Ref. Ares(2023)7884193 - 20/11/2023

HORIZON-EUROHPC-JU-2021-COE-01

EXCELLERAT P2

The European Centre of Excellence for Engineering Applications

Project Number: 101092621

D2.8 Report on the m-AIA Application Use Case

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

The EXCELLERAT P2 project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101092621. The JU receives support from the European Union's Horizon Europe research and innovation programme and Germany, Italy, Slovenia, Spain, Sweden and France.

Work Package:	2	Use-Case Execution
Author(s):	Ansgar Niemöller, Thomas Hösgen, Matthias Meinke	RWTH
Approved by	Executive Centre Management	20.11.2023
Reviewer	Marvin Hubl	SSC
Reviewer	Dennis Grieger	USTUTT
Dissemination Level	Public	

Date	Author	Comments	Version	Status
2023-10-26	Ansgar Niemöller, Thomas Hösgen, Matthias Meinke	Initial draft	V0.1	Draft
2023-11-04	Ansgar Niemöller	Second draft	V0.2	Draft
2023-11-16	Ansgar Niemöller	Finalised version	V1.0	Pending approval

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

List of abbreviations

ACARE	Advisory Council for Aeronautical Research in Europe
AI	Artificial Intelligence
CAA	Computational aeroacoustics
CFD	Computational fluid dynamics
DOF	Degrees of Freedom
FWH	Ffowcs Williams-Hawkings
HPC	High-performance computing
LES	Large-eddy simulation
m-AIA	multi-physics AIA
MPI	Message passing interface
NACA	National Advisory Committee for Aeronautics
SMC	Small metal chevron
STL	Stereolithography
UC	Use Case

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

Executive Summary

This document presents the progress made in the m-AIA Application Use Case UC-3 within reporting period 1 covering the first year of the EXCELLERAT P2 project. Based on the detailed roadmap of the workflow development defined in deliverable D2.1, the workflow of the use case is summarised and the achieved progress with respect to the defined workflow, objectives and success criteria is presented.

In summary, the workflow development for UC-3 has progressed according to the schedule defined in deliverable D2.1. Work has been performed on the individual tasks planned for the first year of the project.

That is, progress has been made concerning the development of the individual workflow components, the workflow automation, the exascale readiness of the m-AIA code and the implementation of an optimisation approach.

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

Table of Contents

1 Introduction ... 7
2 Objectives of the Use Case ... 8
3 Workflow Description ... 8
4 Progress achieved ... 10
5 Conclusion ... 13
6 References ... 13

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

Table of Figures

Figure 1: Baseline (left) and typical chevron nozzle (right) 7

Figure 2: Computational setup with the LES or CFD, and the CAA domain for the jet noise prediction..... 7

Figure 3: Coupling of the LES or CFD with the CAA solver and the FWH method 8

Figure 4: Workflow components for the constrained shape optimisation of chevron nozzles .. 9

Figure 5: Chevron geometry generator: close-up of STL data for the SMC006 baseline and two variations of the chevron geometry. 10

Figure 6: Strong scaling for a coupled CFD/CAA chevron jet application with m-AIA......... 11

Figure 7: m-AIA coupled CFD/CAA simulation: baseline nozzle without chevrons showing flow structures (bottom) and the acoustic field close to the nozzle (top). 11

Figure 8: Flow field around NACA 2412 airfoil for two angles of attack. 11

Figure 9: Convergence of lift to drag coefficient ratio for NACA 2412 airfoil at two angles of attack. 11

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

1 Introduction

In this use case, a shape optimisation of chevron nozzles is performed, which is representative for problems with expensive objective function evaluations in the important technical field of noise reduction. It includes a constraint of minimum jet thrust loss, which has to be weighted with the goal of noise reduction. The accurate prediction of the emitted sound is based on a turbulence scale resolving CFD method, directly coupled with a computational aeroacoustics solver, which is a typical example of a multiphysics application. The necessary large number of objective function evaluations can only be performed on exascale HPC systems. An AI-based optimiser, developed by the project partner FhG, will be used to efficiently identify the optimal solution. The motivation for the use case is given e.g., by ACARE, who established the Flightpath 2050, a new goal for more rigorous noise reduction by 65 percent relative to the capabilities of typical new aircraft in 2000. Despite the progressive introduction of high-bypass-ratio aircraft engines and chevron nozzles, which possess a sawtooth-like shape at the engine's trailing edge, jet noise still is a significant source of aircraft engine noise.

Unlike the optimisation of the aerodynamic performance or structural weight, noise reduction is still to a large extent an unsolved problem. One of the challenges connected to noise reduction is a reliable and accurate prediction of the sound pressure level in the far field, which is often generated by intricate flow phenomena. Therefore, turbulence scale resolving simulations have to be performed in many cases to obtain the correct sound pressure level. For example, the tip gap vortex in an axial fan can generate noise, which cannot be predicted with methods based on Reynolds averaged solutions. At present, shape optimisations are not possible for such cases due to the computational expensive evaluation of the objective function.

A typical chevron nozzle shape is shown in Figure 1 and the setup for the flow and the acoustic field prediction is depicted in Figure 2.

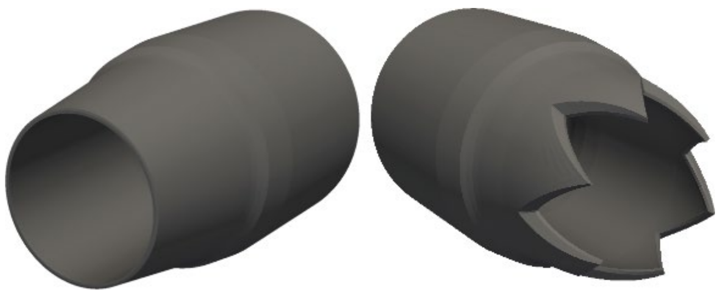

Figure 1: Baseline (left) and typical chevron nozzle (right).

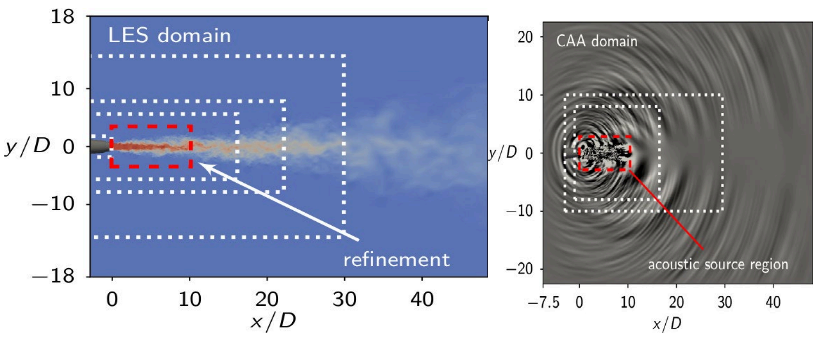

Figure 2: Computational setup with the LES or CFD, and the CAA domain for the jet noise prediction.

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

2 Objectives of the Use Case

It is the objective of this use case to perform a shape optimisation for the shape of the chevrons at the nozzle trailing edge to minimise the jet noise. During aircraft take-off and cruise flight, however, the use of chevron nozzles can lead to a substantial and unwanted loss of thrust. Therefore, the optimisation has to be performed under the constraint that the thrust is minimally affected. Since the objective function evaluation is computationally extremely expensive, HPC clusters of exascale have to be used efficiently. Therefore, an advanced workflow has to be designed, which can fully exploit the available computing resources. The goal is to implement all necessary workflow components such as the automatic generation of chevron shapes, the execution of the simulation for the flow field prediction and the determination of the complex objective function. Additionally, highly resolved simulations should be performed, which enable the identification of the essential noise source locations and to determine the effect of the various chevron parameters on the generation of the acoustic waves. For the execution of the use case, the workflow couples a high-fidelity 3-D aeroacoustics solver with state-of-the-art AI optimisation algorithms provided by the project partners FhG and data analytic tools to systematically perform thrust-constrained noise minimisation of chevron nozzle shapes.
For UC-3 the following success criteria are used during the course of the project to verify its success:

- Development of a workflow, which can automatically perform shape optimisations.
- Efficient usage of HPC hardware during the workflow execution.
- Identification of chevron shapes with minimized noise emission.
- Successful analysis of noise generation mechanisms based on large scale simulation runs on exascale hardware.
- Analysis of noise source mechanisms based on large scale simulations.

3 Workflow Description

The accurate prediction of the overall sound pressure level in the far field requires several numerical methods which must be combined to obtain the final result for the objective function. The CFD solver for the prediction of the turbulent flow field is directly coupled to the CAA solver for the determination of sound pressure levels in the acoustic near field, which then delivers unsteady data for a FWH method, which can compute the overall sound pressure level in the far field. The coupling between the CFD and CAA solver is schematically explained in Figure 3.

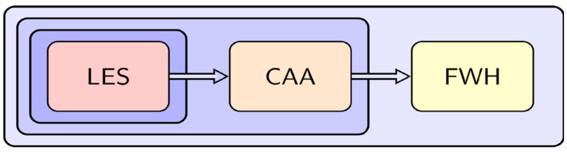

Figure 3: Coupling of the LES or CFD with the CAA solver and the FWH method.

A single simulation is performed in various stages to minimize the required computational effort. First, the CFD solver is executed alone, until a fully developed flow field and a sufficiently converged time averaged flow field is obtained. Subsequently, the CFD and CAA solver are executed in a fully coupled manner until the unsteady acoustic pressure signal is obtained in the near sound field. Finally, the history of the unsteady pressure signal on a closed control surface is used to predict the noise in the acoustic far field. The overall sound pressure level is determined at reference locations for the evaluation of the objective function. In addition, the flow field is postprocessed to compute the nozzle thrust.

The full workflow to perform chevron shape optimisations is schematically depicted in Figure 4:

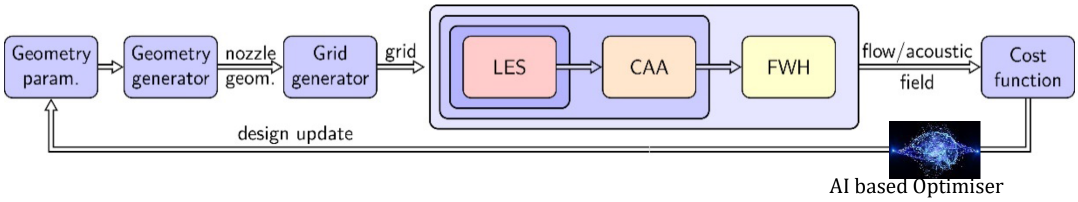

Figure 4: Workflow components for the constrained shape optimisation of chevron nozzles.

In more detail, the workflow components for the jet noise prediction are:

- AI-based optimiser provides a chevron nozzle design, i.e., new chevron geometry parameters.
- Geometry generator generates STL surfaces for the mesh generation for the specific chevron nozzle shape.
- The grid generation generates refined meshes for the LES/CFD and CAA simulation.
- The LES/CFD simulation is performed to determine the time averaged mean flow field.
- The coupled CFD/CAA simulation is performed to predict the acoustic pressure in the near field.
- FWH method to predict the acoustic far-field.
- Postprocessing of the FWH and CFD results for the cost or objective function evaluation.

Regarding an optimisation process, the simultaneous execution of the many required individual runs poses a significant challenge for the coordination of the computational resources (various hardware platforms, queueing systems). Since the runs are expected to last for different lengths of time an intelligent, optimised coordination of the simulations is necessary. Since the AI-based optimisation algorithm is directly fed with data from the simulation (in situ) the simultaneous provision of large CPU/GPU resources is challenging. To enable a subsequent analysis of the simulation results, it is necessary to efficiently store the data on long-term storage systems which is challenging due to the large amount of data.

In summary, multiple simulations with different phases or configurations, various output files, and large data volume have to be performed. This requires a reliable and fault tolerant automation and high efficiency without the necessity of user interaction or supervision.

The exascale execution profile is defined by many small scale runs for the shape optimisation, i.e., requiring on the order of $O(1000)$ runs. In addition, a few large scale runs ($O(10)$) will be performed for advanced data analytics to identify physical mechanisms. Pre-exascale hero runs ($O(1)$) are planned for higher Reynolds number, coaxial jet configurations and very high-fidelity results to demonstrate exascale readiness.

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

4 Progress achieved

In UC-3 an optimisation workflow for chevron nozzles with the target of noise mitigation is developed. To achieve this, the individual workflow components required for an aeroacoustic optimisation workflow have been defined such that ultimately the full approach will be able to perform the targeted chevron nozzle optimisation to reduce jet noise. Furthermore, the HPC hardware requirements for all workflow steps with respect to the exascale execution profile with different types of simulation runs have been analysed.
Based on the workflow analysis, work has been performed on the development of the individual workflow components, which are schematically depicted in the full optimisation workflow in Figure 4. As an input for the grid generator and the subsequent simulation runs the chevron nozzle shape needs to be generated according to some chevron geometry parameters, that will ultimately be provided by the AI-based optimiser. Therefore, a geometry generator to generate parametrised chevron nozzle geometries has been developed. The Matlab/Octave tool generates a new chevron geometry in form of STL data as a direct input for m-AIA. For now, the geometry parameters include the chevron tip angle, i.e., the immersion of the chevron in the jet radial direction, and a shape factor to modify the roundness/sharpness of the chevron shape in the jet circumferential direction. In Figure 5 chevron nozzle geometries are compared to the SMC006 baseline geometry. That is, the chevron tip angle of the baseline configuration is increased from 5 to 10 degrees in a first step. In a second step, the triangular chevron shape is modified such that more rounded chevron lobes are obtained.

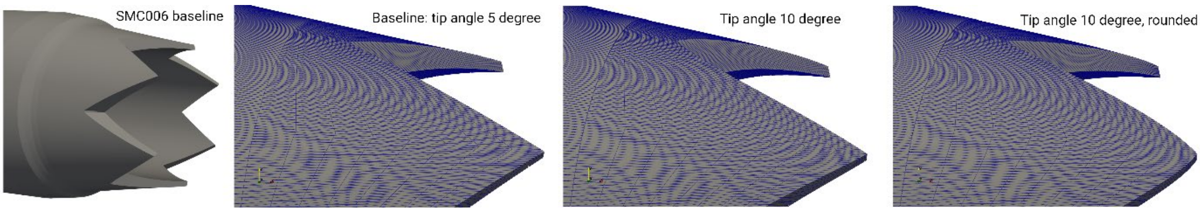

Figure 5: Chevron geometry generator: close-up of STL data for the SMC006 baseline and two variations of the chevron geometry.

Regarding the parallel hierarchical Cartesian grid generator integrated in the m-AIA simulation framework several improvements regarding the generation of complex grids have been achieved. That is, by balancing the grid cells after each step of the grid generation among all MPI ranks the process is significantly sped up, while at the same time less compute nodes are required. Considering that each chevron nozzle configuration will require multiple grid files to be created for the different m-AIA simulation steps, these runtime savings will quickly add up reducing resource usage, and especially a faster overall time to solution in the optimisation workflow can be achieved. For the computation intensive simulation steps with the m-AIA code, significant improvements in performance and parallel efficiency for large-scale multi-physics simulations have been achieved. Repeated testing of an aeroacoustic application that was scaled up to utilize the full HAWK HPC system of about 500,000 compute cores allowed the identification and elimination of performance issues and emerging bottlenecks, which were not visible for smaller scale runs or less complex simulation setups. For example, a critical issue related to a specific inter-process communication was discovered. Appropriate changes were introduced into the critical part of the communication modules in m-AIA, which eliminated the observed performance issues at large scale.

The strong scalability of a realistic coupled CFD/CAA chevron jet application with m-AIA on HAWK is shown in Figure 6. The predicted flow field and the acoustic field for a baseline nozzle without chevrons is visualised in Figure 7. With about 300 million CFD cells and $1*10^9$ CAA DOF this setup corresponds to a smaller scale run according to the exascale execution profile defined in WP2 for UC-3. As evident the code shows excellent strong scalability when

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

going from 2048 up to 262144 MPI processes, i.e., the maximum allocation size on HAWK, achieving about 86 simulation timesteps per second compared to 0.68 timesteps for the baseline.

In summary, the progress made on the m-AIA code is key for its future exascale readiness, showing excellent strong scalability and efficient resource usage for a realistic use case on the full HAWK system.

In addition to the work on the individual workflow components, work has been performed on the workflow automation of the m-AIA simulation pipeline for aeroacoustic predictions, to automatically perform all the individual preprocessing steps, simulation runs and postprocessing steps that are required for a single function evaluation embedded in the optimisation loop, i.e., a single jet noise prediction for a specific chevron nozzle geometry.

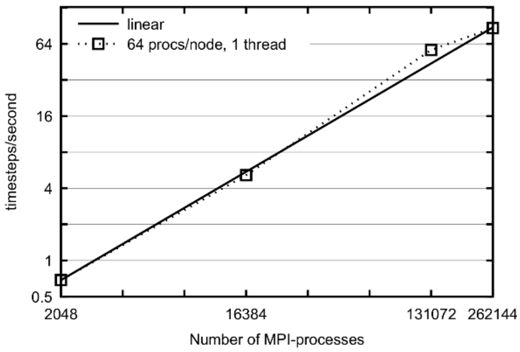

Figure 6: Strong scaling for a coupled CFD/CAA chevron jet application with m-AIA.

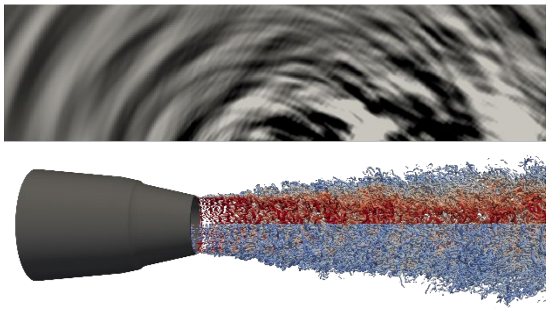

Figure 7: m-AIA coupled CFD/CAA simulation: baseline nozzle without chevrons showing flow structures (bottom) and the acoustic field close to the nozzle (top).

Regarding the targeted AI-based optimisation, RWTH started the collaboration with FhG to work on the optimisation approach with a problem analysis, definition of requirements and plan with next steps. In a first step a simplified small-scale optimisation test setup consisting of a NACA profile was determined, for which the optimisation workflow will be tested at low computational cost. The NACA airfoil will be optimised for reduced drag subject to certain constraints, such as, lift and volume of the airfoil. Initial tests at low inflow Mach- and Reynolds numbers were conducted. In Figure 8 exemplary results for a NACA 2412 airfoil at two angles of attack (AoA) of 0 and 5 degree are shown. The computational grids consist of approximately 10^5 cells. As evident by the ratio of lift and drag coefficient in Figure 9 the simulations quickly converge in about 10,000 time steps rendering the setup suitable for testing the optimisation workflow at low computational cost.

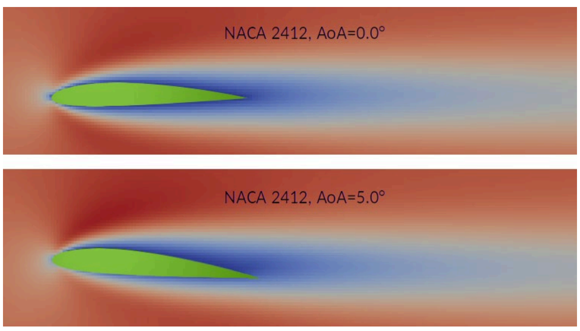

Figure 8: Flow field around NACA 2412 airfoil for two angles of attack.

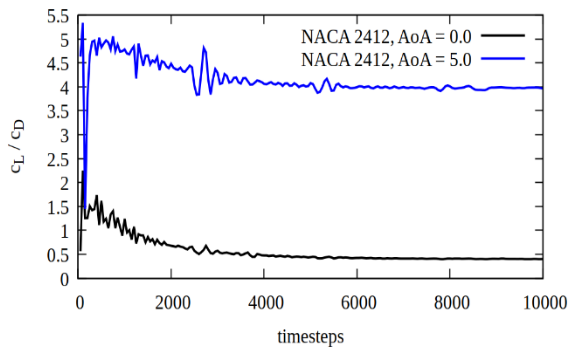

Figure 9: Convergence of lift to drag coefficient ratio for NACA 2412 airfoil at two angles of attack.

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

In the next steps, these CFD-based simulation models, starting with the NACA airfoil testcase, will be integrated into a surrogate-based sequential closed-loop optimisation procedure allowing for a high degree of automation. Bayesian optimisation (BO) with Gaussian Process Regression (GPR) [1] prior as the surrogate machine learning model is one example of surrogate-based optimisation strategy. This method is going to be developed and tested, first on the NACA airfoil testcase and later on the nozzle shape optimisation task. The GPR-based BO workflow consists of the following steps.

1. Formulate the original optimisation problem including the objective function f to be optimised for the optimisation variables $\boldsymbol{x}$ and constraints forming the feasible set $\boldsymbol{\mathcal{X}}$.
2. Perform initial objective function evaluations by running the m-AIA-enabled CFD simulations to compute the objective function f using a space-filling design method for the optimisation variables like Latin Hypercube Sampling (LHS).
3. Start the main optimisation loop using the GPR-based BO procedure, which sequentially updates the data and the GPR surrogate model, i.e., $\hat{f} \approx f$ with $\hat{f}(\boldsymbol{x}) \sim \mathcal{GP}\big(m(\boldsymbol{x}), k(\boldsymbol{x}, \boldsymbol{x}')\big)$, and solves an auxiliary optimisation problem to trade-off exploration and exploitation to efficiently leverage the surrogate model to propose the next simulation to be run, i.e., efficient simulation planning.

A general original problem formulation is stated in the following equation with $\boldsymbol{x}^*$ indicating the true optimum and $\boldsymbol{p}$ denoting preset resp. given parameters, like flow conditions characterised by the problem-specific Reynolds and Mach number.

$$\boldsymbol{x}^* = \underset{x \in \boldsymbol{\mathcal{X}}(\boldsymbol{p})}{\arg\max}\, f(\boldsymbol{x}|\boldsymbol{p})$$

The original problem formulation is reformulated to the auxiliary problem given in the next equation, which is sequentially updated using the surrogate model $\hat{f}$ and solved, hence the index $n+1$, to obtain the next queried simulation with inputs $\boldsymbol{x}_{n+1}$. The updated data in step n is denoted as $D_n = \{(\boldsymbol{x}_i, y_i = f_i + \text{noise}) | 1 \leq i \leq n\}$ and the corresponding prior knowledge as $\boldsymbol{\Pi}_n$, which is used to select the prior mean $m(\boldsymbol{x})$ and kernel function $k(\boldsymbol{x}, \boldsymbol{x}')$ for the GPR model.

$$\boldsymbol{x}_{n+1} = \underset{x \in \widehat{\boldsymbol{\mathcal{X}}}(\boldsymbol{D}_n, \boldsymbol{\Pi}_n)}{\arg\max}\; \alpha\big[\hat{f}(\boldsymbol{x}|\boldsymbol{p}, \boldsymbol{D}_n, \boldsymbol{\Pi}_n, \boldsymbol{\theta})|\boldsymbol{\theta}_\alpha\big]$$

In that sense, exploration means to run simulations, which are the most informative to decrease the uncertainty of the surrogate model. In contrast to that, exploitation refers to the execution of simulations, which are already predicted by the surrogate model to the most promising to be optimal. The trade-off of exploration and exploitation by optimising an acquisition function α with hyperparameters $\boldsymbol{\theta}_\alpha$, e.g., Upper Confidence Bound (UCB) in the simplest case, instead of directly optimising the surrogate model $\hat{f}$ renders this procedure as a global optimisation method. In general, the efficiency of surrogate-based optimisation workflows depends on a carefully designed and tuned surrogate model, i.e., prior mean $m(\boldsymbol{x}|\boldsymbol{\theta}_m)$ and prior kernel function $k(\boldsymbol{x}, \boldsymbol{x}'|\boldsymbol{\theta}_k)$ and it hyperparameters $\boldsymbol{\theta} = [\boldsymbol{\theta}_m, \boldsymbol{\theta}_k]$, and a proper definition resp. handling of optimisation bounds, like box constraints $\boldsymbol{x}_{\min} \leq \boldsymbol{x} \leq \boldsymbol{x}_{\max}$, resp. the feasible set $\widehat{\boldsymbol{\mathcal{X}}}$ in general. When using a stochastic surrogate model like the GPR it is possible to incorporate prior knowledge $\boldsymbol{\Pi}_n$ about the physics-based optimisation problem, like smoothness assumptions of the objective f, to handle noisy expensive-to-evaluate black-box objective functions such as CFD-based simulations in an informed manner. In combination with automatic differentiation applied on the GPR, gradient-based Non-Linear Program (NLP) solvers can be used to solve the corresponding auxiliary problem to realise an efficient and

Copyright © 2023 Members of the EXCELLERAT P2 Consortium

highly automated optimisation procedure.

5 Conclusion

RWTH will focus on the development of the optimisation loop using the Bayesian optimiser and the AI based optimiser provided by the project partner FhG. Once the full optimisation workflow and a reliable and fault tolerant automation of the complete aeroacoustic prediction process have been established, it will be possible to perform constrained shape optimisations. First, the approach will be tested and validated using the defined simplified small-scale test setups. Collaboration with FhG will then concentrate on the integration of the AI-based optimiser, while also more sophisticated aeroacoustic test cases will be targeted.
Furthermore, the simultaneous execution and the interleaved scheduling of individual runs poses a significant challenge. Since the individual runs for the determination of the objective function are expected to last for different run-times an intelligent, optimised coordination of the simulations will be necessary.

6 References

[1] Rasmussen, C. E.: Gaussian Processes for Machine Learning. In: Rasmussen, C. E., Williams, C. K. I., The MIT Press, Cambridge, Massachusetts, USA, ISBN 0-262-18253-X, 2006.