D4.9 Uncertainty Quantification

HiDALGO2

CENTRE OF EXCELLENCE

Date: February 28, 2025

EuroHPC
Joint Undertaking

Document Identification			
Status	Final	Due Date	28/02/2025
Version	1.0	Submission Date	28/02/2025

Related WP	WP4	Document Reference	D4.9
Related Deliverable(s)	D3.7, D3.8, D4.3, D4.4 and D4.5	Dissemination Level (*)	PU
Lead Participant	USTUTT	Lead Author	Flavio Cesar Cunha Galeazzo (USTUTT)
Contributors	PSNC, SZE, MTG, FAU, UNISTRA, ICCS	Reviewers	Angela Rivera Campos (MTG)
			László Környei (SZE)

Keywords:
uncertainty quantification, simulation, methodology, technology, surrogate tools, environmental simulations

Disclaimer for Deliverables with dissemination level PUBLIC
This document is issued within the frame and for the purpose of the HiDALGO2 project. Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Poland, Germany, Spain, Hungary, France, Greece under grant agreement number: 101093457. This publication expresses the opinions of the authors and not necessarily those of the EuroHPC JU and Associated Countries which are not responsible for any use of the information contained in this publication. This deliverable is subject to final acceptance by the European Commission.
This document and its content are the property of the HiDALGO2 Consortium. The content of all or parts of this document can be used and distributed provided that the HiDALGO2 project and the document are properly referenced.
Each HiDALGO2 Partner may use this document in conformity with the HiDALGO2 Consortium Grant Agreement provisions.
(*) Dissemination levels: **PU**: Public, fully open, e.g. web; **SEN:** Sensitive, restricted under conditions set out in Model Grant Agreement; **EU-C: European Union Classified**, the unauthorised disclosure of this information could harm the essential interests of the Consortium.

Document Information

List of Contributors	
Name	**Partner**
Ayyala Somayajula, Venkata Naga Satya Sai Ravi Kiran	FAU
Michał Kulczewski	PSNC
Bartosz Bosak	PSNC
Flavio C. C. Galeazzo	USTUTT
Davide Padeletti	USTUTT
Zoltán Horváth	SZE
Mátyás Constans	SZE
László Környei	SZE
Angela Rivera Campos	MTG
Luis Torres	MTG
Christophe Prudhomme	UNISTRA

Document History			
Version	**Date**	**Change editors**	**Changes**
0.1	31/10/2024	Flavio Galeazzo (USTUTT)	TOC for approval
0.2	10/11/2024	Marcin Lawenda (PSNC), Harald Köstler (FAU)	ToC, timeline and responsibilities approved
0.3	05/02/2025	All partners	Contents of chapters 1 to 5
0.4	19/02/2025	Flavio Galeazzo (USTUTT), Angela Rivera Campos	Internal review

		(MTG), László Környei (SZE), Marcin Lawenda (PSNC)	
0.5	24/02/2025	All partners	Implementation of internal review
0.6	25/02/2025	Flavio Galeazzo (USTUTT)	Formatting
0.7	25/02/2025	Flavio Galeazzo (USTUTT)	Formatting
0.95	27/02/2025	Rahil Doshi	Quality assurance check
1.0	28/02/2025	Marcin Lawenda	Final check and improvements

Quality Control		
Role	Who (Partner short name)	Approval Date
Deliverable Leader	Flavio Galeazzo (USTUTT)	25/02/2025
Quality Manager	Rahil Doshi (FAU)	27/02/2025
Project Coordinator	Marcin Lawenda (PSNC)	28/02/2025

Table of Contents

List of Figures

List of Acronyms

Abbreviation / acronym	Description
AI	Artificial Intelligence
API	Application Programming Interface
CDO	Climate Data Operators
EC	European Commission
DAS	Deep Active Subspace
Dx.y	Deliverable number y belonging to WP x
GFS	Global Forecast System
IFS	ECMWF's Integrated Forecasting System
LBM	Lattice Boltzmann method
ML	Machine Learning
MTW	Material Transport in Water
mUQSA	multipurpose Uncertainty Quantification and Sensitivity Analysis toolbox
PDE	Partial Differential Equation
POD	Proper Orthogonal Decomposition
UAP	Urban Air Project
UB	Urban Building

UQ	Uncertainty Quantification
RES	Renewable Energy Sources
RL	Reinforced Learning
ROM	Reduced Order Model
SA	Sensitivity Analysis
WF	Wildfires
WP	Work Package

Executive Summary

This document, Deliverable D4.9 "Uncertainty Quantification" outlines the efforts towards uncertainty quantification (UQ) within the HiDALGO2 project.

UQ evaluates how uncertainties in model inputs or parameters affect the outputs simulations or experiments. The different HiDALGO2-relevant UQ methods are discussed in Chapter 2. Typically, an uncertainty quantification study requires a large number of simulations to achieve statistically significant results, which demands a large computational effort. To reduce the computational footprint, the use of surrogates comes into place, which are reduced-order models that can reproduce the behaviour of the full simulation while being computationally less expensive. Surrogates and other UQ tools are described in Chapter 3. An overview of how each pilot use case in HiDALGO2 is incorporating UQ is given in Chapter 4.

Two pilots have partially completed UQ and SA analysis, which provided valuable insights about their simulations and the workflow. The other pilots have identified their major sources of uncertainty and are working on the results.

This document is part of Work Package 4 (WP4) and interact with other tasks as WP4.3 (AI for Global Challenges) and WP3.4 (Ensemble Scenarios). The uncertainty quantification studies in this deliverable use some components already described in other project deliverables (e.g. D3.7 [1], D4.3 [2]).

1 Introduction

Uncertainty quantification (UQ) evaluates how uncertainties in model inputs or parameters affect the outputs of interest in computational and experimental contexts. It encompasses two primary methodologies: forward and inverse uncertainty quantification. In forward uncertainty propagation, input uncertainties are first characterized, then propagated through a computational model, and the resulting outputs are statistically analysed. In contrast, inverse uncertainty quantification begins with uncertain responses, often experimental observations, which are propagated backward through the model to infer the associated input uncertainties. This inverse approach is frequently used for tasks like parameter estimation and model calibration. The focus of this work is on forward uncertainty propagation.

Uncertainty in model inputs can be categorized as aleatory or epistemic. Aleatory uncertainty represents irreducible variability inherent in the system or environment and is typically described using probabilistic methods, such as the statistical distribution of population height. Epistemic uncertainty, on the other hand, arises from a lack of knowledge or incomplete information about inputs and is often expressed using interval-based approaches or subjective probability distributions. While aleatory uncertainty is intrinsic to the system and cannot be reduced, epistemic uncertainty may diminish with improved data or understanding.

An uncertainty quantification study typically involves a large number of simulations to achieve statistically significant results, which can be computationally intensive. To reduce the computational cost, one effective approach is to use surrogates—reduced-order models that replicate the behaviour of the full simulation.

In summary, uncertainty quantification is a powerful tool to study how uncertainties in model inputs affect the results of simulations, however with a large computational footprint.

1.1 Purpose of the document

This Deliverable D4.9 “Uncertainty Quantification” is prepared in the context of WP4 and provides information about the methods and technologies developed for uncertainty quantification. Also, the uncertainty quantification studies identified for the different pilot test cases are described in this document.

1.2 Relation to other project work

The uncertainty quantification studies described in this deliverable use tools and methods developed in tasks T3.3 (Ensemble Scenarios) and T4.3 (Artificial Intelligence for Global Challenges). These tools and methods are part of the deliverable series D3.7 [1], D3.8 (future reference) "Ensemble Scenarios for Global Challenges", and D4.3 [2], D4.4 and D4.5 (future reference) "Advances in HPDA and AI for Global Challenges".

1.3 Structure of the document

This document is structured in 3 major chapters and a conclusion.

The current **Chapter 1** introduces and defines uncertainty quantification and explains why HiDALGO2 simulations need it. **Chapter 2** delineates the methodological principles of uncertainty quantification studies. **Chapter 3** presents the methods and technologies used in the pilot use cases of HiDALGO2 for uncertainty quantification studies. **Chapter 4** discusses the results or the plans of uncertainty quantification studies in each HiDALGO2 pilot use case. Finally, the **Conclusion** summarizes the report.

2 Methodologies for Uncertainty Quantification

This chapter describes the methods used in uncertainty quantification studies in HiDALGO2.

Computer simulations have become indispensable tools in a wide range of fields, including engineering, physics and climate science, fields related to HiDALGO2. By solving complex mathematical models, these simulations provide critical insights into systems that are often too expensive, risky, or impractical to study experimentally. However, the predictive power of any simulation is inherently limited by uncertainties. These uncertainties may have various sources, including imprecise input parameters, numerical approximations, and incomplete understanding of the modelled processes. As simulations play an increasingly vital role in decision-making processes, studying these uncertainties is crucial to ensure the reliability, accuracy, and robustness of their predictions. This is where uncertainty quantification (UQ) fits in.

UQ is the systematic process of identifying, characterizing, propagating, and analysing uncertainties in computational and experimental models. By rigorously addressing these uncertainties, UQ provides a framework for quantifying confidence in simulation results, which is essential for risk assessment, optimization, and design under uncertainty.

For a comprehensive review of UQ methods in engineering, we refer readers to Haldar and Mahadevan [3].

2.1 The Role of Uncertainty Quantification in Computer Simulations

In computer simulations, models aim to approximate reality through a combination of physical laws, empirical data, and engineering assumptions. While these models provide valuable predictions, they are rarely perfect representations of the real world. There are inherent uncertainties in the modelling process, including:

- Input Uncertainties: Parameters such as material properties, boundary conditions, or environmental factors may be imperfectly measured or inherently variable. For instance, wind velocity in atmospheric models is subject to natural variability.
- Model Uncertainties: Simplifications or approximations in the mathematical representation of a system may exclude certain physical processes or interactions.
- Numerical Uncertainties: Discretisation errors, convergence tolerances, and other numerical artefacts can affect simulation outputs.

- Experimental Uncertainties: Experimental data used for calibration or validation often carry their own uncertainties.

Without addressing these uncertainties, the outputs of simulations may appear precise but could be misleading or unreliable. Uncertainty quantification provides tools to understand the impact of input uncertainties on simulation outputs. This enables decision-makers to quantify risks, optimize designs, and prioritize areas where additional data or modelling effort is most needed.

2.2 Methodologies in Uncertainty Quantification

UQ methods can be broadly classified into two main approaches: forward uncertainty quantification and inverse uncertainty quantification, where each serves distinct purposes and involves different techniques. In HiDALGO2 we are focused on forward UQ.

2.2.1. Forward Uncertainty Quantification

In forward UQ, the primary goal is to propagate uncertainties in model inputs through the simulation to assess their impact on the outputs. The process typically begins by characterizing input uncertainties using either probabilistic methods, where probability density functions derived from the data are employed, or non-probabilistic techniques such as intervals or fuzzy sets.

The uncertainties are then propagated through the computational model using various methods. Monte Carlo Simulations use random sampling for reiterated model evaluations, offering high accuracy at an also high computational cost. Alternatively, Stochastic Collocation strategically selects sample points to approximate output uncertainty with less computational demand, while Polynomial Chaos Expansion represents outputs as a series of orthogonal polynomials. Finally, the resulting outputs are statistically analysed to compute metrics such as mean, variance, and confidence intervals, providing a comprehensive summary of the system's behaviour under uncertainty.

Forward UQ is particularly useful for sensitivity analysis, where the aim is to identify which input uncertainties have the most significant impact on the outputs. It also supports robust design optimization by ensuring that designs perform reliably across a range of uncertainties.

2.2.2. Inverse Uncertainty Quantification

Inverse UQ method, in contrast, focuses on deducing input uncertainties from observed data. This is often used for parameter estimation, model calibration, and data assimilation. The process involves defining observations (typically experimental or field data) and prior knowledge of the model parameters. Then the Inverse Problem is solved, in order to determine the input parameter distributions that best explain the observed data. Bayesian inference is a common approach, where a prior distribution is updated using likelihood functions to obtain a posterior distribution of the inputs. Finally, the uncertainty in inputs is quantified, as the posterior distribution provides insight into the uncertainty associated with the inferred inputs. Both data noise and modelling assumptions should be accounted for.

Inverse UQ is especially suited for calibrating complex models, especially when direct measurements of input parameters are unavailable or infeasible. However, inverse problems are often ill-posed, meaning they may have non-unique or unstable solutions. Advanced techniques such as regularization and Bayesian frameworks help address these challenges.

2.3 Classifications of Uncertainty

A key aspect of UQ is understanding the nature of uncertainties, which can be broadly divided into two categories:

- **Aleatory Uncertainty**, also known as irreducible uncertainty, arises from inherent variability in the system or environment. For example, the random fluctuations in wind velocity or the natural variability in material properties represent aleatory uncertainty. This type of uncertainty is considered intrinsic to the system and cannot be reduced, though its effects can be modelled probabilistically.
- **Epistemic Uncertainty**, or reducible uncertainty, originates from a lack of knowledge or incomplete information about the system. Examples include insufficient data on material behaviour or uncertainty about the validity of a model under new conditions. Unlike aleatory uncertainty, epistemic uncertainty can be reduced through better data acquisition, refined models, or additional experiments.

In practice, real-world problems often involve a combination of aleatory and epistemic uncertainties. Separating and addressing these uncertainties is critical to ensure robust and interpretable UQ results.

2.4 Challenges in Uncertainty Quantification

UQ faces several challenges in its implementation in computational simulations. One major hurdle is **computational cost**, as many UQ methods, such as Monte Carlo simulations, require a large number of model evaluations, which becomes especially difficult for high-fidelity simulations with long runtimes. This issue is exacerbated in high-dimensional input spaces, where the computational demand increases exponentially with the number of uncertain parameters, which is often referred to as “the curse of dimensionality”. Additionally, the complexity of models, such as those involving non-linear, stochastic, or multi-physics phenomena (and which are focus of HiDALGO2), introduces further difficulties, as these systems may demand advanced techniques for both uncertainty propagation and inverse analysis. **Scarcity of data** compounds these challenges, particularly for epistemic uncertainties, where limited information makes it hard to construct reliable probability distributions or intervals for input parameters. Finally, achieving **validation and verification** is critical for ensuring trustworthy UQ results, requiring rigorous comparison with experimental data and careful verification of the numerical methods employed. Addressing these challenges is essential for the applicability of UQ in practical scenarios.

2.5 Advances in Uncertainty Quantification

To address these challenges, recent advancements in UQ have focused on enhancing computational efficiency and improving the integration of UQ with machine learning and optimization. The use of **surrogate models** (reduced-order models) significantly reduces the computational cost, while approximating the behaviour of high-fidelity simulations. Examples include Gaussian process models, neural networks, machine learning (ML) models and polynomial chaos expansions. Also, the use of **adaptive sampling** techniques such as adaptive Monte Carlo or sparse grid methods, that dynamically allocate computational resources to regions of the input space that contribute most to uncertainty, can save computational resources. Also, the use of **parallel computing** makes it possible to perform large-scale UQ studies more efficiently and in a timely manner.

2.6 Summary

Uncertainty Quantification enables simulations to account for and communicate their inherent uncertainties. By characterizing, propagating, and analysing uncertainties, UQ enhances the reliability and interpretability of simulation results. Despite the challenges of computational cost, high-dimensional input spaces, and data scarcity, advancements in surrogate modelling, adaptive sampling, and high-performance computing are pushing the boundaries of the use of UQ in simulation workflows as the ones in HiDALGO2.

3 Technologies for Uncertainty Quantification

This chapter describes the various tools and technologies used in the UQ studies in HiDALGO2.

3.1 Uncertainty Quantification Tools

3.1.1 Urban Air Project (UAP)

In the Urban Air Project (UAP) pilot, the main UQ tools are natively implemented into the solutions. To reduce the code dependencies and improve the performance of the overall solution, UQ methods are constructed and state-of-the-art solutions (of Python, or R-packages) are (re-)implemented in the UAP software.
One of the CFD solvers used in UAP, RedSim (Reduced Simulations), has inherently the possibility to run reduced order models (ROMs). These ROMs are considered surrogate models for the forward UQ. In RedSim the Proper Orthogonal Decomposition (POD) for time-stepping has been implemented. The POD module of RedSim uses offline the OpenMPI + CUDA implementation of RedSim to collect the snapshot matrix *X* from the simulation results for several parameters, then still offline the singular value decomposition (SVD) of *X*, and construct a *U* projection matrix from the left singular vectors corresponding to the largest singular values and store *U* for the online phase. In the online phase of the POD module of RedSim the time-stepping is computed in the projected space of the image set of *U*.
Concerning ensemble generation, the MathSO portal has an ensemble generating feature with support for parameter sweeping for selected variables for instances that are implemented in the MathSO portal. This functionality applies for the UQ of variables with a small number of degrees of freedom (e.g. when the inlet wind is characterized by some wind direction and speed values only).
For the evaluation of ensemble run results for the sake of UQ, some scoring methods are used, in particular the CRPS (Continuous Ranked Probability Score, see [17],[18]). The CRPS is a widely used metric for evaluating probabilistic forecasts. It measures the difference between the forecast Cumulative Distribution Function (CDF) and the actual outcome—typically represented as a step function—and is defined as the integrated squared error over all possible values. For a discrete forecast distribution, the CRPS can be computed using the formula

$$\mathrm{CRPS}(F, y) = \sum_{i=1}^{N} p_i |x_i - y| - \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} p_i p_j |x_i - x_j|,$$

where $x_1, x_2, \dots, x_N$ are the possible forecast outcomes, $p_1, p_2, \dots, p_N$ are their probabilities, and y is the observed value.

For CRPS generalizes the concept of absolute error from deterministic forecasts to probabilistic ones, with lower scores indicating a forecast that more closely matches the observed outcome.

3.1.2 Urban Building (UB)

In the Urban Building (UB) pilot, we decided to integrate Uncertainty Quantification (UQ) directly into the simulation workflow rather than treating it as an external, post-processing step. This decision is motivated by the large number of uncertain parameters, arising from design assumptions, physical model parameters, construction workmanship, occupant behaviour, and environmental conditions, that can significantly affect simulation results. We leverage two primary tools to manage these uncertainties:

1. OpenTURNS

OpenTURNS (Open-source Treatment of Uncertainties, Risks, and Statistics, see Deliverable 3.7 [1] and [11]) is an open-source C++/Python library providing a rich set of functionalities for UQ, sensitivity analysis, and probabilistic modelling.

- Forward Propagation: We use OpenTURNS to define our uncertain inputs (e.g., parameter distributions, correlations) and to run sampling-based ensemble simulations (Monte Carlo, quasi-random, etc.).
- Sensitivity Analysis: Built-in methods (e.g., Sobol indices, FAST) help us identify which parameters dominate the variance in simulation outputs.
- Python Wrappers: Although implemented in C++, OpenTURNS has a Python API that simplifies integration into existing workflows and HPC environments.

2. Feel++

Feel++ (Finite Element Embedded Language in C++) [12] is our in-house high-performance framework for simulations based on Partial Differential Equations (PDEs) (see Deliverable 3.7 [1]). It features:

- Finite Element Methods: A domain-specific embedded language closely mirroring the mathematical formulation of PDEs.
- Parallel Execution: MPI-based parallelism to manage ensemble runs (i.e., repeated simulations with varied input parameters) for large-scale simulations.
- Planned UQ Integration: While we already perform ensemble runs in Feel++, we intend to extend the framework with native UQ functionalities—such as sampling and statistical post-processing—that integrate seamlessly with OpenTURNS or other libraries.

By combining OpenTURNS (for defining uncertainties and processing outputs) with Feel++ (for large-scale PDE simulations), we aim to streamline the entire UQ workflow in a single, HPC-enabled environment. A more detailed description of these tools and their planned integration can be found in Deliverable 3.7 [1].

3.1.3 Renewable Energy Sources (RES)

The mUQSA platform [10] is developed by the Poznan Supercomputing and Networking Center with the main aim to comprehensively support Uncertainty Quantification and Sensitivity Analysis of computational models. The mUQSA provides a user-friendly graphical interface for both scenario definition and results analysis, as well as the full automation of uncertainty quantification and sensitivity analysis execution on HPC resources. This tool is described in detail in Deliverable 3.7 [1].

3.1.4 Wildfires (WF)

The WF pilot use case is under integration with the built-in QCC-PilotJob mechanism to launch simulations associated with possible weather conditions as well as different ignition points compatible with the observations, in order to estimate the uncertainty in the evolution of the fire spread and the smoke dispersion. New statistical functions to analyse the spatial uncertainty are foreseen with the help of CDO. CDO (Climate Data Operators) is a command line suite or library for manipulating and analysing climate data from the Max Plank Institute. More details about the use of QCC-PilotJob in the WF pilot are found in Deliverable 3.7 [1].

3.1.5 Material Transport in Water (MTW)

The waLBerla framework [13] is an in-house developed software, based on lattice Boltzmann methods (LBM), which is designed for large-scale flow simulations on EuroHPC clusters. It enables multiple simulations by varying particle shapes and input parameters to analyse their impact on flow properties. The simulation results are used to evaluate variations in velocity, temperature, and heat flux. The simulation outputs are also used to train surrogate models which aid in faster flow predictions. For surrogate modelling, we plan to explore Deep Ensembles [14] and Monte Carlo Dropout [15] (MC Dropout). Deep Ensembles improve uncertainty estimation by training multiple independent neural networks and aggregating their predictions. This method effectively captures epistemic uncertainty through model disagreement and can also address aleatoric uncertainty when explicitly modelled. Monte Carlo Dropout (MC Dropout) approximates Bayesian inference by keeping dropout active during inference and performing multiple forward passes. It efficiently captures epistemic uncertainty without requiring multiple trained models.

3.2 Surrogate Tools

3.2.1 Relexi

One of the possible solutions to the very high computational demand of a UQ study is the use of surrogates. Machine learning (ML) models are one alternative, where the user can substitute one part of the workflow or the whole simulation by a neural network.

Among other initiatives to introduce AI capabilities in HiDALGO2 [2], an interface to the CFD OpenFOAM[1] to the tool Relexi [4][1][5][6][7] is being developed. Relexi is a Reinforcement Learning (RL) framework, with RL meaning an intelligent agent that uses ML and optimal control strategies to maximize a reward signal. The framework currently supports the high-order HPC flow solver FLEXI [8]. Relexi is built upon TensorFlow and its RL extension TF-Agents, while the communication and data handling is done using SmartSim package with its SmartRedis communication client [9].

SmartSim is made up of two parts: (i) the SmartSim workflow library, that launches the ML infrastructure using e.g. PyTorch or TensorFlow on HPC systems alongside the user workloads and (ii) SmartRedis, which connects the HPC workloads to the ML infrastructure at runtime. SmartRedis supports applications written in Fortran, C, C++ and Python, and sends the data to a remote SmartSim infrastructure to execute the ML models and scripts on GPU or CPU.

OpenFOAM is currently used by the HiDALGO2 pilots UAP and WF. If the results of the coupling and the use of the ML models are positive with OpenFOAM, other solvers could also be integrated into Relexi.

3.2.2 EasySurrogate

Yet another possible tool to provide a surrogate for the part of a workflow or complete simulation is EasySurrogate[2]. EasySurrogate is a toolkit dedicated to simplifying the construction of surrogate models – swift approximations of complex (multiscale) simulations. A key feature is its inclusion of Deep Active Subspace (DAS) surrogates, which leverage neural networks alongside active subspace principles to effectively

[1] https://www.openfoam.com/
[2] https://github.com/wedeling/EasySurrogate

reduce dimensionality. Furthermore, EasySurrogate incorporates quantized softmax networks – a neural-network-based bootstrapping technique that resamples observed data. This method is particularly beneficial for creating surrogate models of subgrid-scale elements, commonly encountered in simulations of turbulent flows

EasySurrogate is coupled with EasyVVUQ (part of the mUQSA toolkit) to train a model based on UQ ensembles. The feasibility of this use is analysed by a pilot RES use case.

4 Uncertainty Quantification for Global Challenges

This chapter presents how individual HiDALGO2 use cases incorporate uncertainty quantification into their workflow.

4.1 Urban Air Project (UAP)

The Urban Air Project (UAP) pilot simulate the wind flow and gaseous air pollutants' concentrations in cities (5–8-kilometre horizontal length and 1-2 kilometre vertical domains) under high resolution (1-2-meter spatial resolution at ground level and unsteady in time for 1 day to 1 year) and then computes society-relevant indicators like wind-comfort measures, annual exceedance numbers of air-quality limit values, urban planning-related indicators and more. The UAP simulations use data from global and regional weather and air quality forecasts, local weather and air quality measurements, urban geometry 3D models, and emission sources. The input data in UAP has various degrees of freedom, depending on the modelling level: in a coarse model the weather data used for boundary conditions are modelled by some real numbers (inlet wind speed, wind direction) while for a detailed model by a scalar field with many or large number of variables (e.g. weather data of a regional weather model, or the initial state on a large-scale mesh).

These data sources bring uncertainties into each component of the UAP simulation: epistemic uncertainties for the initial and boundary conditions from the global and regional weather simulations, aleatoric ones from the geometry models (e.g. position of buildings, physical parameters of the pollutants).

To quantify the uncertainties of the UAP variables, currently the UAP uses the classical methodology of UQ, namely:

- create ensembles for the parameter or variable under UQ,
- execute the simulation with a surrogate model for the ensemble individuals, and then
- evaluate the ensemble run by using a suitable scoring.

Concerning the ensemble generation in the cases of variables with small degree of freedom, like the coarse model of the wind boundary condition, a sampling with Sobol or other Monte-Carlo based methods is used. For a large-scale variable (e.g. initial state of the state variables in a large-scale simulation) a finite number of perturbations of the variable are used. In the latter case either the perturbations are given with the external data (e.g. all 50 perturbations with the ECMWF model values) or UAP generate itself the perturbations. In the latter case, the methodology of generating singular vectors of a linearized forecast propagator shall be implemented (for the detailed methodology in the case of operational weather forecasts of ECMWF, see **Error! Bookmark not defined.**[16]; for the UQ of UAP initial state an analogue of this approach will be constructed and implemented during the next phase of HiDALGO2).

For a surrogate model, the POD and its variants are used, see the description in Section 3.2. In HiDALGO2, a novel interpolation technique and a deep-learning model are currently being developed based on the POD simulations, to address the curse of dimensionality in the linear algebra computations of the POD.

As a preliminary result of the UQ for the air quality computation of UAP, see Figure 1. This result is taken by the OpenFOAM computations of the main value on a fine mesh for the city of Győr; in the ensemble 50 perturbations of the simulated boundary condition stand. The corresponding CRPS value was 13.4 for a variable of size 10^2, which shows that the result is moderate and both the simulation and the UQ methodology have to be improved.

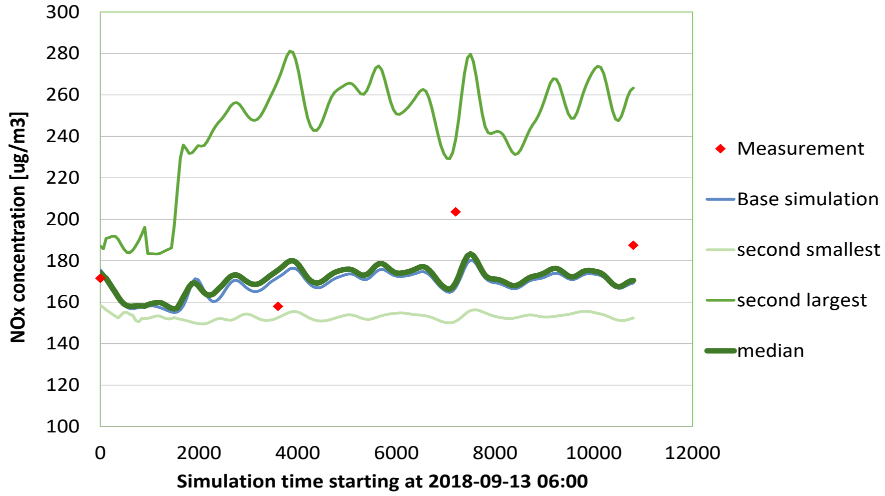

Figure 1. Pollutant concentration values according to the base and some perturbed simulations for UAP in a case study for the city of Győr.

4.2 Urban Building (UB)

The primary objective of the UB pilot is to simulate energy consumption and thermal comfort in urban buildings. The complexity of physical processes, the variability of occupant behaviour, and the uncertainties in building envelope properties make UQ a critical component for reliable simulation results and robust decision-making. Studies have shown that occupant behaviour often has a larger impact on energy calculations than local variability in weather data, emphasizing the need to handle these uncertainties carefully.

Although a full UQ study is ongoing (with preliminary steps already in progress), the main categories of uncertain parameters we have identified include:

- Aleatoric Uncertainties

These uncertainties have known or assumed probability distributions:

  – Weather Conditions: Outdoor temperature, solar irradiance, and wind speed are derived from historical data or typical meteorological year datasets. Variations can also come from synthetic weather data.

  – Occupant Schedules: Randomness in presence or usage patterns (e.g., working hours, weekend behaviour) introduces stochastic variation in building energy demand.

- Epistemic Uncertainties

These stem from a lack of knowledge or incomplete information regarding model parameters:

  – Material Properties: Thermal conductivity, heat capacity, or infiltration rates can deviate significantly from nominal "catalogue" values due to workmanship issues and real-world construction practices.

  – Model Calibration Parameters: Empirical coefficients in heat transfer or occupant behaviour models may be based on simplified assumptions and thus have limited accuracy.

4.2.1 Planned Approach

We employ Feel++ to conduct ensemble runs over distributions of the above uncertain inputs. Each simulation is assigned to an MPI group, enabling parallel execution without external schedulers. This strategy is essential for large-scale, city-wide energy analyses.

With OpenTURNS, we manage sampling (Monte Carlo, quasi-random, etc.) and perform statistical post-processing to compute quantities of interest (e.g., means, confidence intervals, Sobol indices). This workflow ensures a robust treatment of uncertainties and clarifies which parameters most affect energy consumption and thermal comfort.

The ongoing work include:

- Data Collection: We continue to gather more data on materials, occupant behaviour, and building usage to refine or reduce epistemic uncertainties.
- Calibration and Validation: We will calibrate our simulation models against real-world data from actual buildings, thus reducing knowledge gaps and increasing model fidelity.

- Integrated Tooling: Plans include extending Feel++ for direct UQ capabilities—such as advanced sampling algorithms and automated post-processing—thus minimizing the need for external scripts. This tight integration will streamline HPC workflows and enable near-real-time feedback loops in the future.

4.2.2 Broader Context and Additional Resources

Accurate UQ at both building and city scales underpins effective policy-making, retrofit decisions, and design strategies. Bayesian approaches are increasingly being explored, for example, in Westermann & Evins [19], where Bayesian Deep Learning offers uncertainty-aware surrogate models for building energy simulations.

As we progress with these UQ activities, we will document findings and methodological enhancements in subsequent deliverables. Quantifying uncertainties allows us to provide more reliable and actionable predictions of energy performance, directly informing sustainable and efficient urban development and future coupling with Urban Air Pilot or Renewable Energy Sources pilots.

4.3 Renewable Energy Sources (RES)

RES pilot activities concentrate on forward uncertainty quantification, i.e. estimating the impact of input uncertainties on the simulation results and relationship between parameters (sensitivity analysis). The inverse analysis is particularly important to verify and validate the model, though it is not considered at this moment.

UQ can also deliver more application-oriented results. An example is the RES-damages module, which estimates the potential impact of extreme weather events on overhead electrical networks. A case study was conducted in a Polish city, focusing on excessive wind speeds, and the detailed results are available in [10]. The workflow employed three nested domains, with the outermost domain covering the entire Poland at a resolution of 3.6 km, middle domain covering the voivodeship (province) with 800m resolution, the innermost domain discretised at a detailed 100-meter spacing covering the whole city, and the 33 eta vertical levels (vertical resolution between surface and top of the model domain). A custom overhead electrical network was created for this study, based on the real electrical network coverage. This high level of detail enabled the simulation to generate meteorological results for each individual component of the infrastructure. Five distinct electrical lines located in the outskirts and city centre were analysed. Ensembles were generated to simulate various wind conditions, including direction and speed by the means of the mUQSA toolkit. The wind sector was divided into eight directions: N, NE, E, SE, S, SW, W, NW, while the wind speeds were assigned uniformly from 1 to 20 m/s. The initial conditions were taken from the Global

Forecast System (GFS) forecast data, and then only wind speed and direction were updated with the generated values.

Figure 2 illustrates the wind speeds at which electrical lines are impacted under various conditions. The x-axis represents the five electrical lines studied, while the y-axis indicates wind speed in m/s. The mean wind speed is relatively low, especially for two sites, indicating their lower vulnerability to wind conditions. However, the remaining three sites are more susceptible to the impact of wind speed and direction. Only one site shows minimal impact, while the others are more vulnerable. Electrical line number 5 is particularly endangered due to the surrounding terrain, which amplifies the wind speed beyond the initial boundary conditions.

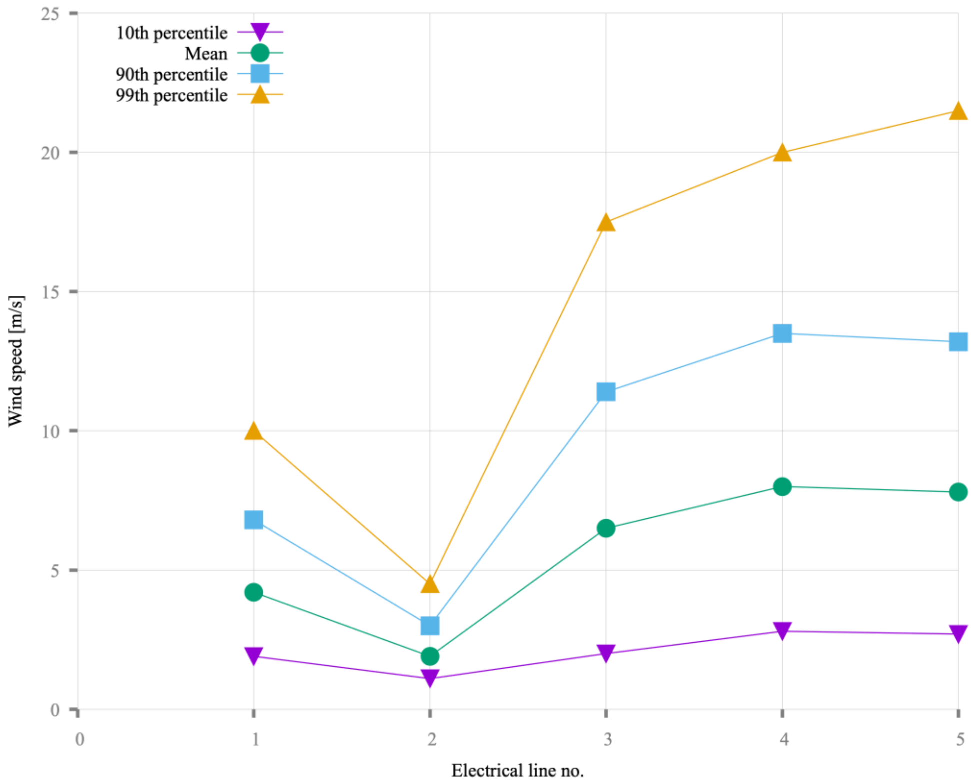

Figure 2. RES: Wind speed at different sites

Sensitivity analysis (SA) investigates how analysed uncertainty parameters influence simulation results. In RES-damages, SA is used to determine whether wind direction or wind speed has a greater impact on wind speed at different sites. Figure 3 illustrates the influence of wind speed and direction.

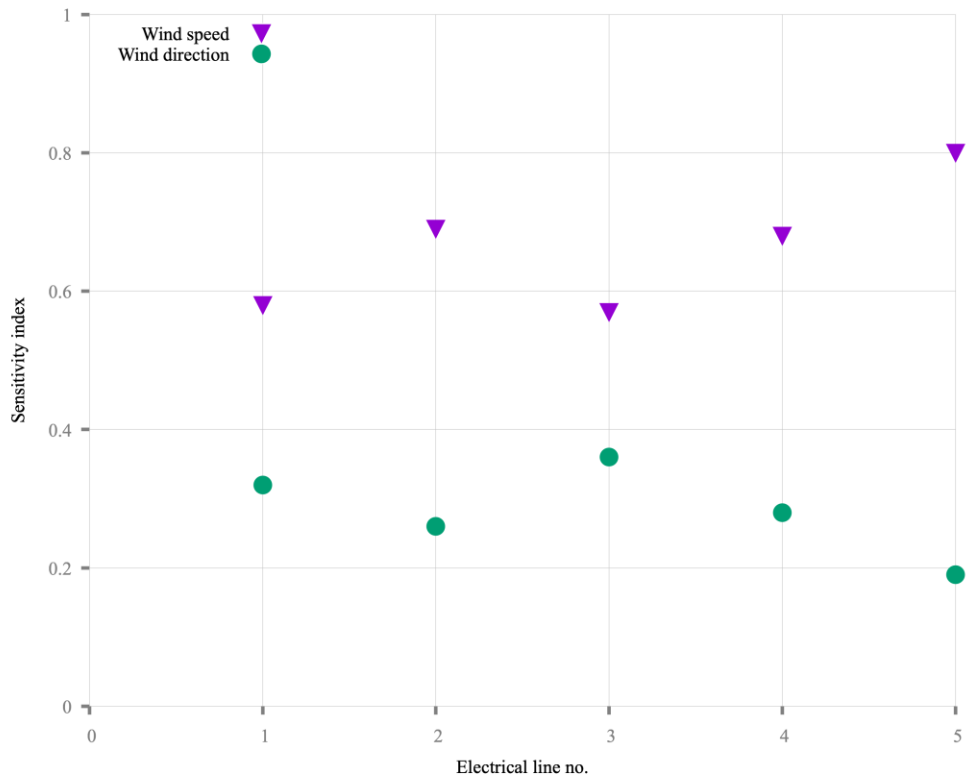

Figure 3. RES: Impact of wind speed and direction at different sites

Overall, wind speed emerges as the most significant factor for all five sites. However, the analysis reveals variations in influence. At some sites, such as site number 3, wind direction is nearly as influential as wind speed, while at others, like site number 5, wind direction has minimal impact on the probability of damage. As expected, higher wind speeds increase the vulnerability of electrical lines to damage, regardless of wind direction. With a certain margin of error, further analysis could potentially exclude wind direction from the uncertainties, as it has a moderate to minor impact on the final wind speed at different sites.

The UQ and SA can support renewable energy source owners and electrical network operators in preparing for extreme weather events. By identifying the most vulnerable sites, particularly those susceptible to excessive wind speeds, these tools provide valuable insights.

4.4 Wildfires (WF)

Within the WF Pilot several sources of uncertainty are considered:

- weather forecast uncertainty,
- vegetation and field data structure,
- fire spread models: aleatory and epistemic uncertainty.

Uncertainty associated with weather conditions

This uncertainty has several possible sources:

- Reliability of the meteorological boundary conditions. Whether using observed data or predictions from a meteorological model, there is a random uncertainty in the wind, temperature and relative humidity conditions used associated with errors in the measuring instruments, in the spatialisation of the data, and in the prediction of their future evolution.
- As the prediction time increases (i.e., predictions further into the future), the reliability of the results decreases. For this reason, we use ensembles of predictions with several models to identify trends and probabilities.
- High-precision prediction models based on downscaling processes lose reliability as more nested domains are coupled to achieve high resolutions. This is a specific issue in high-resolution predictions with mesoscale models like WRF, which we use in the HiDALGO2 wildfire pilot.

Uncertainty associated with data and data structure

Vegetation structure is another important factor contributing to uncertainty. For the representation of vegetation, catalogues of forest fuel models are used to represent approximately the average conditions of fire behaviour in that type of vegetation. This form of representation does not take into account the inter-annual differences in the same vegetation mass due to the different meteorology that has influenced its growth. In addition, some characteristics are determined by parameters that condition the transition between different types of fire, for example in wooded areas the transition process from surface fire to crown fire is simulated by the height of the base of the crown, which in reality varies from one tree to another.

Vegetation moisture also contributes to uncertainty. In practice, this is assumed to be uniform over relatively large areas, but in reality, this is much more complex, as there may be areas within the canopy where moisture is locally higher or lower, especially in critical locations. In addition, categorical deterministic rules are applied to estimate moisture content instead of probability functions, which adds uncertainty to the simulations.

Uncertainty associated with fire spread models.

As explained in D3.7 Ensemble scenarios for global challenges, spread models are another important factor of aleatory and epistemic uncertainty due to both the modelled transmission method and the spatial resolution used, that can be partially avoided using data assimilation techniques.

There are also phenomena that are poorly simulated by the models and whose simulation is still in the research phase, as explained in D3.7, such as flying embers, fire whirls, eruptive fire behaviour, and pyrocumulus clouds. These phenomena, their appearance, and development are difficult to predict and even harder to locate geographically, making them the most significant source of uncertainty for firefighting operations.

To address the analysis of these uncertainties, we foresee two types of tests to be performed:

- Ensemble of scenarios to address the quantification of uncertainties associated with the operational simulation of forest fires.
- Uncertainties associated with topographic conditions, fuel type and humidity, generation of secondary fires associated with sparks or crown fires.

4.5 Material Transport in Water (MTW)

In the current pilot, UQ plays a crucial role in enhancing the robustness of both numerical and surrogate models. Accurately assessing uncertainties ensures that model predictions remain reliable and applicable across varying conditions.

For this pilot, there are two primary approaches to quantifying uncertainties. The first approach involves numerical simulations where variations in particle shapes are used to estimate flow fields, velocity, and temperature distributions. This requires running multiple simulations on EuroHPC clusters, allowing for a detailed analysis of how particle shape influences key output parameters such as velocity, temperature, and heat flux distribution. Understanding the sensitivity of these flow properties to geometric variations is essential for optimizing system performance and ensuring predictive accuracy in real-world scenarios.

The second approach focuses on surrogate modelling for flow prediction. UQ is an integral part of developing surrogate models, ensuring that they can efficiently predict flow properties while accounting for various sources of uncertainty. These uncertainties arise from both aleatoric and epistemic sources.

Aleatoric uncertainties include measurement noise in training data, such as errors in velocity, pressure, and temperature measurements, as well as inherent variability in physical systems, such as turbulence and unsteady flow behaviour. These uncertainties follow probabilistic distributions and cannot be eliminated but can be modelled to improve prediction reliability.

Epistemic uncertainties, on the other hand, arise from limitations in deep learning architectures, such as incomplete representations of the underlying physics. Additionally, data scarcity can lead to incomplete training datasets, missing flow configurations, and unaccounted boundary conditions, further impacting model reliability. These uncertainties can potentially be reduced by refining models or incorporating additional data sources.

The process of UQ in surrogate modelling follows several critical steps. First, uncertainties in input parameters, such as inlet velocity, pressure, and boundary conditions, are identified and characterized. Inlet velocity distributions represent aleatoric uncertainty, while boundary conditions and model parameters often introduce epistemic uncertainty due to incomplete knowledge. Then, surrogate models, including U-Net [20] and Fourier Neural Operators (FNOs) [21], are trained on diverse datasets to capture a wide range of flow scenarios. Once trained, input uncertainties are propagated through the model to assess their impact on predicted outputs, with probabilistic distributions or confidence intervals quantifying the resulting uncertainties. Finally, validation and calibration are performed by comparing surrogate model predictions with high-fidelity numerical simulations or experimental data, ensuring the robustness and reliability of the model.

By systematically integrating UQ into surrogate modelling, this approach ensures that predictions remain robust and actionable, even under uncertain conditions. This framework not only enhances the applicability of surrogate models in complex flow simulations but also improves their generalizability across different operational scenarios.

5 Conclusions

This deliverable defines Uncertainty Quantification in HiDALGO2 and highlights the importance of UQ for the pilot use cases. By incorporating UQ methods, HiDALGO2 analyses and enhances the quality of the complex simulation models used by the pilots. The methodologies and tools outlined in this document show the project's commitment to advancing UQ studies.

The deliverable also illustrates how UQ is integrated into HiDALGO2's various pilot use cases, and their use in analysing and improving the quality of the predictions. The Urban Air Project (UAP) and Renewable Energy Sources (RES) pilots have partially completed UQ and SA analysis. For example, the UQ for the RES pilot provided valuable insights about the most vulnerable sites of the electrical network to excessive wind speeds. The other pilots have already identified their major sources of uncertainty and are working on the results.

This work will be further developed and expanded upon in the upcoming report, “D4.10 Uncertainty Quantification,” scheduled for release in M47.

References

[1] Flavio C. C. Galeazzo et al., "HiDALGO2 D3.7 Ensemble Scenarios for Global Challenges", http://dx.doi.org/10.13140/RG.2.2.28421.15841

[2] Giorgos Stamou et al., "HiDALGO2 D4.3 Advances in HPDA and AI for Global Challenges", http://dx.doi.org/10.13140/RG.2.2.10065.95848

[3] Haldar A and Mahadevan S, Probability, Reliability, and Statistical Methods in Engineering Design, Wiley, New York, 2000

[4] Kurz, M., Offenhäuser, P., Viola, D., Resch, M., & Beck, A. (2022). Relexi—A scalable open source reinforcement learning framework for high-performance computing. Software Impacts, 14, 100422.

[5] Kurz, M., Offenhäuser, P., Viola, D., Shcherbakov, O., Resch, M., & Beck, A. (2022). Deep Reinforcement Learning for Computational Fluid Dynamics on HPC Systems. Journal of Computational Science, 65, 101884.

[6] Kurz, M., Offenhäuser, P., & Beck, A. (2023). Deep reinforcement learning for turbulence modeling in large eddy simulations. International Journal of Heat and Fluid Flow, 99, 109094.

[7] Beck, A., & Kurz, M. (2023). Toward discretization-consistent closure schemes for large eddy simulation using reinforcement learning. Physics of Fluids, 35(12), 125122.

[8] Krais, N., Beck, A., Bolemann, T., Frank, H., Flad, D., Gassner, G., Hindenlang, F., Hoffmann, M., Kuhn, T., Sonntag, M., & Munz, C.-D. (2021). FLEXI: A high order discontinuous Galerkin framework for hyperbolic–parabolic conservation laws, Computers & Mathematics with Applications, 81, 186-219.

[9] Partee et al., "Using Machine Learning at Scale in HPC Simulations with SmartSim: An Application to Ocean Climate Modeling", Journal of Computational Science, Volume 62, 2022, 101707, ISSN 1877-7503

[10] Kulczewski et al., "Fostering uncertainty quantification in Global Challenges with mUQSA toolkit", Parallel Processing and Applied Mathematics. PPAM 2024. Lecture Notes in Computer Science (in publish)

[11] Baudin, M., Dutfoy, A., Iooss, B., & Popelin, A. (2016). OpenTURNS: An Industrial Software for Uncertainty Quantification in Simulation. Handbook of Uncertainty Quantification. https://openturns.github.io/

[12] Prud'homme, C., Chabannes, V., Saigre, T., Trophime, C., Berti, L., Samaké, A., Van Landeghem, C., Szopos, M., Giraldi, L., Bertoluzza, S., & Maday, Y. (2024). feelpp/feelpp: Feel++ Release V111 preview.10 (v0.111.0-preview.10). Zenodo. https://doi.org/10.5281/zenodo.12742155

[13] Chair for System Simulation, “waLBerla (widely applicable Lattice Boltzmann from Erlangen)”. Zenodo, Oct. 30, 2023. doi: 10.5281/zenodo.10054460.

[14] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in Advances in Neural Information Processing Systems (NeurIPS), vol. 30, pp. 6402–6413, 2017.

[15] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in Proceedings of the 33rd International Conference on Machine Learning (ICML), vol. 48, pp. 1050–1059, 2016.

[16] ECMWF, IFS Documentation CY49R1 - Part V: Ensemble Prediction System. DOI: 10.21957/956d60ad81

[17] Jordan, A., Krüger, F., & Lerch, S. (2019). Evaluating Probabilistic Forecasts with scoringRules. Journal of Statistical Software, 90(12), 1–37. https://doi.org/10.18637/jss.v090.i12

[18] CRAN: scoringRules: Scoring Rules for Parametric and Simulated Distribution Forecasts. https://cran.r-project.org/web/packages/scoringRules/ (Last retrieve: 15.02.2025.)

[19] Westermann, P., & Evins, R. (2020). Using Bayesian deep learning approaches for uncertainty-aware building energy surrogate models. arXiv preprint, https://arxiv.org/abs/2010.03029

[20] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Cham, Switzerland: Springer International Publishing, 2015, pp. 234–241.

[21] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, A. Stuart, and A. Anandkumar, “Fourier neural operator for parametric partial differential equations,” arXiv preprint arXiv:2010.08895, 2020.