DOI: 10.13009/EUCASS2025-523
11TH EUROPEAN CONFERENCE FOR AERONAUTICS AND SPACE SCIENCES (EUCASS)

Machine Learning Methods for Design and Operation of Chemical Space Propulsion Engines - Results of the ESA Project ADMIRE

Günther Waxenegger-Wilfing[⋆‡†], *Kai Dresia*[⋆], *Jeannine Schmacka*[⋆], *Michael Börner*[⋆], *Jan Martin*[⋆], *Till Hörger*[⋆], *Justin Hardi*[⋆], *Jan Deeken*[⋆§], *Annafederica Urbano*[¶], *Jose Felix Zapata Usandivaras*[¶], *Andreas Gernoth*[*]
[⋆]*German Aerospace Center (DLR), Institute of Space Propulsion, 74239 Hardthausen am Kocher, Germany*
[‡]*University of Würzburg, Institute of Computer Science, 97074 Würzburg, Germany*
[§] *RWTH Aachen University, Institute of Jet Propulsion and Turbomachinery, 52062 Aachen, Germany*
[¶] *Fédération ENAC ISAE-SUPAERO ONERA, Université de Toulouse, 31055 Toulouse, France*
[*] *European Space Research and Technology Centre, European Space Agency, 2201 AZ Noordwijk, Netherlands*
guenther.waxenegger@dlr.de
[†]Corresponding author

Abstract

Machine learning methods have demonstrated remarkable performance across various domains. In engineering, they enable the creation of data-driven models that are highly accurate and computationally efficient. These models not only facilitate complex design optimization, but also allow real-time deployment under resource-constrained conditions, an essential factor for applications ranging from autonomous driving to spacecraft control. As part of the ESA project ADMIRE, the potential of machine learning techniques for the design and operation of chemical space propulsion systems has been thoroughly investigated. The suitability of models such as neural networks, as well as their prerequisites, have been critically evaluated. In particular, an assessment of the available versus required (experimental) data has been conducted. Furthermore, promising application areas have been identified and quantitatively assessed through various test cases. These test cases cover key aspects such as cooling channel design, prediction of thermoacoustic instabilities, injection processes, and engine control.

1. Introduction

Machine learning methods have the potential to make the development and operation of space propulsion systems more efficient, faster, and safer. In addition to the associated cost savings, the expected increases in performance are essential, especially for future reusable systems. The main reason for the great opportunities associated with the use of machine learning is the following. Many engineering problems typically need accurate models and simulations to evaluate the implications of design variables and constraints. Furthermore, the computational cost needs to be moderate to enable optimization loops or even real-time applications. Machine learning models like neural networks represent a convincing way to fulfill both aspects. In the ESA project ADMIRE the suitability of state-of-the-art machine learning methods within typical design and development processes for space propulsion systems was investigated.

2. Assessment of machine learning methods

2.1 Modelling and simulation

The main disadvantage of high-fidelity computational fluid dynamics (CFD) or finite element method (FEM) calculations is that they are not suitable for design space exploration and extensive sensitivity analysis due to their large calculation effort.[6] By constructing surrogate models using samples of the computationally expensive calculation, one can alleviate this burden. However, it is crucial that the surrogate model mimics the behaviour of the simulation model as closely as possible and generalizes well to unsampled locations while being computationally cheap to evaluate. Machine learning models like neural networks have been successfully applied as surrogate models in several domains, e.g. fluid mechanics, structural mechanics, acoustics, and combustion.

DOI: 10.13009/EUCASS2025-523

2.2 Control and monitoring

Key technologies for the successful operation of reusable space transportation systems are the control and condition monitoring of the engines.[4,12] Space transportation systems that land again with retro-thrust require additional deep thrust throttling and restart capabilities. Ongoing work to integrate machine learning methods into control and monitoring systems spans a wide range of industries, aiming to improve system performance, adaptability, and fault diagnosis.[15] Furthermore, machine learning techniques are increasingly being used for prognosis, enabling the estimation of a system's remaining useful life based on historical and real-time data. This supports predictive maintenance strategies, where maintenance actions are planned based on the predicted condition of components rather than fixed schedules.

2.3 Data requirements

The amount of training data required for machine learning depends on several factors, such as the complexity of the problem, the quality of the data, and the type of model being used.[1,7] In general, the more training data you have, the better the model will perform. However, there is no set rule for how much data is required. Some problems may require only a few hundred data points, while others may require millions. As a general guideline, it is recommended to have at least several thousand data points per class or category in classification problems, and enough data to cover the range of variations in the problem. It's important to note that having a large amount of data does not necessarily guarantee good performance. The quality of the data is equally important, and it is important to ensure that the data is representative of the problem being solved and does not contain biases or errors. In addition to the amount and quality of data, the training process itself is also important. Hyperparameters, such as the learning rate and batch size, can have a significant impact on model performance, and it is important to experiment with different settings to find the optimal combination for ones problem.

2.4 Synthetic data

For some problems there is the possibility to increase the number of data samples with the help of so-called synthetic data. This data is created using algorithms or simulations and is designed to mimic the characteristics of real-world data. Synthetic data can be useful in situations where there is limited or no access to real-world data, or where the real-world data is insufficient. It is important to note that synthetic data should be validated to ensure that it accurately represents the characteristics of the real-world data it is meant to mimic. Otherwise, the machine learning model trained on synthetic data may not perform well when applied to real-world scenarios. In case the synthetic data contain a minor modelling error but some real data exist, approaches like transfer learning can be used. Transfer learning is a machine learning technique where a model is first pre-trained on a large dataset, and then fine-tuned on a smaller dataset. The idea behind transfer learning is that the pre-trained model has already learned useful features and patterns from the initial dataset, which can then be applied to the new dataset.

3. Test cases

Since the requirements are different for different learning methods and application topics, different application topics have been investigated: cooling channel design, injector design, combustion instability forecasting, and engine control design.

3.1 Cooling channel design

Several methods exist to study the regenerative cooling of liquid rocket engines thrust chambers. A simple approach is to use semi-empirical one-dimensional correlations to estimate the local heat transfer coefficient. However, one-dimensional relations are not able to capture all relevant effects that occur in asymmetrically heated channels like thermal stratification or the influence of turbulence and wall roughness. Especially when using methane as the coolant, the prediction is challenging and simple correlations are not sufficient.

The starting point of the test case was a neural network, which was developed by Waxenegger-Wilfing et al.[19] and later improved by Dresia et al.[3] using data from computational fluid dynamics to predict hot gas wall temperatures along cooling channels. The proposed neural network was a fully connected, feedforward network, fully connected meaning every neuron of one layer is linked with all neurons of the next layer. The network consisted of 4 hidden layers with 408 neurons per layer. Figure 1 shows such an exemplary network. The input parameters shown are all

DOI: 10.13009/EUCASS2025-523

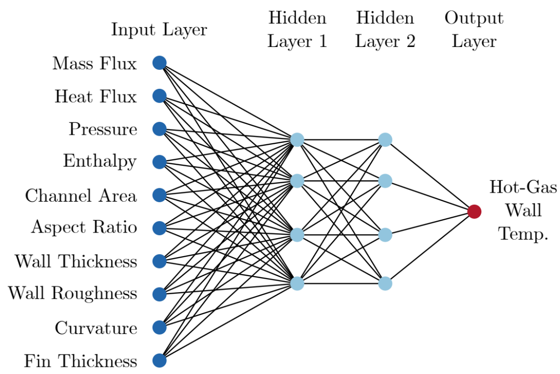

Figure 1: Schematic neural network architecture for wall temperature prediction

parameters used for this case as well: coolant mass flux, heat flux, pressure, enthalpy, channel area, aspect ratio, wall thickness, wall roughness, curvature, and fin thickness.

This baseline neural network was then used for transfer learning with experimental data from different application cases. First, data from the HARCC (High Aspect Ratio Cooling Channel) experiments with methane at the P8 test bench in Lampoldshausen were used.[8] Second, data from experiments with a nitrous oxide cooled 22 N thruster that has been operated at the M11 test bench in Lampoldshausen were investigated.[9] In both cases, the transfer learned network was able to adapt without the need of extensive training data by using only a very small set of experimental data points. Thus, the transfer learning approach could be concluded successfully.

The HARCC combustion chamber segment is divided into four sectors around the circumference, each containing different cooling channel geometries. Table 1 and Table 2 show the mean average error and the maximum absolute error for two HARCC test cases.

Test case / quadrant	Mean average deviation [K]	Maximum absolute error [K]
T1 / Q1	-444.99	605.51
T1 / Q2	-330.49	492.96
T1 / Q3	-30.87	81.26
T1 / Q4	-97.68	135.42
T2 / Q1	-221.83	264.81
T2 / Q2	1.43	38.67
T2 / Q3	-72.46	108.99
T2 / Q4	-83.6	101.41

Table 1: Performance of baseline neural network

Test case / quadrant	Mean average deviation [K]	Maximum absolute error [K]
T1 / Q1	-19.96	115.92
T1 / Q2	6.26	79.35
T1 / Q3	25.87	34.46
T1 / Q4	5.58	42.96
T2 / Q1	19.84	58.36
T2 / Q2	-6.93	23.68
T2 / Q3	-58.19	74.5
T2 / Q4	-58.42.6	71.58

Table 2: Performance of refined neural network

In Figure 2 the results for the predictions with both the baseline neural network and the refined neural network are shown.

DOI: 10.13009/EUCASS2025-523

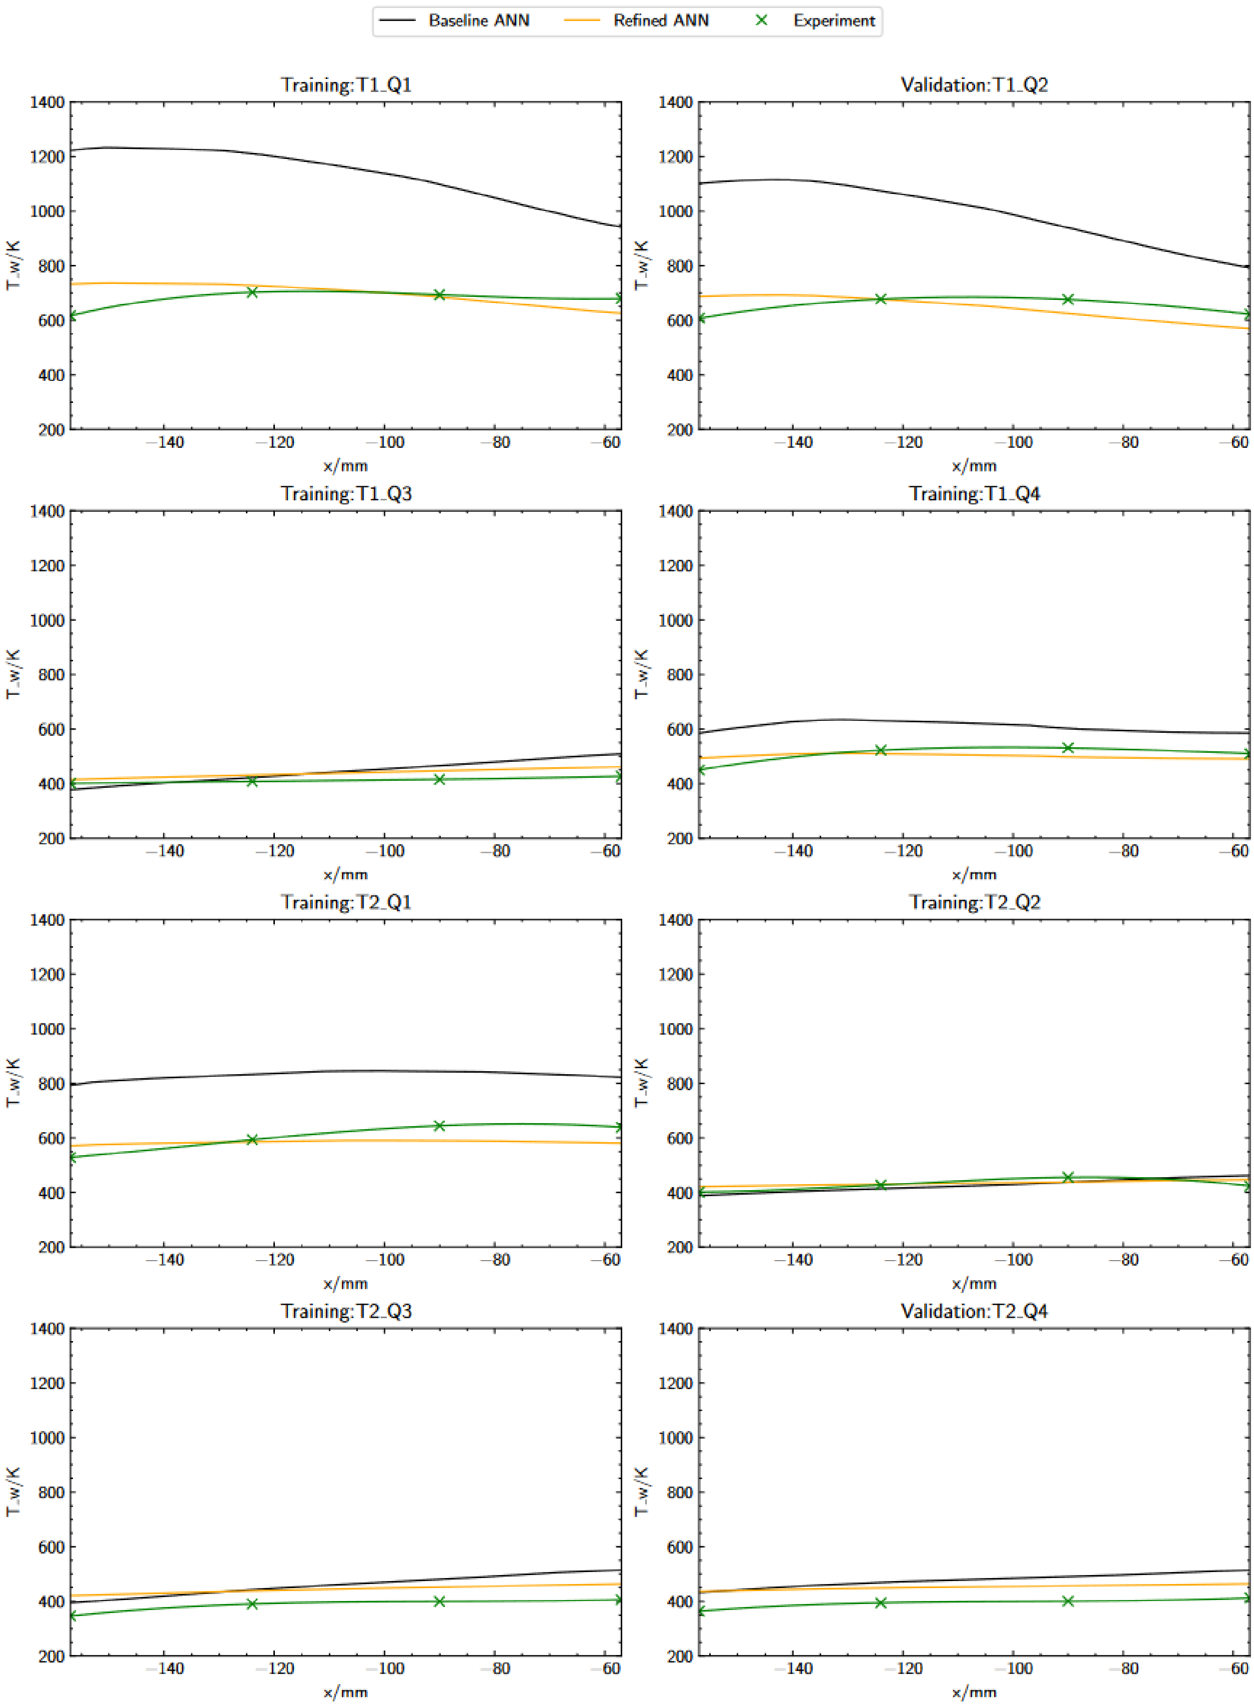

Figure 2: Predictions for HARCC test cases with baseline and refined neural network

DOI: 10.13009/EUCASS2025-523

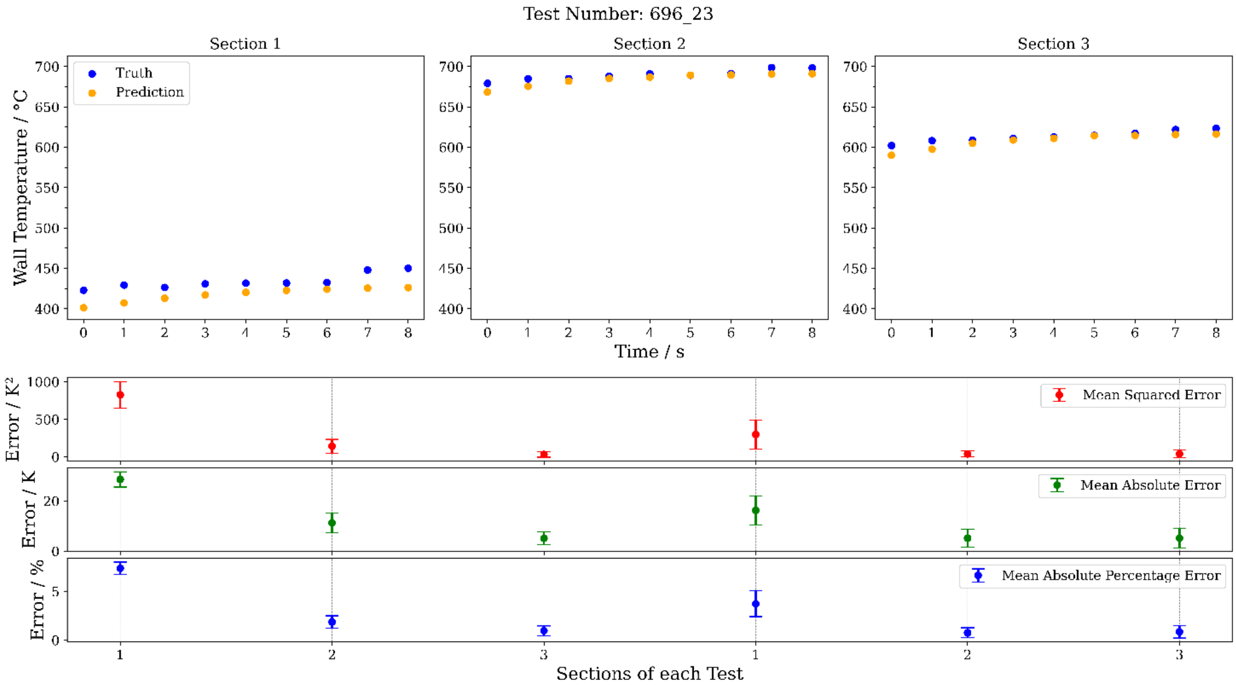

Figure 3: Predicted wall temperatures and associated errors for a test case with nitrous oxide cooling

Figure 3 shows a typical result for a test case with nitrous oxide cooling. Predictions have also been made for a representative combustion chamber geometry and heat flux profile of an upper stage engine using methane as coolant. However, the transfer learning (with HARCC data) caused the neural network to specialize in the new case instead of a generalization. It was noticed that this capability of a neural network to efficiently fine-tune to a specific test case could provide a new possibility for adapting to a specific combustion chamber instead of creating a generalized tool.

3.2 Injector design

Surrogate models that map the influence of design variables on the performance of injection systems with a high degree of precision have great potential to accelerate the design of new combustion chambers.[16,17] However, models such as neural networks are bound to learn the structural relations of the data they are trained upon. Consequently, models obtained from low-fidelity data will reproduce the inherent errors of low-resolution samples.

In this work, we intended to address this conundrum by adopting a multi-fidelity strategy. For such, two large-eddy simulation (LES) datasets of shear-coaxial injectors, operating on a gaseous oxygen (GOx) and gaseous methane (GCH4) propellant couple, were created. The number of high-fidelity samples was limited due to their high computational cost. The application of transfer learning techniques while making use of both the high and low fidelity datasets, enabled the creation of data-driven emulators with satisfactory performance on the high-fidelity test dataset.

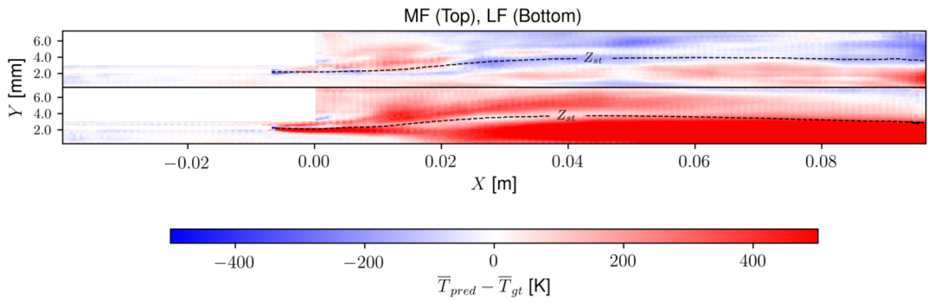

Figure 4: Difference between predicted time-averaged temperature field and ground-truth value for both the multi-fidelity (MF) and low-fidelity (LF) models

DOI: 10.13009/EUCASS2025-523

The change between multi-fidelity (MF) and low fidelity (LF) is shown in Figure 4, where the difference between the estimated, and the ground-truth field, for the same sample, is displayed. A strong red area, below the stoichiometric line, is shown for the LF model. Comparatively, the MF model error figures do not display this region, indicating that transfer learning provided the sought correction.

Furthermore, the same exact network and methodology as above was used to finetune the model to the LUMEN injector.[14] It is important to remember that there are many differences between the two cases: thermodynamic conditions, perfect vs real gas, geometrical details of the injectors. Figure 5 shows the difference between reference and predicted temperature fields for both the original model (top) and the finetuned one (bottom) on a validation case. It clearly shows the strength and possibilities offered by the developed transfer learning methodology.

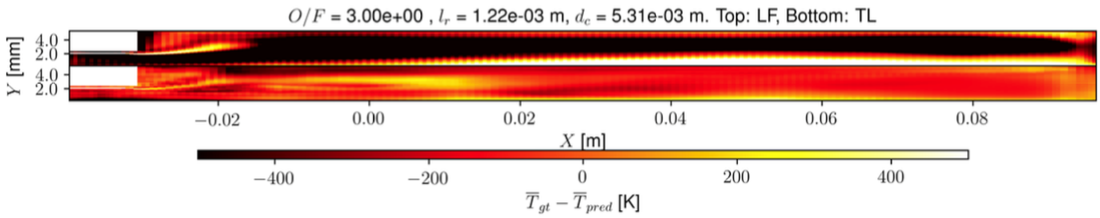

Figure 5: Temperature difference between reference solution, original and finetuned models

3.3 Combustion instability forecasting

Having a reliable forecast of the dynamic pressure in the combustion chamber can be an important tool in instability forecasting. The early detection or prognosis of instabilities could protect test combustion chambers from the damage caused by combustion instabilities during the development phase.[20]

In this test case we further developed models that forecast the mean RMS (Root Mean Square) value of all dynamic pressure sensors within the combustion chamber into the future.[13] The available dataset consisted of ignition-sequences of the BKA, which was operated at the DLR test bench P8 in Lampoldshausen.[11] Combustion instability was frequently observed during these ignition tests. For each ignition sequence, data from several sensors has been gathered and analyzed. This included (dynamic and static) pressure sensors, thermometers, fuel rate sensors and so on. Two different neural network categories were trained and evaluated for the given task. First, we investigated Feedforward Neural Networks (FNNs), consisting of Convolutional and Dense Layers. These are the most basic and straight-forward neural network models, and perform well on a wide variety of tasks. Second, we investigated Recurrent Neural Networks (RNNs), consisting of stateful LSTM layers. These networks contain connections that form cycles, therefore giving them a kind of memory of previous inputs, making them particularly suitable for time series forecasting. Furthermore, autoregressive models were investigated, but discarded as they did not show any significant improvements. Figure 6 shows the predictions for an exemplary test ignition sequence. The mean RMS value is predicted 150 ms into the future.

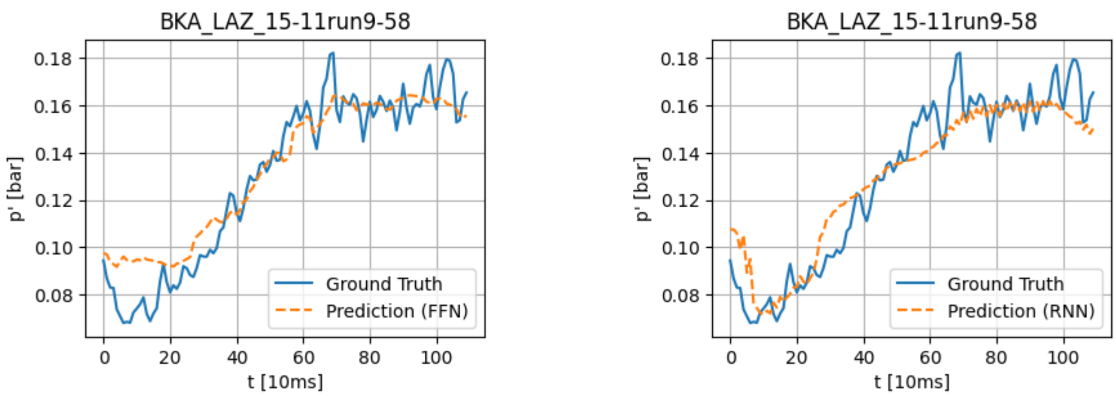

(a) Trained FNN

(b) Trained RNN

Figure 6: Prediction of the mean RMS of the dynamic pressure

DOI: 10.13009/EUCASS2025-523

One can see that the trend is predicted quite well. However, depending on the point in time and test case, the prediction of the exact value may show larger deviations. Table 3 and Table 4 show the mean absolute error and the maximum absolute error for all test ignitions.

Test case	Mean absolute error [bar]	Maximum absolute error [bar]
T1	0.015	0.052
T2	0.018	0.180
T3	0.075	0.388
T4	0.139	0.413
T5	0.009	0.028

Table 3: Trained FNN

Test case	Mean absolute error [bar]	Maximum absolute error [bar]
T1	0.023	0.108
T2	0.025	0.478
T3	0.050	0.363
T4	0.129	0.491
T5	0.009	0.031

Table 4: Trained RNN

For the last case, the prediction is very accurate, for the third and fourth test case, there are maximum deviations in the range of 0.4-0.5 bar. Thus, the predictions are still too inaccurate to be reliably used for combustion instability forecasting. This is mostly due to the following facts. First, the available training dataset is limited, especially once short or ill-suited experiments are sorted out. Second, the experimental data originates from the ignition phase, which has much more complex dynamics than the phases thereafter, in which combustion instabilities have already been successfully forecasted with greater accuracy using machine learning methods.

3.4 Engine control design

An already existing 22 N experimental nitrous oxide/ethane thruster was used as test case and is shown in Figure 7. The thruster operates with a self-pressurized propulsion system. Due to the evaporation of the propellant, the tanks cool down during thruster operation. Thus, the vapor pressure in the tank changes continuously and for a steady operation of the thruster, a permanent adjustment of the mass flow is needed. The control objectives were given by regulating the mixture ratio and the combustion chamber pressure via flow control valves.

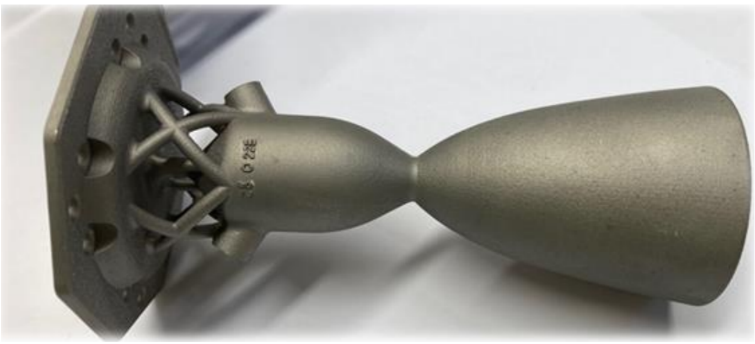

Figure 7: Thruster used for testing

Based on a simulation model with EcosimPro / ESPSS, deep reinforcement learning was used to train the controller represented by a neural network.[2,5,10,18] First, the trained controller was tested in the simulation environment. Two different scenarios were evaluated. First, the controller had the task of setting a specific target operating point in terms of pressure and mixture ratio, when starting from a different operating point. Second, it had to follow a time dependent target. The Control frequency was 10 Hz in both cases. Figure 8 shows a single simulation of a set-point control task.

DOI: 10.13009/EUCASS2025-523

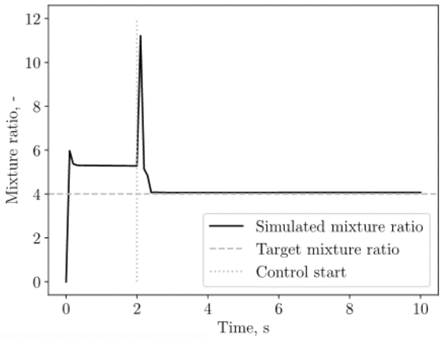

(a) Mixture ratio

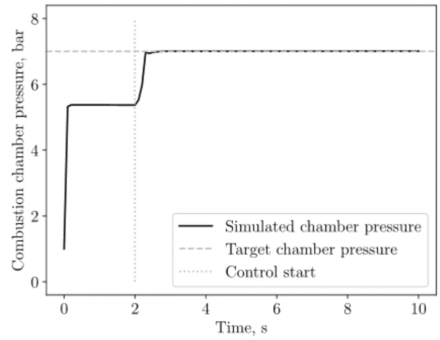

(b) Chamber pressure

Figure 8: Simulated single set-point control for mixture ratio and chamber pressure

In Figure 9 a trajectory of six different set-points is given. The set-point changes every 10 s in a 'step response-like' manner. Both controlled variables (chamber pressure and mixture ratio) are changed. First the chamber pressure is kept constant while the mixture ratio is changed. At 40 s the mixture ratio is kept constant, while the chamber pressure is changed. It can be seen, that the target values of chamber pressure and mixture ratio were met nearly perfectly.

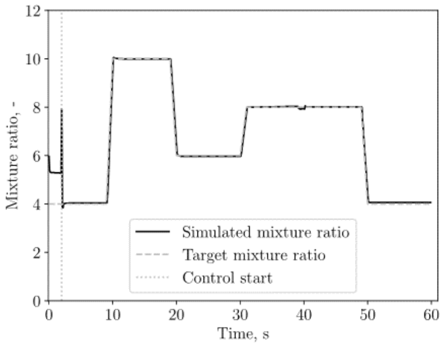

(a) Mixture ratio

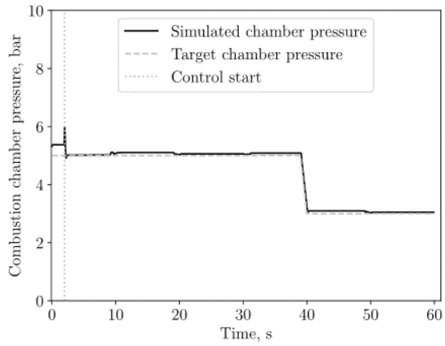

(b) Chamber pressure

Figure 9: Simulated trajectory with step response like changes in mixture ratio and chamber pressure

Although the controller worked very well in the simulation environment, the performance at the test facility was worse and not yet satisfying. As can be seen in Figure 10 the neural network was able to settle the right values, however strong superimposed fluctuations in both controlled variables and long settling times are not acceptable for real applications.

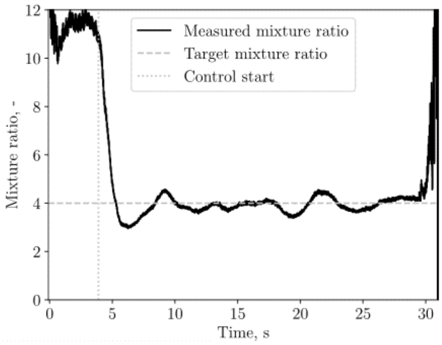

(a) Mixture ratio

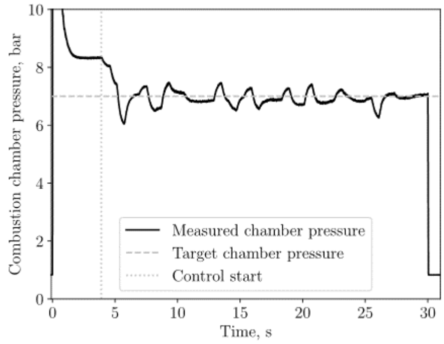

(b) Chamber pressure

Figure 10: Real single set-point control for mixture ratio and chamber pressure

DOI: 10.13009/EUCASS2025-523

4. Conclusions and outlook

Two main conclusions can be drawn from the results. First, machine learning methods can support the design process and development of new engines by, for example, enabling efficient optimization based on powerful surrogate models or extending the lifetime of test combustion chambers through improved online prediction of combustion instabilities. In addition, the use of trained controllers promises increased performance for propulsion systems with limited computing resources. Second, the biggest challenge is the small amount of training data that is naturally available in the development of (new) rocket engines. Solutions such as the use of synthetic data or transfer learning approaches improve the situation but are not always possible and must be adapted to the particular real application.

Follow-up activities should further increase usability so that machine learning methods can prove that their application can improve the development and operation of engines, save costs by reducing development times, and ensure that the engine control meets the increased requirements resulting from reuse. These activities include:

- Improve transfer learning approaches that enable fine-tuning surrogate models (e.g. NN-based heat transfer predictors) from one engine design to another with minimal additional data.
- Investigate transfer learning methods that align simulation output with test-stand or flight data.
- Explore transfer learning across related tasks, for example, from steady-state operation to transient startup.
- Develop domain adaptation approaches to handle distributional changes as engines move between test stands, ground hotfire, and flight application.
- Investigate the lifelong adaptation of control policies to enable controllers to maintain performance as the engine ages.

A recurring theme across these activities is bridging gaps-between different engines, between simulation and test-stand data, and between evolving operating conditions. By leveraging techniques such as transfer learning, domain adaptation, and continual learning, practitioners can develop models that generalize across engine variants, adapt to new regimes without retraining from scratch, and continuously improve.

5. Funding

This work has been carried out under a funded contract from the European Space Agency under contract number 4000139505/22/NL/MG. The views expressed in this publication can in no way be taken to reflect the official opinion of the European Space Agency.

References

[1] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. Springer, New York, 2006.

[2] Kai Dresia, Antonius Adler, Anirudh Saraf, Robson Hahn, Eldin Kurudzija, Tobias Traudt, Jan Deeken, and Günther Waxenegger-Wilfing. Rocket engine control with neural networks: Experimental results of the lumen turbopump test campaign. In *AIAA SCITECH 2023 Forum*, page 0137, 2023.

[3] Kai Dresia, Eldin Kurudzija, Jan Deeken, and Günther Waxenegger-Wilfing. Improved wall temperature prediction for the lumen rocket combustion chamber with neural networks. *Aerospace*, 10(5):450, 2023.

[4] Kai Dresia, Tobias Traudt, Till Hörger, Jonas Dauer, Jan C Deeken, and Günther Waxenegger-Wilfing. Overview of future rocket engine control systems. In *Proceedings of the 9th Edition of the Space Propulsion Conference*, 2024.

[5] Kai Dresia, Günther Waxenegger-Wilfing, and Michael Oschwald. Nonlinear Control of an Expander-Bleed Rocket Engine using Reinforcement Learning. In *Proceedings of the Space Propulsion 2020+1 Conference*, Virtual Event, 2021.

[6] Michael Frank, Dimitris Drikakis, and Vassilis Charissis. Machine-Learning Methods for Computational Science and Engineering. *Computation*, 8(1):15, March 2020.

DOI: 10.13009/EUCASS2025-523

[7] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Adaptive Computation and Machine Learning. The MIT Press, Cambridge, Massachusetts, 2016.

[8] Jan Haemisch, Michael Oschwald, and Dmitry Suslov. Experimental analysis of heat transfer processes in cooling channels of a subscale combustion chamber at real thermal conditions for cryogenic hydrogen and methane. In *Proceedings of the Space Propulsion 2018 Conference*, 2018.

[9] Till Hörger, David Heyer, Kai Dresia, and Günther Waxenegger-Wilfing. Transfer learning for wall temperature prediction in regeneratively cooled 22n thrust chambers. In *Proceedings of SPAICE2024: The First Joint European Space Agency/IAA Conference on AI in and for Space*, pages 383–387, 2024.

[10] Till Hörger, Lukas Werling, Kai Dresia, Günther Waxenegger-Wilfing, and Stefan Schlechtriem. Experimental and simulative evaluation of a reinforcement learning based cold gas thrust chamber pressure controller. *Acta Astronautica*, 219:128–137, 2024.

[11] Sebastian Klein, Michael Börner, Justin S Hardi, Dmitry Suslov, and Michael Oschwald. Injector-coupled thermoacoustic instabilities in an experimental lox-methane rocket combustor during start-up. *CEAS Space Journal*, 12:267–279, 2020.

[12] Eldin Kurudzija, Kai Dresia, Jan Martin, Tobias Traudt, Jan C Deeken, and Günther Waxenegger-Wilfing. Virtual sensing for fault detection within the lumen fuel turbopump test campaign. In *Proceedings of the 9th Space Propulsion Conference*, 2024.

[13] Ushnish Sengupta, Günther Waxenegger-Wilfing, Jan Martin, Justin Hardi, and Matthew P Juniper. Forecasting thermoacoustic instabilities in liquid propellant rocket engines using multimodal bayesian deep learning. *International Journal of Spray and Combustion Dynamics*, 14(3-4):218–228, 2022.

[14] Tobias Traudt, Wolfgang Armbruster, Christopher Groll, Robson Henrique Dos Santos Hahn, Kai Dresia, Michael Börner, Sebastian Klein, Dmitry Suslov, Max Axel Müller, Jan Haemisch, et al. Lumen: Liquid upper stage demonstrator engine-a versatile test bed for rocket engine compo-nents: Hot-fire test results. In *2024 IAF Space Exploration Symposium at the 75th International Astronautical Congress, IAC 2024*, 2024.

[15] Andrea Urgolo, Chiara Gei, Kai Dresia, Manuel Freiberger, Eldin Kurudzija, Heike Neumann, Ingo Pill, Federico Pittino, Gleb Radchenko, and Günther Waxenegger-Wilfing. Feature selection and virtual sensing based mixture ratio estimation for liquid propellant rocket engine control systems. In *9th EDITION OF THE 3AF INTERNATIONAL CONFERENCE ON SPACE PROPULSION*, page 416, 2025.

[16] José F Zapata Usandivaras, Michael Bauerheim, Bénédicte Cuenot, and Annafederica Urbano. Deep convolutional autoencoders for the time–space reconstruction of liquid rocket engine flames. *Proceedings of the Combustion Institute*, 40(1-4):105382, 2024.

[17] Jose Felix Zapata Usandivaras, Michael Bauerheim, Bénédicte Cuenot, and Annafederica Urbano. Data-driven multifidelity surrogate models for rocket engines injector design. *Data-Centric Engineering*, 6:e2, 2025.

[18] Günther Waxenegger-Wilfing, Kai Dresia, Jan Deeken, and Michael Oschwald. A reinforcement learning approach for transient control of liquid rocket engines. *IEEE Transactions on Aerospace and Electronic Systems*, 57(5):2938–2952, 2021.

[19] Günther Waxenegger-Wilfing, Kai Dresia, Jan Christian Deeken, and Michael Oschwald. Heat Transfer Prediction for Methane in Regenerative Cooling Channels with Neural Networks. *Journal of Thermophysics and Heat Transfer*, 34(2):347–357, 2020.

[20] Günther Waxenegger-Wilfing, Ushnish Sengupta, Jan Martin, Wolfgang Armbruster, Justin Hardi, Matthew Juniper, and Michael Oschwald. Early detection of thermoacoustic instabilities in a cryogenic rocket thrust chamber using combustion noise features and machine learning. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(6), 2021.