furthermore, the data is segmented into ten sequences, from startup to shutdown. Sequences measure different operation points within the test facility or show transition states. For each operation point, all relevant variables were adjusted through their entire range to examine the effects and the system behavior.
The low-frequency data measures the pressure distributions on wall and blade surfaces, temperature variations, and operation parameters such as tailboard or throttle position. Analyzing the data using a Principal Component Analysis (PCA) shows clear clusters in those variables for each operation point as well as clear differences among startup, shutdown, measuring, and transitioning between operation points.
In addition to the low-frequency data, strain data from the investigated blade surface were measured at a high frequency. For this purpose, four points on the suction side and four on the pressure side were determined by a modal analysis for ideal sensor placement. These sensor positions are intentionally redundant in pairs to observe sensor failures or occurring anomalies. The measured strain data shows differences between the redundant sensors, which are most likely caused by the high Mach numbers in the TGK, but the cause can only be determined by a second measurement campaign, which makes the strain data currently unfit for the proposed model.
The data does not contain any known anomalies. Further measurements are scheduled to create a validation dataset with known anomalies.

Anomaly Detection
We propose using a Multi-Convolution Heads Variational Autoencoder (MCHVAE) for anomaly detection. Our proposed model is a convolution-based structure inspired by the Multi-Scale Temporal Variational Autoencoder (MSTVAE), which has been shown to effectively capture both intra-series temporal dependencies and inter-series correlations, while requiring less computational power than recurrent neural network (RNN)-based VAEs. Figure 1 shows the architecture of the model. Since our dataset does not contain anomaly labels, we evaluated our model's performance using a public dataset and

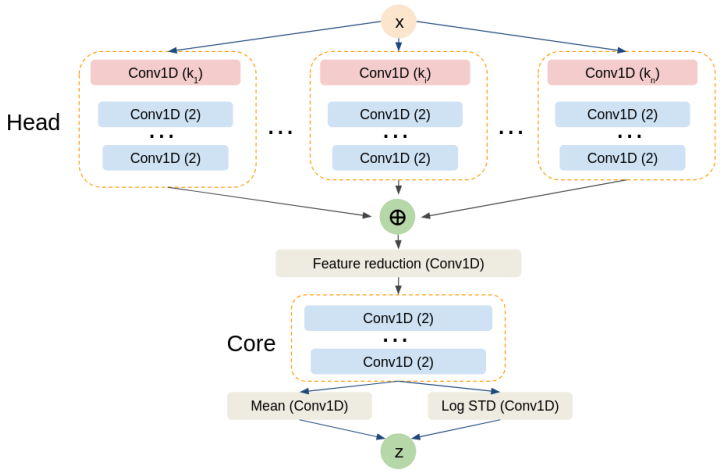


Figure 1: Architecture of the MCHVAE encoder

protocol from Liu et al (4). Our proposed model ranked third among deep learning models and was the top-performing VAE-based approach. Additionally, the MCHVAE shows great potential for our use case, as it can effectively detect anomalies and correct sensor signals.

Conclusion and further work

Our proposed model demonstrates promising results in the benchmark tests we conducted and is highly applicable to our use case, specifically for predicting anomalies in experimental measurements and for signal correction. Since our current data lacks anomaly labels, we scheduled future measurements to gather data with anomalies, which will enable a more effective evaluation of our model's performance. This will enable us to analyze how to better integrate the unique structure of our data into the model for reliable online anomaly detection and signal correction. Additionally, we are exploring methods to not only detect anomalies but also identify the most likely anomalous sensor, which will aid in improved troubleshooting.

References:

(1) Gori, V., Veneri, G., Ballarini, V.: Continual learning for anomaly detec- tion on turbomachinery prototypes - a real application. pp. 1–7 (07 2022). https://doi.org/10.1109/CEC55065.2022.9870234.

(2) Pawełczyk, M., Fulara, S., Sepe, M., Luca, A., Badora, M.: Industrial gas turbine operating parameters monitoring and data-driven prediction. Eksploat- acja i Niezawodnosc - Maintenance and Reliability 22, 391–399 (06 2020). https://doi.org/10.17531/ein.2020.3.2

(3) Hergt, A., Pesch, M., Grund, S., Flamm, J.: Low reynolds number effects in com- pressor blade design. International Journal of Gas Turbine, Propulsion and Power Systems 15, v15n4tp08 (08 2024). https://doi.org/10.38036/jgpp.15.4_v15n4tp08

(4) Liu, Q., Paparrizos, J.: The elephant in the room: Towards a reliable time-series anomaly detection benchmark. In: NeurIPS 2024 (2024)

Note : References are in Vancouver style

Optimal Drone Control with Liquid Neural Networks

Tudor-Mihai Avarvarei (ONERA - ISAE), Nicolas Drougard (ISAE) and Mario Cassaro (ONERA)

Presentation type
On-going collaborative work

Keywords
autonomous landing, quadrotor flight, liquid-time constant networks, learning, control

Introduction

Vision-based landing has emerged as a critical capability in the field of autonomous aerial systems, offering a potential solution for precise, GPS-independent navigation in complex or unstructured environment during the last phases of the flight. From a control perspective, this task requires the integration of visual data into real-time guidance and control loops, presenting challenges due to the nonlinearity, latency, and uncertainty inherent in the perception systems (1).

Existing approaches have primarily focused on the vision-related aspects, particularly on tasks such as runway detection and segmentation (2) (3). While significant progress has been made in this domain, the control component of the landing process has largely relied on classical solutions, including pose estimation techniques and traditional control algorithms. This separation between perception and control has left room to explore more integrated, learning-based alternatives.

In this context, Liquid-Time Constant Networks (LNNs) (4) present an intriguing research approach. Although less expressive than deep non-linear architectures, LNNs offer benefits in terms of interpretability (*i.e.* smaller network and lesser hyperparameter), and training (*i.e.* smaller dataset) — features especially desirable in safety-critical scenarios like autonomous landing, (5). Notably, LNNs have already shown promising results in end-to-end control tasks, such as driving, where they map raw visual inputs directly to control commands, (6).

The current work builds on these insights by investigating the control capabilities of LNNs in the context of vision-based landing. The focus is not on improving vision models for perception and control, but rather on understanding how LNNs can learn control strategies from visual inputs. As a first step, this study considers a simplified setting involving a drone, where an LNN is trained through supervised learning using a dataset tailored to perform basic flight control tasks. The goal is to assess the feasibility and limitations of LNNs in learning flying strategies and to lay the groundwork for more advanced applications in visual-based autonomous landing.

Problem Formulation

To investigate the control capabilities of various neural network architectures, a simplified drone control task is considered. The objective is to navigate a drone through way-points with a minimal power consumption strategy. While the task itself is relatively trivial, it provides a suitable environment for evaluating how different models might learn and generalize control strategies.

The dataset to be used in this study is sourced from the TU Delft drone control repository[1]. It was generated by solving an energy-optimal control problem, where the drone is driven from randomly sampled initial states to a final hovering position. This dataset is based primarily on the work of Ferede et al. (7), who proposed a framework for generating optimal control trajectories for quadrotors. The control system inputs (observations) are the drone's states, *i.e.* its pose and velocities. The outputs (controls) are the four motor commands, one per propeller, controlling their rotational speed.

Based on this framework, the planned work involves training and evaluating several neural network architectures for the control task such as the state-of-the-art classic neural networks (NN), recurrent neural networks (RNN) or continuous-time RNNs (CT-RNN). This setup is intended to provide insight into the learning and control capabilities of each architecture, with a particular focus on evaluating the potential of LNNs.

[1] https://github.com/tudelft/optimal_quad_control_SL

Implementation and Preliminary Results

Focusing on both LNN prediction accuracy and overall closed-loop dynamic behavior, performances are assessed through command time histories comparison against ground truth, in the first place, and through LNN on-the-loop full flight simulation consequently. Results in Figure 1 demonstrate the high level of accuracy reached by the proposed LNN architecture in learning a control dynamics.

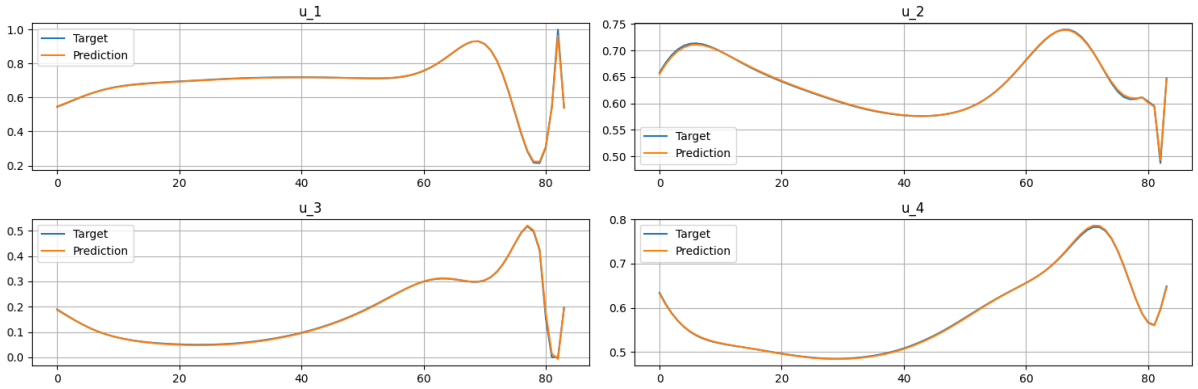


Figure 1: Motors commands time histories, LNN prediction (orange) *vs.* ground truth (blue)

The implementation of the LNN on-the-loop full flight simulation is still a work in progress. The preliminary objective is to control the drone from a random initial state to a stable hovering at a waypoint, while flying at minimal energy consumption. Pertinent results and discussion will be presented in the final version of the paper.

References:

(1) S. Kraïem, M. Cassaro, C. Seren et al., A vision-based control solution for autonomous approach and landing of an airliner. In Proceedings of the 2024 CEAS EuroGNC conference. Bristol, UK. (June 2024)

(2) Y. Li, Y. Xia, G. Zheng et al., YOLO-RWY: A Novel Runway Detection Model for Vision-Based Autonomous Landing of Fixed-Wing Unmanned Aerial Vehicles, Drones, 8(10) (2024)

(3) W. Dai, Z. Zhai, D. Wang et al., YOMO-Runwaynet: A Lightweight Fixed-Wing Aircraft Runway Detection Algorithm Combining YOLO and MobileRunwaynet, Drones, 8(10) (2024), p. 330

(4) R. Hasani, M. Lechner, A. Amini et al., Liquid Time-constant Networks, In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35 (12) (feb 2021).

(5) K. Kumar, A. Verma, N. Gupta and A. Yadav, Liquid Neural Networks: A Novel Approach to Dynamic Information Processing, 2023 International Conference on Advances in Computation, Communication and Information Technology, Faridabad, India, (2023), pp. 725-730

(6) R. Hasani, M. Lechner, A. Amini et al. Closed-form continuous-time neural networks. Nat Mach Intell. 4 (2022), p.992–1003

(7) R. Ferede, G.H.C.E. de Croon, C. de Wagter and D. Izzo, ”End-to-end neural network based optimal quadcopter control”, Robotics and Autonomous Systems, vol. 172, (Feb 2024)

Improving Fault Diagnosis in Complex Systems: Investigating xAI Techniques for Abnormal Sensor Isolation

Baptiste Lambert (Institute of Data Science - German Aerospace Center, Mälzerstraße 3-5, Jena 07745 Germany)

Presentation type
Result of research work

Keywords
xAI, anomaly detection, multivariate time series

Introduction

Anomaly detection (AD) in multivariate real-world data is a critical task across various domains, including industrial monitoring, healthcare, and cybersecurity. The ability to detect deviations from expected patterns enables early identification of system failures, fraudulent activities, or potential risks. However, conventional AD models often function as black boxes, providing little to no insight into the underlying reasons for detected anomalies. This lack of interpretability poses challenges for decision-making and corrective actions, especially in high-stakes environments where understanding the cause of failure is as important as detecting the anomaly itself.

In high-dimensional systems with thousands of variables, one of the main challenges operators face is identifying which part of the system is responsible for an anomaly. Most anomaly detection (AD) models only provide timestamps of detected anomalies, offering little insight into their root causes. To address this limitation, we explore explainable artificial intelligence (xAI) methods to determine whether they can help answer a key question: Which variables contribute the most to the AD model's prediction at the time of the anomaly? Additionally, to ensure generalizability, we restrict our investigation to model-agnostic xAI methods that do not require domain-specific knowledge about the monitored system.

Preliminary Results

Before exploring various xAI methods, the first challenge is selecting a suitable dataset for anomaly detection—one that includes clear labels indicating the origin of each anomaly. To keep the problem simple and easily explainable, we focus on meteorological conditions (wind speed, temperature, humidity, precipitation, atmospheric pressure, and rain cover) in the Atacama Desert, Chile, from 2000 to 2024. This region, the driest in the world, has experienced three unprecedented flash floods between 2015 and 2024. We define these floods as anomalies, with hourly accumulated precipitation serving as the clear variable responsible for them.

As stated in the introduction, our focused is to explore model-agnostic xAI methods. Specifically, we investigate the following three approaches:

- Window Shapley Additive Explanations (WSHAP) (1)
- (Standard) Shapley Additive Explanations (SHAP) (2)
- Counterfactual Explanations (CFE) (3)

Since CFE is an optimization problem, we consider two distinct approaches for solving it: a gradient descent-based method (CFE-Grad) and a genetic algorithm (CFE-GA).

To evaluate the performance of these methods, we use an explainability score defined as:

$$1 - \frac{1}{n} \sum_{i=1}^{n} \|S_i - T_i\|$$

where S_i represents the probability assigned by the xAI method to variable i as being responsible for the anomaly, and T_i denotes the true probability of variable i causing the anomaly. Table 1 presents the average performance of the different xAI methods, calculated using 10 distinct anomaly detection techniques.

	SHAP	WSHAP	CFE – Grad	CFE – GA
Average Score	0.782	0.872	0.890	0.677
Run Time/Inference	35.9k	5k	73.4k	80k

Table 1: Average score of the xAI methods using 10 different anomaly detection models. The last row shows the average run time of the xAI methods relatively to the cost of one prediction by the anomaly detection model.

Although CFE-Grad is the most accurate method, its computational cost makes it impractical compared to WSHAP in this scenario. However, SHAP is known for its computational complexity increasing exponentially with the number of features—a challenge that WSHAP aims to mitigate. Further investigation is required to understand WSHAP's scalability and efficiency in high-dimensional datasets.

The methods and results will be presented in greater detail during the workshop, along with additional examples using other datasets.

References:

(1) Nayebi A, Tipirneni S, Reddy CK, Foreman B, Subbian V. WindowSHAP: An Efficient Framework for Explaining Time-series Classifiers based on Shapley Values. J Biomed Inform. 2022;104438.

(2) Lundberg SM, Lee S-I. A Unified Approach to Interpreting Model Predictions. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. Red Hook, NY: Curran Associates Inc.; 2017. p. 4768-77.

(3) Wachter S, Mittelstadt B, Russell C. Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. Harv J Law Technol. 2018;31(2):841-87.

Recent works at ONERA on the runway detection with Neural Network models for autonomous vision-based landing of an airliner

Sofiane Kraïem, Augustin Fuchs, Tudor Avarvarei, Gustav Öman Lundin, Mario Cassaro (DTIS, ONERA, Université de Toulouse)

Presentation type
Result of research work

Keywords
Pose Base Vision Servoing, Neural Network model, Yolo, Runway Detection.

Introduction

Enhancing autonomy and safety in civil aviation, particularly in GPS-denied contexts, can significantly benefit from exteroceptive sensors. Recent advances in vision-based solutions and computer vision using Neural Networks (NN) for detection and segmentation (1), (2), (3), have shown promise in estimating aircraft pose relative to a desired runway. Additionally, dataset generation tools have recently been available, enabling the training of NN models for runway detection in images (4), (5). In this presentation, we propose to detail our recent progress made on vision-based control strategies for autonomous landing of an airliner using NN models.

Overall Pose Base Visual Servoing strategy

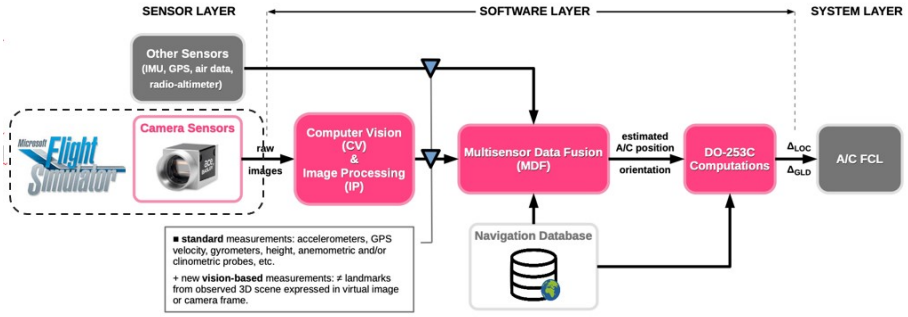


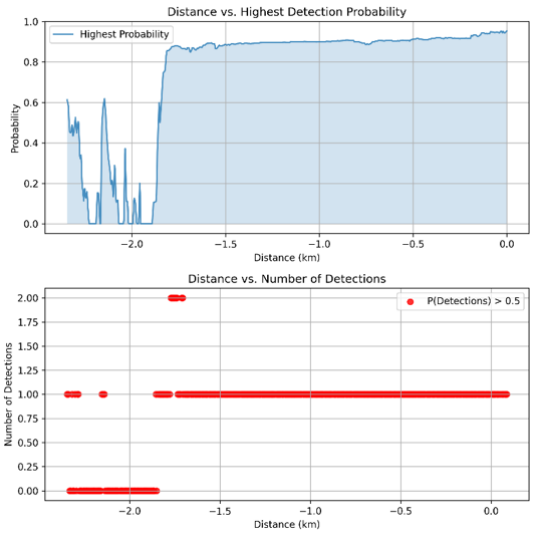
Figure 1: PBVS guidance law block diagram.

The proposed Pose Base Visual Servoing (PBVS) strategy, depicted in Figure 1, is an extension of our previously developed solution for autonomous airliner landing (6). This enhancement involves incorporating Artificial Intelligence (AI)-based measurements into the estimation of the aircraft pose. Specifically, pixel coordinates of detected runway corners serve as AI-based measurements, leveraging their corresponding real-world positions available in the onboard navigation dataset. Our proposed control strategy draws inspiration from Instrument Landing Systems (ILS), which estimate an aircraft's lateral and vertical deviations from the runway using visual measurements.

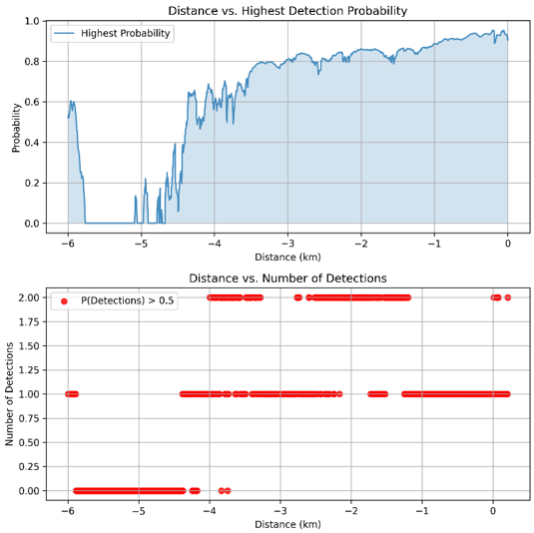
Computer Vision and Image processing algorithms

The presentation details the computer vision algorithms performing the derivation of AI-based measurements from image flows and their integration into the multi-sensor data fusion filter. Two critical aspects to integrate AI-based measurements in the estimation process are the inference times and accuracy of both detection and segmentation of NN models. Impacts of model's size and choice of datasets to obtain

satisfactory runway detection performances are discussed as illustrated in Figure 2.



(a) YOLOv8s (trained on LARD and BARS) with inference time of 0,005s



(b) YOLOv11s (trained onBARS) with inference times of 0,01s

Figure 2: Comparison of two YOLO models performances: number of runway detected and the associated confidence in function of the distance to the runway

Conclusion

The presentation gives an overview of the recent works made at ONERA on the runway detection using NN models for autonomous vision-based landing solutions. Both results in open-loop and closed-loop performed on a real-time simulation platform are presented to validate the control strategy and highlight its feasibility (7).

References:

(1) Li Y, and al. Yolo-rwy: A novel runway detection model for vision-based autonomous landing of fixed-wing unmanned aerial vehicles. Drones, 8(10):571, 2024.

(2) Dai W, and al. Yomo-runwaynet: A lightweight fixed-wing aircraft runway detection algorithm combining yolo and mobilerunwaynet. Drones, 8(7):330, 2024.

(3) Yu G, and al. Real-time vision-inertial landing navigation for fixed-wing aircraft with cfc-ckf. Complex & Intelligent Systems, 10(6):8079–8093, 2024.

(4) Wenhui Chen, and al. Bars: a benchmark for airport runway segmentation. Applied Intelligence, 53(17):20485–20498, 2023

(5) Ducoffe M, and al. Lard–landing approach runway detection–dataset for vision-based landing. arXiv preprint arXiv:2304.09938, 2023.

(6) Kraïem S, and al. A vision-based control solution for autonomous approach and landing of an airliner. In CEAS EuroGNC 2024, 2024

(7) Kraïem S, and al. A real-time simulation environment for aircraft vision-based autonomous landing system design. In 10th ICARA, pages 260–264. IEEE, 2024.

Ensuring the Safe Deployment of Neural Networks in Rocket Engine Control Systems

Jonas Dauer (DLR, Institute of space propulsion), Günther Waxenegger-Wilfing (DLR, Institute of space propulsion; University of Würzburg, Institute of Computer Science), Kai Dresia (DLR, Institute of space propulsion)

Presentation type
Result of research work

Keywords
Reinforcement learning, Safety, Propulsion

Introduction

The requirements for the control of a rocket engine are getting more and more complex. Reusability, throttling and multi-restart are just some aspects. Significant nonlinearities and safety-constraints must be handled. In addition, controllers for rocket engines have to make very quick decisions on resource-constrained hardware. Although neural networks have the ability to meet the mentioned requirements they are rarely used to control a safety-critical system. In the field of rocket engines in particular, neural networks have not been used at all to date.

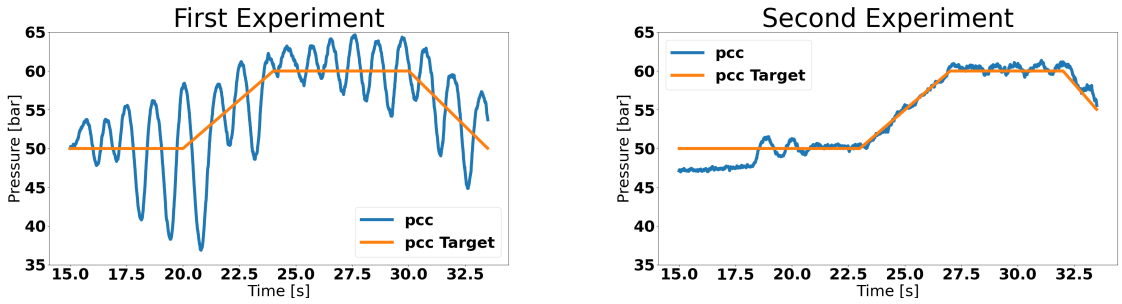
The Problem

This work will introduce an intelligent controller for the LUMEN (Liquid Upper stage demonstrator Engine) (1) engine. The controller is trained with the help of reinforcement learning (RL). Since the creation of real training data would be far too expensive and reinforcement learning learns according to the trial and error principle, it is inevitable that the agent is trained in a simulation. Previous work has already shown that training an intelligent agent in a simulation works very well (2). However, transferring the agent from the simulation to reality is often problematic. This problem is often described as the sim-to-real gap. The sim-to-real gap in reinforcement learning refers to the discrepancy or difference between the performance of an RL agent in a simulated environment and its performance when deployed in the real world. Addressing the sim-to-real gap in reinforcement learning for rocket engine control is critical for ensuring the safe and trustworthy transfer of trained agents from simulation to real-world applications.

Our Approach

We propose an algorithm designed to assess whether, and to what extent, a trained agent is capable of successfully operating within a given environment. The core of this method is an ensemble of neural networks that predicts the cumulative reward expected over the next n time steps. This predictive mechanism enables the system to estimate how well the agent is likely to perform in the near future, based on its current state and policy. To determine the agent's actual competence, the predicted rewards generated by the ensemble are compared to the rewards that are actually observed during interaction with the environment. Discrepancies between predicted and observed rewards serve as indicators of performance degradation or mismatch between training and deployment conditions. Additionally, the ensemble's standard deviation is used as a proxy for uncertainty in the predictions, offering insight into how confident the model is in its forecast. High uncertainty, especially when paired with low prediction accuracy, may signal that the agent is operating in a previously unseen or significantly altered environment. This approach allows for a systematic and interpretable evaluation of agent reliability, particularly in real-world deployment scenarios where safety and robustness are critical.

It is shown in two ways how it can be verified that the intelligent controller can be transferred. Firstly, deviations that did not occur during training are integrated into the simulation. It will be shown that the occurrence of such disturbances can be recognized. On the other hand, we will also apply the method to real-world data from actual LUMEN rocket engine tests and show that the method can be used to verify whether the trained agent's performance aligns with real-world behavior. By comparing the agent's actions in the simulation with data from the LUMEN engine real test runs, we can assess whether it is safe and reliable to transfer the trained agent from the simulated environment to the real-world rocket engine system. This application to real LUMEN data ensures the practical applicability and robustness of the intelligent controller in real-world conditions.
Figure 1 presents two real world trajectories that are compared to evaluate the performance of the algorithm. Figure 1a illustrates a trajectory executed using a controller that failed to transfer successfully from simulation to the real system. The underlying issue was an inaccurate modeling of certain valve dynamics within the simulation, which resulted in a substantial sim-to-real gap. Following the identification and correction of these modeling errors, the agent was retrained under the updated simulation conditions. The outcome of this retraining is shown in Figure 1b. As evident from the results, the controller's oscillatory behavior was significantly reduced, allowing the sequence to be executed successfully in the real system.



(a) Combustion chamber pressure during the first experiment. The pressure fluctuated significantly during the experiment.

(b) Combustion chamber pressure during the second experiment. The fluctuations have decreased significantly.

Figure 1: Combustion chamber pressure during the two experiments. The sequences vary slightly, which explains the differences in the trajectories.

We utilize the developed algorithm to demonstrate that the second policy is more capable of solving the problem, as the adjustments made to the simulation have led to a significantly improved alignment with the real-world system. As previously described, we use the error and the standard deviation of the ensembles to determine the width of the sim-to-real gap.

Conclusion

In summary, this work demonstrates that the trained reinforcement learning agent can be successfully transferred from simulation to real-world applications. By introducing deviations in the simulation and comparing the agent's behavior with real-world data, we can confirm that the intelligent controller is capable of safely and effectively operating in the real-world environment, ensuring high performance and reliability for rocket engine control.

References:

(1) Traudt, T.; Börner, M.; Suslov, D.; Saraf, M.; Groll, C.; Klein, S.; Deeken, J.; Hardi, J. and Schlechtriem, S., LUMEN: Test Platform for Rocket Engine Technologies. Project Overview and Upcoming Steps in the Development", in International Symposium on Space Technology and Science (ISTS), Kurume, Japan, 2023.

(2) Dresia, K.; Waxenegger-Wilfing, G.; Dos Santos Hahn, R. H.; Deeken, J.; Oschwald, M., Nonlinear Control of an Expander-Bleed Rocket Engine using Reinforcement Learning", in Space Propulsion 2020+1 Conference, Virtual Event, 2021.

Investigating the limits of ad-hoc extension of single-intruder collision avoidance policies to multi-intruder scenarios.

Thibault Roux, Filipo S. Perotto, Gauthier Picard - ONERA-DTIS-EDO

Presentation type
Result of research work

Keywords
Deep Reinforcement Learning, Airspace Deconfliction, Autonomous Collision Avoidance, Multi-Agent Systems.

Introduction

The rapid densification of airspace, due to the increasing number of flights and to the deployment of other technologies like unmanned aerial vehicles (UAVs), presents significant challenges for air-traffic management and requires the development of robust collision avoidance systems to mitigate the risk of near mid-air collisions between aircraft. This presentation reproduces the results published in (1), where we analyze the possibility of tackling these problems using deep reinforcement learning (DRL) and multi-agent DRL methods, which have the potential to compute good maneuvers in complex multi-aircraft scenarios. We focus particularly on the ACAS-X system.

ACAS-X

The Airborne Collision Avoidance System (ACAS-X) (3) (4) (2) is a next-generation solution designed to enhance aviation safety by overcoming the limitations of the currently deployed equipment, Traffic Collision Avoidance System II (TCAS-II). ACAS-X reduces both the risk of mid-air collisions and the frequency of unnecessary maneuvering alerts, enabling collision avoidance through cost tables computed by dynamic programming over a stochastic process that models an airspace encounter.

However, ACAS-X is inherently limited by its design, both because it relies on a simplified, discretized state and action space (5), and because it assumes a single-intruder scenario (1), which is extrapolated for the multi-intruder case. While these assumptions allow for efficient computation via dynamic programming, they become dangerous in dense airspace environments with multiple intruders, where the state space grows exponentially with the number of aircraft. Thus, this method has its safety compromised in the context of multiple aircraft operating in a dense airspace.

Let's consider a scenario where N aircraft operate at a fixed altitude in a 2D plane, which corresponds to the ACAS problem limited to its horizontal dimension. The objective is to minimize the occurrence of collisions while simultaneously reducing the number of necessary maneuvers. Once per second, each aircraft selects a heading adjustment from the set $A = \{0°/s, \pm 1.5°/s, \pm 3°/s\}$, corresponding to no change, a weak turn, or a strong turn to the left or right. A *Near Mid-Air Collision* (NMAC) is occurring when the separation between two aircraft falls to 500 feet or less (about 150 meters). Figure 1 shows the information available to the aircraft to make their decisions. Each aircraft i has access to its own velocity v_i, as well as the velocity v_j, distance $\rho_{i,j}$, relative heading $\theta_{i,j}$, and relative bearing $\psi_{i,j}$ to every other aircraft. It also has information about its own last action, a_i.

Other previous studies in the literature have extended ACAS-X methods to multiple intruders but always by decomposing the problem into single-intruder subproblems, which leads to suboptimal solutions. To achieve solutions that generalize to an unknown number of intruders, deep reinforcement learning (DRL) emerges as a promising approach. By exploiting function approximation through deep neural networks, DRL can be very effective in environments with large or continuous state and action spaces, making it well suited for the intricacies of dense airspace collision avoidance. Deep multi-agent reinforcement learning (MARL) has demonstrated its effectiveness in collision avoidance for robotics, allowing agents to learn coordination directly from raw sensor data (7). Similarly, MARL has been successfully applied to UAV swarm control, where decentralized agents must cooperate in dynamic environments (6). These achievements position MARL as a relevant approach for addressing multi-agent

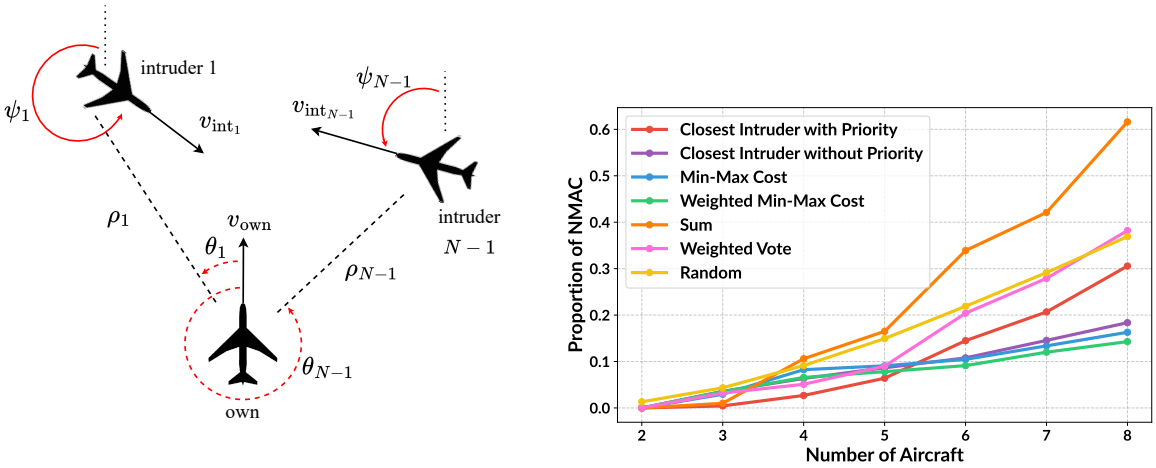


Figure 1: Illustration of the simplified 2D collision avoidance problem with several intruders.

collision avoidance in dense airspace.

Conclusion

We compared the proportion of Near Mid-Air Collisions for seven different ad-hoc strategies, with regard to the number of agents in the simulation, computed through 50000 random simulations. Each different strategy is a manner of aggregating the solution proposed to avoid each different intruder. These experimental results indicate that ad-hoc methods to face multi-intruder scenarios are not completely effective, asking for more robust strategies. Deep MARL is an interesting candidate, and the next step of this research is to investigate this hypothesis.

References:

(1) Roux, Perotto, Picard. Towards Scalable Collision Avoidance in Dense Airspaces with Deep Multi-Agent Reinforcement Learning. Proc. of ALA 2025. IFAAMAS.

(2) EUROCONTROL. 2022. Airborne Collision Avoidance System (ACAS) guide.

(3) Kochenderfer, Holland, Chryssanthacopoulos. 2012. Next-Generation Airborne Collision Avoidance System. Lincoln Lab Journal 19, 1 (2012), 17–33.

(4) Holland, Kochenderfer, Olson. 2013. Optimizing the Next Generation Collision Avoidance System for Safe, Suitable, and Acceptable Operational Performance. Air Traffic Control Quarterly 21, 3 (2013).

(5) Perotto. 2024. Using Deep RL to Improve the ACAS-Xu Policy: Concept Paper. In Proc. of the 2nd Int. Conf. on Cognitive Aircraft Systems: ICCAS. INSTICC, SciTePress, 110–117.

(6) Chen et al. 2024. Global-and-Local Attention-Based Reinforcement Learning for Cooperative Behaviour Control of Multiple UAVs. IEEE Transactions on Vehicular Technology 73, 3 (2024), 4194–4206.

(7) Long et al. 2018. Towards Optimally Decentralized Multi-Robot Collision Avoidance via Deep Reinforcement Learning. arXiv:1709.10082 [cs.RO].

AI for Engineering and Operation

Contents

- Deep Learning-Based Deterministic and Probabilistic Wind Nowcasting for Wake Vortex Warning Systems at Vertiports *Sooseok Lee, Frank Holzäpfel, Anton Stephan, Dieter Moormann*.
- Lost in Translation: Human Preferences in Multi-Criteria Decision Making *Sarah Ternus, Boris Djartov*.

Deep Learning-Based Deterministic and Probabilistic Wind Nowcasting for Wake Vortex Warning Systems at Vertiports

Sooseok Lee (Institut für Physik der Atmosphäre, Deutsches Zentrum für Luft- und Raumfahrt / Technologien für Kleinflugzeuge, Deutsches Zentrum für Luft- und Raumfahrt / Institute of Flight System Dynamics, RWTH Aachen University)
Frank Holzäpfel (Institut für Physik der Atmosphäre, Deutsches Zentrum für Luft- und Raumfahrt)
Anton Stephan (Institut für Physik der Atmosphäre, Deutsches Zentrum für Luft- und Raumfahrt)
Dieter Moormann (Institute of Flight System Dynamics, RWTH Aachen University)

Presentation type
On-going collaborative work | Result of research work

Keywords
Wind Nowcasting, Wake Vortex Prediction, Wake Vortex Warning System

The development of Urban Air Mobility (UAM) and vertiport operations has presented new challenges in predicting environmental hazards. One significant concern is wake vortices generated by large aircraft during landing or takeoff, which can impact electrical Vertical Take-Off and Landing (eVTOL) aircraft like air taxis operating at low altitudes near vertiports located near the runways of major airports. Given their relatively low mass and limited momentum compared to larger aircraft, VTOL vehicles are more vulnerable to such disturbances due to the little time and altitude available for recovery during encounters. Since the trajectory of wake vortices is highly influenced by wind strength and direction, it becomes an important factor to accurately predict wind patterns. In order to ensure safety and maintain the operational reliability of vertiport systems, it is essential to investigate under what wind conditions wake vortices can be transported to adjacent vertiports.

To address this issue, the Flight Safety of Small Aircraft under Adverse Meteorological Conditions (SMART-MET) project, led by DLR, is currently developing a dedicated Wake Vortex Warning System for Vertiports (WVWS-V), with an initial focus on Rome Fiumicino Airport. This system integrates three main components: (1) wind nowcasting using both statistical persistence methods [1] and various machine learning techniques, (2) wake vortex evolution simulation using the P2P model [2], and (3) encounter assessment based on wake-induced perturbation criteria. Within the project, wind nowcasting plays a fundamental role in accurately predicting vortex behavior.

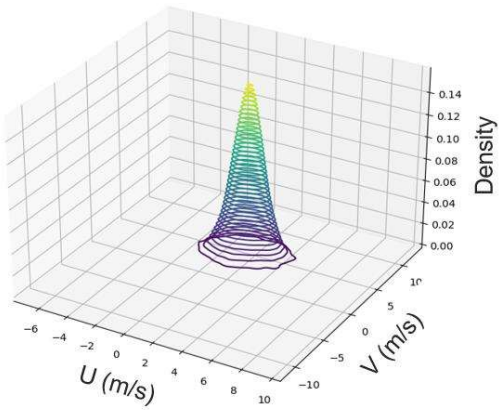


Figure 1: An example of the probabilistic nowcasting of wind in the u and v directions 16 minutes after the reference time. The superposed distribution is derived from the mean and deviation structure encoded in the dependent u–v matrix.

To be specific, this study addresses two aspects of short-term wind prediction: deterministic and probabilistic nowcasting, demonstrating the combined use of deep learning and probabilistic modeling for safety from wake vortices. Deterministic wind nowcasting is achieved using multiple machine learning models such as Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and Convolutional Neural Network (CNN), while integrating them into a hybrid model that incorporates cyclic time encodings to improve performance over longer prediction horizons. This study not only presents the results of comparing several models with different architectures, but also introduces an improved ensemble model that combines 2D CNN, LSTM, and a technique for capturing dynamic patterns using Multi-Resolution Analysis (MRA).

For probabilistic nowcasting, this introduces a Mixture Density Network (MDN) with superposed multivariate Gaussian distributions over wind components, capturing their empirical distribution more accurately than conventional models. A weighted Mahalanobis distance approach was developed to extract defined confidence intervals from mixture outputs, resulting in narrower prediction ranges of wind speed and direction compared to traditional statistical persistence models.

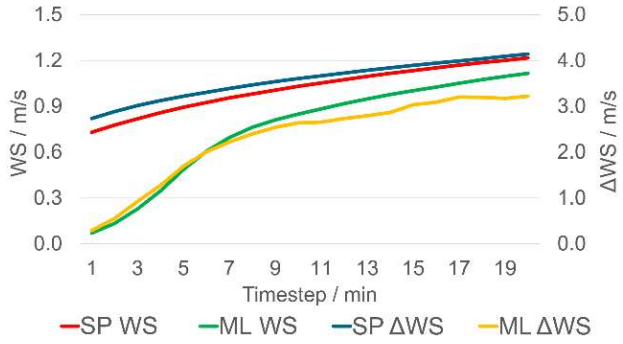


Figure 2: Model comparison results.
The primary Y-axis shows the Mean Absolute Error (MAE), while the secondary Y-axis represents the range between the minimum and maximum values required to achieve a 95% confidence interval.
The results indicate that the machine learning (ML) model outperforms the statistical persistence (SP) model in both accuracy and confidence interval compactness.

This deterministic and probabilistic wind nowcasting used anemometer data from 2018 to 2023 provided by UrbanV, with 2018 to 2022 used for training and 2023 for testing and evaluation. For benchmarking and comparison, a statistical persistence model [1] was employed. Overall, the proposed machine learning models demonstrated superior performance in both deterministic and probabilistic aspects. Notably, the results varied depending on the architecture used, with the hybrid model introduced for the first time in this study demonstrating the most outstanding performance. Nevertheless, there remain areas for improvement even within the hybrid model, which are also discussed in this work. Furthermore, the study demonstrated how the predicted wind conditions can lead to the intrusion of wake vortices into the flight paths of small aircraft. These results will provide practical solutions for safe integration of air taxis near busy runways, supporting future aviation safety systems.

References:

[1] Frech M, Holzäpfel F, Gerz T, Konopka J. Short-term prediction of the horizontal wind vector within a wake vortex warning system. Meteorol Appl. 2002 Mar;9(1):9-20.

[2] Holzäpfel F. Probabilistic Two-Phase Wake Vortex Decay and Transport Model. J Aircr. 2003 Mar-Apr;40(2):323-331.

Lost in Translation: Human Preferences in Multi-Criteria Decision Making for Alternate Airport Selection

Sarah Ternus (DLR, Institute of Flight Guidance, Human Factors Department), Boris Djartov (DLR, Institute of Flight Guidance, Human Factors Department)

Presentation type
On-going collaborative work

Keywords
Multi-Criteria Decision Making, Alternate Airport Selection, Preferences

Introduction

Expressing preferences in decision-making is a fundamental challenge, particularly when these preferences must be translated into mathematical formulations for computational systems. This problem is one of the key aspects within Multi-Criteria Decision Making (MCDM), as it requires stakeholders to specify their priorities among conflicting objectives. These methods assume that decision-makers can clearly articulate and quantify their preferences. Yet in practice, humans may not naturally express preferences in ways that align with MCDM requirements. Their reasoning often reflects training, experience, and situational judgment rather than predefined weights or consistent rankings. This paper explores the process of preference elicitation from pilots in the context of alternate airport selection and examines its compatibility with MCDM frameworks. By investigating how preferences are communicated, we aim to understand the challenges involved in integrating human decision-making with formal decision support systems.

Multi-Criteria Decision-Making

Decision-making is a fundamental process in both daily life and complex systems, involving the evaluation of multiple options based on different criteria. In most cases, decision-making is about navigating trade-offs—balancing competing priorities to arrive at the best possible outcome given the available information and constraints. Whether in personal choices, business strategies, or technical systems, selecting one option often means sacrificing certain advantages in favor of others.
Aviation provides a clear illustration of these trade-offs. One example is the Dynamic Alternate Airport Selection (DAAS) problem, where pilots or automated systems must choose an alternate airport in response to in-flight changes, such as weather conditions, fuel constraints, or mechanical issues. The decision involves weighing factors like proximity, weather at potential landing sites, airport capacity, and fuel consumption. Another example is air traffic control, where controllers must continuously manage airspace efficiency while ensuring safety (3)(4). Routing aircraft in order to optimize fuel efficiency and scheduling while also reducing collision risks, requires controllers to balance safety, efficiency, and regulatory constraints.

Preference Communication in Alternate Airport Selection

To better understand how human preferences manifest in real-world aviation scenarios, data was analyzed from recorded sessions in an A320 simulator. These sessions involved emergency situations such as hydraulic failures, engine fires, or fuel leaks, in which flight crews applied the structured FOR-DEC process—Facts, Options, Risks & Benefits, Decision, Execution, and Check. All conversations between pilots during these scenarios were transcribed for qualitative analysis, with a focus on how preferences and trade-offs were discussed during the selection of an alternate airport. The dialogues were originally in German and were translated by the authors.
The analysis of these transcribed dialogues reveals that preferences in alternate airport selection are rarely presented as explicit rankings of criteria or are rarely directly expressed. Instead, they emerge within the flow of discussion, through comparative reasoning, expressed tendencies, or diverging priorities between pilots. Some of the recurring patterns which were identified are:

Pilots frequently comparing specific attributes of options—such as crosswind components, weather, or infrastructure—highlighting in terms of what is ”better” or ”worse” in given context.

> *“Pisa has over 30 km/h CWC. Definitely not an option with the information we have. Light rain is probably still better than showers [...].”*

At certain times, pilots explicitly ask their crew member to express a directional preference.

> *A: ”What are you leaning toward?” B: “I’m leaning toward low visibility. We can’t really compensate for [the wind], [...] and in poor visibility we can at least do [...].”*

When a rationale is offered, it often reflects personal experience and intuition. Therefore, in some cases, participants emphasize different priorities, such as company operations, availability of maintenance, or safety, leading to negotiation.

> *A: “[Billund] is in the middle of nowhere. The operational impact for the company is just terrible. What do we want in Billund?” – B: “To get the aircraft down safely.”*

These exchanges show that preferences are shaped by specific values and experiences and not always aligned or resolved. In the end, the Captain decides what alternate airport is chosen and takes the First Officers recommendations into consideration. Overall, preferences seem to be situational, often implicit, and dynamically negotiated. They surface through practical reasoning, rather than being stated in abstract or structured form.

Preference expression difficulty

The illustration of how human preferences are communicated in practice highlights the inherent difficulty of translating them into structured inputs for MCDM frameworks. Traditional MCDM methods often rely on explicit preference elicitation techniques, such as pairwise comparisons or direct ratings which assume stable and clearly defined criteria weights. However, the simulator dialogues reveal that human reasoning frequently departs from these assumptions, favoring contextual, intuitive, and often inconsistent judgments (1).

Empirical studies and theoretical models (e.g., latent variable models) show that preference consistency is difficult to achieve in practice. Decision-makers may not exhibit transitivity, may weigh hidden or situational criteria, and may not differentiate clearly between personal preferences and stakeholder priorities. Moreover, even expert judgments, which are often treated as more reliable, are susceptible to biases and contextual variance (5)(6). While experts may bring valuable experience, their decisions do not necessarily reflect greater information use, but rather different weighting strategies (10). This raises the question of how best to integrate expert insights into MCDM frameworks, and whether such methods can account for nuanced, situational reasoning observed in real decision-making contexts (1).

Ultimately, the analysis of pilot behavior reinforces a broader insight: human preferences are not fixed inputs but evolving constructs. Capturing their richness requires MCDM models that are not only mathematically rigorous, but also psychologically informed, context-aware, and able to accommodate negotiation and ambiguity.

Future Work

This initial analysis opens up several directions for future work on integrating human preferences into decision support systems for aviation. Future work could hereby focus on developing MCDM models that integrate psychological realism, contextual sensitivity, and mechanisms for handling ambiguity and negotiation, to better align with the complexity of human decision-making. Another option is to investigate possibilities to translate how humans express preferences, often informally, contextually, and dynamically, into structured, machine-readable formats suitable for MCDM. Furthermore, it could be explored how preference-aware systems can visualize trade-offs, highlight conflicting priorities within crews, or adapt over time to individual decision styles. Evaluating how such systems affect trust, decision quality, and workload will be essential to ensure their usefulness and acceptance in operational settings.

References:

(1) Djartov B, Oswald F, Jakobi J. Through the psychological lens: unveiling biases in multi-criteria decision-making. In: Ahram T, Casarotto L, Costa P, editors. Human Interaction and Emerging Technologies (IHIET 2024). AHFE (2024) International Conference. AHFE Open Access; 2024. p. –. AHFE International, USA. doi:10.54941/ahfe1005476.

(2) Hwang CL, Yoon K, Hwang CL, Yoon K. Methods for multiple attribute decision making. Multiple attribute decision making: methods and applications a state-of-the-art survey. 1981:58-191.

(3) Djartov B, Mostaghim S, Papenfuß A, Wies M. Description and first evaluation of an approach for a pilot decision support system based on multi-attribute decision making. In2022 IEEE Symposium Series on Computational Intelligence (SSCI) 2022 Dec 4 (pp. 141-147). IEEE.

(4) Djartov B, Mostaghim S, Papenfuß A, Wies M. A Learning Classifier System Approach to Time-Critical Decision-Making in Dynamic Alternate Airport Selection. In2024 IEEE Congress on Evolutionary Computation (CEC) 2024 Jun 30 (pp. 1-8). IEEE.

(5) Odu GO. Weighting methods for multi-criteria decision making technique. Journal of Applied Sciences and Environmental Management. 2019 Sep 11;23(8):1449-57.

(6) Borgstede M, Eggert F. Squaring the circle: From latent variables to theory-based measurement. Theory & Psychology. 2023 Feb;33(1):118-37.

(7) Thurstone LL. A law of comparative judgment. InScaling 2017 Jul 5 (pp. 81-92). Routledge.

(8) Day J. What is an expert?. Radiography. 2002 May 1;8(2):63-70.

(9) Weinstein BD. What is an expert?. Theoretical medicine. 1993 Mar;14:57-73.

(10) Shanteau J. How much information does an expert use? Is it relevant?. Acta psychologica. 1992 Oct 1;81(1):75-86.

AI for Sensing and Image Processing

Contents

- Enhancing Composite Micrograph Analysis with Semantic Segmentation *Jonas Naumann, Jonas Appels, Philipp Sämann, Timo de Wolff, Christoph Brauer* .
- AIRforschen: AI-Enhanced Meteorological Data Collection from General Aviation for Improved Weather Forecasting *Sreerag V Naveenachandran, Julia Fligge-Niebling, Christian Thiel* .
- onboard semantic compression for remote sensing payloads *Protim Bhattacharjee, Peter Jung* .
- Synthetic Images for Training Deep Learning Perception Models *Joachim Rüter*.

Enhancing Composite Micrograph Analysis with Semantic Segmentation

Jonas Naumann (DLR, Institute of Lightweight Systems, Stade), Jonas P. Appels (DLR, Institute of Lightweight Systems, Stade), Philipp Sämann (DLR, Institute of Lightweight Systems, Stade), Timo de Wolff (TU Braunschweig, Institute of Analysis and Algebra), Christoph Brauer (DLR, Institute of Lightweight Systems, Stade)

Presentation type Result of research work

Keywords Carbon Fiber Composites, Micrograph Analysis, Semantic Segmentation

Introduction

Hydrogen-powered aircraft are a key element for future CO2-efficient air transport. This requires special transport tanks for liquid hydrogen (LH2), which must be integrated into the aircraft structure [1]. A critical requirement is tightness at cryogenic temperatures. The materials face significant challenges due to repeated temperature cycles between LH2 (-252.87 °C) and ambient conditions, especially during maintenance when the tank is fully drained. Suitable material combinations and manufacturing methods are needed to withstand these conditions. Carbon fibre-reinforced polymers (CFRP) offer high strength and low weight [2] but are prone to cracking under thermal cycling [3], which threatens LH2 tightness [4]. Therefore, relevant material and process parameters affecting crack susceptibility must be identified. To study this, small cylindrical specimens undergo repeated immersion in liquid nitrogen (LN2), with microscopy after each cycle to monitor material changes. While visual inspection of micrographs is common, it is time-consuming and yields qualitative results. Machine learning offers a promising alternative for automating tasks like microcrack detection and evaluating the local fibre volume ratio (FVR)—a key indicator of mechanical performance [5]. Local FVR variations may compromise strength and promote cracks, making their analysis in micrographs highly relevant.

Contribution

In this work, we discuss the use of semantic segmentation to inspect CFRP micrographs with a special focus on crack formation. The contribution of this work is fourfold:

1) We present a multi-step labelling pipeline for the CFRP micrographs that we used to create the database required for supervised machine learning. The image database consists of micrographs of seven planar and cylindrical specimens that differ in terms of mechanical and thermal manipulation, direction of the cross section and fibre architecture.

2) We introduce deep learning models for semantic segmentation that, when trained on the training split consisting of more than 3,000 patches of size 512×512 px, successfully distinguish between the different visual components of the CFRP specimens (see Figure 1). Thereby, they provide starting points for an

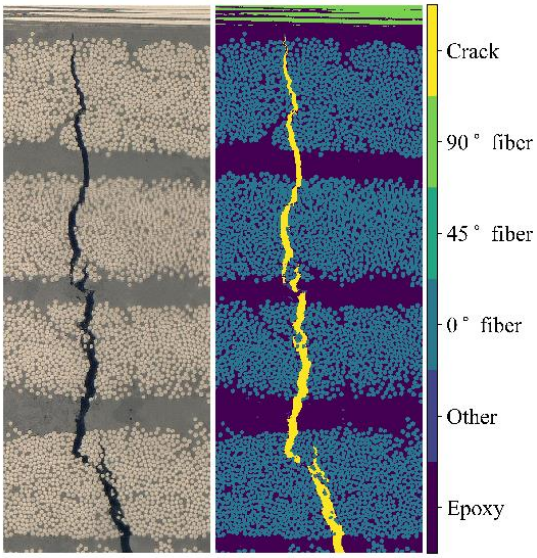


Figure 1: An exemplary micrograph excerpt with a wide crack from the test data split (l) and the corresponding segmentation model output (r).

advanced analysis and drastically reduce the time required for manual inspection. Our best model is based on the U-Net [6] and InternImage [7] and achieves a mean intersection over union (mIoU) of 0.9071 on the test set.

3) We demonstrate two applications of the model's outputs: The first one simplifies and accelerates the search for microcracks by highlighting coarse regions of the micrograph which contain at least one pixel with a crack probability exceeding a specified threshold. A threshold of 20 % yields a crack pixel recall of 97 % which represents a significant improvement compared to the model's crack pixel recall of 78.16 %.

4) The other application tool enables a quantitative investigation of local FVR variations. This is achieved by estimating the latter pixelwise by the ratio of 0° fibre and epoxy pixels in a specified neighbourhood in the segmentation mask. The results are visualized as heatmaps. Comparing the estimated, global FVRs for four specimens to the true FVRs measured with a thermogravimetric analysis (TGA) yields an average deviation of 2.16 %pt..

Conclusion

In conclusion, the proven accuracy of the segmentation model and the tools based on it allows us to state that they greatly simplify, accelerate, and quantify the analysis of fibre composite micrographs.

References:

(1) How to store liquid hydrogen for zero-emission flight. Last accessed on 20 Mar. 2025. https://www.airbus.com/en/newsroom/news/2021-12-how-to-store-liquid-hydrogen-for-zero-emission-flight.

(2) Wiedemann, M. System Lightweight Design for Aviation. Cham: Springer, 2023.

(3) Grogan, D., Leen, S., Semprimoschnig, C., & Ó Brádaigh, C. "Damage characterisation of cryogenically cycled carbon fibre/PEEK laminates". Composites Part A: Applied Science and Manufacturing 66 (2014): 237-250.

(4) Grogan, D. "Damage and permeability in linerless composite cryogenic tanks". PhD thesis, National University of Ireland, Galway, 2015.

(5) Schürmann, H. "Wichtige Kenngrößen der Einzelschichten und des Laminats". Konstruieren mit Faser-Kunststoff-Verbunden. Berlin, Heidelberg: Springer, 2005.

(6) Ronneberger, O., Fischer, P. & Brox, T. "U-Net: Convolutional Networks for Biomedical Image Segmentation". Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. Munich, Germany, October 5-9, 2015. Navab, N., Hornegger, J., Wells, W. & Frangi, A. (eds). Springer. pp. 234–241.

(7) Wang, W., Dai, J., Chen, Z., Huang, Z., Li, Z., Zhu, X., Hu, X., Lu, T., Lu, L., Li, H., Wang, X. & Qiao, Y. "InternImage: Exploring Large-Scale Vision Foundation Models with Deformable Convolutions". IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Vancouver, BC, Canada, June 17-24, 2023. pp. 14408-14419.

AIRforschen: AI-Enhanced Meteorological Data Collection from General Aviation for Improved Weather Forecasting

Sreerag V Naveenachandran, Dr. Julia Fligge-Niebling, apl. Prof. Dr. Christian Thiel

Presentation type

On-going collaborative work

Keywords

AI, anomaly detection, data compression, weather forecasts, sensor data, aviation, aerospace, crowd-sourcing, weather services.

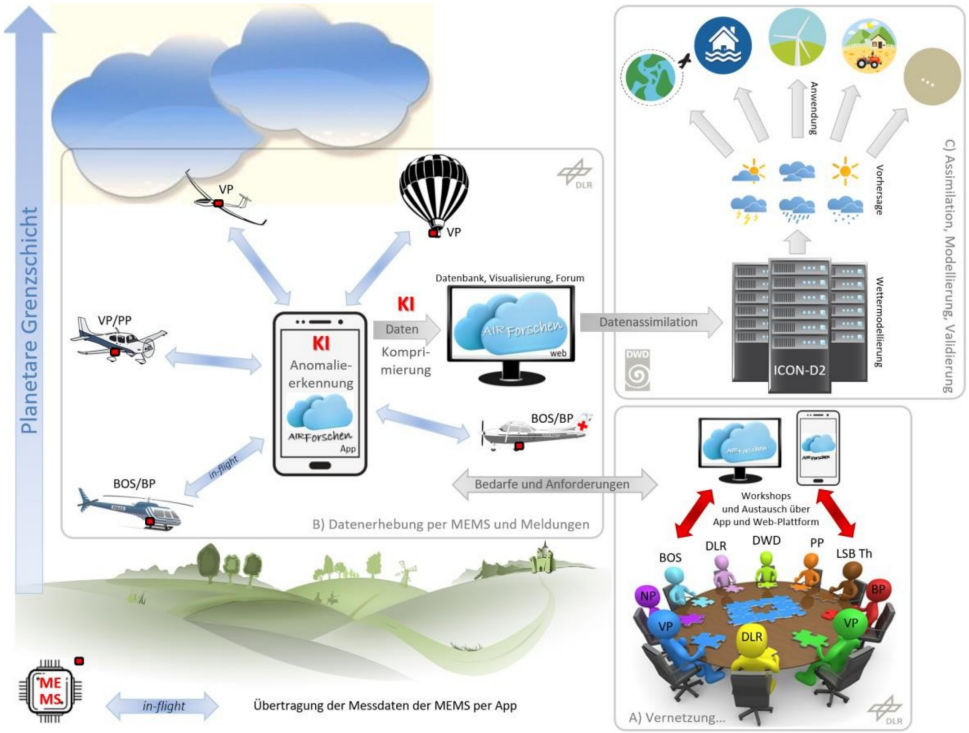


Figure 1: AIRforschen structure (KI: artificial intelligence, MEMS: microelectromechanical systems, VP: club pilot, PP: private pilot, BP: professional pilot)

Introduction

Accurate weather forecasting is vital for safe and efficient traffic management, optimizing renewable energy systems, supporting precision agriculture, and preparing for extreme weather events. Modern forecast systems rely on advanced numerical meteorological models, such as the DWD's ICON model and Météo-France's ARPEGE, which depend heavily on the frequency and spatial distribution of observational data. These data are gathered from various sources, including ground-based networks, weather satellites, commercial aviation, and remote sensing technologies. Although these sources provide valuable meteorological data, significant gaps remain, particularly within the planetary boundary layer (up to 2 km), due to technical limitations, typical flight altitudes, and restricted coverage. Studies indicate that increasing the measurement density in sparsely monitored regions can greatly enhance the accuracy of the forecast. (1) (2).

To address these limitations, AIRforschen aims to improve data collection by incorporating data from general aviation pilots. These pilots include authorities with security tasks (e.g., police, customs), BOSs, professional pilots registered with DAeC (over 100,000 members in Germany) and club pilots. This initiative builds on the successful Voluntary Observing Ship Programme (VOS) by NOAA. Using a network of these pilots equipped with sensor kits, AIRforschen aims to bridge existing data gaps and improve meteorological models with more frequent and spatially diverse observations.

AI-based anomaly detection & data compression

Anomaly detection (3) is crucial in maintaining data reliability by identifying irregular sensor readings and potential faults in real time. Deep learning methods, such as diffusion models or autoencoders, excel in learning compact representations of high-dimensional data, making them particularly effective in detecting anomalies. These algorithms can automatically detect & exclude or impute inaccurate readings, or trigger timely replacement of defective sensors.

Beyond anomaly detection, AI-based techniques enable efficient data compression, reducing storage and transmission demands while preserving data integrity. However, deploying these AI-driven solutions on pilots' smartphones presents hardware constraints, necessitating careful optimization of computational resources to ensure seamless operation.

Conclusion

The AIRforschen project marks a significant advance in weather forecasting by integrating data from general aviation aircraft into existing meteorological models. Using voluntary contributions from aviation stakeholders, it establishes a scalable, cost-effective, and sustainable approach to atmospheric data collection. AI-driven anomaly detection and data compression further enhance data quality, ensuring efficient collection and transmission. This campaign fills critical observational gaps, particularly within the planetary boundary layer, improving forecast accuracy. The data collected will be publicly accessible, benefiting both the scientific community and the general public. Additionally, AIRforschen data will be integrated into DWD's meteorological models, and in collaboration with third-party partners, prototypes of advanced weather services for the transport and mobility sectors will be developed.

Enhanced weather forecasting will support key industries such as transportation, energy, and urban development. Real-time weather information will improve traffic management, optimize electric vehicle performance, and strengthen infrastructure resilience through better maintenance planning and disaster preparedness. In addition, accurate meteorological data will help utilities optimize power distribution and mitigate disruptions caused by extreme weather conditions.

References:

(1) Pinto, J.O. et al. (2021) 'The status and future of small uncrewed aircraft systems (UAS) in operational meteorology', Bulletin of the American Meteorological Society, 102(11), pp. E2121–E2136. doi:10.1175/bams-d-20-0138.1.

(2) Dong, J., M. Xue, and K. Droegemeier, 2011: The Analysis and impact of simulated high-resolution surface observations in addition to radar data for convective storms with an ensemble Kalman filter. Meteor. Atmos. Phys., 112, 41-61, https://doi.org/10.1007/s00703-011-0130-3.

(3) A. B. Nassif, M. A. Talib, Q. Nasir and F. M. Dakalbab, "Machine Learning for Anomaly Detection: A Systematic Review," in IEEE Access, vol. 9, pp. 78658-78700, 2021, doi: 10.1109/ACCESS.2021.3083060.

Onboard semantic compression for remote sensing payloads

Protim Bhattacharjee, Peter Jung (Institute of Optical Sensor Systems, German Aerospace Center (DLR), Berlin)

Presentation type
Result of research work

Keywords
compressed learning, semantic compression, learning structured matrices, remote sensing payloads.

Introduction
Earth observation plays a crucial role in the assessment, protection, and sustainable usage of resources on the planet. Remote sensing platforms such as satellites, aircrafts, balloons, UAVs, and HAPs are used to conduct surveys of the Earth's surface and beyond to produce actionable data for socio-economic activities. In most cases, an imaging payload, like an camera or a spectrometer, is responsible for acquiring the images which are further analyzed. For perspective, the current Copernicus-Sentinel (1) mission archive is estimated to be around 91 PB. These huge datasets upon acquisition need to be downlinked from the remote platforms to the ground station. Though the no. of platforms available for acquisition is increasing everyday, the corresponding increase in the downlink capabilities is not at-par. There exists a downlink bottleneck that needs to be addressed with data-centric compression methods and system architectures. Moreover, often one is not interested in the imagery but in the final analysis, downstream task as it is known in remote sensing literature, that needs to be concluded. Thus, instead of downlinking the entire image one may just downlink the necessary features required for the downstream task. As most analysis is currently done using an encoder-decoder type neural architecture, the straightforward way would be to implement the encoder onboard the remote platform to generate the features from the sensed imagery. However, the SOTA neural network models require intensive memory and power, which is limited on the considered platforms. To address this challenge we propose an end-to-end learnable semantic compression pipeline based on compressed learning where a learned linear compressor is implemented onboard and the compressed features are downlinked for the downstream task. Figure 1 describes the entire pipeline. The complete semantic communication pipeline consists of a learned low-complexity compression matrix that acts on the noisy camera output to generate onboard a vector of observations that is downlinked through a communication channel, processed through an unrolled network and then fed to a deep learning model performing the necessary downstream tasks; image classification is studied.

GLODISMO + Classifier (2)
The proposed compression pipeline consists of three main learnable parts, a compression matrix, a recovery network, and a classifier. All the three parts are learned in an end-to-end manner via backpropagation. The compression matrix and the recovery model together form the GLODISMO (3) sub-network, which is followed by a classifier sub-network, and together they form the proposed model. GLODISMO allows learning of sampling distributions with discrete structural constraints by applying the Gumbel-softmax trick during backpropagation. This allows learning of a distribution of structured binary matrices from which compression matrices can be sampled that are implementable on constrained hardware. To address distortions that may be introduced by the camera hardware and the communication channel, a few layers of an unrolled reconstruction network, specifically NA-ALISTA (4) with a wavelet sparsity is used to add robustness to the complete pipeline. To ensure semantic flow through the network, the compression matrix, the unrolled algorithm and the classifier are jointly learned in an end-to-end manner. Training is carried out with two different losses, classification loss with and without the reconstruction loss.

Further to ensure efficient compression of the compressed vector, a learnable entropy encoder is

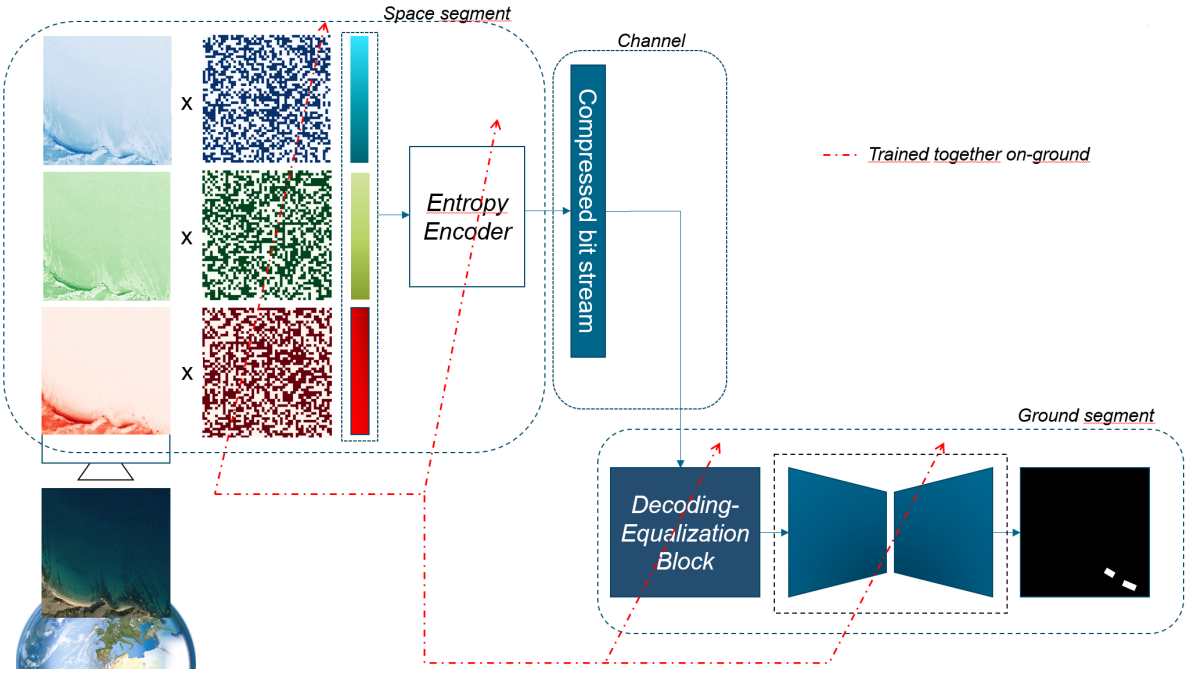


Figure 1: Proposed compressed learning based semantic communication pipeline

added after the compressor. This allows for a trade-off study between the rate-distortion and semantic loss. The end-to-end training ensures that the best codes are learned with respect to the required downstream task.

Conclusion

In this work an end-to-end compressed learning framework is designed and studied for application in onboard operations of remote sensing platforms. The framework is able to learn low-complexity structured binary matrices that can be implemented on low cost off the shelf hardware. The parameters of the compression matrix, entropy encoder, reconstruction algorithm, and the classifier are tied together such that the compression pipeline is learned according to a given application reducing the bandwidth requirements for downlinking. It is shown that addition of a recovery loss along with the task dependent losses improves the downstream performance in noisy settings at low compression ratios. As the framework is plug-and-play in nature, different sensor constraints, such as spectral sampling, noise robust algorithms and downstream tasks can be included with minimal compute or memory overheads.

References:

(1) Copernicus Data Space Ecosystem [Internet], [cited 2025 Apr 9]. Available from: https://dashboard.dataspace.copernicus.eu/

(2) Bhattacharjee P, Jung P. Compressed learning-based onboard semantic compression for remote sensing platforms. Proc SPIE. 2024 Nov 19;13196:1319608.

(3) Sauder J, Genzel M, Jung P. Gradient-based learning of discrete structured measurement operators for signal recovery. IEEE J Sel Areas Inf Theory. 2022;3(3):481–92.

(4) Behrens F, Sauder J, Jung P. Neurally augmented ALISTA. In: Proceedings of the International Conference on Learning Representations (ICLR); 2021.

Synthetic Images for Training Deep Learning Perception Models

Joachim Rüter (German Aerospace Center (DLR) - Institute of Flight Systems - Unmanned Aircraft)

Presentation type
Result of research work

Keywords
Synthetic Data, Deep Learning, Sim-to-Real Gap, Game-Engine

Introduction

Future autonomous unmanned aircraft systems (UAS) require robust environmental perception capabilities, for example, to detect emergency landing sites or to plan safe flight paths. Deep learning algorithms have shown state-of-the-art performance for many perception tasks. However, they rely heavily on large and diverse datasets for training and evaluation. In aviation, such datasets are often unavailable, particularly for edge cases and adverse weather conditions. Furthermore, manual annotation of these datasets is labor-intensive, costly, and prone to inconsistencies.

Meanwhile, visual simulation environments, such as modern game engines, have become increasingly realistic. These environments allow to extract visual sensor data and automatically generate corresponding annotations across a wide variety of scenarios, at significantly lower cost and much higher speed. While synthetic data is promising for training deep learning models, models trained solely on synthetic data often struggle to generalize well to real-world settings. This issue is commonly referred to as the sim-to-real gap.

This work presents research into the effectiveness of training deep learning perception models for aviation use-cases using synthetic images generated with modern game engines. The study explores to what extent such models can be trained on synthetic data and still perform effectively on real-world images.

Method

This study investigates the performance of deep learning perception models trained on synthetic images for multiple use-cases. The generalization capability of the investigated models is evaluated using separate real-world test datasets. The experiments further explore influences on the generalization challenges.

The first use-case explores drogue detection during air-to-air refueling. This scenario presents an interesting combination of conditions that may result in a good generalization from synthetic to real-world images. Specifically, the high altitude and the limited variability of participating objects reduce the occurrence of shadows, reflections, distractions, and occlusions, resulting in visually simple images. Furthermore, the limited number of objects in sight, as well as the absence of complex ground textures, contribute to a more controlled and reproducible environment.

The second use-case explores semantic segmentation from a low-altitude perspective of a UAS. The use-case includes common objects such as buildings, roads, vehicles, vegetation, and people. The presentation describes the data generation process and also discusses key challenges in generating synthetic images and provides potential solutions to improve realism. Example images of the real-world dataset and our synthetic counterpart are shown in Figure 1.



Figure 1: Comparison of exemplary scenes from the real-world dataset (1) and our synthetic counterpart (2)

Results and Conclusion

The experiments show promising results for using synthetic images to train deep learning perception models. Especially for the drogue detection use-case, model trained on synthetic data achieve good real-world performance. This suggests that synthetic data can serve as a promising substitute when matching real-world datasets are unavailable.

For the semantic segmentation use-case, the experiments reveal that synthetic images alone are not enough to train competitive deep learning models. However, combining synthetic with real-world images can reduce the amount of real-world images needed while maintaining the same performance. This allows to reduce the manual annotation burden and reduce the cost of real-world data generation. Additionally, the results show that combining synthetic and real-world images can even lead to better performance than training on real-world images alone.

Overall, the findings highlight the value of synthetic images and show promising directions for future research. Additionally, the presentation offers initial insights into factors that contribute to the sim-to-real gap, such as the model architecture (4) and shows potential application fields such as using synthetic images to increase the coverage of an Operational Design Domain (ODD) (5).

References:

(1) Marcu A, Licaret V, Costea D, Leordeanu M. Semantics through Time: Semi-supervised Segmentation of Aerial Videos with Iterative Label Propagation. Asian Conference on Computer Vision. 2020.

(2) Hinniger C, Rüter J. Synthetic Training Data for Semantic Segmentation of the Environment from UAV Perspective. Aerospace. 2023.

(3) Rüter J., Schmidt, R. Using Only Synthetic Images to Train a Drogue Detector for Aerial Refueling. International Conference on Modelling and Simulation for Autonomous Systems. 2023.

(4) Rüter J, Durak U, Dauer JC. Investigating the Sim-to-Real Generalizability of Deep Learning Object Detection Models. Journal of Imaging. 2024.

(5) Rüter J, Maienschein T, Schirmer S, Schopferer S, Torens C. Filling the Gaps: Using Synthetic Low-Altitude Aerial Images to Increase Operational Design Domain Coverage. Sensors. 2024.

AI Engineering Methods and Approaches

Contents

- Data-Driven AI Surrogate Model for Rapid 3D Flow Approximation in Axial Fans *Krishna Srinitha Vithanala*.
- Contrail Optimization for North-Atlantic Traffic *Xavier Pucel*.
- Coupling of a Genetic Algorithm and a Thermoelectric Network Model for Radioisotope Thermoelectric Generator (RTG) Optimization *Stylianos Kyrimis, Eliott Lallemand, Christian Stiewe, Pawel Ziolkowski, Eckhard Müller, Johannes De Boor*.
- Designing Satellite Assembly Lines from Experts' Objectives *Anouck Chan, Thomas Polacsek*.
- Towards Smarter Engineering Workflows Using Large Language Models *Stanislaus Reitenbach, Martin Siggel, Martin Bolemant*.
- Variational autoencoder with weighted samples for high-dimensional non-parametric adaptive importance sampling *Julien Demange-Chryst, Jérôme Morio, Francois Bachoc*.

Data-Driven AI Surrogate Model for Rapid 3D Flow Approximation in Axial Fans

Krishna Srinitha Vithanala, *Marcel Aulich*, *Christian Voß*
German Aerospace Center (DLR), Cologne, Germany

Presentation type
On-going collaborative work

Keywords
Artificial Intelligence (AI), Transformer Architecture, Axial Fans, Surrogate Modeling, Turbomachinery, Deep Neural Networks (DNN), Computational Fluid Dynamics (CFD), Reynolds-averaged Navier–Stokes (RANS) Simulations

Introduction

Computational Fluid Dynamics (CFD) simulations are essential for the design and optimization of various engineering systems, but they require significant computational resources and time. This work introduces a data-driven AI surrogate model designed to quickly approximate 3D CFD results for axial fans, significantly reducing the computational cost of obtaining 3D flow solutions. The AI model uses existing CFD data to understand the complex relationships between geometric designs, boundary conditions, and flow solutions. This enables the model to predict flow patterns for new, previously unseen geometries.

Application

The application of this AI model is in predicting flow for industrial axial fans. The model is trained using unstructured CFD data from turbomachinery flow applications to approximate the low Mach number, incompressible flows characteristic of these fans.

AI Model Architecture

The AI model is trained using a Design of Experiments (DOE) database of geometric surface points (XYZ coordinates), rotational speed, and inlet/outlet boundary conditions. The model predicts key flow variables, including velocities, pressure, density and wall shear stresses.

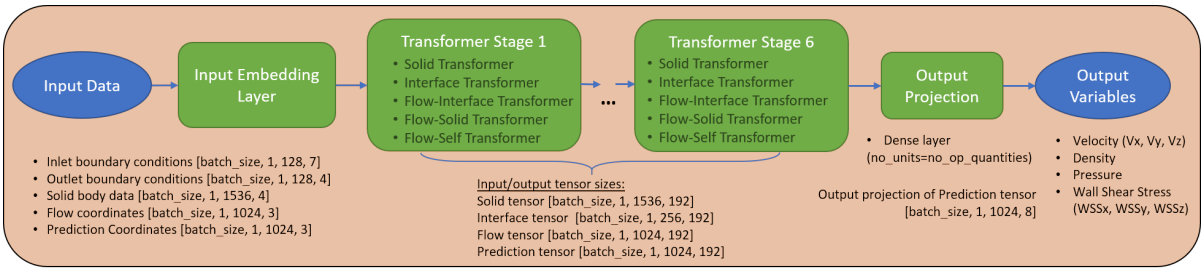


Figure 1: AI model architecture

The AI model utilizes a transformer-based architecture to predict 3D flow fields. This architecture incorporates multi-head attention mechanisms (2) to effectively capture spatial dependencies within the input data. The model's input layer processes raw input data through a positional encoding layer, converting it into a high-dimensional sine-cosine representation. This approach enhances the representation of input data to facilitate training by scaling numerical values to a range of -1 to 1 for numerical stability and representing values across different scales for improved function approximation.

Dataset and Hardware

A Design of Experiments (DOE) dataset, comprising 900 unique axial fan configurations, was used for model training. A separate test dataset containing 320 optimized fan designs was used to evaluate the model's ability to generalize to new geometries. The AI model was trained using an NVIDIA RTX A6000 GPU with 48GB memory and CUDA 12.6. TensorFlow and JAX were used to build and train the model, efficiently handling large CFD datasets. The training and validation data were generated from RANS simulations performed in STAR-CCM+ (3).

Results

The AI model's performance was evaluated by comparing its predictions to CFD simulations. Qualitative comparisons of flow fields showed that the model captures essential flow features. Quantitative analysis for train, validation (0.5% DOE) and a test (optimized fan designs) demonstrated an agreement between the model's predictions and CFD results for key performance indicators, including mass flow rate, and area-averaged pressure.

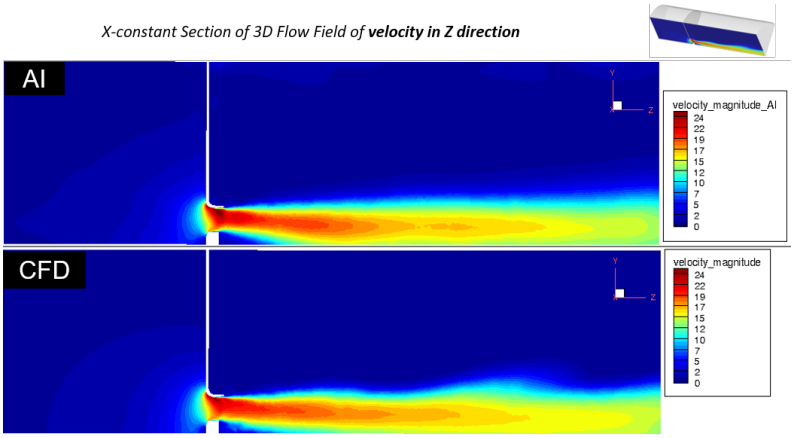


Figure 2: Comparison of Velocity-Z Field Predictions between AI and CFD for a Validation Design at OP1

Figure 2 presents a comparison of flow field predictions between the AI model (top) and a CFD simulation (bottom) for a validation design at OP1 (free blowing point). This design, while part of the Design of Experiments dataset, was not included in the AI model's training. A qualitative assessment indicates that the AI approximation captures the overall flow pattern reasonably well, particularly the high-velocity regions near the impeller and the gradual dissipation of velocity downstream.

Conclusion

This research successfully developed and validated an AI surrogate model for 3D flow approximation in axial turbomachinery. The model demonstrates its capacity to approximate flow fields for unseen fan designs using unstructured CFD data by generalizing to a wide range of geometries and operating conditions. However, discrepancies remain between AI predictions and CFD results, particularly in pressure and velocity distributions. To address these limitations, future research will focus on enhancing accuracy by incorporating turbulence quantities into the loss function, improved metrics, and expanding the training dataset with diverse boundary conditions. Sampling strategies are being optimized to reduce errors in sensitive regions, such as impeller surfaces. Hybrid approaches, integrating physical laws like parts of the Navier-Stokes equations, are being explored to enhance reliability and align the model closely with underlying physics. Future plans involve expanding the model to handle unsteady flow predictions and integrating physical laws to improve reliability, making it a valuable tool for efficient and accurate turbomachinery design.

References:

(1) Aulich M, Goinis G, Voß C. Data-Driven AI Model for Turbomachinery Compressor Aerodynamics Enabling Rapid Approximation of 3D Flow Solutions. Aerospace. 2024; 11(9):723. https://doi.org/10.3390/aerospace11090723

(2) Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. In: Advances in Neural Information Processing Systems. 2017;30:5998–6008.

(3) Siemens Digital Industries Software. Simcenter STAR-CCM+. 2024. Available from: https://www.plm.automation.siemens.com/global/en/products/simcenter/STAR-CCM.html

(4) Vithanala KS, Aulich M, Voß C. Data-Driven AI Surrogate Model for Rapid 3D Flow Approximation in Axial Fans. ASME Proceedings 2025. As part of this ongoing work, this paper provides detailed findings and methods discussed in this research.

Contrail Optimization for North-Atlantic Traffic

Xavier Pucel (ONERA)

Presentation type

Ongoing ressearch work

Keywords

Air Traffic Management, Trajectory planning, Optimization

Introduction

The contribution of air traffic to climate change involves not only CO_2 effects, but also condensation trails (contrails) [1, 2]. This paper addresses the problem of anticipating contrail-prone conditions and mitigating them through air-traffic regulation, a question addressed by [3] with individual flights, in [4] focusing mainly on separation rules, and in [5] for central Europe traffic. Focusing on the north-Atlantic airspace, we study the ability to mitigate contrail formation by modifying the trajectories of all flights in the region before the flights take off.

Anticipating contrail formation is difficult [6], as it depends on atmospheric conditions as well as the behaviour of each aircraft. Nevertheless, it is known that Ice Super-Saturated Regions (ISSRs) are prone to producing contrails when an aircraft flies through them. The problem of forecasting ISSRs is still under active research [7], but is currently possible to some extent and precision.

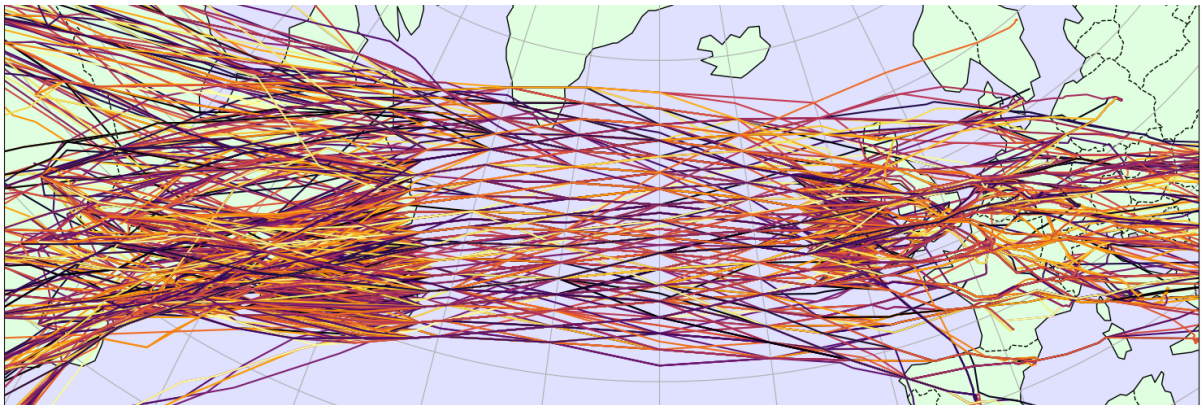


Figure 1: Daily flight plans crossing the North Atlantic Airspace.

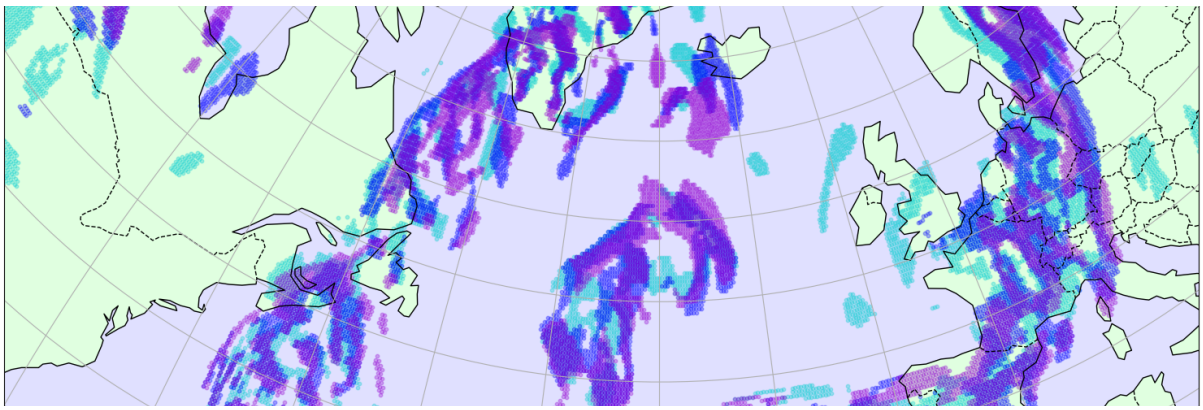


Figure 2: Ice Super-Saturated Regions in North Atlantic Airspace at 300 (turquoise), 250 (blue) and 230 (violet) HPa.

The North-Atlantic airspace is a free flight airspace, however the current practice is to use a small

set of waypoints with round longitude and latitude values as waypoints. This is visible in Figure 1: the zones with regulated flight routes exhibit a dense, maze-like network of flight plans, while in the free flight zone all flight plans neatly follow a lattice of integer latitudes, and multiple of 10 longitudes. As a result, it seems feasible to act upon the entirety of the traffic in order to minimize the distance travelled inside ISSRs. Figure 2 illustrates ISSRs at some point in time for different pressure values.

The ATM-lab is a software developed at ONERA that computes forward and reverse flight simulations and provides precise evaluations for climate related indicators [8, 9]. However, an optimisation algorithm that needs to evaluate millions of alternatives requires a simpler model with faster evaluation time.

Problem statement and computational approach

Given a set of flight plans and a forecast of atmospheric conditions, the problem consists in computing collective actions to modify these flight plans, and evaluating their capacity to mitigate their climate impact.

The computation of collective actions consists in designing actions (such as closing a part of the airspace), and evaluating its impact on the flight plans.

We split the evaluation part in two steps. The first only looks at the distance travelled inside ISSRs, as it can be computed efficiently for use inside an optimisation algorithm. The second step uses the ATMlab to precisely simulate and evaluate the climate impact of the air traffic with and without the collective actions taken, and evaluates the impact of these actions.

References:

[1] Klöwer M, Allen M, Lee D, Proud S, Gallagher L, Skowron A. Quantifying aviation's contribution to global warming. Environmental Research Letters. 2021;16(10):104027.

[2] Lee DS, Fahey DW, Skowron A, Allen MR, Burkhardt U, Chen Q, et al. The Contribution of Global Aviation to Anthropogenic Climate Forcing for 2000 to 2018. Atmospheric environment. 2021;244:117834.

[3] Barten KPAM. Contrail Mitigation through Flight Planning. 2017.

[4] Rodionova O, Sbihi M, Delahaye D, Mongeau M. North Atlantic Aircraft Trajectory Optimization. IEEE Transactions on Intelligent Transportation Systems. 2014 Oct;15(5):2202-12.

[5] Vergnes F, Bedouet J, Olive X, Sun J. Environmental Impact Optimisation of Flight Plans in a Fixed and Free Route Network. In: ICRAT 2022. Tampa, United States; 2022. .

[6] Gierens K, Matthes S, Rohs S. How Well Can Persistent Contrails Be Predicted? Aerospace. 2020;7(12):169.

[7] Riggi-Carrolo E, Dubot T, Sarrat C, Bedouet J. AI-Driven Identification of Contrail Sources: Integrating Satellite Observation and Air Traffic Data. Journal of Open Aviation Science. 2023 Oct;1(2).

[8] Bedouet J, Huynh N, Kervarc R. GAMME, a Meta-Model to Unify Data Needs in Simulation Modeling (WIP). In: Proceedings of the Symposium on Theory of Modeling & Simulation - DEVS Integrative M&S Symposium. DEVS 13. San Diego, CA, USA: Society for Computer Simulation International; 2013. p. 1-6.

[9] Joulia A, Dubot T. Analysis and Assessment of U-space Tactical Conflict Management Services. In: AIAA AVIATION 2022 Forum; 2022. p. 3543.

Coupling of a Genetic Algorithm and a Thermoelectric Network Model for Radioisotope Thermoelectric Generator (RTG) Optimization

Stylianos Kyrimis, Eliott Lallemand, Christian Stiewe, Pawel Ziolkowski, Eckhard Müller, Johannes de Boor (DLR)

Presentation type

Result of research work

Keywords

Radioisotope Thermoelectric Generator (RTG), Genetic Algorithm (GA), thermoelectric network model, Lunar exploration, Deep space exploration

Introduction

Lunar exploration requires a continuous supply of electrical power, independent of solar irradiation availability, which can be provided by radioisotope thermoelectric generators (RTG) [1]. RTG contain a radioisotope heating unit (RHU) enclosed in an aeroshell core and coupled with thermoelectric modules (TEM) [2], as seen in Fig. 1b. The latter utilize the temperature difference developed between their hot side, T_h, adjacent to the RTG core, and their cold side, T_c, connected to radiation fins attached to the outer RTG surface, to create a thermoelectric voltage in the TEM through the Seebeck effect, allowing them to convert part of the released heat into electricity. Design and operation of RTG is a complex multi-disciplinary challenge, as they are comprised of multiple components that affect performance. The surfaces of the aeroshell not connected to TEM are covered by thermal insulation, driving most of the source heat through the TEM. Excessive insulation however, could undesirably increase the system mass which should be minimized. Further, the materials and geometry of the TEM affect the heat-to-electricity conversion efficiency as well as the temperature at the TEM hot side which should not exceed the TEM's safe operating temperature to prevent device failure. To address these challenges and optimize the RTG for lunar operation, suitable multi-physics and multi-parametric tools are necessary to model and investigate the RTG's performance under realistic mission scenarios.

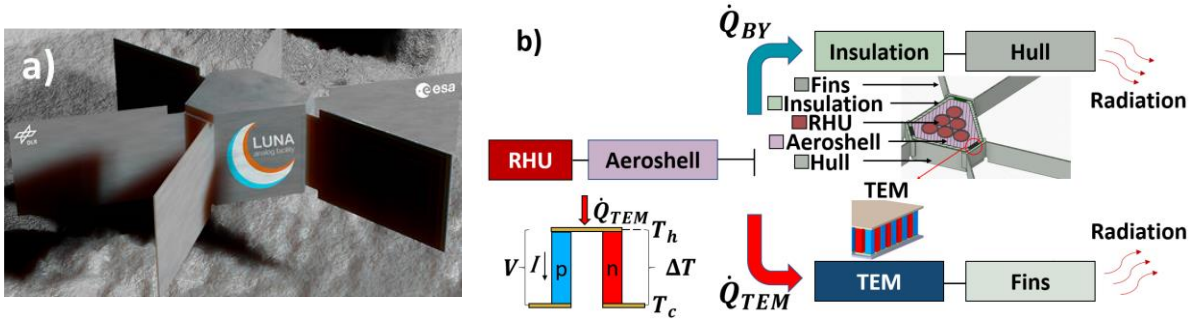


Figure 1: a) Artistic representation of the RTG on the lunar surface and b) TE network model approach.

TE network model

Here, we developed a thermoelectric (TE) network model to represent the RTG design, seen in Fig. 1b, suitable for rapid design investigations. Our network model simplifies the RTG components as thermal resistors and TE elements, thus calculating the heat transfer within the RTG by assuming two parallel paths, the heat moving through the TEM, $\dot{Q}_{TEM}$, and the parasitic heat loss through the insulation, $\dot{Q}_{BY}$. The two heat paths allow the determination of the temperature differences within the RTG components, while solving the TE heat transfer equations on the TEM provides an estimate of the RTG's performance. In our model, the RTG components are represented by a combination of

parameters, e.g. cross-section area, leg length and fill factor of the TEM and length, width and height of the radiator fins, which can be optimized to achieve pre-defined goals.

Genetic Algorithm

Our TE network model is coupled with a Genetic Algorithm (GA) which performed a detailed parametric study on the key component dimensions to identify the RTG design which maximizes the specific power, i.e. electrical power output over the system mass, while adhering to the safe operating temperatures of the TEM. Over 300 generations, the GA created a population of RTG designs by assigning values to the component parameters within pre-defined limits. The GA then utilised the TE network model to solve each RTG design and calculate its specific power and TEM temperature, ranking each solution based on highest specific power. The best designs of each generation were selected and altered through cross-over and mutation to create the children of the subsequent generation, gradually improving the result until a plateau is reached. The optimisation process was tracked in Fig. 2a, where, by finding the best compromise between power and mass, the specific power gradually improves until reaching a maximum after about 170 generations.

Results

The optimized RTG design, which uses a Bi_2Te_3 TEM, achieves a specific power of 1.42 W/kg, with an electrical power output of 13.9 W_{el} and a conversion efficiency of 6.9%, as seen in Fig 2b. In this configuration, 86% of the source heat passes through the TEM. This is a significant improvement over the baseline design currently in development by the European Space Agency (ESA), which aimed for an electrical output of 10 W with a 5% conversion efficiency [2]. Our model can also adapt the RTG design to incorporate new TEM which can operate at higher temperatures and thus allow for higher performances. As such, our tools can enable designs with higher specific powers and conversion efficiencies, guiding the efforts for the development of next generation RTGs.

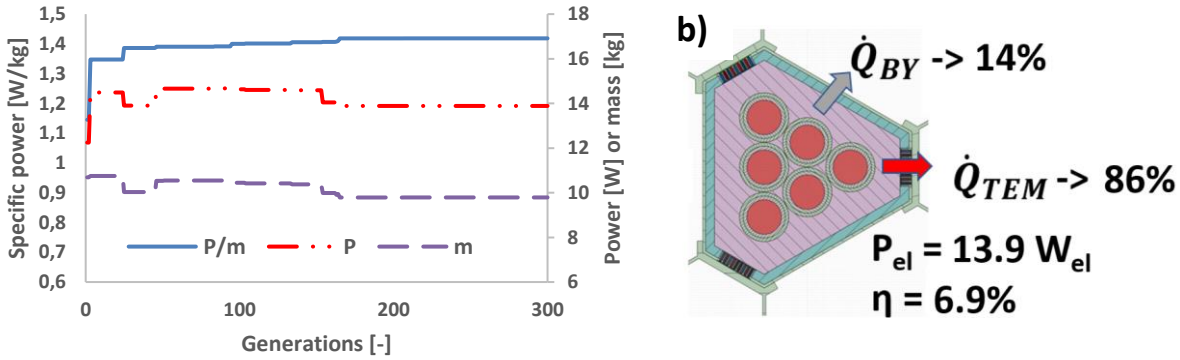


Figure 2: a) Specific power, power and mass produced by the GA and b) TE performance of the optimized RTG.

Conclusion

A TE network model was developed that approached the RTG as a series of thermal resistors and TE elements, assuming heat transfer by two competitive paths. By coupling it with a GA, we were able to optimize its design and operation, reaching an RTG with a specific power of 1.42 W/kg, power output of 13.9 W_{el} and a conversion efficiency of 6.9% while maintaining the safe temperature limits of the TEM. Our tools can thus be used to design and improve next generation RTGs.

References:

(1) Zillmer A.J., The history of the United State's Flight and Terrestrial RTGs, The Technology of Discovery, 2023, 7-33

(2) Ambrosi R.M. et al., European Radioisotope Thermoelectric Generators (RTGs) and Radioisotope Heater Units (RHUs) for Space Science and Exploration, Space Sci. Rev 2019 (Vol. 215, Issue 8, pp. 55).

Designing Satellite Assembly Lines from Experts' Objectives

Anouck Chan (ONERA), Thomas Polacsek (ONERA)

Presentation type
Result of research work

Keywords
Industry 4.0, Industry 5.0, Assembly line, Automatic calculus, Goal-requirement modelling.

Introduction

The aerospace industry is undergoing a transformation with the rise of satellite constellations for telecommunications, Earth observation, and meteorology. These projects increasingly rely on mass production, requiring scalable and adaptable assembly lines. The complexity of the satellites and production environments presents significant design challenges, particularly when precise assembly and testing are required. Assembly line design involves various stakeholders with conflicting goals, including safety, cost, worker conditions etc. To address this, structured methodologies are needed to guide early-stage decision-making. This work addresses two research questions: (1) How can the objectives and constraints of the different stakeholders be identified? (2) How can these be used to assess the feasibility and quality of potential assembly line designs? We propose a methodology that integrates operations research with *goal-oriented requirements engineering (GORE)*. This approach generates a high-level design of an assembly line that aligns with expert knowledge and diverse stakeholder goals. A real industrial case involving the production of three satellite types illustrates our method.

Background and Related Work

Satellite Assembly Line

An assembly line consists of stations equipped with jigs, fixtures, and tools (including machines, robots, environments) to perform specific operations. Each satellite type requires a specific sequence of assembly operations, tools, and routes between environments. In our case study, we aim to design an assembly line capable of manufacturing three satellite types with different delivery constraints.

Related Work

Early-stage assembly line design has been widely explored using models like *Resource-Constrained Project Scheduling Problem* [1]. Multi-objective optimisation approaches often use Pareto fronts and may be combined with robust strategies [2]. However, most of these works assume predefined and equally weighted objectives. In contrast, our context requires identifying relevant and sometimes conflicting goals, evaluating their satisfaction conditions, and managing trade-offs specific to our industrial case. To address this, we combine operations research models with GORE and *Participatory Enterprise Modelling* [3]. GORE models stakeholder goals as a directed graph (*goal model*), while participatory modelling captures expert knowledge through collaborative sessions.

Methodology

Our methodology starts with eliciting stakeholder objectives via participatory modelling. Goals are organised into a hierarchical goal model, where links capture the influence of one goal's satisfaction on another, allowing satisfaction propagation through the graph. Unlike in traditional operations research, where objectives and constraints are distinct, GORE defines a *goal* as a desired system state. Thus, a goal may act either as an optimisable objective or a hard constraint in the operations research model. A subset of goals is then selected as constraints for optimisation. Using operations research techniques, we model and solve the assembly line design problem to satisfy these goals. The resulting design is evaluated by mapping satisfied goals back onto the goal model, enabling satisfaction propagation and trade-off visualisation. This process can be repeated with different goal sets. Results are reviewed with domain experts to refine goal definitions, adjust preferences, and resolve conflicts.

Application

We applied our methodology to real datasets provided by a satellite manufacturer. These datasets describe the assembly processes for three satellite types, involving 147 distinct resources and spanning an 18-month planning horizon. Each satellite requires between 40 and 90 operations to complete. Figure 2a

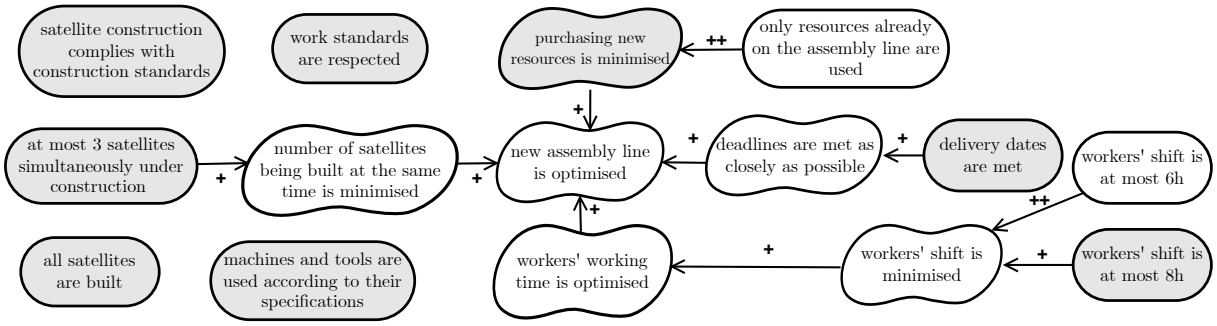
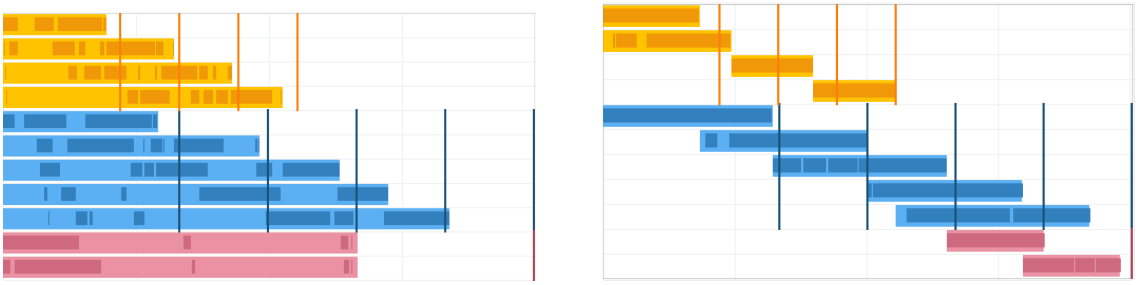


Figure 1: Extract of the asssembly line goal model



(a) Simple assembly line design

(b) Assembly line with experts' objectives

Figure 2: Schedule of assembly line

shows the baseline assembly line configuration, generated without considering expert goals. Each horizontal line represents a satellite under construction; darker segments indicate active operations, lighter segments correspond to idle time in clean rooms. Colours distinguish satellite models, and vertical lines mark delivery deadlines. An extract of the goal model is shown in Figure 1, where each node represents a goal and arrows indicate positive relationships. The full goal model is composted of 58 distinct elements. From this model, we selected an initial subset of goals and translated them into constraints for a constraint programming algorithm. These constraints can be enabled or disabled to explore various combinations of stakeholder objectives. The algorithm takes the satellite dataset as input and produces a detailed assembly line configuration, including operation sequences, resource allocation, and timing. We implemented our models using mostly the solver IBM ILOG CP Optimizer 20.1 [4]. Applying the goal set coloured in grey in Figure 1, we obtained the assembly line shown in Figure 2b.
Note that when we add the goal *worker shift is at most 6 hours*, the solver reports no feasible solution. In this case, stakeholders must make a choice between competing goals for the final assembly line design.

Conclusion

This paper presents a method for designing high-level satellite assembly lines while incorporating diverse and sometimes conflicting stakeholder goals. Objectives are derived from expert knowledge and structured into a goal model (1). Based on the selected goals, a constraint programming model generates a valid assembly configuration (2) that can be further studied by industrial experts.

References

[1] Hartmann S, Briskorn D. An updated survey of variants and extensions of the resource-constrained project scheduling problem. Eur J Oper Res. 2022;297(1):1-14.

[2] Saif U, Guan Z, Wang B, Mirza J. Pareto Lexicographic α-Robust Approach and Its Application in Robust Multi Objective Assembly Line Balancing Problem. Frontiers of Mechanical Engineering. 2014 Sep;9(3):257-64.

[3] Stirna J, Persson A. Enterprise Modeling - Facilitating the Process and the People. Springer; 2018.

[4] Laborie P, Rogerie J, Shaw P, Vilím P. IBM ILOG CP optimizer for scheduling: 20+ years of scheduling with constraints at IBM/ILOG. Constraints. 2018 03;23:210-50.

Towards physics-informed convolutional neural networks for steady turbulent duct flows

Gaurav Bokil (DLR Institute of Electrified Aircraft Propulsion), Sebastian Merbold (DLR Institute of Electrified Aircraft Propulsion), Stefanie deGraaf (DLR Institute of Electrified Aircraft Propulsion)

Presentation type
Result of research work

Keywords
Physics-informed machine learning, convolutional neural network, turbulent flow, surrogate modelling

Introduction

Numerical simulations have served as the foundation of engineering design since their emergence in the 1970s. Discretization techniques such as Finite Difference Methods (FDM), Finite Element Methods (FEM), and Finite Volume Methods (FVM) are at the forefront of this field. In aerospace applications, accurately modelling fluid flow at high velocities is essential. Consequently, Computational Fluid Dynamics (CFD) simulations necessary for the design of aircraft components become computationally intensive. At present, the electrification of aircraft propulsion systems is being explored as a sustainable alternative to traditional combustion engines. Propulsion configurations primarily reliant on hydrogen fuel cells and batteries are particularly sensitive to operational temperatures and temperature gradients. With fuel cells functioning at around 50% efficiency, effectively managing heat rejection presents a considerable challenge for the Thermal Management System (TMS). Therefore, the design of cooling ducts and heat exchangers that can efficiently dissipate heat is of utmost importance. To create cooling ducts and heat exchanger designs tailored for specific applications, sensitivity analyses and geometry optimization must be performed across an extensive design space. Nevertheless, the computational expense associated with executing CFD simulations serves as a significant bottleneck. To address this issue, a surrogate modelling approach has been selected as a viable alternative.

Methodology

A CNN framework is developed to predict velocity and pressure using geometry representation as input. It consists of two U-nets for outputting velocity components and pressure. Geometry is represented by binary masks and distance function fields, along with boundary masks for enforcing Dirichlet boundary conditions. A sixth-order central difference scheme computes spatial gradients, forming a loss function that includes conservation laws. Numerical diffusion ensures stability, and the Adam optimizer trains the RANS-CNN unsupervised. This method's performance is evaluated through three case studies: an S-shaped duct using literature data, a ducted heat exchanger compared to experimental data, and a forward facing step problem as a benchmark for CFD codes.

Conclusion

The overall results observed demonstrate that the proposed approach of utilising:

- higher-order finite difference scheme for computing gradients,
- numerical stabilisation, and
- strong boundary enforcement

enabled the CNN model to simulate the RANS equations. The work aimed to develop a framework that produces accurate results which can be further improved via provided labelled data.
This work focuses on developing cooling ducts for electric aircraft thermal management. The study aims to create a physics-informed CNN to simulate fluid flow through ducts at high Reynolds numbers. Results confirm the RANS-CNN method's effectiveness. The next step involves creating a RANS-based surrogate model for ducted heat exchangers to aid in designing optimal cooling ducts.

Variational autoencoder with weighted samples for high-dimensional non-parametric adaptive importance sampling

Julien Demange-Chryst [1], *Jérôme Morio* [1], *François Bachoc*[2]

[1] *ONERA/DTIS, Université de Toulouse, F-31055 Toulouse, France*

[2] *Institut de Mathématiques de Toulouse, UMR5219 CNRS, 31062 Toulouse, France*

Presentation type

~~Collaboration proposal~~ — ~~On-going collaborative work~~ — Result of research work

Keywords

Variational autoencoder, importance sampling, high dimension, rare event estimation.

Introduction

Importance sampling is a well-known uncertainty quantification method, classically used as a variance-reduction technique for Monte Carlo integration including rare event estimation (1), or for generating points from a target probability distribution known up to a constant (2). The common denominator of every importance sampling procedure is that they all require to estimate a target probability distribution g^* with weighted samples, and obviously, the accuracy of the algorithm depends on the quality of the estimation of the distribution. Moreover, we also need to be able to not only sample from the built auxiliary distribution, but also to have access to its PDF values.

Generative models for high-dimensional density approximation

Existing density approximation methods for importance sampling, namely non-parametric and parametric models, strongly suffer from the curse of dimensionality and/or require prior knowledge on the target distribution to be accurate. To overcome these issues and to combine both flexibility and robustness faced to the dimension, we propose to use generative models. Indeed, generative models are recent probabilistic tools for approximating high-dimensional complex distributions of observed data and generating new samples. More precisely, as shown in Figure 1, given a latent distribution p on a latent space $\mathcal{Z}$, it consists in building a generative function G satisfying:

$$G(\mathbf{Z}) \sim g^*, \text{ with } \mathbf{Z} \sim p. \tag{1}$$

The main advantage of this family of models is that the generative function G can be as expressive as wanted. In particular, it can be parametrized by a deep neural network.

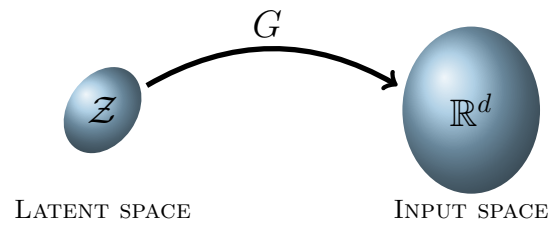


Figure 1: Principle of a generative model.

Among the various generative models, we propose in this work to use a variational autoencoder (3). Three main arguments motivate this choice. First, since it is a generalization of a classical autoencoder, its learning procedure provides a lower dimensional latent representation of the target distribution which makes the estimation procedure more robust faced to the dimension. Second, the resulting approximating distribution is very flexible because it is an infinite Gaussian mixture. At last, a variational autoencoder provides an explicit expression of the resulting PDF, which is necessary for its use in importance sampling.

Contribution

In the present communication (4), we extend the existing framework of a variational autoencoder to the case of weighted samples by introducing a new objective function. The resulting importance sampling auxiliary distribution is close to an infinite mixture of Gaussian distributions. Then, its flexibility makes it as expressive as a non-parametric model, and despite the very high number of parameters to estimate, it is much more efficient in high dimension than the classical Gaussian or Gaussian mixture families. Moreover, in order to add even more flexibility to the model and to be able to learn multimodal distributions, we consider a learnable prior distribution for the variational autoencoder latent variables. We also introduce a new pre-training procedure for the variational autoencoder to find good starting weights of the neural networks to prevent as much as possible the posterior collapse phenomenon to happen.

At last, we explicit how the resulting distribution can be combined with importance sampling. Indeed, the existing procedure (5) to compute the PDF values of the resulting distribution of a variational autoencoder leads to a biased and non-convergent importance sampling estimator. In order to keep an unbiased and consistent estimator, we introduce a new way to compute the PDF values. Then, we show how to integrate the whole suggested procedure into existing reliability algorithms, such as the cross-entropy algorithm (1), for rare event estimation. Finally, we illustrate the practical interests of the previous efforts on two multimodal rare-event-estimation problems. The code to reproduce the numerical experiments is publicly available at: `https://github.com/Julien6431/Importance-Sampling-VAE.git`.

References:

(1) Reuven Y Rubinstein and Dirk P Kroese. The cross-entropy method: a unified approach to combinatorial optimization, Monte Carlo simulation, and machine learning, volume 133. *Springer*, 2004

(2) Jean-Michel Marin, Pierre Pudlo, and Mohammed Sedki. Consistency of adaptive importance sampling and recycling schemes. *Bernoulli*, 25(3):1977 – 1998, 2019.

(3) Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. *Article accepted in the 2nd International Conference on Learning Representations 2014*, 2014.

(4) Julien Demange-Chryst, François Bachoc, Jérôme Morio, and Timothé Krauth. Variational autoencoder with weighted samples for high-dimensional non-parametric adaptive importance sampling. *Transactions on Machine Learning Research*, 2024.

(5) Hechuan Wang, Monica F Bugallo, and Petar M Djuric. Adaptive importance sampling supported by a variational auto-encoder. In *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 619–623. IEEE, 2019.

Other AI Related Contributions

Contents

- Development of a Benchmark for Deep Reinforcement Learning Based Control of Liquid Propellant Rocket Engines *Vincent Bareiß, Kai Dresia, Günther Waxenegger-Wilfing*.
- Deterministic and Probabilistic Decision Models for GSaaS-based Satellite Communication Resource Management *Hénoïk Willot, Gauthier Picard, Jean-Loup Farges, Philippe Pavero*.
- Using AI for day-to-day work at DLR department Space Launcher System Analysis *Ingrid Dietlein*.
- Measuring and Addressing Information Leakage in Concept Bottleneck Models *Raffael Schön, Baptiste Abeloos, Stéphane Herbin*.
- Predicting damage location information based on single actuator-sensor-path information using regularized multilinear regression *Tim Behrens, Maria Moix-Bonet, Daniel Schmidt, Peter Wierach.*

Development of a Benchmark for Deep Reinforcement Learning Based Control of Liquid Propellant Rocket Engines

Vincent Bareiß (German Aerospace Center (DLR), Institute of Space Propulsion), Kai Dresia (German Aerospace Center (DLR), Institute of Space Propulsion), Günther Waxenegger-Wilfing (German Aerospace Center (DLR), Institute of Space Propulsion; University of Würzburg, Institute of Computer Science)

Presentation type
Collaboration proposal

Keywords
Liquid Propellant Rocket Engines, Deep Reinforcement Learning, Robust Reinforcement Learning, Sim2Real gap, Benchmarking.

Introduction

Driven by new trends in the space industry, such as reusability, the propulsion subsystem of modern launch vehicles is faced with new challenges requiring more sophisticated control architectures [1]. Future challenges include precise throttling and multiple re-light capabilities, requiring increasingly complex objective functions. Such objective functions must not only take account not only the directly controlled variables but also secondary control objectives—for example, minimizing wear on engine components by reducing bearing temperatures. Deep Reinforcement Learning (DRL)-based techniques have shown promising results in solving such new control challenges [2]. Due to their black-box nature, however, the development of trust in such AI-based control systems remains a key challenge for their adoption [1].
In this work, we present the design of a benchmark based on the Liquid Upper Stage Demonstrator Engine (LUMEN) to evaluate a DRL-based control system's performance and its ability to adapt to modeling errors as well as faults. By incorporating experimental data gathered during hot run tests of LUMEN we increase confidence in the benchmark results.

The LUMEN Control Problem

The developed benchmark is based on the digital simulator model of LUMEN, an expander-bleed cycle liquid propellant rocket engine using cryogenic liquid methane (LNG) and liquid oxygen (LOX) [3]. The LUMEN simulator is based on EcosimPro/ESPSS and has been developed and tested at the DLR Institute of Space Propulsion. The simulator has been extensively tested and validated using experimental data gathered at the P8.3 test facility at DLR Lampoldshausen. In this work, we consider the control task of optimal trajectory tracking for the combustion chamber pressure and mixture ratio. Two control valves are used to regulate the oxidizer and fuel turbine mass flow, which in turn control the combustion chamber pressure and mixture ratio. Due to the structure of the expander-bleed engine cycle, the mixture ratio and combustion chamber pressure are highly coupled [2]. The resulting LUMEN control task is a 2×2 MIMO problem. The applicability of DRL to the LUMEN control problem has previously been studied and tested by Dresia et al. in two test campaigns at DLR Lampoldshausen [1].
Three key challenges have emerged when applying DRL to the LUMEN control problem. First, how to develop and train a DRL agent that can regulate previously unseen trajectories efficiently, generating smooth actuator commands. Second, how to increase the agent's ability to adapt to changes in the dynamics due to modeling errors when applied to a physical system, crossing the Sim2Real gap. Third, how to train an agent that can adapt to substantial variations in the dynamics function, action space, or state space caused by faults. In addition, detecting conditions under which the agent is no longer able to adapt to the changed dynamics, as presented in the talk by Jonas Dauer, is of great interest.

Structure of Benchmark

We structure the benchmark as a series of test cases to evaluate a control system across different performance aspects. Each test case is represented as an OpenAI Gym environment, encapsulating the LUMEN simulator model.

The nominal tracking performance is evaluated in two test cases. For both test cases, the nominal simulator model is used. In test case 1, prior knowledge of the tracked reference trajectory is available during training. In test case 2, the agent is evaluated using a randomized reference trajectory. These two test cases serve as a baseline for a given control system prior to the introduction of variations in the model dynamics.

The agent's capability to adapt to variations in the system's dynamics is evaluated in three test cases. In test case 3, a multiplicative parametric uncertainty is applied to the model. Prior knowledge of the uncertainty distribution and affected parameters is provided at training time. In test case 4, a similar parametric uncertainty is applied to the simulator dynamics, with no data on distribution and affected parameters given. Finally, in test case 5, the agent's capability to adjust to slow, continuous changes in the dynamics, as caused by, e.g., aging of components, is evaluated.

Test cases 6 and 7 evaluate the fault tolerance of the developed control system. Two faults are simulated based on hot run data gathered at the P8.3 test facility at DLR Lampoldshausen. In test case 6, a binary signal notifying the agent of a fault in a certain component is provided at both training and evaluation time. This represents an external fault detection system such as [4]. In test case 7, no external fault detection is provided.

Table 1 shows an overview of the test cases.

Table 1: Structure of Benchmark

Testcase	Description of Testcase
Testcase 1	Tracking in nominal case with a-priori known reference trajectory
Testcase 2	Tracking in nominal case with unknown trajectory
Testcase 3	Domain adaptation with prior knowledge of target domain
Testcase 4	Domain adaptation without prior knowledge of target domain
Testcase 5	Slow continuous change in model dynamics during episode
Testcase 6	Fault with external fault detection
Testcase 7	Fault without external fault detection

Conclusion

We have presented the design of a benchmark for Deep Reinforcement Learning based control of a liquid propellant rocket engine. The benchmark is based on the LUMEN control problem, a 2x2 MIMO optimal tracking problem with a high degree of coupling. The challenges of the control problem have been highlighted and transferred into seven test cases that will be implemented in the benchmark. The benchmark will be shared, and we invite collaborators to propose solutions.

References:

[1] Dresia, K, Traudt, T, Horger, T, Dauer, J, Deeken, J, Waxenegger-Wilfing, Overview of Future Rocket Engine Control Systems. In Proceedings of the 9th Edition of the Space Propulsion Conference 2024 .

[2] Dresia, K, Waxenegger-Wilfing, G, Dos Santos Hahn, R, Deeken, J, Oschwald, Nonlinear Control of an Expander-Bleed Rocket Engine using Reinforcement Learning. In Space Propulsion 2020+1 2021

[3] Traudt, T, Armbruster, W, Groll, C, Dos Santos Hahn, R, Dresia, K, Borner, M, Klein, S, Suslov, D, Kurudzija, E, Haemisch, J, Müller, M, Deeken, J, Hardi, J, Schlechtriem, LUMEN, the Test Bed for Rocket Engine Components: Results of the Acceptance Tests and Overview on the Engine Test Preparation. In Proceedings of the 9th Edition of the Space Propulsion Conference 2024 .

[4] Kurudzija, E, Dresia, K, Waxenegger-Wilfing, G, Behler, H, Auer, D, Fröhlke, K, Neumann, H, Frank, A, Laurent, J, Fabreguettes, L. "Automation of Testing and Fault Detection for Rocket Engine Test Facilities with Machine Learning". International Journal of Energetic Materials and Chemical Propulsion 2023.

Deterministic and Probabilistic Decision Models for GSaaS-based Satellite Communication Resource Management

Hénoïk Willot, Jean-Loup Farges, Gauthier Picard (DTIS, ONERA, Université de Toulouse), Philippe Pavero (Airbus Defence and Space)

Presentation type
Result of research work

Keywords
Earth observation, communication resources, GSaaS, Integer Linear Programming.

Introduction

Conventionally, Earth Observation (EO) space missions possess individualized resources. Each mission is equipped with a satellite constellation, ground stations, and a command center for satellite communication, uploading observation blueprints, and downloading collected data. In recent years, the development of third-party communication stations adopting the Ground Station as a Service (GSaaS) model has been observed (1). This approach allows customers to reserve communication resources from ground segment service providers rather than constructing their own facilities.

To facilitate coordination among disparate EO missions, a multi-agent federation layer has been proposed in the DOMINO-E project (2). This layer aims to streamline client access to multiple satellite constellations and communication sites for the assembly and retrieval of extensive acquisition data in less time than traditional, uncoordinated methods. The communication dispatcher, coined Satellite Communication and Resource Management System (SCRMS), is central to the federation layer, orchestrating communications for N satellites across multiple ground stations, each with its own antenna.

This work aims at developing and assess methods to solve this allocation problem. We study two cases, one where all the scheduled communication contacts are accepted by the GSaaS and one where acceptations are following probability distributions, hence are not guaranteed.

Selection problem

SCRMS is in charge of providing opportunities for communicating with ground to all satellites of all federated missions. It manages communication between N satellites, seeking for communication windows for uplink and downlink data transfer, and sites including one or several ground stations. Each ground station controls a specific antenna. Missions provide orbits of their satellites and ground station networks publish coordinates and visibility masks of their stations.

Propagating orbits and computing the events with respect to sites coordinates and visibility masks leads for each satellite i to L_i potential contacts with the ground. Each contact $l \in \{1, \dots L_i\}$ is characterized by its site on ground, its time window $[\underline{u}_{i,l}, \overline{u}_{i,l}]$, allowing to deduce its duration $d_{i,l}$. The decision variable associated to a contact is $x_{i,l} \in \{0, 1\}$, valued to 1 if contact l of satellite i is chosen in the plan.

The selection of contacts allow to partially fulfill the required communication needs of the satellite. Each satellite i has K_i needs to fulfill, defined by a communication duration $D_{i,k}$, to fulfill within a time window $[\underline{t}_{i,k}, \overline{t}_{i,k}]$, with a list of bookable sites $\mathcal{S}_{i,k}$ and a communication band (S band for uploading the observation plan to the satellite and maintenance operations, and X band for downloading the measurements). A need can also be specialized with an observation area, in that case it is called localized need. The monitoring of this observation area (e.g. in case of wildfire) requires communications when the satellite is near, before entering the area if the need's band is S and after leaving the area if it is X band. These communications shall occur in a time window of length Δ_k respectively before and after flying over the area.

A contact can be selected to partially fulfill a need k if its time window is contained in the need's one, the site belongs to the bookable sites and the communication band is right (and belong to the authorized time window for a localized need). Let $\mathcal{Z}_{i,l}$ be the set of needs a contact l is able to fulfill.

Our decision problem is then finding a selection of contacts such that the resulting communication plan verifies for each satellite i and each of its needs k $\sum_{l \in \{1,...L_i\}/k \in \mathcal{Z}_{i,l}} d_{i,l} x_{i,l} \geq D_{i,k}$.

If the problem is feasible, setting all $x_{i,l}$ variables to one would respect this constraint. However, there is a financial cost associated to contact selection and selecting some contacts together induces a risk of not being able to actually use them because of conflict or mutual jamming. The total cost of an allocation is obtained by adding the individual cost of the selected contacts, which depends on the cost model of the station of the contact which can be pay per use (a cost per minute time the duration), pay per pass (a fixed cost for using the contact) and pay per pass with commitment (a certain number of contacts are pre-paid each day, fixed cost for the rest). Two contacts l of satellite i and m of satellite j can generate an amount $b_{i,l,j,m}$ of conflict and jamming if their time windows overlap and if they share the same site (not necessarily the same antenna). The total jamming cost is obtained by summing the quantities generated by each pair of contacts. We consider here a lexicographic ordering between the criteria, minimizing first the cost (resp. jamming) and at equal value we then minimize the jamming (resp. cost).

Solving methods

Deterministic case. In this scenario, all the contacts, when selected, are accepted by the GSaaS. We tested several approximate solving methods such as greedy allocations, round-robins, local and large neighborhood searches. We also implemented an exact formulation using Mixed Integer Linear Optimization (MILO). In the constraints of the problem, the expression of the cost and of the jamming are the only non linear constraints but can be linearized without approximation at a small cost.

Probabilistic case. In this scenario, each GSaaS has different probability levels different probability of acceptance depending on the contact cost model. If the GSaaS has a billing policy with multiple prices, e.g. for early or last minute reservation, a refused contact can be asked again the next day with the next cost model, if available. This uncertainty makes the fulfillment of the needs requirement hard to achieve. We propose to handle this case by finding a plan which expectation (with a margin based on its standard deviation) fulfills the requirements. We modified the MILO of the deterministic case to integrate this constraint, using a McCormick envelope approximation for the standard deviation.

Results and perspectives

The methods performances are compared on a realistic 60 days scenario, with a federation of 3 satellite constellations of 2, 4 and 4 satellites respectively, with evolving needs among this period. The methods must plan the next 10 days every day. Depending on the station policy and cost, the contacts can be booked 7, 5, 3 or 1 day earlier. The results obtained are good, especially for the MILO which runs faster than the imposed time limit of 5 minutes for each plan. We plan to extend this work by taking into account, in the uncertain scenario, the missed communication time, due by the contacts rejection, by increasing the needs for the next day. We also aim to investigate finer models of uncertainties, like the load of GSaaS services due to the other missions and satellites. A first approach could consist in learning the correlation between the rate of acceptance and the number of overflying satellites. This work is part of the DOMINO-E (https://domino-e.eu/) project which received funding from the European Union's Horizon Europe Program for Research and Innovation under Grant Agreement n°101082230.

References:

(1) Carcaillon E, Bancquart B. Market perspectives of Ground Segment as a Service. 71st International Astronautical Congress (IAC). 2020.

(2) Farges JL, Perotto FS, Pralet C, Picard G, de Lussy C, Guerra J, et al. Going Beyond Mono-Mission Earth Observation : Using the Multi-Agent Paradigm to Federate Multiple Missions. In : 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-24); 2024.

Using AI for day-to-day work at DLR department Space Launcher System Analysis

Ingrid Dietlein (DLR Bremen)

Presentation type

Result of research work

Keywords

Large Language Models, LLM, AI, coding.

Introduction

The department Space Launcher System Analysis located at Bremen as part of the DLR Institute for Space Systems is dedicated to the pre-design of launcher concepts with a specific focus in reusable launch systems. This department has more than a quarter of a century of experience in this domain and it is unique in its kind in Germany if not Europe as it houses and combines all relevant discipline competencies involved with the multi-disciplinary pre-design of a multi-stage launch system. These disciplines involve engine cycle analysis, aerodynamics, flight mechanics, structural analysis, trajectory optimization for ascent, descent and combined ascent-descent, mass and cost estimation. Apart from the design work itself, other tasks not directly related to the design process are recurrent for employees in this department. This involves the documentation of the work, the presentation to team members and clients or cooperation partners and the summarizing of major results in scientific articles and publications. In addition, tools used within the design process have to be maintained, debugged and modified to expand their capabilities. From time to time, new tools have to be coded to extend the department capabilities. New code has to be tested and documented.
On an individual basis, some team members have already integrated Large Language Models (LLM) like ChatGPT in their daily work. The presentation proposed to this workshop intents to share gained experience on the use of LLMs in the day-to-day work and initiate an exchange on AI use for various use cases.

Most relevant use cases for LLMs for day-to-day work at DLR Space Launcher System Analysis department

Apart from more trivial tasks related to text generation, translation, research of information the field of coding is probably currently the most beneficial use case inside the author's department:
Following list provides a number of use cases within this field:

- First draft responding to a coded function, e.g. on one dimensional heat transfer, calculating orbital parameters from velocity and position vector,
- Providing code for visualization
- Support debugging and code optimization
- Providing programming language reference

Limits of use of LLMs for the day-to-day work

While publicly available LLMs was found to be helpful during quotidian work it is necessary to emphasize the limitations related to LLMs such as ChatGPT. It was found essential that results from LLMs have to be checked with respect to reliability irrespective of the use case. This double-check is even more important when using LLMs for code generation or other code tasks as cases occurred that provided codes contained bugs and formulation errors which make code testing highly relevant.

Envisaged future uses of LLMs

Following ideas to expand the use of LLMs inside the departments are envisaged to develop in the near future:

- Expand the use of LLMs for code debugging and optimization and maintenance,
- Checking code with existing documentations on theory behind the code,
- Generate code documentation for inhouse tools,
- Use auto-completion for certain documents

Conclusion

LLMs are already used by the majority of employees of the author's department to assist them during quotidian work. It was found to be very beneficial for coding and writing purposes but with the caveat that output from LLM prompts are to be handled with a sound portion of skepticism. The answers cannot be trusted blindly and require double checks by the user.
In summary, the experience with LLMs were judged sufficiently beneficial to envisage expanding to further use cases in the near future.

Measuring and Addressing Information Leakage in Concept Bottleneck Models

Raffael Schön (DTIS,ONERA), Baptiste Abeloos (DTIS,ONERA), Stéphane Herbin (DTIS,ONERA)

Presentation type
On-going collaborative work

Keywords
Concept Bottleneck Models (CBMs), Information Leakage, Explainable AI.

Introduction Concept Bottleneck Models (CBMs) are an interpretable by design approach typically leveraged for image classification. A CBM models its decision process such that it first predicts a set of human understandable concepts (1) as an intermediate step, based on which the final class is predicted. The model can therefore be split up in a concept predictor and a class predictor, which are chained together. Due to this separation three training methods exist for training a CBM, namely independent, sequential and joint training. Despite their interpretability advantages, a significant challenge persists: *Information Leakage*, where unintended information in the concept layer influence predictions, undermining trustworthiness and interpretability (2), (3). Existing evaluation methods are indirect, such as pruning concepts or comparison against oracle models. Thereby limiting the precise understanding of how Information Leakage affects a prediction as well as hindering the development of methods to mitigate it.

Measuring Information Leakage To address this gap, we introduce the *Irrelevant Concept Contribution (ICC)* metric, designed to directly quantify the degree of Information Leakage in a model prediction. The core principle of this method relies on prior information in the training dataset about the relationship between concepts and classes; i.e. which concepts are relevant to a class. In order to determine the ICC, we count the contributions from concepts that are not relevant to class predictions. The resulting ICC score offers an interpretable measure of the reliance on unintended information, with higher scores indicating greater leakage.

Architectures and Training Objectives To mitigate information leakage, we propose the *RelevNet CBM* architecture, which restricts connections from concepts to classes solely to relevant concepts, thereby exhibiting no Information Leakage by design. Furthermore, we propose an additional training objective for Concept Bottleneck Models, which repurposed the ICC metric as a loss function (ICCL). Leveraging ICC as a training objective encourages the trained CBM to rely on relevant concepts for its predictions.

Experiments To evaluate our proposed training objective, we train multiple CBMs using different activation functions—with and without ICCL. Additionally we also train a RelevNet. For each model we train it three times with the different Concept Bottleneck training methods, namely independent, sequential and joint training. We perform our experiments on the CUB (4) dataset and report for each training test accuracy as well as the ICC score. The results of our experiments are shown in Figure 1.

Conclusion This work introduces the **ICC metric** as a practical tool for quantifying information leakage in CBMs, along with the RelevNet and a CBM training objective to eliminate reliance on irrelevant concepts. Our findings demonstrate that Information Leakage in sequentially trained CBMs can be mitigated by introducing our training objective, while retaining the same level of accuracy. Furthermore, we show that RelevNet an architecture with significant constrains achieves top of the line accuracy results while possessing no Information Leakage inherently. These advances enable more trustworthy Concept Bottleneck models, which are especially crucial in safety-critical domains. Future work will compare the RelevNet to State of the Art approaches and extend our experiments to more datasets.

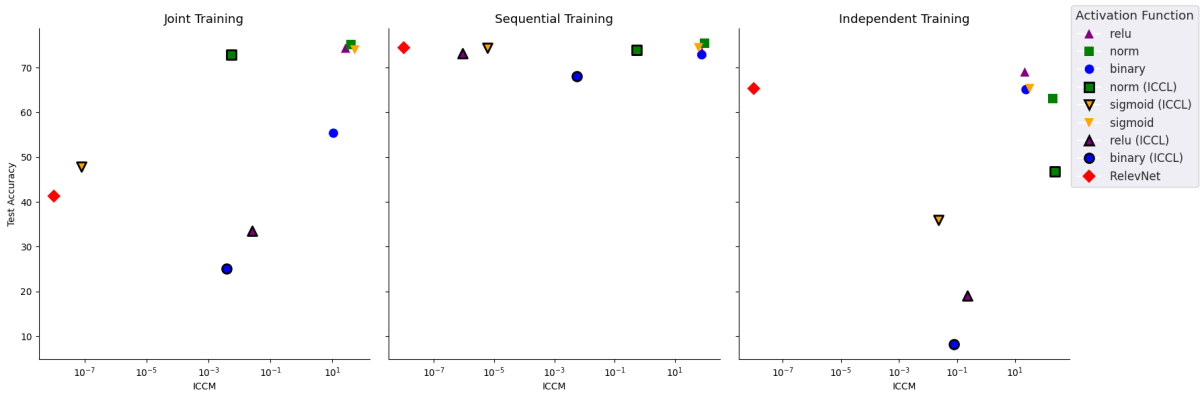


Figure 1: Scatterplots display how different training strategies affect test accuracy and ICC on CUB dataset. The highlighted symbols are CBMs trained with additional ICCL. On the Y-axis, the test accuracy is displayed, while on the X-axis the ICC value is displayed as a log scale.

References:

(1) Koh PW, Nguyen T, Tang YS, Mussmann S, Pierson E, Kim B, Liang P. Concept Bottleneck Models. In: Proceedings of the 37th International Conference on Machine Learning; 2020. p. 5338–5348.

(2) Margeloiu A, Ashman M, Bhatt U, Chen Y, Jamnik M, Weller A. Do Concept Bottleneck Models Learn as Intended? In: International Conference on Learning Representations 2021 Workshop on Responsible AI; 2021. Available from: https://arxiv.org/abs/2105.04289.

(3) Mahinpei A, Clark J, Lage I, Doshi-Velez F, Pan W. Promises and Pitfalls of Black-Box Concept Learning Models. In: Proceedings of the 38th International Conference on Machine Learning; 2021. [arXiv preprint arXiv:2106.13314].

(4) Wah C, Branson S, Welinder P, Perona P, Belongie S. The Caltech-UCSD Birds-200-2011 Dataset. California Institute of Technology; 2011. CNS-TR-2011-001. Pasadena, CA, USA.

Predicting damage location information based on single actuator-sensor-path information using regularised multilinear regression

Tim Behrens (DLR-SY), Maria Moix-Bonet (DLR-SY), Daniel Schmidt (DLR_SY), Peter Wierach (DLR-SY)

Presentation type

Result of research work

Keywords

structural health monitoring, ultrasonic guided wave, machine learning, probability reconstruction algorithm, damage localisation

Introduction

In the aviation industry maintenance is a cost which researcher of structural health monitoring plan to reduce by enabling a fast damage detection and localisation. Using guided ultrasonic waves, probability reconstruction algorithms are a method to locate a damage. They work by calculating the probability for each actuator-sensor-path on whether there is a damage on the path or not. By superposition of each path and its damage probability, a damage localisation is done. The disadvantage of this is, that the damage localization resolution is limited by the number of paths crossing each other. (1–3)

To overcome this, the hypothesis of this investigation is that the information of a path can not only be used to determine whether a damage is present, but that additional information about the location within the path can be calculated as well. This way a localization resolution can be higher than by only relying on the path density.

Distance Prediction

The goal is to evaluate the capability to predict the distance of a damage based on the information of one single actuator-sensor-path. The signal that is available for this is a time signal at the sensor position in a pristine structure state and a second on after a damaged was induced. In the signal processing several features are computed for each of five time windows. With these features a regularised multiple linear regression machine learning approach is carried out. The dataset is split into 80% trainings and 20% test data.

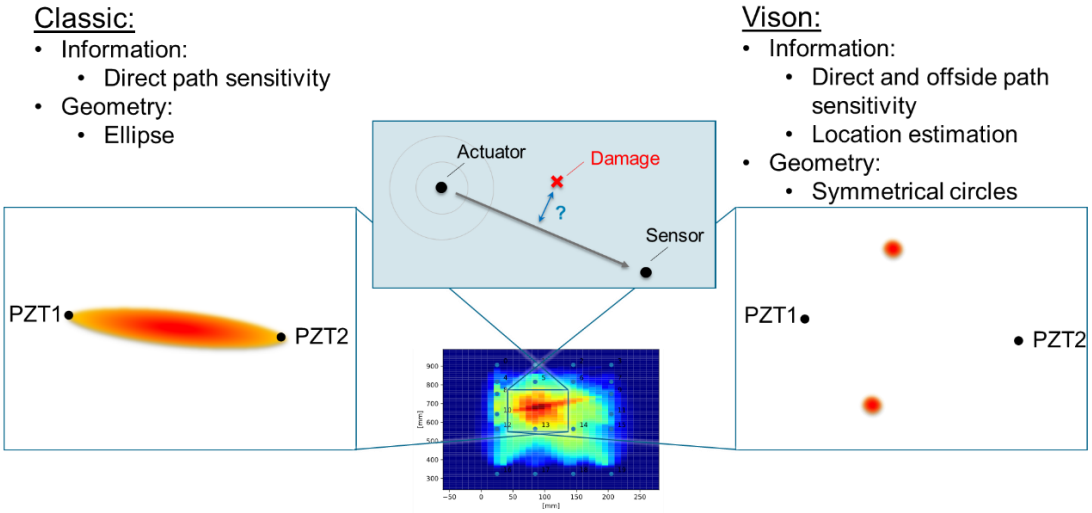


Figure 1: Envisioned concept for an improved damage localization via probability reconstruction algorithm

Figure 1 displays the principle of the investigated damage localization, the probability reconstruction algorithms as well as the envisioned more accurate localization. The predicted distance of this research is sketched in the top.

Conclusion

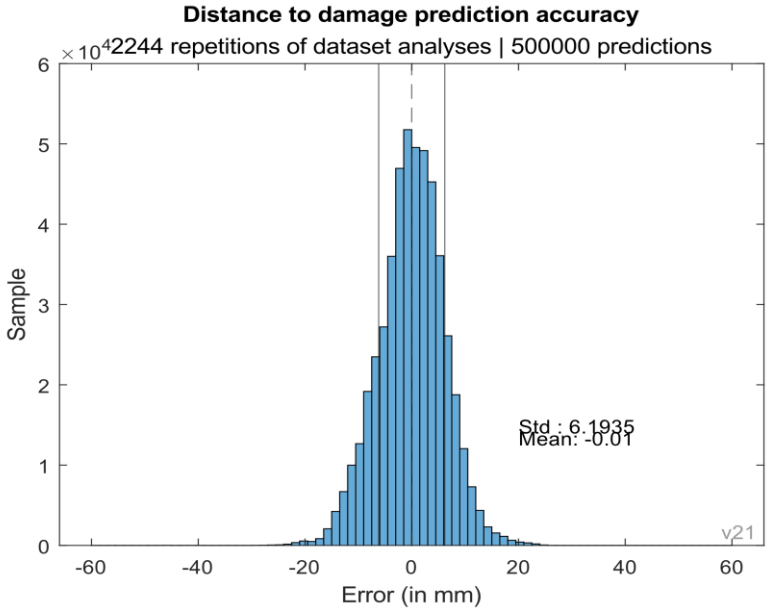


Figure 2: Distance to damage prediction accuracy of 500000 predictions

A standard deviation of under 7 mm in predicting the damages distance (seen in Figure 2), shows that a single path can give a geometric information of the location of the damage. Thinking about the combination of multiple paths in a structural health monitoring network this offside path distance prediction shows a promising step into an accurate damage localisation with probability reconstruction algorithms.
Also, the prediction process provides a tool to evaluate the quality of parameters in structural health monitoring. One can study different features and can evaluate those based on the standard deviation of the error. It is the same with other process parameters like the regularisation or the time windows.

References:

1. Moll J, Kexel C, Kathol J, Fritzen C-P, Moix-Bonet M, Willberg C et al. Guided Waves for Damage Detection in Complex Composite Structures: The Influence of Omega Stringer and Different Reference Damage Size. Applied Sciences 2020; 10(9):3068. doi: 10.3390/app10093068.
2. Hay TR, Royer RL, Gao H, Zhao X, Rose JL. A comparison of embedded sensor Lamb wave ultrasonic tomography approaches for material loss detection. Smart Mater. Struct. 2006; 15(4):946–51. doi: 10.1088/0964-1726/15/4/007.
3. Moix-Bonet M, Wierach P, Loendersloot R, Bach M. Damage Assessment in Composite Structures based on Acousto Ultrasonics – Evaluation of Performance. SARISTU Final Conference 2015. doi: 10.1007/978-3-319-22413-8_28.