AeroFusion: Data Fusion and Uncertainty Quantification for Entry Vehicles

Steven Snyder*, T. J. Wignall†, Justin S. Green‡, Satvik G. Kumar§, Michael W. Lee¶,
Tenavi Nakamura-Zimmerer‖, James B. Scoggins**, R. Anthony Williams††
NASA Langley Research Center, Hampton, VA, 23681

AeroFusion is a NASA Langley initiative to incorporate advances in data science into the aerodynamic modeling process to improve efficiency. The effort can largely be categorized in three components: reduced-order modeling techniques, surrogate modeling techniques, and uncertainty quantification. By combining various methods from these categories, AeroFusion aims to reduce the development cost of aerodynamic models, both in terms of time and money.

I. Introduction

MARS Lander Aerodynamic Model Data Fusion Using Machine Learning with Embedded Uncertainty Quantification, or simply AeroFusion, is an Early Career Initiative (ECI) project funded by the Space Technology Mission Directorate (STMD) at NASA. The goal of AeroFusion is to leverage advances in data science to improve the aerodynamic database generation process at NASA. Specifically, we employ reduced-order methods and surrogate modeling techniques to quantify uncertainty more rigorously throughout the entire database creation process. The intended outcome is a high-level Python software package and workflow for modelers to use without having to implement the procedures from lower-level packages.

The application area of focus is currently entry, descent, and landing vehicles; specifically targeting use on novel configurations that will enable landing payloads of 20 metric tons, including humans, on Mars. For entry vehicles, the primary purpose of the aerodynamic models is to provide data to the trajectory simulations used to develop the guidance, navigation, and control systems for the vehicle and provide targeting and landing ellipse prediction during flight operations. In this paper, we will examine the AeroFusion developments on an analytically defined, axisymmetric heat shield for the Orion vehicle. We use this configuration in our case study because data collection for this vehicle has been ongoing for more than a decade, there is an existing baseline aerodynamic model for comparison, and the results for this particular geometry have been cleared for uncontrolled export. We note that aerodynamic models are important for the assessment of control and stability objectives for aircraft as well, so improvements to aerodynamic modeling processes have wider benefits for NASA than just the entry vehicles on which we currently focus. Improvements to the aerodynamic modeling process could also help enable methods for certification by analysis [1].

The current paradigm for the aerodynamic modeling process of the Orion vehicle [2, 3] involves extensive testing, both computational and experimental, often in the form of one-factor-at-a-time sweeps across large portions of the state space. Expert judgment is often relied upon to fuse different data sources, such as computational and experimental tests. For example, the procedure described in Section III of [3] consists of ad hoc methods of pruning and modifying data before applying smoothing splines. Computational fluid dynamics (CFD) data was discarded in the trim area (expected operating region) in order to force the dominance of the wind tunnel data, which was considered to be of higher fidelity. To extend past the region where wind tunnel data was captured, the CFD data was further modified to preserve the shape or trend of the data whilst anchoring it to the end point of the wind tunnel data.

In contrast, AeroFusion seeks to provide readily accessible data-driven tools for data and model fusion. Throughout the design and development process, there may be changes to the shape of the vehicle, including the backshell angle or thermal protection system sizing. Data from previous designs can be incorporated when analyzing new designs to reduce overall development costs, and flight test data can be incorporated to improve the aerodynamic models for

*Research Aerospace Engineer, Dynamic Systems and Control Branch, AIAA Member
†Aerospace Engineer, Configuration Aerodynamics Branch
‡Aerospace Engineer, Atmospheric Flight and Entry Systems Branch, and AIAA Member
§Pathways Intern, Configuration Aerodynamics Branch, AIAA Student Member
¶Research Aerospace Engineer, Configuration Aerodynamics Branch, AIAA Member
‖Research Aerospace Engineer, Flight Dynamics Branch
**Research Aerospace Engineer, Aerothermodynamics Branch, AIAA Member
††Research Computer Scientist, Atmospheric Flight and Entry Systems Branch

future flights. By making use of existing and low-fidelity data, the amount of high-fidelity data that must be collected can, in theory, be reduced. This can provide significant time and cost savings. Streamlining the model development process results in models that are equally (or more) predictive, reduces the engineering judgment required to quantify uncertainties, and provides easy fusion of prior information.

Also, the current modeling process examines each force and moment independently, not considering correlations within the uncertainties. The correlations may be implicitly captured in some of the constraints applied to derived values, such as an independent constraint on the ratio of the lift and drag forces, but they are not considered outright. This implicit correlation of outputs is an example of expert judgment that may be explicitly determined by the data.

II. Methods

A. Reduced-order Modeling for Data Generation and Sensitivity Analysis

A reduced-order model (ROM), in the context of AeroFusion, is an inexpensive mathematical representation of a complex, high-fidelity system. These formulations aim to capture the principal physics within these source models so that engineers can quickly study the dominant effects using minimal computational resources. It is worth noting that a ROM is different from a model with simplified physics; where the former takes the governing equations as they are and makes assumptions about the coherent structures of the dynamical system, the latter simplifies the governing equations themselves but allows the resulting coherent structures to be completely emergent. For example, a Reynolds-Averaged Navier-Stokes (RANS) simulation simplifies the physics of the Navier-Stokes equations by performing Reynolds averaging and applying a turbulence closure relation to the fluctuating field. This closure, while empirical, makes the equations solvable in such a way that relatively coarse domain meshes can yield relatively accurate results. Conversely, for reduced-order models, the number of independent parameters is reduced by coupling physical and/or temporal structures to one another, based on trustworthy system observations, and then propagating those coupled structures (in lieu of all uncoupled structures, which may emergently couple via the governing equations) through the solution space. In [4], aerodynamic trends are investigated for the Orion crew module under different RANS turbulence models. In [5], compressible fluid dynamics are simulated with low-cost ROMs such that more Orion flight conditions can be considered and parameter sensitivities will be quantifiable.

Due to their coupling of state variables, ROMs are inherently of lower fidelity than the source models, but they provide a computationally efficient means of producing a large amount of additional data around the source model. Because of this, ROMs enable a computationally tractable way to (a) perform sensitivity analyses and uncertainty quantification on numerical experiments, (b) provide dense, lower-fidelity data that can be combined with the sparse, higher-fidelity source data and multi-fidelity modeling approaches to create more confident data-fused models, and (c) (locally) extrapolate higher-fidelity studies to provide sufficient data to draw conclusions at a reduced cost.

One of the largest challenges with producing ROMs for planetary lander vehicles is properly capturing the physics of a fluid flow influenced by reacting flows, shock waves, and a turbulent wake. The flow physics are so complex that simplified physics models fail to capture the critical flow regions which exist throughout the parameter space. Standard compressible flow ROMs are likewise so sensitive that their limited stability makes their application for AeroFusion's purpose intractable. To address this, the AeroFusion team developed a multi-basis, affine compressible Galerkin ROM with linear stabilization functionality. A particle swarm optimization technique is used to provide the stabilization. In [5], this ROM formulation is demonstrated on a difficult flow regime with planned efforts to apply it directly to the Orion lander. Being able to capture these critical flow regions enables this type of ROM to be used in the development and analysis of aerodynamic models for planetary lander vehicles, providing all of the benefits of ROMs mentioned above.

Rigorous quantification of uncertainty in CFD simulations is difficult to perform. While spatial and temporal grid resolution studies can be conducted, and multiple solvers and models can be applied to similar flow problems, the interconnectedness of the simulation parameters, combined with the variable nature of the quantities of interest (e.g., distributed loads versus integrated forces and moments), makes an encompassing, rigorous uncertainty quantification procedure difficult to develop, as well as computationally expensive to conduct. In practice, often the procedure ends up relying on engineering judgment, posterior fidelity estimates, conservative design margins, and extra factors of safety. Due to their lower computational cost for simulating fluid flows, ROMs can be used to efficiently assess the sensitivity of the flow to various physical and computational parameters. In [6], a boundary penalty method is introduced to a Galerkin Proper Orthogonal Decomposition (POD) ROM so that a global sensitivity analysis method called Morris screening can be applied to a benchmark lid-driven cavity problem in a computationally efficient manner. We plan to apply an extension of this method to lander vehicles to assess sensitivity to physical parameters, such as surface roughness, that are too computationally intensive to study with high-fidelity simulations. By quantifying this sensitivity, the engineer can be more informed and compute more rigorous uncertainty bounds on the CFD data.

By combining aspects of physics-based fluid ROMs with interpolation-based methods, emergent physics can be identified at parameter points that were not used during the ROM's creation. Three different techniques are examined on a benchmark regularized lid-driven cavity problem in [7]: spanning ROMs, basis concatenation ROMs, and basis interpolation ROMs. The spanning and basis concatenation ROMs were demonstrated to be capable of accurate fluid flow simulations over both short- and long-time integration when the Reynolds numbers (parameter points) used to construct them are sufficiently close together. As the Reynolds numbers (parameter points) of the reference data were increased, stability over short-time integration was preserved, but stability over long-time integration became less consistent. Basis interpolation was seen to be less robust than the other methods, though it generally has a lower memory requirement. In all three cases, the parametric ROMs were able, to some degree of accuracy, to represent the fluid physics at parameter points where there was no explicit collection of data. Though these emergent physics can manifest from the ROM, the empirical basis on which the ROM is based mandates that these physics are taken with a grain of salt. Such ROMs can thus be used to generate lower-fidelity space-filling data that can be combined with multi-fidelity modeling methods, such as those outlined in Section II.B.

For most databases, CFD data is provided in the form of averaged quantities of interest. If the simulation is formulated as steady-state to begin with, the reported value is usually just the values computed at the final iteration; if the simulation is formulated to be time-accurate, then the reported value is usually an average over a significant, non-transient window of iterations, i.e., a time average. Running CFD simulations long enough to perform meaningful time averaging typically requires thousands of additional iterations after initial transient departure. For buffet and flutter analysis, where the quantities of interest are presented as functions of time or frequency, fine frequency resolution requires even more iterations. To reduce the cost of these analyses, only a subset of the statistically stationary CFD iterations can be performed by the full solver. Then, these snapshots can be used to train ROMs that allow efficient extension of these runs and provide sufficient data for time averaging or buffet and flutter analysis at a reduced cost. By separating a portion of the higher-fidelity CFD snapshots for testing, the accuracy loss incurred by the reduced-order modeling can be quantified.

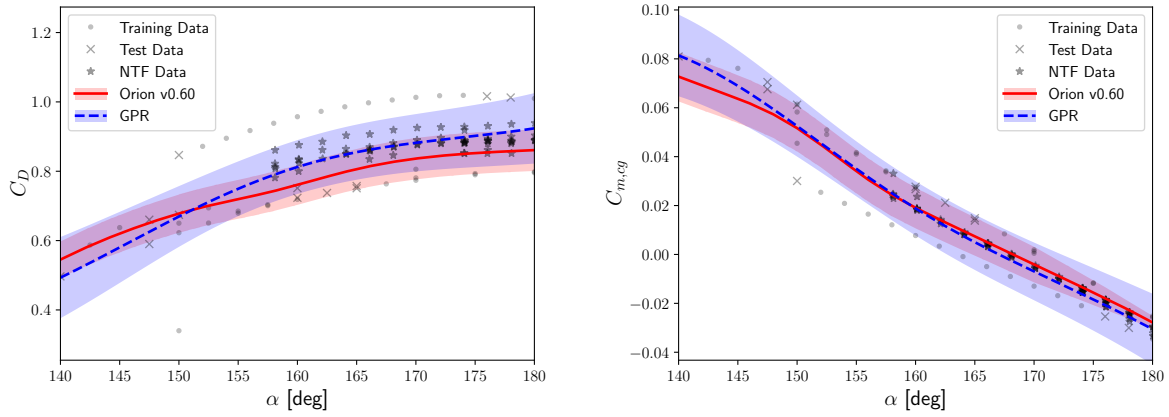
B. Surrogate Modeling with Multiple Fidelity Levels

A surrogate model, sometimes called a data-fit model or response surface model, seeks to fit data from physical experiments or high-fidelity models to provide a similar estimate of the quantities of interest while being faster to evaluate. It is not uncommon for aerodynamic surrogate models to take the form of a complex lookup table allowing proper blending between flight regimes, based on the existing data. For our trajectory simulations, it is only the net aerodynamic forces and moments acting on the vehicle that are of concern. As such, we create a surrogate model that captures the main effects of the fluid physics (i.e., the net forces and moments) without expending the resources to resolve the entire state of the fluid. Such models are typically constructed by performing a number of experiments, possibly physical (such as wind tunnel tests) or computational (such as CFD simulations), to measure the net forces and moments at various vehicle states. From this finite number of data points, regression analysis or some interpolation scheme is implemented to predict the quantities of interest (net forces and moments) as a function of the vehicle states.

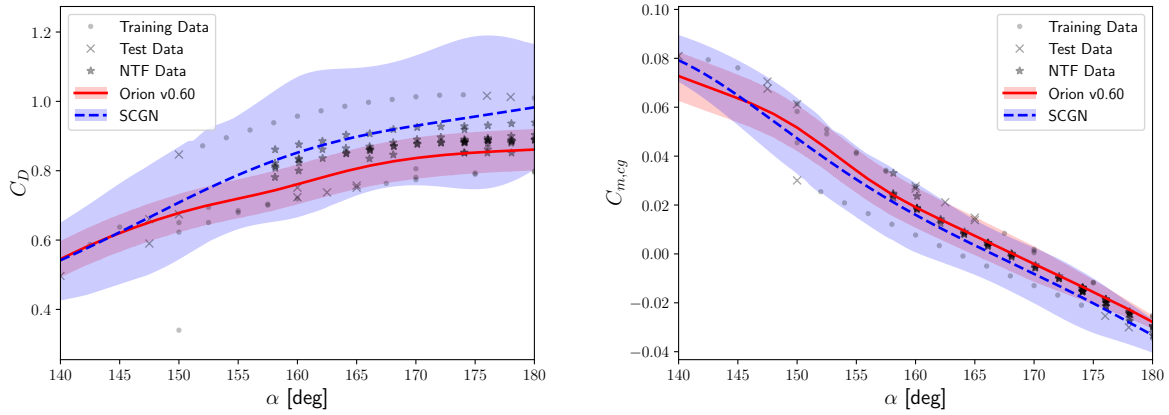
Regardless of the testing procedure used, the aerodynamic data collected will have some randomness associated with it, whether that is due to measurement noise, slight variations in conditions, or other sources of uncertainty. Because of this, it is insufficient to provide a single average estimate of the quantities of interest if a robust risk analysis of the mission is to be performed. Instead, an appropriate range of values should be given so that tests can be performed over the range of conditions that the vehicle might encounter. Traditionally, a single point estimate is provided, and a term-based interval is added on top of that where the modeled terms are based on engineering judgment and experience. The uncertainty estimates for these terms are built up by combining multiple sources of uncertainty in the data, such as turbulence modeling, grid resolution, wind tunnel repeatability, roughness effects, or other physical phenomena.

There are many analyses that depend only on the output (net forces and moments) of the aerodynamic model and not on the entire fluid state, e.g., control system analysis, structural load analysis, or trajectory analysis. Modeling the entire fluid flow would quickly become computationally prohibitive in simulations intended for these output-only analyses. Because surrogate models efficiently capture the main effects of the fluid physics, they are excellent candidates for such aero-dependent analyses. For example, in trajectory analyses (described further in Section II.D), Monte Carlo simulations are performed typically using thousands of trajectories, each with hundreds of thousands of calls to the aerodynamic model. Computational efficiency is key.

Beyond computational efficiency, it is difficult to compare the performance of various surrogate models. As part of AeroFusion, a toolbox of different surrogate modeling methods has been assembled. Different methods have different pros and cons. While a "nominal" estimate for the predicted forces and moments is important, so too is the distribution



(a) GPR



(b) SCGN

Fig. 1 Visual comparison of model predictions for drag coefficient C_D and pitching moment coefficient $C_{m,cg}$ at a constant Mach of 0.5 and various values of angle of attack α. Shaded region represents a 90% prediction interval.

of the predictions. Some of the models considered do not provide a straightforward means of deriving a prediction interval. However, these models can be combined with other methods by applying a secondary modeling technique to the residuals of the primary model. Figure 1 shows a visual comparison of two methods from the toolbox, Gaussian Process Regression (GPR) and Structured Covariance Gaussian Network (SCGN), against a baseline model, labeled Orion v0.60. Each of the models depicted here expresses a prediction distribution as well as a "nominal" function. When comparing models, either "nominal" models or prediction distributions, there is a question as to what metrics are appropriate for the comparison. In [8], a number of the modeling methodologies from the AeroFusion toolbox are implemented, and different metrics are discussed for comparing the resulting models. Relative to the existing baseline model, the models produced by the newer methods were developed in significantly less time. Less engineering judgment went into the production of the models as well since there was no need to sift through the data by hand to extract uncertainty quantification bounds. However, the engineering judgment that went into the term-based uncertainty quantification in the baseline model is able to provide assessments of known unknowns that are not present in the dataset. Likely, a combination of data-driven and judgment-based modeling methods will be needed to achieve reduced modeling costs in conjunction with the desired coverage of unknowns.

The first method shown in Fig. 1 is GPR, which is a method of interpolating that assumes the underlying process is Gaussian. One of the challenges of GPR is that the computational cost to evaluate the model increases with the number of data points used during training. In [9], a GPR formulation is provided which allows data of multiple fidelity levels to

be fused together into an ensemble of models. During the development of a vehicle, multiple configuration changes may occur, resulting in multiple computational or physical tests being conducted with varying degrees of fidelity for the final design. Due to constraints on collecting new datasets for each configuration, engineers typically must use existing data and make engineering assumptions based on limited updates to the data. One of the primary benefits of this method is that the probabilistic aerodynamics model that is created can be used during the design phase and later updated as flight data becomes available. The evaluation complexity of this version of GPR scales based on the largest subset of data when split by fidelity level.

The other approach depicted in Fig. 1, SCGN, was developed for nonlinear regression and uncertainty quantification based on a pair of neural networks that parameterize the mean and covariance functions of a multivariate Gaussian process. The pair is trained simultaneously so as to maximize the log-likelihood of observing the given data. A derivation and analysis of this new SCGN approach is given in [10]. It is shown that an SCGN can provide uncertainty descriptions that are similar to a traditional GPR model while significantly improving scalability with respect to data set size. An efficient algorithm to generate sample functions is also presented. The sample functions produced by SCGN are continuous and recover the point-wise covariance of the model. The evaluation of these functions is also computationally efficient, especially compared to GPR models that have large training sets. These properties are desirable for use in trajectory simulations, in which sample functions must be evaluated many times.

C. Data Fusion

Data fusion is the process of integrating multiple, potentially heterogeneous, data sources to produce more useful knowledge than that provided by the individual data sources. For aerodynamic testing procedures, it is common to compare CFD data with results from physical experiments to verify and validate the results. However, CFD data allows for flow visualization while it is difficult to capture many measurements associated with the fluid during physical testing. Typically, one only has access to measure the integrated forces and moments applied to the vehicle, along with a handful of discrete pressure and temperature measurements. When differences between quantities of interest — such as the lift and drag coefficients — exist, it is difficult to determine why the data sources disagree, due to the lack of full-field flow data. To remedy this, AeroFusion is investigating machine learning methods that are trained on discrete data akin to that available in a wind tunnel test and constrained by certain fluid flow properties in order to approximate flow field snapshots. The constructed snapshots have been found to agree with the true flow fields at large scales, even for complex incompressible fluid flows. In Fig. 2, results are shown for flow over a cylinder at a Reynolds number of $20,000$. More details on this method and its results can be found in [11]. The goal of this tool is to enable researchers to understand why quantities of interest differ between their experimental and computational datasets. Doing so helps improve the estimation of the uncertainty of the data. Sensitivity analysis performed on the flow reconstruction method can identify areas where those sparse physical measurements will provide the most information. It can also be used to identify if higher fidelity sensors would have a meaningful impact on the reconstruction. A study on the sensitivity of the neural networks used to reproduce flow fields is provided in [12].

D. POST2

The Program to Optimize Simulated Trajectories II (POST2) is a flight mechanics software that is used widely across NASA, industry, and academia that was first developed in the 1970s [13, 14]. A wide variety of atmospheric entry problems have successfully utilized POST2, including multiple Mars missions [15–20]. POST2 is the workhorse integrated simulation framework for assessing the flight performance of both human and robotic spacecraft at varying solar system destinations. Within the modular architecture of POST2, many simulation models can be incorporated to calculate vehicle-specific data, e.g., aerodynamics, different vehicle configurations, or different mission requirements.

The state of the art (SOTA) for aerodynamic modeling with POST2 involves an aerodynamicist and a flight mechanic. The aerodynamicist is tasked with developing an appropriate aerodynamic database (surrogate model) from the available data. This data can be from physical tests such as wind tunnels, or it may come from CFD numerical simulations. Once the aerodynamic database is generated, the flight mechanic uses POST2 to ingest the aerodynamic database and run analyses of the integrated system. The results of these analyses are then reviewed by each of the corresponding subject matter experts to confirm the behavior of the vehicle is appropriate.

For the work described herein, a framework was developed in order to utilize the methods described previously within POST2 for aerodynamic databases. This framework consists of a localized web socket server written in Python to allow POST2 to communicate with the AeroFusion models that utilize machine learning (ML) and uncertainty quantification (UQ). The socket-server interface is especially useful because it allows the strengths of POST2 and ML+UQ within Python to be leveraged simultaneously. Inputs to the server include Mach number, angle of attack, and angle of sideslip. The outputs consist of the aerodynamic force and moment coefficients. Outside of the inputs

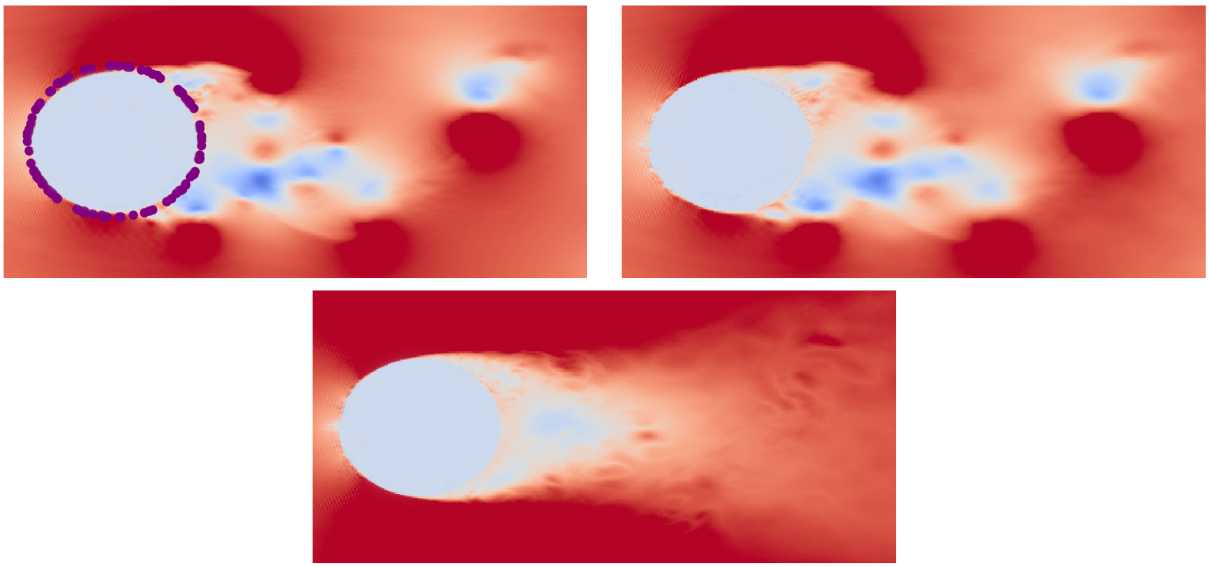


Fig. 2 Flow reconstruction results [11]. Top left: True flow with sensor positions (purple dots). Top right: Reconstructed flow based on the discrete measurements. Bottom: Normalized reconstruction error (red denotes a low value).

and outputs to the Python server, POST2 is agnostic to which model is used. Once the AeroFusion models are vetted against the SOTA for aerodynamic databases, the web socket server framework can be utilized to compare integrated performance of the two databases. We note that there is a natural trade off between the computational time of a trajectory evaluation and the complexity of the aerodynamic database. Increased complexity may bring about increased accuracy; however, it likely also increases the time needed to analyze the trajectories. A balance needs to be struck between these competing objectives.

For example, Fig. 3 shows a comparison of the Mach number at mortar fire between the baseline Orion API and a preliminary AeroFusion model. (The mortar fire is used to release the drogue parachute.) From the POST2 results, information about the corresponding distributions can be analyzed, including the mean and quantiles. A significant statistical difference can be seen here in the mean value of the Mach at mortar fire; however, the training data used to generate these two models were not the same. Future comparisons will provide more meaningful insights when the same underlying data can be used for each model.

III. Fusing the Pieces

Given the various methods discussed in Section II, how do we combine them into an efficient workflow to develop aerodynamic models for new vehicles? Some initial data will always be required, so our proposed workflow will start with some high-fidelity CFD data taken over a sparse subset of the parameter space. From this data, ROMs can be computed at the parameter points to assess the sensitivity of the data. The knowledge gained from the sensitivity analysis can be used to drive the collection of additional high-fidelity data (either computational or experimental) where it will be most useful. As noted in Section II, ROMs can even be used for time savings on producing enough data for time averaging. By combining the ROMs with interpolation-based methods (such as those mentioned in [7]), the parameter space can be filled out with ROM-based lower-fidelity data. The combination of high-fidelity CFD data and the lower-fidelity ROM-based data can then be collected to construct surrogate models (such as those in Section II.B) for use in the trajectory simulations. Note that while fast relative to the higher-fidelity CFD solvers, the ROMs still spend time determining fluid state information that is not necessary for a trajectory simulation.

This creates a first-iteration model for trajectory analysis. As vehicle design changes based on programmatic needs and these early results, we can then repeat that process for the updated design. The relatively quick turn around time between iterations should reduce the overall development time of the aerodynamic model while remaining agile enough for nontraditional vehicle designs. As each iteration of the vehicle becomes available, multi-fidelity methods (such as the one detailed in [9]) can be employed to make use of all of the data available thus far. When experimental results are added to the computational data, flow reconstruction methods (such as the one discussed in [11]) can be used to improve

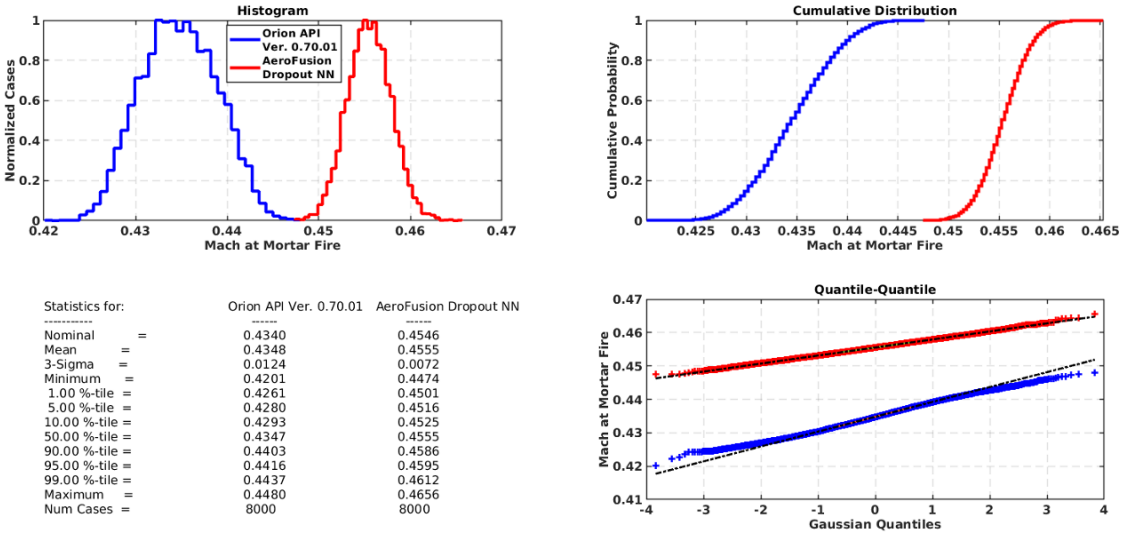


Fig. 3 Comparison of a preliminary AeroFusion model with the baseline Orion model.

our understanding of the experimental results and assist in identifying why these different data sources might disagree. A schematic of the proposed workflow is shown in Fig. 4.

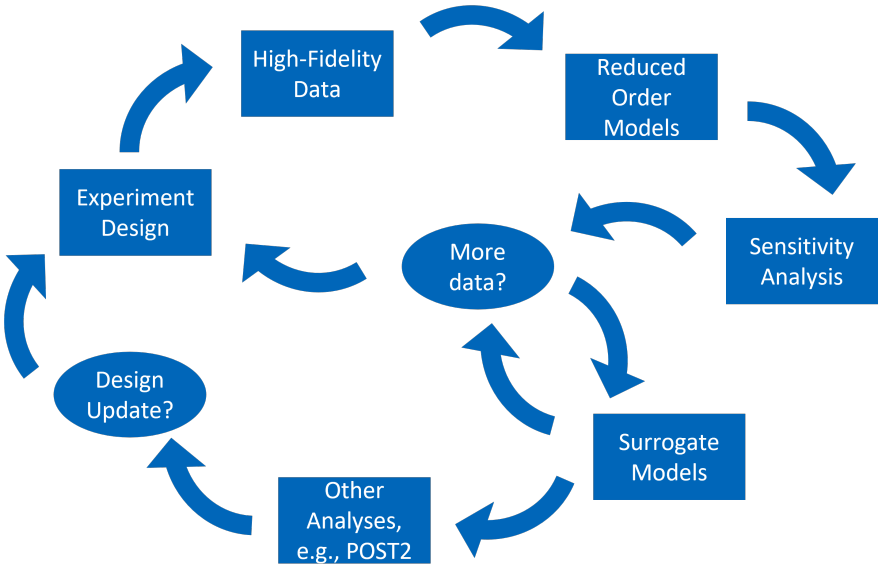


Fig. 4 Schematic of proposed workflow.

IV. Summary

The goal of AeroFusion is to improve the aerodynamic modeling process by making more efficient use of data and requiring less engineering judgment. To achieve this, we recommend the use of simplified physics models [4] and reduced-order models [5] to capture trends. These more computationally-friendly models can then be combined with sensitivity analysis [6] and interpolation methods [7] to provide uncertainty quantification and low-cost, volume-filling data. For use in other analyses that do not require fluid state information, we recommend using this low-cost data to construct surrogate models that can capture uncertainties [9, 10]. Some modeling methods can even account for different levels of fidelity while operating more efficiently than simply pooling all available data together [9]. Different surrogate modeling techniques may be more or less beneficial for certain tasks. Because of this, we discuss different metrics for

comparing these probabilistic models [8]. Finally, we suggest flow reconstruction from discrete measurements [11] as a means of comparing data from disparate sources, such as wind tunnel test data and CFD data. Sensitivity analyses on these flow reconstruction methods can even provide input on optimal sensor placement [12].

By no means do we consider the efforts to improve the aerodynamic modeling process presented here to be complete. There is still room for improvement within the various individual methods mentioned and the connection of these component technologies. Full implementation of the envisioned workflow on a vehicle of interest has yet to occur, but it is the goal of this paper to spark interest within the community on enhancing this key process for the aerospace industry.

Acknowledgments

We would like to thank the Space Technology Mission Directorate (STMD) Early Career Initiative (ECI) program for funding this work.

References

[1] Mauery, T., Alonso, J., Cary, A., Lee, V., Malecki, R., Mavriplis, D., Medic, G., Schaefer, J., and Slotnick, J., "A Guide for Aircraft Certification by Analysis," Tech. Rep. CR-20210015404, NASA, 2021.

[2] Bibb, K., Brauckmann, G., Walker, E., and Robinson, P., "Development of the Orion Crew Module Static Aerodynamic Database, Part I: Hypersonic," *AIAA Applied Aerodynamics Conference*, Honolulu, HI, 2011.

[3] Bibb, K., Walker, E., Brauckmann, G., and Robinson, P., "Development of the Orion Crew Module Static Aerodynamic Database, Part II: Subsonic/Supersonic," *AIAA Applied Aerodynamics Conference*, Honolulu, HI, 2011.

[4] Buchmeier, J., Ramezanian, D., and Bermejo-Moreno, I., "Parametric Studies of Subsonic and Supersonic Reentry Phases of the Orion Crew Module through Numerical Simulations," *AIAA SciTech Forum*, 2023. (Submitted).

[5] Ramezanian, D., Lee, M., and Bermejo-Moreno, I., "Application of an Affine Nonlinear Galerkin Reduced-order Model to Compressible Fluid Flows," *AIAA SciTech Forum*, National Harbor, MD, 2023. (Submitted).

[6] Hanes, H., Lee, M., Ramezanian, D., and Smith, R., "Low-cost Quantification of Fluid Flow Parameter Sensitivity using Reduced-order Modeling," *AIAA SciTech Forum*, National Harbor, MD, 2023. (Submitted).

[7] Edwards, C., Lee, M., Ramezanian, D., and Smith, R., "Quantifying Emergent Fluid Dynamics Using Reynolds-Interpolated Reduced-Order Models," *AIAA SciTech Forum*, National Harbor, MD, 2023. (Submitted).

[8] Wignall, T. J., Nakamura-Zimmerer, T., Scoggins, J. B., Snyder, S., Kumar, S. G., and Colbert, B. K., "Comparisons of Performance Metrics and Machine Learning Methods on an Entry, Descent, and Landing Database," *AIAA SciTech Forum*, National Harbor, MD, 2023. (Submitted).

[9] Scoggins, J. B., Wignall, T. J., Nakamura-Zimmerer, T., and Bibb, K., "Multihierarchy Gaussian Process Models for Probabilistic Aerodynamic Databases using Uncertain Nominal and Off-Nominal Configuration Data," *AIAA SciTech Forum*, National Harbor, MD, 2023. (Submitted).

[10] Nakamura-Zimmerer, T., Stringer, M. T., Colbert, B. K., and Scoggins, J. B., "Structured Covariance Gaussian Networks for Orion Crew Module Aerodynamic Uncertainty Quantification," *AIAA SciTech Forum*, National Harbor, MD, 2023. (Submitted).

[11] Lebedev, Y., Lee, M., and Zare, A., "Construction of a Fluid Flowfield from Discrete Point Data using Machine Learning," *AIAA SciTech Forum*, National Harbor, MD, 2023. (Submitted).

[12] Hanes, H., Lebedev, Y., Smith, R., and Zare, A., "Convergent Uncertainty Quantification in Fluid Dynamics using Reduced-order Models and Machine Learning," *AIAA SciTech Forum*, National Harbor, MD, 2023. (Submitted).

[13] Brauer, G. L., Cornick, D. E., and Stevenson, R., "Capabilities and Applications of the Program to Optimize Simulated Trajectories (POST)," Tech. Rep. CR-2770, NASA, Feb. 1977.

[14] Williams, R. A., Lugo, R. A., Marsh, S. M., Hoffman, J. A., Shidner, J. D., and Aguirre, J. T., "Enabling Thread Safety and Parallelism in the Program to Optimize Simulated Trajectories II," *AIAA SciTech Forum*, National Harbor, MD, 2023. (Submitted).

[15] Braun, R. D., Powell, R. W., Engelund, W. C., Gnoffo, P. A., Weilmuenster, K. J., and Mitcheltree, R. A., "Mars Pathfinder Six-Degree-of-Freedom Entry Analysis," *Journal of Spacecraft and Rockets*, Vol. 32, No. 6, 1995, pp. 993–1000. doi: 10.2514/3.26720.

[16] Desai, P. N., Schoenenberger, M., and Cheatwood, F. M., "Mars Exploration Rover Six-Degree-of-Freedom Entry Trajectory Analysis," *Journal of Spacecraft and Rockets*, Vol. 43, No. 5, 2006, pp. 1019–1025. doi:10.2514/1.6008.

[17] Desai, P. N., Prince, J. L., Queen, E. M., Schoenenberger, M., Cruz, J. R., and Grover, M. R., "Entry, Descent, and Landing Performance of the Mars Phoenix Lander," *Journal of Spacecraft and Rockets*, Vol. 48, No. 5, 2011, pp. 798–808. doi:10.2514/1.48239.

[18] Striepe, S., Way, D. W., Dwyer, A. M., and Balaram, J., "Mars Science Laboratory Simulations for Entry, Descent, and Landing," *Journal of Spacecraft and Rockets*, Vol. 43, No. 2, 2006, pp. 311–323. doi:10.2514/1.19649.

[19] Maddock, R. W., Cianciolo, A. M. D., Korzun, A. M., Litton, D. K., Zumwalt, C. H., and Karlgaard, C. D., "InSight Entry, Descent, and Landing Postflight Performance Assessment," *Journal of Spacecraft and Rockets*, Vol. 58, No. 5, 2021, pp. 1530–1537. doi:10.2514/1.A35023.

[20] Way, D. W., Dutta, S., Zumwalt, C., and Blette, D., "Assessment of the Mars 2020 Entry, Descent, and Landing Simulation," *AIAA SciTech Forum*, San Diego, CA, 2022.