

Distributionally Robust Bayesian Optimization For Two-Phase Convergent-Divergent Nozzle Shape Design

Matteo Battisti¹, Ahmed Harbi², Andrea Cioffi³, Roberto Agromayor³, Fredrik Haglind³,
 Giacomo Persico², Anne Barros¹, Yiping Fang¹

1. Chair on Risk and Resilience of Complex Systems, Laboratoire Génie Industriel, CentraleSupélec, Université Paris-Saclay, France.

2. Laboratory of Fluid Machines, Department of Energy, Politecnico di Milano, Via Lambruschini 4a, Milano.

3. Department of Civil and Mechanical Engineering, Technical University of Denmark Koppels Allé, Building 404, 2800 Kgs. Lyngby, Denmark.

E-mail: matteo.battisti@centralesupelec.fr; ahmed.harbi@polimi.it; ancio@dtu.dk; roagr@dtu.dk, fhag@dtu.dk, giacomo.persico@polimi.it, anne.barros@centralesupelec.fr, yiping.fang@centralesupelec.fr

Two-phase convergent-divergent nozzles play a critical role in heating, ventilation, air conditioning, and refrigeration systems. Designing efficient nozzle geometries is challenging due to geometrical, operational, and epistemic uncertainties arising from manufacturing tolerances, variable operating conditions, and incomplete understanding of the underlying physics. These uncertainties limit the effectiveness of conventional nozzle design methods. This study proposes a Distributionally Robust Bayesian Optimization (DRBO) approach for the shape optimization of a convergent-divergent nozzle coupled with homogeneous equilibrium flow models. To address distributional uncertainty during adaptive sampling, ambiguity sets are characterized using maximum mean discrepancy (MMD) and Wasserstein distance metrics. The study compares the motive nozzle isentropic efficiency obtained using General DRBO approach, where the uncertainty reference distribution is provided to the method, and the Data-Driven DRBO approach, where the uncertainty reference distribution is updated using data from previous iterations. Both approaches are evaluated for the shape optimization of the motive nozzle in an ejector refrigeration system using CO₂ as the working fluid. The results suggest that the robust optimized nozzle achieves higher isentropic efficiency compared to the nominal and non-robust optimized designs under uncertain conditions.

Keywords: Distributionally Robust Bayesian Optimization, Two-phase flow, Convergent–divergent nozzles, Computational fluid dynamics

1. Introduction

The urgent need to reduce greenhouse gasses has become increasingly clear in recent years; for this reason the European Commission committed to making the EU climate neutral by 2050 (DG, 2019). To support this transition, it is essential to improve the use of renewable energy sources and the efficiency of energy conversion systems, which require significant research and innovation in clean technologies.

In this context, Heating, Ventilation, Air Conditioning, and Refrigeration (HVAC&R) practices represent a major opportunity for emissions reduction through improved efficiency and sustainable working fluids. An important step in achieving carbon neutrality is the transition from high global warming potential (GWP) working fluids, to en-

vironmentally friendly refrigerants, which are advantageous because of their availability, low cost, and minimal environmental impact (Balali et al., 2023). Among natural refrigerants, CO₂ is recognized as an efficient long-term solution due to its negligible GWP, non-flammability, non-toxicity, and favorable thermophysical properties. Another advantage of CO₂ is that its thermodynamic cycle allows for smaller and more compact components, making it beneficial to various applications (Kim et al., 2004).

A fundamental component for the overall performance of CO₂ HVAC&R systems is the expansion device. Traditional expansion devices, like simple expansion valve, do not recover any of the supplied mechanical energy and therefore impose a significant energy loss. To address this issue,

alternative expansion technologies have been explored, like two-phase ejectors, which can recover some of the energy lost during the expansion process (Elbel and Lawrence, 2016).

Despite its advantages, the complexity of two-phase ejector flows presents challenges in achieving reliable and efficient designs. Their performances are highly dependent on an appropriate motive nozzle design, since flow expansion and phase change occur entirely through the nozzle. However a limited understanding of non-equilibrium two-phase flow phenomena reduce the predictive accuracy of the flow models and ejector performances are sensitive to inlet operating conditions and geometry variations, causing a efficiency reduction in off-design condition. Conventional deterministic optimization approaches, as well as standard stochastic optimization methods, typically assume that these uncertainties follow a known and fixed probability distribution. This assumption is rarely satisfied in two-phase CO₂ flows, because of the lack of data or epistemic uncertainty on the flow models.

To the best of the authors' knowledge, no previous studies have addressed the optimization of CO₂ motive nozzles under uncertainty, nor has robust optimization been systematically applied to CO₂ ejector systems, representing a significant gap in the literature on ejector refrigeration technologies. To address this gap, this paper explores the application of Distributionally Robust Bayesian Optimization (DRBO) (Kirschner et al., 2020) techniques to address these challenges. Bayesian optimization is a model-based sequential optimization framework for expensive-to-evaluate black-box functions, successfully applied in various fields (Shahriari et al., 2015).

The objective is to showcase the benefit of the application of DRBO in the robust design optimization of a two-phase ejector nozzle for CO₂ heat pump applications ultimately contributing to the design of robust ejectors accounting for uncertainties in two-phase flow behavior.

The rest of the paper is structured as follows: Section 2 summarizes the problem description, Section 3 describes the computational framework, Section 4 presents the results, and Section 5 con-

cludes with a summary of findings and suggestions for future work.

2. Problem description

In this section, we present the formulation of the ejector nozzle design problem.

2.1. Uncertainties formulation

The problem presents two main sources of uncertainties: geometrical and operational. Geometric uncertainties arise from manufacturing tolerances and imperfection or damage to components caused by harsh working conditions. Operational uncertainties, on the other hand, arise from the variation of the inlet condition of the nozzle due to external environmental changes. This variability propagates within the component, causing changes in expected performances (Razaaly et al., 2019).

Characterization of these uncertainties is challenging, especially for geometric ones. In fact, collecting a sufficiently large sample of actually manufactured components to build a statistically converged data set can be time consuming. To solve this problem, we decided to follow the Wunsch and Hirsch (2021) approach by deriving input uncertainties from ISO 2768-1, 1989 (Tolerances, 1989). The variability ranges of operational conditions, pressure and temperature at the inlet, are based on reasonable values derived from the little information available on the topic in the literature (Ringstad et al., 2020)

Due to the just mentioned lack of detailed statistical data, we adopt a independent uniform distributions as a modeling assumption of the uncertainties. However, these distributions should not be interpreted as accurate representations of the true uncertainty, which is unknown and potentially deviates from the assumed one. Precisely because of this, we employ the DRBO framework to ensure that the optimized design remains robust against deviations from our assumption. The range of geometric tolerances are reported in Table 1, while the operational ranges are in Table 2.

The uncertainty model is formulated to represent realistic manufacturing and operational constraints; however, while the motive nozzle is

Table 1. Nominal values and tolerances for geometry variable

Name	Nominal (mm)	Tol (mm)
Throat Height	0.24	$\pm 2.00 \cdot 10^{-2}$
Outlet Height	0.39	$\pm 2.00 \cdot 10^{-2}$
Inlet Height	10.00	± 1.00
Divergent Length	56.15	± 1.50
Convergent Length	27.35	± 1.50
Width	3.00	± 0.50

Table 2. Nominal values and tolerances for inlet operating conditions.

Name	Nominal	Tol
T_{in} (K)	310.15	± 4.00
P_{in} (MPa)	9.10	± 0.50

treated as a representative case study, the objective of this work is not limited to optimizing a single technical configuration. The proposed framework is intended to be generalizable and applicable to other turbomachinery components facing similar challenges, where robust design has evident technical relevance.

2.2. Problem formulation

As explained in Section 1, the nozzle is a critical component of the two-phase ejector system, responsible for accelerating refrigerant flow and facilitating phase change. The design of the nozzle significantly impacts the overall performance and efficiency of the ejector system. As a test case, we consider the Nakagawa et al. (2009) nozzle. It is a well-known two-phase ejector nozzle described in the literature of CO₂ refrigeration systems. The geometry of the nozzle is defined by a set of design variables: inlet diameter, outlet diameter, convergent and divergent length, throat diameter, and width of the nozzle. The nominal values are reported in Table 1. The performance of the nozzle is evaluated on the basis of the isentropic efficiency metric, which is defined as:

$$\eta = \frac{(h_{01} - h_2)}{(h_{01} - h_{2s})} \quad (1)$$

where h_{01} is the total enthalpy at the nozzle inlet, h_2 is the actual enthalpy at the nozzle outlet, and h_{2s} is the enthalpy at the nozzle outlet for an isentropic expansion process.

In an ejector system under any operating condition, the nozzle design must be able to provide the required refrigerant mass flow rate for the cycle. To ensure this, we define a constraint function to minimize the deviation of the mass flow rate of the nozzle, $\dot{m}$, from its reference value, $\dot{m}_{ref}$, under any operating conditions. We use the Absolute Percentage Error (APE) as a metric to represent the overall mass flow rate deviation. The APE is defined as:

$$APE_{mass} = \frac{|\dot{m} - \dot{m}_{ref}|}{\dot{m}_{ref}} \times 100\% \quad (2)$$

We impose $APE_{mass} < 3\%$ to restrict the deviation.

3. Computational Framework

This section presents the one-dimensional ejector nozzle model, the surrogate modeling approach, and the optimization algorithm for the robust design problem.

3.1. 1D HEM formulation

To predict the performance of the nozzle, we implemented a one-dimensional (1D) Homogeneous Equilibrium Model (HEM). When phase change occurs within the nozzle, the resulting two-phase flow is modeled by assuming that the two phases are in mechanical and thermodynamic equilibrium in every section, with a single velocity, pressure, temperature and chemical potential describing the fluid mixture.

The governing conservation equations for mass, momentum, and energy for the mixture are solved along the axial direction of the nozzle. Spatial discretization is performed using a finite-volume approach. The source terms account for the geometry variation and losses due to viscous friction:

$$\frac{\Delta(\rho v A)}{\Delta z} = 0 \quad (3)$$

$$\frac{\Delta[(p + \rho v^2)A]}{\Delta z} = \frac{p A_w \sin \theta}{\Delta z} - \frac{\tau_w A_w \cos \theta}{\Delta z} \quad (4)$$

$$\frac{\Delta [\rho v A (h + v^2/2)]}{\Delta z} = 0 \quad (5)$$

where ρ is the volume-averaged mixture density, v is mixture velocity, p is mixture pressure, h is static enthalpy, A is the cross-sectional area, A_w is wall wetted area, θ is the opening angle of the nozzle, and Δz is the axial distancing between the control volume.

The model was validated under both subcooled and supercritical inlet conditions using two cases from Nakagawa et al. (2009): case b at an inlet pressure of 6.1 MPa and case b at an inlet pressure of 9.1 MPa. Using the one-dimensional shear stress model proposed by Richardson (1958), we obtain an average Normalized Mean Absolute Error (NMAE) of 17.7% for the static pressure distribution predictions. This is the same result for 1D HEM models with Richardson in (Tammone et al., 2024).

The flow model is implemented in Python, facilitating its integration with optimization algorithms and surrogate modeling techniques.

3.2. Surrogate modeling

The model described in the previous section is computationally efficient compared to 2D or 3D CFD simulations. However the computational cost arises from the high number of evaluations required by the DRBO framework. In particular, the need to repeatedly evaluate performance under varying uncertainty realizations and to identify worst-case distributions within the ambiguity set, significantly increases the total number of model queries.

To mitigate this issue, a surrogate metamodel is utilized to approximate the objective function. Furthermore, this approach ensures the scalability of the framework, making it suitable even when higher-fidelity models are considered. This ensure it can be seamlessly extended to 2D or 3D CFD simulations in future developments.

In this study we chose as surrogate model a Gaussian Process Regression (GPR) for its ability to accurately capture the complexity of multi-dimensional response functions (Duvenaud, 2014). A GPR is a distribution over functions such that any finite set of function values

$f(x_1), f(x_2), \dots, f(x_N)$ has a joint Gaussian distribution (Seeger, 2004). A GPR model, before data acquisition, is specified by its mean function $\mu(x)$ and the covariance function, also called kernel.

The structure that can be captured is entirely determined by its kernel. To better capture the data structures, we decided to implement the automatic model construction method proposed by (Duvenaud, 2014). The method defines a grammar to build kernels and then searches the kernel structure that best fits the data. To accomplish this task, a greedy algorithm is used to iteratively add or remove kernels based on their Bayesian Information Criterion (BIC):

$$BIC(M) = \log p(D|M) - \frac{1}{2}|M|\log N \quad (6)$$

where $p(D|M)$ is the marginal likelihood with optimal kernel hyperparameters, $|M|$ is the number of kernel parameters, and N is the number of data points. The reader is referred to (Duvenaud, 2014) for a detailed description of the grammar construction.

Our implementation of the surrogate model is based on the `GPFlow` (Matthews et al., 2017) python library, which provides a flexible framework for building and training Gaussian Process models in Python.

3.3. Optimization algorithm

Given the uncertain nature of the distributions of the problem and the computational cost associated with the high number of fluid simulations, we employ the Distributionally Robust Bayesian Optimization (DRBO) algorithm introduced by Kirschner et al. (2020), to obtain a nozzle design that remains optimal under the worst-case distribution of uncertain parameters.

The DRBO process aims to optimize a black-box objective function $f : \mathcal{X} \times \mathcal{C} \rightarrow \mathbb{R}$ over the set $\mathcal{X} \subset \mathbb{R}^m$ of the design variable x and the set $\mathcal{C} \subset \mathbb{R}^n$ of uncontrolled contexts c , which in our case is described in Section 2.1. The method can be described as a Bayesian Optimization (BO) process that uses the solution to a Distributionally Robust Optimization (DRO) problem as its acquisition function solution. Next, we briefly describe

the BO and DRO acquisition function. More details can be found in (Kirschner et al., 2020; Tay et al., 2022; Kuhn et al., 2019).

The Bayesian Optimization algorithm seeks the global maximum of the black-box function f by sequentially querying f through observations $y_t := f(x_t, c_t)$ for a certain number of iterations $t = 1, \dots, T$. The selection of the next query point is guided by a so called acquisition function. Each query to the surrogate is considered costly, meaning that the maximum should be found in as few queries as possible (Frazier, 2018). The black box-function is modeled using a GPR. The GPR is updated at each iteration with the new query points.

As an acquisition function, the DRO approach aims to track the optimal solution that maximizes the expected function value under the worst-case distribution of an external contextual parameter:

$$\max_{x \in \mathcal{X}} \inf_{q_t \in U_t} \mathbb{E}_{c \sim q_t} [f(x, c)] \quad (7)$$

Here, the uncontrolled context c is considered a random variable that is sampled from a distribution $p_t \in U_t \subset \mathbb{R}^{|C|}$. The nominal distribution p_t is a prior knowledge of the distribution of c . However, since we do not have a clear indication about the true distribution of the uncertainties, the DRO procedure considers the setting where c is subject to distribution shift. This means that there exists an unknown true distribution p_t^* such that $c \sim p_t^*$. In Eq. (7) this is modeled by U_t which is called ambiguity set and it is defined to contain all possible p_t^* . To handle distributional shifts, we define an ambiguity set U_t as a ball of radius ϵ_t around the nominal distribution p_t , such that $U_t = \{q \in \mathcal{P}(\mathcal{C}) : d(p_t, q) \leq \epsilon_t, \mathbf{1}^\top q = 1\}$. The DRO acquisition function then identifies the design x that maximizes performance under the most unfavorable distribution q within this ambiguity set.

The definition of the nominal distribution p_t plays a important role in the optimization strategy. Depending on how the nominal distribution p_t is defined, two different DRBO approaches can be formulated: General DRBO and Data Driven DRBO (Kirschner et al., 2020). In the General DRBO the nominal distribution is defined at the

beginning of the optimization process and never changed through out the process, while in Data Driven DRBO the nominal distribution is updated dynamically according to the new (c) contexts observed at each iteration.

In addition to the choice of the nominal distribution update strategy, the definition of the ambiguity set itself requires selecting a suitable probability metric. In this study, we consider maximum mean discrepancy (MMD) (Muandet et al., 2017) and Wasserstein distance (Kuhn et al., 2019) as distribution distances obtaining four different DRBO settings: Wasserstein General (W-General), Wasserstein Data Driven (W-Data Driven), MMD General (MMD-General), MMD Data Driven (MMD-Data Driven).

The DRBO algorithm is implemented in Python drawing on the material provided by (Tay et al., 2022). Gaussian process regression was performed using the GPFLOW library (Matthews et al., 2017), convex optimization tasks were handled with CVXPY (Diamond and Boyd, 2016), while the main Bayesian Optimization loop was implemented using Trieste python library (Picheny et al., 2023).

4. Results

In this section, we present the results obtained from the implementation of the DRBO algorithm.

4.1. Surrogate model performance

We decided to build a surrogate model using GPR to approximate the 1D HEM model described in Section 3.1. We built two distinct surrogate models, one to predict the isentropic efficiency and another one to predict the mass flow rate. The first is used to return the observation of the objective function, while the second is used to evaluate the constraint of the optimization problem. For both models, we followed the same procedure for data generation and model training.

The initial Design of Experiments (DoE) was made up of 100 samples generated using Latin Hypercube Sampling (LHS) to ensure well-distributed coverage of the input space. Each sample needs to be evaluated using the 1D HEM model presented in Section 3.1 to obtain the isen-

tropic efficiency and mass flow rate values. However, such data set presents the issue of containing too many sample points with low efficiency values, making it difficult for the surrogate model to accurately learn the relationship between inputs and outputs in high efficiency regions. To address this, an adaptive sampling strategy was employed, where the initial GPR model was trained on the original data set and then used to identify and sample additional points in regions of high efficiency. This was possible by imposing a lower bound constraint that forced the sampler to select points with an efficiency greater than 40%. The initial DoE was iteratively infilled until a total of 180 samples were collected, ensuring a more balanced representation of the input-output relationship across the entire efficiency spectrum. The kernel of the model was learned along the other hyperparameters applying the method described in Section 3.2 to the initial DoE and the final data set.

The performance of the surrogate models was then evaluated using a new LHS test set of 100 samples. We evaluated the performance of the models using the coefficient of determination (R^2) and the Mean Absolute Error (MAE). The results shows a R^2 and MAE scores of 0.85 and 0.043 for isentropic efficiency and 0.99 and $8.7 \cdot 10^{-4}$ for mass flow rate. Since the isentropic efficiency is bounded in the interval $[0, 1]$ and the nominal mass flow rate of the Nakagawa nozzle is 2.3×10^{-2} , we considered these error levels sufficient to proceed with the subsequent optimization procedure in this current work. The investigation of advanced adaptive sampling strategies to further improve model fidelity is left for future work.

4.2. Optimization Results

The results of the General and Data Driven DRBO algorithm with both the Wasserstein (W-General, W-Data Driven) and MMD (MMD-General, MMD-Data Driven) distances are presented in this section. They are compared against the nominal geometry of Nakagawa and a non-robust optimized geometry (O-NR), which is obtained as the solution of a BO loop with the Expected Constrained Improvement (ECI) acquisi-

tion function (Gardner et al., 2014). The designs are first compared at the nominal conditions, and then an uncertainty quantification (UQ) analysis is conducted on the nominal optimal profiles to assess and compare their performances. The analysis is carried out using the 1D HEM model and the uncertain parameters are sampled through an LHS of 100 points.

Nominal conditions: The nominal Nakagawa geometry at nominal operating conditions yields an efficiency of 48.3%, while the O-NR design yields a value of 50.7%. The Wasserstein robust optimization results show efficiencies of 66.8% and 53.2% in the General and Data Driven approaches, respectively. While the MMD robust optimized designs show efficiencies of 53.7% and 59.4% in the General and Data Driven approaches, respectively. We notice that the robust optimization approach outperforms both the nominal and non-robust optimized design. The profiles of the nominal designs are shown in Table 3.

The results suggest that the highest performances are achieved by the shortest nozzles. This is due to the accumulation of frictional losses along the nozzle wall, that cause shorter nozzles to operate more efficiently by minimizing the wetted surface area. In fact the primary determinant of nozzle performance is found to be wall friction, which reduces both the maximum achievable mass flow rate and the total pressure of the flow. This total pressure loss arises because the mechanical energy dissipated by friction is converted irreversibly into heat and cannot be recovered as useful work. In addition to these effects, wall friction also delays the near wall flow, reducing the effective cross-sectional area and displacing the location of maximum velocity downstream of the geometric throat as seen for all the cases.

Uncertain conditions: the UQ analysis is conducted by running 100 simulations for each design. The results of the analysis are presented in Figure 1, which shows the Probability Density Function (PDF) of the efficiency for the nozzle design solutions. The optimization results indicate that mean efficiency values of the robust optimized results are better than those of the O-NR and the Nominal geometry. Similarly, we observe




Fig. 1. PDF comparison between the robust optimized and baseline profiles.

Table 3. Nominal designs

	Throat Height	Outlet Height	Divergent Length	Convergent Length	Inlet Height	Nozzle Width
Nominal	1.20e-04	1.95e-04	5.61e-02	2.74e-02	5.00e-03	3.00e-03
O-NR	1.44e-04	1.97e-04	5.73e-02	2.97e-02	5.67e-03	2.43e-03
MMD-Data Driven	1.43e-04	1.95e-04	4.06e-02	2.93e-02	5.54e-03	2.36e-03
MMD-General	1.44e-04	1.95e-04	5.05e-02	3.01e-02	6.00e-03	2.48e-03
W-Data Driven	1.44e-04	1.95e-04	3.35e-02	2.99e-02	2.82e-03	2.47e-03
W-General	1.44e-04	1.95e-04	3.09e-02	3.01e-02	2.75e-03	2.08e-03

Table 4. Efficiency statistics for the optimized and nominal Nakagawa designs

Design	Mean(%)	5%-q(%)	95%-q(%)
Nominal	46.99	42.43	50.62
O-NR	50.35	47.07	53.67
W-Data Driven	54.83	48.59	61.70
W-General	56.67	48.59	65.33
MMD-Data Driven	53.36	50.19	57.61
MMD-General	51.26	47.84	54.22

that the 5% and 95% quantiles for the nominal and non-robust designs are worse than those of the robust solutions. The numerical values are reported in Table 4. These results highlight the effectiveness of the DRBO approach in enhancing the robustness of the nozzle design against uncertainties under operating conditions, for both distances, with the best performance achieved by Wasserstein distributional shift characterization.

5. Conclusions and Future Work

In this study we applied the DRBO framework to an optimization of a two-phase ejector nozzle achieving consistent improvements in isentropic efficiency across uncertain operating conditions and geometry imperfections. Overall, all robustly optimized designs outperform the nominal and O-NR design in terms of mean efficiency and quantiles.

This is a preliminary study, and future work can explore several extensions in greater depth. First, higher-fidelity CFD simulators can be included in the framework to better capture the complex fluid dynamics of turbomachinery components. While the current study uses a uniform nominal distribution as a reference, future work could explore non-parametric or data-derived nominal distributions to further refine the ambiguity sets. Future research could also extend the DRBO approach to multi-objective optimization problems.

Acknowledgement

The research presented in this paper has received funding from the Training42Phase Marie Skłodowska-Curie Actions Doctoral Network (MSCA-DN) grant agreement no. 101120366. The financial support is gratefully acknowledged. Disclaimer: Funded by the European Union. Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible.

References

- Balali, A., A. Yunusa-Kaltungo, and R. Edwards (2023). A systematic review of passive energy consumption optimisation strategy selection for buildings through multiple criteria decision-making techniques. *Renewable and Sustainable Energy Reviews* 171, 113013.
- DG, E. C. C. A. (2019). Going Climate-neutral by 2050: A Strategic Long-term Vision for a Prosperous, Modern, Competitive and Climate-neutral EU Economy. Publications Office of the European Union.
- Diamond, S. and S. Boyd (2016). Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17(83), 1–5.
- Duvenaud, D. (2014). Automatic model construction with Gaussian processes. Ph. D. thesis.
- Elbel, S. and N. Lawrence (2016). Review of recent developments in advanced ejector technology. *International journal of refrigeration* 62, 1–18.
- Frazier, P. I. (2018). Bayesian optimization. In *Recent advances in optimization and modeling of contemporary problems*, pp. 255–278. Informs.
- Gardner, J. R., M. J. Kusner, Z. E. Xu, K. Q. Weinberger, and J. P. Cunningham (2014). Bayesian optimization with inequality constraints. In *ICML*, Volume 2014, pp. 937–945.
- Kim, M.-H., J. Pettersen, and C. W. Bullard (2004). Fundamental process and system design issues in CO₂ vapor compression systems. *Progress in energy and combustion science* 30(2), 119–174.
- Kirschner, J., I. Bogunovic, S. Jegelka, and A. Krause (2020). Distributionally robust bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 2174–2184. PMLR.
- Kuhn, D., P. M. Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh (2019). Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations research & management science in the age of analytics*, pp. 130–166. Informs.
- Matthews, A. G. d. G., M. Van Der Wilk, T. Nickson, K. Fujii, A. Boukouvalas, P. León-Villagrà, Z. Ghahramani, and J. Hensman (2017). Gpflow: A gaussian process library using tensorflow. *Journal of Machine Learning Research* 18(40), 1–6.
- Muandet, K., K. Fukumizu, B. Sriperumbudur, and B. Schölkopf (2017). Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning* 10(1-2), 1–141.
- Nakagawa, M., M. S. Berana, and A. Kishine (2009). Supersonic two-phase flow of CO₂ through converging-diverging nozzles for the ejector refrigeration cycle. *International journal of refrigeration* 32(6), 1195–1202.
- Picheny, V., J. Berkeley, H. B. Moss, H. Stojic, U. Granta, S. W. Ober, A. Artemev, K. Ghani, A. Goodall, A. Paleyes, et al. (2023). Trieste: Efficiently exploring the depths of black-box functions with tensorflow. *arXiv preprint arXiv:2302.08436*.
- Razaaly, N., G. Persico, and P. M. Congedo (2019). Impact of geometric, operational, and model uncertainties on the non-ideal flow through a supersonic ORC turbine cascade. *Energy* 169, 213–227.
- Richardson, B. L. (1958). Some problems in horizontal two-phase tow-component flow (thesis). Technical report, Argonne National Lab., Lemont, Ill.
- Ringstad, K. E., Y. Allouche, P. Gullo, Å. Ervik, K. Banasiak, and A. Hafner (2020). A detailed review on CO₂ two-phase ejector flow modeling. *Thermal Science and Engineering Progress* 20, 100647.
- Seeger, M. (2004). Gaussian processes for machine learning. *International journal of neural systems* 14(02), 69–106.
- Shahriari, B., K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas (2015). Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE* 104(1), 148–175.
- Tammone, C., A. Romei, G. Persico, and F. Haglind (2024). Extension of the delayed equilibrium model to flashing flows of organic fluids in converging-diverging nozzles. *International Journal of Multiphase Flow* 171, 104661.
- Tay, S. S., C. S. Foo, U. Daisuke, R. Leong, and B. K. H. Low (2022). Efficient distributionally robust bayesian optimization with worst-case sensitivity. In *International Conference on Machine Learning*, pp. 21180–21204. PMLR.
- Tolerances, G. (1989). Part 1: Tolerances for linear and angular dimensions without individual tolerance indications. *International Organization for Standardization*, 2768–1.
- Wunsch, D. and C. Hirsch (2021). Characterization of manufacturing uncertainties with applications to uncertainty quantification and robust design optimization. *Journal of the Global Power and Propulsion Society* 2021(May), 1–16.