The German Section of the Combustion Institute (CI)

Deutsche Vereinigung für Verbrennungsforschung (DVV)

32.
Deutscher Flammentag
2025

15. - 17.09.2025

Paderborn

Development of a Surrogate-Based Optimization Framework for Hydrogen Burner Design

L. Juris[1*], S. Nerzak[2], J. Willkomm[1], G. Melo[1], M. Gauding[2], H. Pitsch[2], J. H. Schleifenbaum[1]

*luca.juris@dap.rwth-aachen.de

[1] *Chair for Digital Additive Production (DAP), RWTH Aachen University, Campus-Boulevard 73, 52074 Aachen, Germany*

[2] *Institute for Combustion Technology, RWTH Aachen University, Templergraben 64, 52056 Aachen, Germany*

Abstract

Hydrogen's specific combustion properties, such as high consumption speeds, wide flammability limits, and low ignition energy, pose significant challenges for burner design: conventional manual iterations are time-consuming, enable limited insight into the multi-dimensional design space, and risk missing global optimal configurations. A full-factorial exploration, on the other hand, becomes expensive under realistic computational budgets. To overcome these limitations, a simulation driven, surrogate-based optimization framework that couples parametric CAD generation, steady-state Reynolds-Averaged Navier-Stokes (RANS) CFD simulations and an Efficient Global Optimization (EGO) algorithm is introduced. By constructing a Gaussian-Process surrogate and iteratively refining it via Expected Improvement, the framework systematically navigates a four-dimensional burner design space. With this methodology, an improved configuration that achieves a 52 % reduction in mixture inhomogeneity relative to the dataset mean can be identified. Crucially, only 72 high-fidelity simulations were required—a computational saving of 87.5 % compared to a full-factorial approach—demonstrating both the efficiency and effectiveness of the proposed framework for automated hydrogen burner design.

Introduction

Hydrogen is a key enabler for the decarbonisation of energy-intensive sectors and forms a central pillar of the EU's Green Deal and REPowerEU plan, which targets 20 million tonnes of renewable hydrogen by 2030 [1, 2]. One example for decarbonization is the thermochemical energy conversion process by replacing conventional hydrocarbon fuels like methane with hydrogen. Compared to methane, hydrogen exhibits higher laminar flame speeds, broader flammability limits, and lower ignition energy, increasing the risk of flashback and complicating flame stabilization in premixed systems [3, 4]. In addition, the potential for thermal NO_X formation is greater with hydrogen, as the adiabatic flame temperature is approx. 200° C ($\Phi = 1$) higher than with methane [5, 6]. To minimize NO_X formation, modern burner designs rely on lean premixed or partially premixed operation. These configurations have two limitations. First, these conditions are sensitive to thermo-diffusive flame instabilities and therefore exhibit significantly increased reaction rates and super-adiabatic temperatures. Second, in technical applications, highly homogeneous mixing over a short mixing length has to be ensured to minimize the risk of flashback [7, 8]. Modern burner concepts need to be optimized to enable the combustion of hydrogen safely and with low emissions. Different approaches for diffusion and premixed burners have been proposed and special low- NO_X concepts such as micro mixing burners are currently being developed such as in [9]. With the micro mixing concept, the fuel is injected through many small openings to produce a large number of small localized flames. This keeps the local temperature lower and the residence time of the hot combustion products short, which reduces NO_X formation and flashback risk, as the flame front cannot spread into the fuel opening [10]. Similar concepts are pursued by tangentially arranged fuel injectors, which generate a slight swirl. The flames become smaller and the residence time of the hot gas is reduced, thereby lowering NO_X emissions as the swirl increases [11]. The introduced concepts offer potential solutions for safe and efficient hydrogen burners but require complex geometries that are difficult to manufacture with conventional manufacturing processes like milling or casting.

Additive manufacturing (AM), particularly Powder Bed Fusion using a Laser Beam of Metals (PBF-LB/M), enables the layer-wise fabrication of complex internal features without the constraints of conventional tooling or assembly [12]. This design freedom allows for the integration of flow-guiding structures and distributed injection systems directly into the burner geometry. However, PBF-LB/M also imposes limitations such as surface roughness, anisotropic material properties, and geometric deviations—especially in thin-walled or overhanging regions [13]. These effects must be considered early in the design process to ensure manufacturability and consistent performance.

The combination of hydrogen-specific burner requirements and the capability of AM necessitates a novel, simulation-based design approach that integrates combustion dynamics, flow behavior, and manufacturing constraints into a unified optimization process.

This process has to be able to cover the complex, multi-objective challenge due to the interplay of application-specific operating conditions, high-dimensional parameter spaces, and non-linear trade-offs that the design of hydrogen burners presents. Traditional, sequential design and optimization approaches are ill-suited to this context, as they rely heavily on manual iteration and fail to systematically identify high-performance geometries [14]. To address these limitations, automated design frameworks have emerged that integrate parametric modelling, simulation, and optimization algorithms. These frameworks enable efficient and reproducible exploration of the design space, particularly when using inverse design strategies such as generative design [15]. While such methods have been applied in the broader field of hydrogen technologies only a few have been used in burner design [16]. State-of-the-art optimization techniques include gradient-based algorithms, evolutionary strategies, and surrogate-based optimization (SBO). The latter being particularly effective in scenarios involving computationally expensive simulations [17, 18]. Surrogate models, such as Gaussian Processes (GP), approximate the simulation input–output relationships, allowing rapid exploration of a parameterized design space without full CFD simulations for each variant [19]. In early design stages, this mapping uncovers key variables, interactions, and performance trends that isolated case studies or single-point optimizations might miss [20]. Efficient Global Optimization (EGO) pairs a GP surrogate with the Expected Improvement Function (EIF) to iteratively refine designs, balancing exploration and exploitation and converging to global optima with few evaluations in costly, high-dimensional settings [21].

This study aims to develop an additively manufacturable, partially premixed crossflow jet burner for hydrogen combustion, targeting improved mixing quality, thereby ensuring low NO_x emissions. To this end, an automated simulation-driven design approach is pursued, combining parametric geometry generation, cold-flow CFD simulations, and surrogate-based optimization. The proposed framework enables efficient exploration of the design space to identify high-performance burner configurations that fulfil functional and manufacturing constraints.

Methods

This study presents a design automation framework that is specifically tailored for the optimization of hydrogen burner geometries. Figure 1 illustrates the overall architecture of the optimization framework, which establishes an automated loop between the open-source software tools Dakota, Salome, meshing, and OpenFOAM [22, 23, 24]. Dakota generates design parameters, which are converted into parametric geometries in Salome, meshed, and then evaluated via CFD in OpenFOAM. The simulation results are returned to Dakota to inform further sampling and optimization. Each step is described in detail in the following sections.

To reduce computational cost and develop a generalizable design methodology, the framework is applied to a simplified burner configuration based on a jet-in-crossflow setup. This includes multiple fuel inlets injecting transversely into a main oxidizer stream within a confined duct, capturing key mixing mechanisms while avoiding the complexity of full burner assemblies. By isolating a single mixing region from combustion and micro-mixing effects, the setup enables focused analysis of how geometry influences mixture homogeneity at reduced computational cost. The simplified configuration allows for the systematic modulation of several parameters, including the fuel inlet angle α, inlet height h, and inlet diameter d, as well as the quantity of inlets $N_{\mathrm{fuel\_inlets}}$. This parametric flexibility forms the basis for generating a comprehensive dataset used in training and validating the surrogate model.

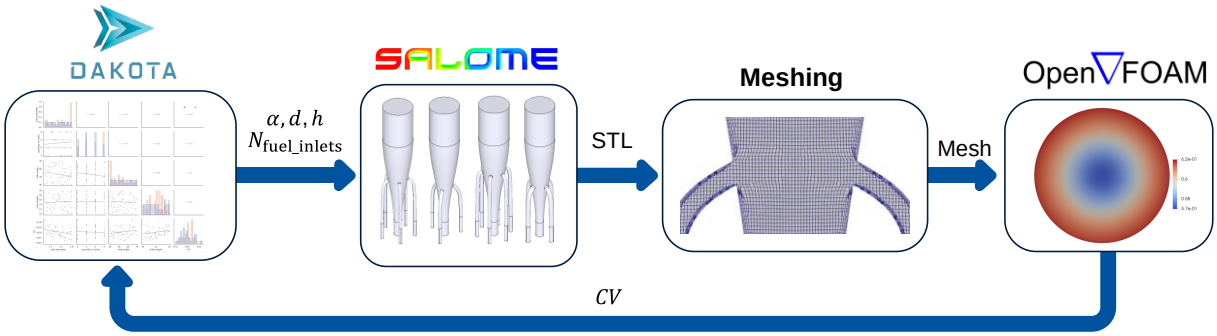

Figure 1. Framework architecture.

Parametric CAD model

The parametric model is built in Salome, which is an open-source tool that supports both graphical and script-based CAD modelling and offers a Python interface for automating geometry generation and variation. In this study, geometric parameters with a significant influence on the mixing behavior are selected. These parameters are systematically varied to populate the design space. In Figure 2, one representative geometry is shown, illustrating the definition and variation of the considered parameters. A parametric Python script automates geometry generation and modification based on parameter sets which are created in Dakota.

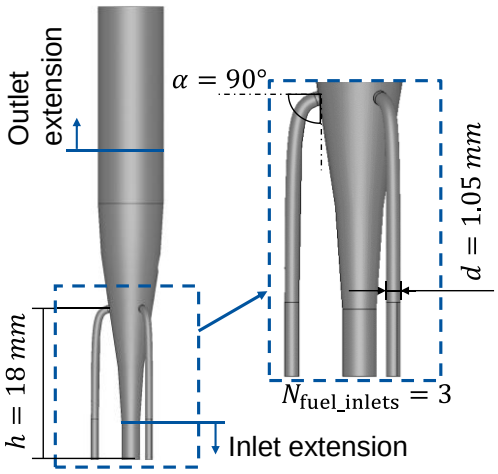

Figure 2. Parametric Salome model.

To improve the numerical stability of the simulation, the geometry generated in Salome includes extensions at both the inlet and outlet. The inlet section is extended by 5 mm to allow the velocity profile to develop before entering the mixing region. Similarly, the outlet is extended by 15 mm to reduce backflow effects and dampen transient behavior, thereby enhancing the robustness of the simulation results.

CFD Simulation

OpenFOAM is an open-source CFD toolbox with modular architecture used for simulating fluid flow and heat and mass transfer. In the presented framework, OpenFOAM is used to simulate the steady-state flow using the RANS equations in combination with the k-ε turbulence model in reactingFoam. The unityLewisFourier scalar transport model is used the which is the standard in the reactingFoam solver. This configuration models a non-reacting hydrogen–air mixing process with a global Reynolds number of 7000 and an equivalence ratio of $\Phi = 0.6$. While the total mass flow rate and equivalence ratio remain constant across all simulations, geometry-dependent variations in the fuel inlet configuration lead to case-specific inlet velocities, turbulence intensities, and k–ε values, which are computed to automate the simulation setup. Table 1 lists all additional boundary and initial conditions that remain constant during the simulations.

Table 1. Simulation setup.

	oxidizer	fuel	boundary condition
inlet	$U_{\text{air}} = 122\ \frac{m}{s}$	$U_{\text{air}} = \text{calculated}$	U = fixed value
	$\rho_{\text{air}} = 1.2\ \frac{kg}{m^3}$	$\rho_{\text{H}_2} = 0.084\ \frac{kg}{m^3}$	
	$Q_{\text{air}} = 6.0367 * 10^{-4}\ \frac{m^3}{s}$	$Q_{\text{H}_2} = 1.5217 * 10^{-4}\ \frac{m^3}{s}$	p = zero gradient
	$T_{\text{air}} = 298\ K$	$T_{\text{H}_2} = 298\ K$	
outlet	-	-	U = zero gradient
			p = total pressure

A consistent mesh is maintained with a 0.14 mm base cell and a 1.2 expansion ratio to balance resolution and cost. Boundary layers are resolved conservatively: the oxidizer uses five layers at a 55 % growth ratio, while the tight, curved fuel inlets use three layers at 30 %, avoiding cell distortion. These settings strike a compromise between physical accuracy and geometric feasibility and are kept constant because manual mesh tuning is not possible in the automated optimization workflow.

The CFD simulations are conducted using the reactingFoam solver in OpenFOAM, which handles compressible, turbulent, reacting flows. A local Euler time integration scheme improves the resolution of unsteady effects in the mixing region compared to a steady state scheme, and an adjustable timestep adapts dynamically to local Courant numbers, enhancing transient accuracy without compromising stability. As the study focuses solely on mixing behavior, the combustion model is deactivated. Each simulation runs for 4000 outer iterations, based on preliminary tests showing that residuals stabilize sufficiently within this range for reliable evaluation. Each CFD simulation, including preprocessing (geometry preparation, meshing) and postprocessing (data extraction, objective function evaluation), requires approximately 60 minutes of wall-clock time on the high-performance computing (HPC) cluster CLAIX-2023 at RWTH Aachen University using one node with 48 processor cores.

Design space exploration and SBO

Dakota is an open-source tool for optimization and design-space exploration. In this framework, it coordinates the workflow by linking geometry generation in Salome with CFD simulations in OpenFOAM and automating the evaluation of design variants.

For the design optimization phase, this study uses the Efficient Global Optimization (EGO) method implemented in Dakota. EGO was selected due to its suitability for simulation-driven workflows with high computational cost, as it enables targeted refinement of the surrogate model while requiring relatively few CFD evaluations. This makes it effective for exploring complex, high-dimensional design spaces under limited computational resources.

For initial design space exploration, the Latin Hypercube Sampling (LHS) method is used. LHS is a stratified sampling technique that ensures a uniform, space-filling distribution of sample points across the parameter space, as compared to simple random sampling [20]. This sample size is chosen based on the Dakota framework's recommendation of selecting between $2x$ and $10x$ samples, where x is the number of design variables, depending on the problem's complexity and nonlinearity [23]. For this study the initial sample size consists of $10x$ samples resulting in 40 initial CFD simulations. In Table 2 the limits for the design variables are shown, which were chosen based on a combination of geometric constraints, manufacturing limitations, and numerical stability considerations.

Table 2. Limits of design variables.

	Inlet angle α [°]	Inlet height h [mm]	Inlet diameter d [mm]	Quantity of inlets $N_{\text{fuel\_inlets}}$
Upper limit	40	10	1.1	3
Lower limit	70	20	1.6	6

The inlet angle is limited to 40–70°: angles below 40° cannot be modelled robustly in Salome across all heights, while angles above 70° cause inconsistent surface patch assignments due to randomized face numbering. The inlet height range of 10–20 mm is chosen to avoid regions of high oxidizer velocities, which occur for h < 10 mm and lead to solver instabilities, and to respect the physical space available in the burner assembly, since h > 20 mm would interfere with adjacent flow paths. The inlet diameter is constrained to 1.1–1.6 mm to prevent manufacturing issues in PBF-LB/M, because diameters < 1.1 mm require extensive post-processing and to maintain sufficient fuel inlet momentum, since diameters > 1.6 mm in combination with many inlets produce excessively low velocities. Finally, the number of fuel inlets is varied from 3 to 6: fewer than three cause asymmetry, while more than six cannot be physically arranged in the diverging duct. Because Dakota's EGO algorithm only handles continuous variables, the quantity of inlets is treated continuous during sampling and rounded to the nearest integer before geometry generation in Salome, preserving automation and ensuring valid discrete configurations.

The objective function is defined based on the spatial distribution of the equivalence ratio (Φ) at the outlet plane (excluding the outlet extension) with the goal to minimize the variation of Φ, which is defined as the ratio of the actual to the stoichiometric fuel–oxidizer ratio:

$$\Phi = \frac{(\dot{m}_{fuel}/\dot{m}_{oxidizer})_{actual}}{(\dot{m}_{fuel}/\dot{m}_{oxidizer})_{stoich}} \tag{1}$$

Values of $\Phi < 1$ correspond to fuel-lean conditions and $\Phi > 1$ to fuel-rich conditions. To reduce numerical fluctuations in the steady-state RANS results, hydrogen and oxidizer mass fractions are averaged over the final outer iterations. Sampling is performed cell-wise on an outlet plane (excluding the outlet extension), and local Φ values are calculated for each cell based on the averages of hydrogen massfraction. To quantify the overall mixture homogeneity, the coefficient of variation (CV) is computed across all outlet surface cells. The CV is defined as the ratio of the standard deviation to the mean of the equivalence ratio:

$$CV = \frac{\sigma_\phi}{\bar{\Phi}}, \tag{2}$$

where $\bar{\Phi}$ is the mean equivalence ratio and σ_ϕ its standard deviation over the outlet plane. A low CV indicates a more homogeneous mixture and thus represents an improved mixing performance. This scalar value serves as the final objective function in the optimization process. CV is chosen because it provides a dimensionless, statistically robust indicator that captures both the magnitude and consistency of local deviations. Alternatively, other metrics such as the variance of the equivalence ratio could replace CV or be incorporated as an additional objective function to further enhance the optimization process.

The accuracy of the surrogate model is assessed both during and after the optimization process. Within the EGO algorithm, surrogate predictions are iteratively validated by selecting new sample points based on the EIF. At each iteration, the design variables set with the highest expected improvement are evaluated using a CFD simulation, and the resulting value is used to update the GP model. This process ensures a continuous refinement of the surrogate model in promising regions of the design space.

The initial sampling was followed by 32 additional simulations during the EGO-based optimization phase, where each newly proposed design was evaluated based on its expected improvement. To assess the predictive accuracy of the surrogate model, the mean absolute percentage error (MAPE) is used in a Leave-One-Out (LOO) cross-validation scheme. For each of the data points, a GP surrogate is trained on the remaining samples, and the excluded point is used for testing. This is repeated for all samples, and the average relative error is computed. MAPE is particularly suitable here, as the objective function takes on small values. Unlike error metrics such as the mean square error (MSE) or R^2, which are sensitive to scale and outliers, MAPE is scale-independent and more robust for low-magnitude targets and varying output scales. The MAPE is calculated according to Eq. 3 where y_i denotes the true objective values, $\hat{y}_i$ the surrogate predictions and N the total number of samples [25]:

$$MAPE = \frac{1}{N}\sum_{i=1}^{N}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \times 100 \quad (3)$$

Results

This section presents the results of the surrogate-based optimization process, including the sampling behavior, surrogate model performance, and the evaluation of the optimal design configuration.

The evaluation of the sample distribution and optimization outcomes shows distinct trends in mixing performance. Figure 3 shows the scatterplot matrix, highlighting scatterplots for each pair of design variables on the off-diagonal elements, histograms of individual variable distributions on the diagonal, and includes Pearson correlation coefficients to quantify the strength and direction of relationships between parameters, providing a comprehensive overview of parameter interactions and sampling behavior.

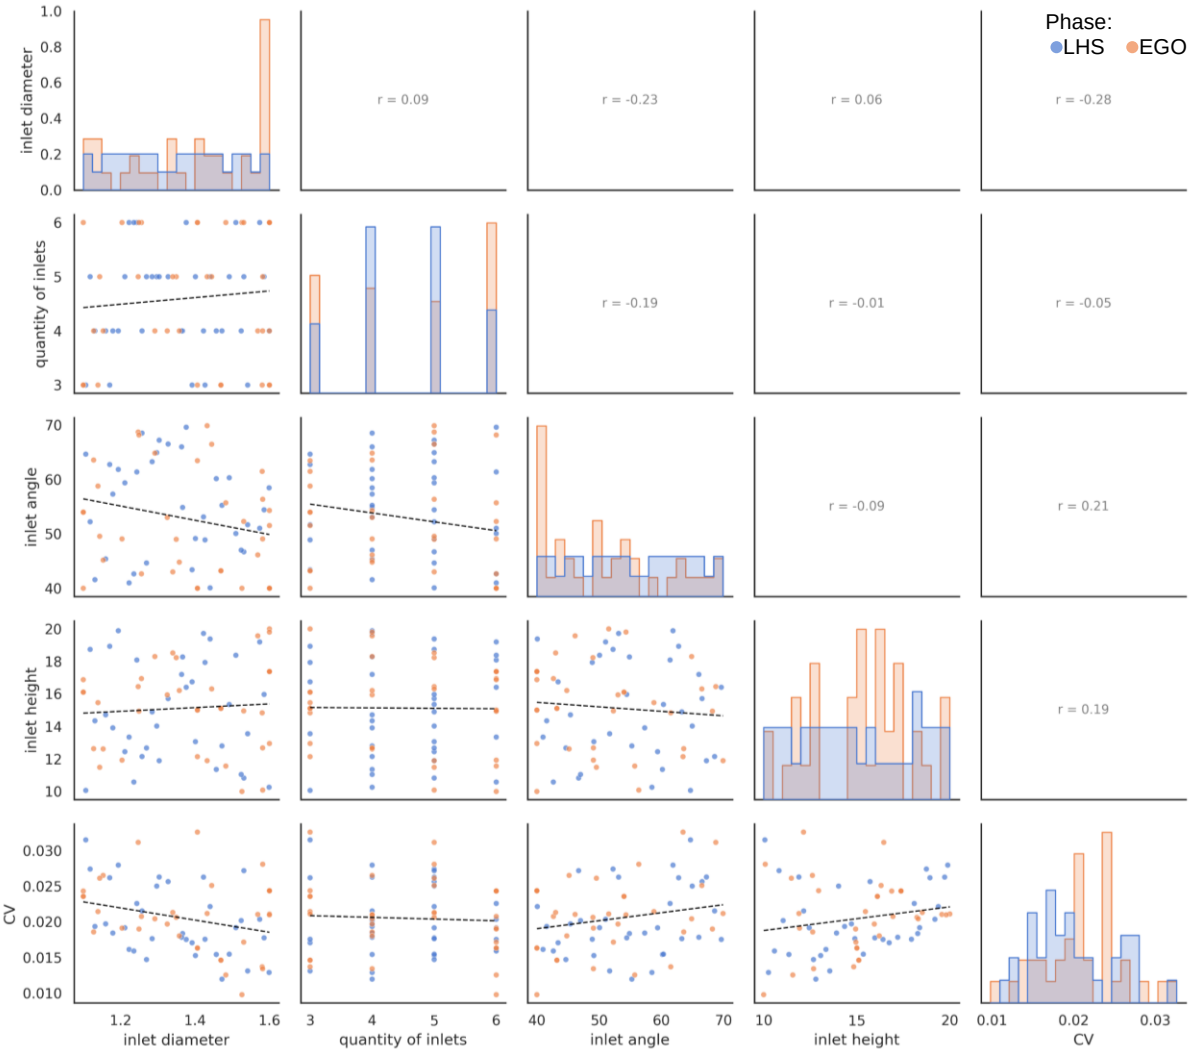

Figure 3. Scatterplot matrix of LHS and EGO samples.

The initial LHS samples, represented by blue points, ensure uniform coverage across diameter and inlet angle, whereas the inlet height exhibits underrepresentation within the range of 16 to 18 mm. Because the number of inlets was sampled continuously and then rounded to integers, the LHS stage produces an uneven distribution of inlet counts which are clustered at 4 and 5 inlets. The EGO phase, denoted by orange points, incorporates an increase in samples with a higher quantity of inlets and larger diameters, which is also indicated by the Pearson correlation (r = -0.28). Notably, EGO adds more points at 1 and 6 inlets to probe underexplored discrete regions. Additionally, EGO favors smaller angles within the range of 40° to 50°, aligning with the observed weak positive correlation (r = 0.21). For the inlet height, supplementary samples are introduced at 12 mm and within the previously underrepresented 16 to 18

mm range to optimize the design space comprehensively. The increase in the quantity of inlets, shows a low linear correlation of r = -0.05.

The prediction accuracy of the surrogate model starts at a mean average percentage error of 19,86% with 40 initial LHS samples and stabilizes at 16,76% after a total of 32 EGO samples are added. This indicates a gradual improvement in the model's reliability for predicting the coefficient of variation (CV) as the EGO phase refined the design space.

The optimization process identifies a design with the following parameters as the best-performing configuration: inlet diameter = 1.53 mm, number of inlets = 6, inlet angle = 40.0°, and inlet height = 10.0 mm. This design achieves a CV value of 0.0098, which is 52.2% smaller than the mean CV of the overall dataset with a value of 0.0205, indicating an improved mixing homogeneity. This parameter set and its corresponding CV value are not a prediction from the surrogate model. It was selected by the EGO algorithm during the optimization process due to its highest expected improvement in that iteration and was subsequently validated through a CFD simulation. Table 3 compares this optimized design with a configuration whose CV value is close to the dataset mean. The mean CV design exhibits a smaller diameter, fewer inlets, larger angles, and higher inlet heights aligning with the observed Pearson correlations.

Table 3. Parameter sets of optimal and mean CV.

	Inlet angle α [°]	Inlet height h [mm]	Inlet diameter d [mm]	Quantity of inlets $N_{\text{fuel\_inlets}}$	CV
Optimal design	40	10	1.53	6	0.0098
Design with mean CV	65	18.3	1.29	4	0.0204

Figure 4 shows a contour plot of the equivalence ratio at the outlet plane of the burner of both designs. The optimized design exhibits a uniform distribution with small local gradients or wall accumulations, whereas the comparison case shows steep gradients of equivalence ratio and in general higher values at the wall. The equivalence ratio in the optimized configuration is bounded between 0.5854 and 0.6073, indicating a higher mixture homogeneity compared to the mean CV value case.

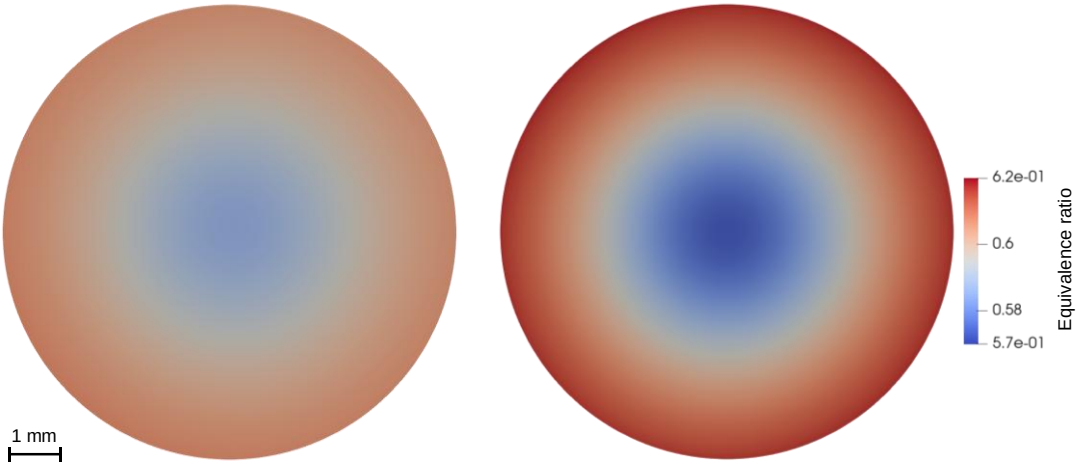

Figure 4. Equivalence ratio at the outlet of optimized (left) and mean CV design (right).

Figure 5 compares the radial profiles of equivalence ratio from the center (radius = 0 mm) to the wall (radius = 4.5 mm) for the optimized design (dashed line) and a mean CV reference design (solid line). These graphs highlight the differences in gradient behavior: in the optimized design, the gradient is shallower throughout the cross-section compared to the mean CV design. In the mean CV case the gradient increases in the near-wall region compared to the optimized case where the gradient decreases near the wall.

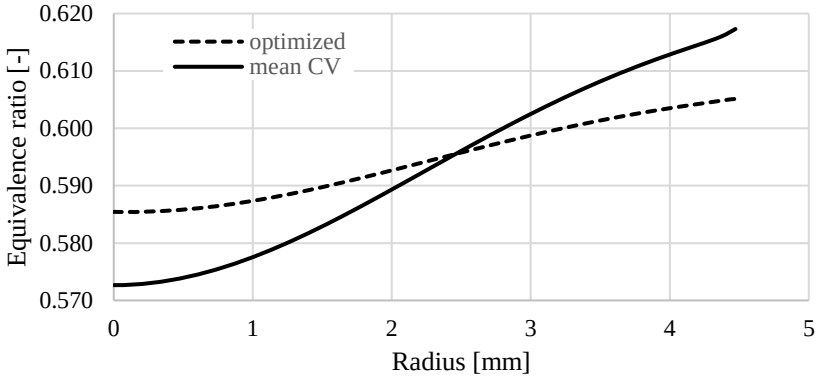

Figure 5. Plot of Equivalence ratio at the outlet of optimized and mean CV design.

In contrast to the surrogate-based approach employed in this study, a full-factorial design of experiments with equivalent variable limits would require 576 simulations to cover the design space. This results from discretising the inlet diameter into six steps of 0.1 mm, the number of inlets from 3 to 6 in integer increments, the inlet angle in 10° steps from 40° to 70°, and the inlet height in 2 mm steps between 10 mm and 20 mm. Such a grid ensures full parameter coverage but relies on relatively coarse resolution for certain variables, potentially overlooking local sensitivities. The surrogate-based optimization, in contrast, reduced the number of required simulations to 72, enabling a more efficient and adaptive exploration of the design space. The SBO framework reduces the computational effort by an estimated 87.5 %, while still enabling the identification of high-performance design candidates.

Discussion

The results demonstrate that the surrogate-based optimization framework can effectively identify burner geometries with improved mixing homogeneity. The optimized configuration exhibits a uniform, low-gradient profile of equivalence ratio across the outlet plane, in contrast to the high-gradient zones observed in the mean CV reference design. This uniform mixing mitigates two critical combustion risks: first, it prevents a lean core with a locally hydrogen-rich boundary layer, which promotes high equivalence-ratio values at the walls and increases flashback risk; second, it reduces localized hot spots caused by uneven equivalence-ratio distributions, thereby decreasing thermal NO_X formation. By achieving a tight equivalence-ratio band with small radial gradients, the optimized design reduces both hazards.

The sampling behavior of the EGO algorithm further underscores its capability to balance exploration and exploitation. Early iterations introduce points in underrepresented inlet-height regions and across a broad range of diameters and inlet counts, reflecting an exploratory phase to map the global design space. Subsequent clustering of samples at larger diameters, higher inlet numbers, and lower angles shows the exploitation of regions predicted by the EIF to yield the improvements in mixing performance. The leading-order effects driving improved mixing performance are primarily associated with larger inlet diameters (Pearson correlation r = -0.28) and smaller inlet angles (r = 0.21), which reduce CV values, while smaller inlet heights (r = 0.19) also contribute by providing a longer mixing length that allows more time for homogenization. In combination with the higher quantity of inlets, despite the low Pearson correlation, lower inlet velocities for hydrogen produce more homogeneous mixtures.

Despite these successes, the framework is subject to several constraints. All key metrics are derived from temporally averaged RANS data, which can obscure localized or transient peaks in equivalence ratio that are critical for flashback initiation and hot-spot-driven NO_x formation and fail to capture time-dependent mixing that could significantly influence performance in practice. Moreover, the fact that the optimal solution lies at the lower limits of inlet angle (40°) and inlet height (10 mm) and at the upper limits of inlet diameter (1.53 mm) and number of inlets (6) suggests that further improvements may exist beyond the chosen parameter ranges. Although the EGO algorithm is designed to efficiently search for global optima, a guarantee for global optimum cannot be provided due to the probabilistic nature of the surrogate model, limited sample budget, and model inaccuracies. While the MAPE of 16.76% indicates a high percentage error, the absolute error remains small given the low CV values. Additionally, the

limited number of EGO samples may not yet fully cover all regions of the design space, potentially missing other high-performing configurations. Finally, the high MAPE and low Pearson correlation observed for the quantity of inlets indicate nonlinear interactions, that the current GP surrogate, particularly given its continuous treatment of a discrete variable, may not fully resolve.

Conclusion and Outlook

This work establishes a surrogate-based optimization framework that effectively integrates parametric geometry generation, steady-state RANS CFD simulations, and an EGO algorithm to systematically explore and optimize hydrogen burner configurations for enhanced mixing homogeneity. While the identified optimal design demonstrates promising performance, achieving a coefficient of variation of 0.0098 (a 52.2 % improvement over the dataset mean), the primary contribution lies in the development of the methodology itself, as the optimized result remains subject to uncertainties inherent to the simulation and modeling assumptions. By automating the exploration of a four-dimensional design space, the framework reduces the required number of high-fidelity simulations from 576 in a full-factorial grid to 72, corresponding to an 87.5 % decrease in computational expense. The optimal design achieves a coefficient of variation of 0.0098, a 52.2 % improvement over the dataset mean, while respecting additive-manufacturing constraints and avoiding problematic overhangs or narrow channels.

Beyond demonstrating substantial gains in mixing performance, the sampling of the EGO algorithm revealed systematic tendencies: early exploration addressed sparsely sampled inlet heights and diameters, whereas later exploitation concentrated on large diameters, high inlet counts, and low angles. This behavior not only demonstrates the EIF balance between exploration and exploitation but also highlights key sensitivity trends, which is information that would remain unexplored in purely manual or grid-based approaches. The surrogate's MAPE reduction from 19.9 % to 16.8 % further indicates that iterative refinement effectively increases predictive accuracy in regions of interest, despite the continuous approximation of discrete variables.

In future work, the framework may be extended and strengthened through several approaches. First, reactive CFD simulations that capture local equivalence-ratio peaks which indicate flashback risks can be incorporated. Second, further investigation into alternative optimization algorithms, as the EGO algorithm does not guarantee a global optimum necessitating exploration of other methods. Third, experimental validation of the optimized geometries manufactured in PBF-LB/M will quantify real-world performance and inform any necessary modifications to the surrogate model or objective definition. Third, expanding or refining the design limits (e.g., angles < 40°, heights < 10 mm, diameters > 1.6 mm, and higher inlet counts) may reveal further improvements, provided manufacturing feasibility is ensured and mesh stability is maintained. To better resolve discrete and strongly nonlinear interactions, the surrogate strategy can be enhanced through mixed-integer GP or multi-fidelity modeling that accommodates integer variables natively. Finally, extending the framework to multi-objective optimization (e.g., balancing mixing homogeneity, NO_x formation, pressure loss, and flashback margins) will support the holistic design of hydrogen burners tailored to industrial applications.

By pursuing these developments, the proposed methodology can evolve into a comprehensive toolchain for the rapid, reliable, and manufacturability-aware design of next-generation hydrogen combustion systems.

Acknowledgement

This work has been funded by the Deutsche Forschungsgemeinschaft, Germany (DFG, German Research Foundation) under the IGNITER project within the DFG-SPP 2419 HyCAM (project number: 523874889). Computations were performed with computing resources granted by RWTH Aachen University under project rwth1837.

Literatur

[1] Europäische Kommission: *Umsetzung des europäischen Grünen Deals:* Auf dem Weg zu einem klimaneutralen Europa bis 2050, Brüssel, 2021.

[2] Europäische Kommission: *Mitteilung der Kommission an das Europäische Parlament, den Europäischen Rat, den Rat, den Europäischen Wirtschafts- und Sozialausschuss und den Ausschuss der Regionen:* REPowerEU-Plan, Brüssel, 2022.

[3] Aniello A., Poinsot T., Selle L. und Schuller T.: *Hydrogen substitution of natural-gas in premixed burners and implications for blow-off and flashback limits* International Journal of Hydrogen Energy, 47, S. 33067–33081, 2022.
[4] Fruzza F., Lamioni R., Tognotti L. und Galletti C.: *Flashback of H2-enriched premixed flames in perforated burners: Numerical prediction of critical velocity* International Journal of Hydrogen Energy, 48, S. 31790–31801, 2023.
[5] Law C. K., Makino A. und Lu T. F.: *On the off-stoichiometric peaking of adiabatic flame temperature* Combustion and Flame, 145, S. 808–819, 2006.
[6] Choudhuri A. R. und Gollahalli, SR: *Characteristics of hydrogen–hydrocarbon composite fuel turbulent jet flames* International Journal of Hydrogen Energy, 28, S. 445–454, 2003.
[7] Zhang L., Wu C., Zhang J., Zhang B. und Sui C.: *Numerical Simulation of the Combustion Characteristics in a Flue Gas Internal Recirculation Burner* ACS omega, 7, S. 42264–42271, 2022.
[8] Pignatelli F.: *Impact of pilot flame and hydrogen enrichment on turbulent methane/hydrogen/air swirling premixed flames in a model gas turbine combustor* Experimental Thermal and Fluid Science, 152, S. 111124, 2024.
[9] Singh G., Schreiner B. D. J., Sun X. und Sethi V.: *A review of hydrogen micromix combustion technologies for gas turbine applications* International Journal of Hydrogen Energy, 127, S. 295–310, 2025.
[10] York W. D., Ziminsky W. S. und Yilmaz E.: *Development and testing of a low NOx hydrogen combustion system for heavy-duty gas turbines* Journal of Engineering for Gas Turbines and Power, 135, S. 22001, 2013.
[11] Kim Y., Kwon H. und Lee M. J.: *Mitigating NOx Emissions in Hydrogen-Fueled Industrial Burners Through Residence Time Control* Case Studies in Thermal Engineering, S. 106523, 2025.
[12] Andreas Gebhardt, Julia Kessler und Laura Thurn, *3D Printing - Understanding Additive Manufacturing*. Hanser Publishers, Munich, ISBN: 978-1-56990-702-3, 2019.
[13] *DIN EN ISO/ASTM 52911-1:2020-05: Additive manufacturing – Design – Part 1: Laser-based powder bed fusion of metals (ISO/ASTM52911-1:2019)*, 52911-1, DIN Deutsches Institut für Normung e. V., Berlin, Mai. 2020.
[14] Reumschüssel J. M., Saldern J. G. R. von, Ćosić B. und Paschereit C. O.: *Data-driven optimization of a gas turbine combustor: A Bayesian approach addressing NO x emissions, lean extinction limits, and thermoacoustic stability* Data-Centric Engineering, 5, 2024.
[15] Willkomm J., Weller V., Bredow A. von, Putz C. und Schleifenbaum J. H.: *Design Automation: Term Definition and Methodical Analysis for Software Selection* RTe Journal, 2025.
[16] Angelo, C., Carla, D., Mariangela, L., Giuseppe, S., & Domenico, L.: *Knowledge Based Engineering for Hydrogen Gas Turbines and Burners Design: A review*, 2022.
[17] Serani A., Scholcz T. P. und Vanzi V.: *A Scoping Review on Simulation-Based Design Optimization in Marine Engineering: Trends, Best Practices, and Gaps* Archives of computational methods in engineering : state of the art reviews, 31, S. 4709–4737, 2024.
[18] Liu X.: *Surrogate-assisted global and distributed local collaborative optimization algorithm for expensive constrained optimization problems* Scientific reports, 15, S. 1728, 2025.
[19] Gauger N. R. *et al., Advances in Evolutionary and Deterministic Methods for Design, Optimization and Control in Engineering and Sciences*. Springer International Publishing; Imprint: Springer, Cham, ISBN: 9783319899886, 2019.
[20] Jiang P., *Surrogate Model-Based Engineering Design and Optimization*. Springer Singapore Pte. Limited, Singapore, ISBN: 9789811507311, 2020.
[21] Jones D. R., Schonlau M. und Welch W. J.: *Efficient Global Optimization of Expensive Black-Box Functions* Journal ofGlobal Optimization, S. 455–492, 1998.
[22] CEA, EDF, Open Cascade, *SALOME PLATFORM*, 9.13, opensource, Rocky 9.6, https://www.salome-platform.org/
[23] Sandia National Laboratories, *Dakota*, 6.22, opensource, Rocky 9.6, https://dakota.sandia.gov/
[24] OpenCFD Ltd, *OpenFOAM*, v2312, opensource, Rocky 9.6, https://www.openfoam.com/
[25] Kim S. und Kim H.: *A new metric of absolute percentage error for intermittent demand forecasts* International Journal of Forecasting, 32, S. 669–679, 2016.