

# SmOOD: Smoothness-based Out-of-Distribution Detection Approach for Surrogate Neural Networks in Aircraft Design

Houssem Ben Braiek  
Polytechnique Montreal  
Montreal, Quebec, Canada  
houssem.ben-braiek@polymtl.ca

Ali Tfaily  
Polytechnique Montreal  
Montreal, Quebec, Canada  
ali.tfaily@aero.bombardier.com

Foutse Khomh  
Polytechnique Montreal  
Montreal, Quebec, Canada  
foutse.khomh@polymtl.ca

Thomas Reid  
Sycodal Inc.  
Montreal, Quebec, Canada  
t.reid@sycodal.ca

Ciro Guida  
Bombardier Aerospace Inc.  
Montreal, Quebec, Canada  
ciro.guida@aero.bombardier.com

## ABSTRACT

Aircraft industry is constantly striving for more efficient design optimization methods in terms of human efforts, computation time, and resources consumption. Hybrid surrogate optimization maintains high results quality while providing rapid design assessments when both the surrogate model and the switch mechanism for eventually transitioning to the HF model are calibrated properly. Feedforward neural networks (FNNs) can capture highly nonlinear input-output mappings, yielding efficient surrogates for aircraft performance factors. However, FNNs often fail to generalize over the out-of-distribution (OOD) samples, which hinders their adoption in critical aircraft design optimization. Through SmOOD, our smoothness-based out-of-distribution detection approach, we propose to code-sign a model-dependent OOD indicator with the optimized FNN surrogate, to produce a trustworthy surrogate model with selective but credible predictions. Unlike conventional uncertainty-grounded methods, SmOOD exploits inherent smoothness properties of the HF simulations to effectively expose OODs through revealing their suspicious sensitivities, thereby avoiding over-confident uncertainty estimates on OOD samples. By using SmOOD, only high-risk OOD inputs are forwarded to the HF model for re-evaluation, leading to more accurate results at a low overhead cost. Three aircraft performance models are investigated. Results show that FNN-based surrogates outperform their Gaussian Process counterparts in terms of predictive performance. Moreover, SmOOD does cover averagely 85% of actual OODs on all the study cases. When SmOOD plus FNN surrogates are deployed in hybrid surrogate optimization settings, they result in a decrease error rate of 34.65% and a computational speed up rate of 58.36 $\times$ , respectively.

## KEYWORDS

Neural network, out-of-distribution detection, surrogate modeling, aircraft design optimization

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ASE '22, October 10–14, 2022, Rochester, MI, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9475-8/22/10...\$15.00

<https://doi.org/10.1145/3551349.3556936>

## ACM Reference Format:

Houssem Ben Braiek, Ali Tfaily, Foutse Khomh, Thomas Reid, and Ciro Guida. 2022. SmOOD: Smoothness-based Out-of-Distribution Detection Approach for Surrogate Neural Networks in Aircraft Design. In *37th IEEE/ACM International Conference on Automated Software Engineering (ASE '22)*, October 10–14, 2022, Rochester, MI, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3551349.3556936>

## 1 INTRODUCTION

The design of aircraft, from conceptualization to commercialization, is expensive. Their lead-time engineering process can take many years and cost billions of dollars from design to manufacture. Multidisciplinary design optimization (MDO) leverages advances in numerical simulation and computer aided design to revolutionize aircraft design [50]. By providing simulation-driven design configurations at an earlier stage [48], next-generation aircraft will have reduced development costs and lead times. Indeed, MDO optimizes different types of objective functions [4] through repetitive assessment of evolving configurations via the use of high-fidelity (HF) mathematical models, which can lead to high computational cost. Even though HF models offer a less expensive alternative to physical testing, HF simulations of aircraft flight dynamics typically take several days to assess thousands of design configurations. Various solutions have been proposed to circumvent the compute-intensive issues of MDO through surrogate modeling [59]. Indeed, surrogate models, also known as metamodels, speed up the optimization process since they assess configurations approximately, but at a much lower expense and faster rate. In the context of aircraft design, developing surrogate models for aircraft design poses a challenge due to the high nonlinearity of some of the aircraft design variables of concern. Due to its accurate predictions and fast inferences, deep learning networks can effectively address the imperfect surrogate modeling issue faced by conventional machine learning models in the field of aircraft design optimization [62]. Neural networks [24] are universal approximators with non-convex learning algorithms that train fast and capture high-nonlinearity input-output mappings. Nevertheless, they fundamentally inherit a closed-world assumption [22], and provide no guidance on how to handle out-of-distribution data (OOD) [66]. Indeed, their generalization capability is guaranteed on novel configurations, but drawn from the same or close distribution as the training set, called in-distribution data (ID). According to the core notions of model

over-parametrization and complexity, modern neural networks can behave worse than their simpler counterparts when exposed to OOD inputs. The resulting corner-case behaviors adversely affect the overall performance assessments for the design configurations, which in turn increases the aircraft design lead time and development budget. However, chasing complete training sets that cover all the facets of the distributions in relation to the quantities of interest, restarts the onset limitation of compute-intensive costs, since the data in aircraft design is typically generated from expensive numerical simulations. Furthermore, the construction of advanced surrogate models using state-of-the-art deep learning techniques [10, 14, 62] requires model engineering efforts, draws on new expertise, and often results in over-optimized models that target individually specific sub-problems.

In response to the aforementioned challenges, we propose *SmOOD*, a smoothness-based OOD detection approach, that allows to codesign the surrogate neural network model for accurate assessments and its OOD detector for selective prediction. Unlike common research on out-of-distribution, *SmOOD* exploits the apriori smoothness property of the simulated system to overcome the main challenges of OOD detection, including complex, high-dimensional input spaces and degeneration of uncertainty estimates beyond the ID regions. This research work makes the following contributions:

- FNNs are investigated as an alternative to GPs in MDO for aircraft design at early stages.
- Pointwise sensitivity profiles are proposed, and their correlations with prediction errors, as well as their discriminative power for ID and OOD samples are examined.
- *SmOOD*, a smoothness-based OOD detection strategy, ensures selective and reliable predictions on the ID inputs that are identified by their proximity to smooth domain regions.
- *SmOOD* is evaluated as a novel criterion for the systematic switch between the surrogate and its HF model counterpart in hybrid optimization settings, and aside from assessing the decrease in errors, overhead costs are also measured.
- The assessment of FNN surrogates, local sensitivity profiles and *SmOOD* approach, was conducted on three aircraft design variables study cases, along with comparisons to baselines, respectively, GPs, uncertainty estimates, and hybrid GP relying on deviations from design neighbors.

**The remainder of this paper is organised as follows.**

Section 2 introduces the relevant concepts and literature in relation with our approach. Section 3 presents our novel smoothness-based OOD detection approach for NN-based surrogate models. Section 4 reports evaluation results, while Section 5 concludes the paper.

## 2 BACKGROUND AND RELATED WORKS

This section provides background information about surrogate modeling for aircraft design and OOD detection strategies.

### 2.1 Surrogate Modeling for Aircraft Design

Surrogate models are data-driven and low-cost substitutes for the exact evaluation of the data points in design space, sometimes dominating the entire optimization process [1], and sometimes serving just as a supplementary aid to speed up the computations [59, 60]. Multiple techniques have been proposed in the literature to build

data-fit surrogates that are trained with regression of high-fidelity simulation data. They usually rely on methods such as gaussian processes [29, 34, 57], proper orthogonal decomposition [25], eigenvalue decomposition [47], artificial neural network [36, 52], and more advanced techniques such as combining GP and neural networks [53]. In the context of aircraft design optimization, the selection of the suitable surrogate modeling technique depends on the complexity of the underlying design problem, the ease of collecting high-fidelity simulation data, and the cost of development and maintenance. **Surrogate modeling solutions are usually presented in conjunction with specific MDO problems from different industries. Hence, advanced techniques pose challenges in regards to the engineering efforts required to adapt them. This paper examines established data-driven modeling methods such as GP and FNN that can be applied directly to a wide variety of MDO problems, including aircraft design optimization. We developed a variety of surrogate models for different quantities of interest (QoI) using mainstream DL frameworks and GP modeling tools.** These QoIs capture certain performance factors of an aircraft and are dependent on the operating flight conditions and the design variables. Our focus is on their deployments to accelerate the investigation of the design space and to find optimum solutions in hybrid surrogate optimization settings [28, 46, 67], where we are able to exploit information coming from both the original model and its surrogate. Therefore, reliable surrogate models with selective predictions (i.e., they are provided only under high-confidence conditions) are essential to ensure the high quality of these accelerated aircraft design optimizations.

### 2.2 Out-of-Distribution Detection Strategies

A trustworthy surrogate model should not only produce accurate predictions based on known context, but also detect unknown inputs and forward them to the High-Fidelity model for safe handling. Indeed, Machine learning (ML) surrogates learn statistically from the available samples of data under the closed-world assumption [66], where the test data is assumed to be drawn i.i.d. from the same distribution as the training samples, known as in-distribution (ID). However, when surrogate models are deployed in concrete aircraft design optimization, the sampled design configurations may represent inputs that are substantially different from the training samples, and should be considered as out-of-distribution (OOD). A wide range of detection strategies [9, 55] has been released to overcome the out-of-distribution challenge. Deep generative models [12, 41, 51] have been leveraged to model efficiently the distribution of inputs  $p(x)$  on the training samples. Then, a membership test is performed on any input  $x$ : if  $p(x)$  is low,  $x$  will be assigned to  $\in \mathcal{D}_{ood}$ , and vice versa. However, this generative modeling faces limitations on large-scale and/or complex input distributions. In contrast, model-dependent OOD detection methods focus on the distribution of model's intricacies, such as hidden features, softmax probability outputs, and uncertainty scores, rather than directly modeling input distributions [23]. Several model-dependent methods [55, 66] have been proposed for classification neural networks. A few research works have focused on regression neural networks through thresholding on uncertainty estimates [33]. Bayesian approaches [7, 16, 27] approximate the posterior distribution of neural

networks parameters through an ensemble of models. Regarding Non-Bayesian ensembling approaches, Kendall and Gal [26] adds the expected prediction’s variance as additional network’s output. Hence, the variance of a prediction is also improved by the learning algorithm, guided by the minimization of the negative log likelihood on the data. Mi et al. [40] proposed training-free uncertainty estimation that leverages the average of output deviations under input or feature map perturbations as a surrogate for uncertainty measurement. One challenge with these uncertainty-based methods is their reliability degeneration on out of distribution inputs on which they may be falsely overconfident [45]. Therefore, OOD examples can be identified by their predicted variances beyond the confidence interval obtained for ID samples. In this paper, we propose a novel OOD detection strategy for regression models under conditions of smoothness that relies on the mapping function’s sensitivity profiling at a given input region to capture potentially OODs.

### 3 APPROACH

In this section, we describe the development steps required to code-sign smoothness-based OOD detectors with FNN surrogates. In Figure 1, the proposed systematic workflow of the SmOOD approach is shown, including the OOD recognition, the training and testing of the inherent regression and classification models.

#### 3.1 Characterization of Out-of-Distribution

In supervised learning, statistical models are commonly estimated via empirical risk minimization (ERM) [17], a principle that considers minimizing the average loss on observed samples of data, as an empirical estimate of the true risk, i.e., the expected true loss for the entire input distribution. Unfortunately, ERM assumes that training and test data are identically and independently distributed (a.k.a. i.i.d. assumption). Distributional shifts often occur in real scenarios for many reasons, such as domain transition, temporal evolution, or selection bias, which degrades the model’s performance since certain captured correlations may not hold on these shifted inputs. Even worse, several research works [3, 15, 37] demonstrated that the optimized models can fail dramatically when involving strong distributional shifts. In many surrogate modeling cases where the HF simulations are expensive, selection bias is almost unavoidable, which ruins the i.i.d. assumption. The design of OOD detection methods for selective prediction is of more critical significance than reducing empirical risk further on the collected training samples. The first step consists of the OOD inputs characterization. In the context of aircraft surrogate performance models, we aim to characterize the foreseeable design configurations that can be sampled during the design optimization, however, they may deviate from the training distribution. As explained above, ERM optimizes the model by minimizing the average loss, which leads to greedily absorbing the correlations and patterns that hold on the majority of training instances. Hence, ERM can produce biased models that are susceptible to outliers, unfair to minor subsets of data, or prone to out-of-distribution samples [35]. Thus, it is common to use held-out validation data as a proxy for unseen data points, aiming at more realistic estimates of the true risk. As shown in Figure 1 (step 3), the validation error rate may reveal the model’s inefficiencies in terms of remaining high-error test inputs on which the learned patterns

could not generalize, and they can be used for approximating the boundaries of the in-distribution. We rely on thresholding over the prediction confidence interval, i.e., the margin of validation errors at a certain level of confidence, and labeling the data point as OOD if its associated error exceeds the estimated margin.

#### 3.2 Computation of Local Sensitivity Profiles

HF mathematical models are typically based on conservation laws and solve a coupled nonlinear system of partial differential equations on a discretized spatio-temporal domain. This allows numerically stable simulations of the aircraft via their approximations in smooth or piecewise smooth dynamics. Hence, their NN-based surrogate counterpart must meet this apriori of smoothness in regards to the predicted quantities of interest in order to be a viable alternative to the HF model. In fact, the smoothness of the FNN’s mapping function affects the model complexity as evidenced by the fact that smooth deep neural networks tend to generalize better than their less smooth counterparts [44]. Nonetheless, rigid smoothness techniques that constrain the learned function excessively by forcing it to be equally smooth throughout the input space, may throw away useful information about the input distribution [54]. As alternatives to these rigid methods, DL practitioners can employ regularization techniques, such as weight decay [39], dropout [61], and early stopping [32], to encourage smoothness of the model and improve generalization. Their hyperparameters should be tuned to achieve the desired level of smoothness that allows the network to allocate capacity as needed to maintain useful diversity, handle input modalities, and capture task-relevant information. In our study of OOD detection for aircraft surrogate performance FNNs, the smoothness is one of the fundamental properties of system design. Apart from regularizing the FNNs to smooth mapping functions, we created pointwise sensitivity profiling to capture the local function smoothness around the neighbor regions of each evaluated data point. On unseen data points, the integrity of the surrogate network can then be ensured by comparing their triggered pointwise sensitivity profiles to the degree of smoothness observed on the optimized DL model on the in-distribution inputs. In the following, we introduce the proposed pointwise sensitivity profiling for FNN that quantify different aspects of the network’s output variations in relation with changes in input variables.

*3.2.1 The expected deviations of the network’s output:* The straightforward way to estimate the sensitivity of the neural network at an individual data point, is to assess the induced output deviations in response to random noise injected into the input features. According to Mi et al. [40], network sensitivity to input perturbations can be used as a surrogate for uncertainty. More precisely, they have proved that sensitivity and uncertainty have a nonnegative correlation in a setting of dense regression networks. Therefore, we include the following statistics on the output deviations under constrained input perturbations in the sensitivity profiling that will be used to separate between OOD and ID inputs.

Let  $\Delta x$  be the perturbation applied to the input of  $f$ , a neural network,  $\delta$  be the maximum threshold of the absolute value of  $|\Delta x|$ , and  $\Delta y$  be the responding output deviation,  $|f(x+\Delta x) - f(x)|$ . Then, the two deviation-based sensitivity profiling metrics are defined

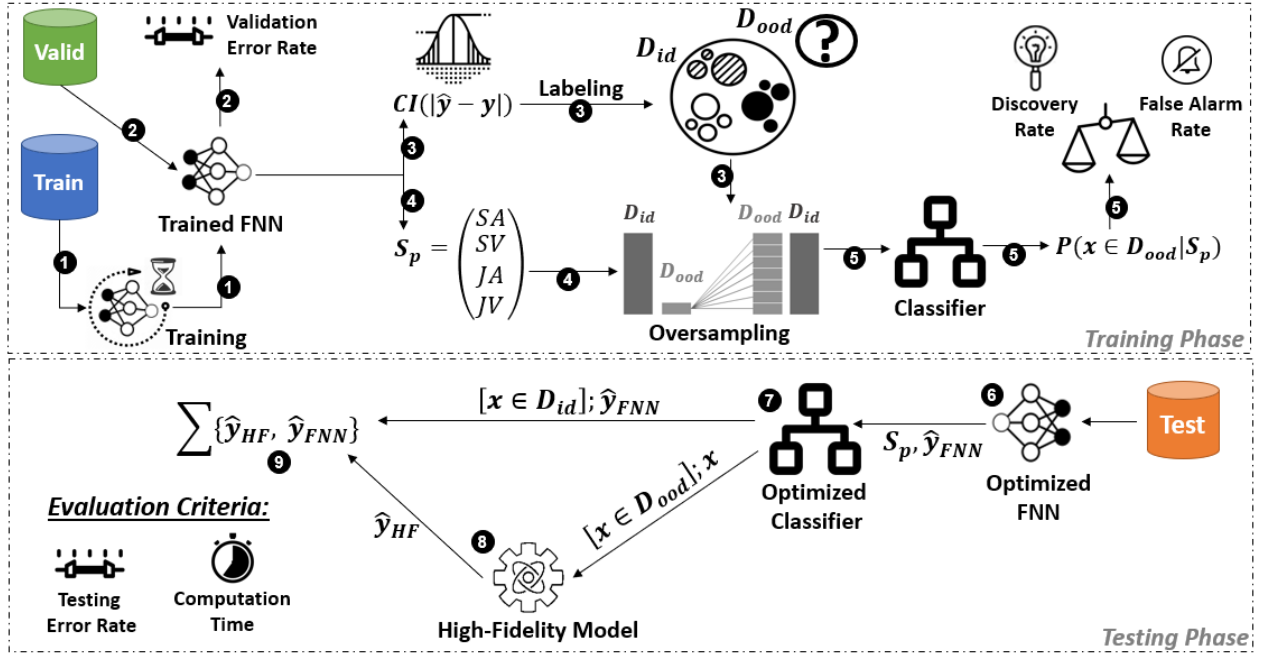


Figure 1: Schema of *SmOOD*: Co-Design of FNN Surrogate and Classification-based OOD Detection Models.

respectively:

$$SA(x) = \text{AVG}(|\Delta y|) = \mathbb{E}_{\Delta x \sim \mathcal{U}(-\delta, \delta)} [|f(x + \Delta x) - f(x)|]$$

$$SV(x) = \text{VAR}(|\Delta y|) = \mathbb{E}_{\Delta x \sim \mathcal{U}(-\delta, \delta)} [ (|\Delta y| - \text{AVG}(|\Delta y|))^2 ]$$

**3.2.2 The input-output jacobian norm:** As the feedforward neural networks are differentiable models, we also investigate its sensitivity at a point through the computation of input-output jacobian norm. Novak et al. [44] have presented extensive experimental evidence that the local geometry of the trained function as captured by the input-output Jacobian can be informative of the prediction confidence at the level of individual test points, and that it varies drastically depending on how close to the training data manifold the function is evaluated. Thus, the computed jacobian norm is likely to be higher at shifted inputs in comparison with the training inputs. Nevertheless, the feedforward networks are typically based on ReLU activations, which makes them not continuously differentiable. The derivative of the mapping function may therefore fluctuate sharply at small scales, and hence, considering the Jacobian of a specific data point will have less meaning than considering the Jacobians of a subset of nearby data points [58]. Thus, we apply random perturbations on the data point to yield samples from its neighborhood. Then, we compute statistics on the vanilla Jacobians estimated at these noisy samples to be included in our network's sensitivity profiling, as follows.

Let  $J(x) = \partial f(x) / \partial x^T$  be the input-output jacobian of a neural network,  $f$ , at the input  $x$ , and  $\|\cdot\|_F$  is the Frobenius norm. Further, we assume that  $\Delta x$  is the perturbation applied to the input of  $f$ , a neural network,  $\delta$  is the maximum threshold of the absolute value of  $|\Delta x|$ , and  $x' = x + \Delta x$ .

$$JA(x) = \|\text{AVG}(J(x + \Delta x))\|_F = \|\mathbb{E}_{\Delta x \sim \mathcal{U}(-\delta, \delta)} [J(x + \Delta x)]\|_F$$

$$JV(x) = \|\text{VAR}(J(x'))\|_F = \|\mathbb{E}_{\Delta x \sim \mathcal{U}(-\delta, \delta)} [(x' - \text{AVG}(J(x')))]^2\|_F$$

As demonstrated in the step 4 of our defined workflow (Figure 1), we define the pointwise sensitivity profiles, denoted  $S_p$ , as a real-valued vector that concatenate the above sensitivity-related statistics within the neighborhood regions of individual data points,  $x$ , which can be formulated as follows,  $S_p(x) = [SA(x) \parallel SV(x) \parallel JA(x) \parallel JV(x)]$ .

### 3.3 Training of ID/OOD Classifier

After being estimated, the pointwise local sensitivity profiles are fed into a ML-based classifier, specifically, a binary classifier. The latter is trained on the sensitivity profiles of validation data points, along with their derived labels, i.e., ID or OOD, on the basis of the expected confidence interval for the validation errors. Nevertheless, the average and the margin of error are two metrics that will drive the development and the selection of surrogate models. Hence, it is reasonable to expect that these measurements are well optimized and only a small fraction of data points would remain with relatively high error over successive validation steps, which results in an imbalanced dataset of sensitivity profiles that is predominantly composed of ID examples with only a low percentage of OOD examples. To deal with class imbalance, we explore several commonly-used oversampling techniques [30] in the literature that produce synthetic data points belonging to the minority class to emulate a semblance of balance to the dataset. These techniques are necessary means of increasing the sensitivity of a classifier to the minority class, allowing the detection of as many OODs as possible from incoming inputs. This is shown in-between the step 4-5 of the *SmOOD* co-design workflow in Figure 1. By interpreting OODs as positive and IDs as negative, we should conduct an in-depth performance assessment of the ML-based classifier with

aim of converging to optimal balance (i.e., the evaluation step 5 in Figure 1) between OOD discovery, referring to the ratio of the truly detected OODs among the actual OOD collection, and OOD false alarms, defined as the ratio of false positive arising from the actual ID samples. In the context of hybrid surrogate design optimization, a high OOD discovery rate ensures trustworthy selective surrogate predictions where only ID samples are submitted to the surrogate model, and hence, the accelerated design optimization converges to high potential and feasible aircraft design configurations. On the other hand, a low OOD false positive rate is imperative to preserve the advantages of surrogate modeling and avoid useless, costly queries to the HF model.

### 3.4 SmOOD: Integration and Evaluation

During the training phase of the SmOOD workflow (Figure 1), the surrogate model is learned first, then its predictive errors are computed on the held-out validation datasets. Next, we leverage the confidence intervals of the estimated errors during the validation to identify the out-of-distribution (OOD) data (i.e., inputs on which the best-fitted model still produces relatively high errors). In parallel, we calculate the local sensitivity profiles for all the validation inputs. Since the best-fitted model was selected for its predictive performance, OOD inputs would have very small sensitivity profiles compared with their counterparts assigned to in-distribution (ID) inputs, and an oversampling procedure is required to amplify synthetically their occurrences in order to analyze the differences in smoothness-related characteristics between both groups of profiles. As a next step, we feed the oversampled sensitivity profiling data into a binary classifier that learns to distinguish between OOD and ID profiles. For optimal classifier selection, we assess both the OOD discovery and false alarm rates using k-fold cross-validation.

For the testing phase in the SmOOD workflow (Figure 1), the test inputs represent aircraft design configurations for which an assessment has been requested. Since SmOOD is a model-dependent OOD detection strategy, we first pass all the test inputs by the optimized FNN to determine their predictions and their corresponding local sensitivity profiles. Then, the optimized classifier serves as a calibrated OOD detector that generates a risk score for each precomputed sensitivity profile, quantifying the likelihood that its source test is indeed an out-of-distribution sample from the FNN surrogate's perspective. The next step depends on a criterion that evaluates whether the risk score falls below a predefined threshold (i.e., a default of 0.5 out of [0, 1.0] score ranges), and two scenarios are then possible. If the criterion is satisfied, the surrogate prediction is returned and no further action is taken. If not, the test sample is sent to the High-fidelity model for simulation and its outcome is returned. Quality of results is assessed using two evaluation criteria. First criterion is the decrease rate achieved in respect to the prediction errors on testing samples, and especially, the further error reduction by the hybrid mode. A reliable OOD detection method must route all high risk samples to the HF model, reducing error rates significantly. Second, the computation time of hybrid mode should be much shorter than the pure HF-driven optimization since this was the main purpose of setting up the surrogate model in the first place. Indeed, a well-calibrated OOD detection method

should spawn a low rate of false alarms because these false positives would cause unnecessary HF simulations, slowing down the progress towards design optimization.

## 4 EVALUATION

In this section, we first introduce the three aircraft performance case studies, as well as our evaluation setup, metrics, and methodology. Next, we evaluate *SmOOD* against standard baseline in terms of predictive performance and computation runtime.

### 4.1 Experimental Setup

This section details the different elements that were set up for SmOOD performance assessment, including aircraft design study cases, inherent models, baseline, as well as evaluation environment, strategy, and metrics.

**4.1.1 Case Studies.** Below, we briefly describe the studied aircraft performance factors [2], the influencer design variables [2], and the design of experiments used for dataset collection.

**Maximum TakeOff Weight (MTOW).** It represents the maximum weight at which the pilot of the aircraft is allowed to attempt to take off given its structural design. MTOW is an important factor in aircraft design. A higher MTOW means the aircraft can take off with more fuel and has a longer range, which makes it more appealing to customers. However, MTOW is subject to several structural constraints during the design optimization process.

**Time To Climb (TTC).** Climbing is the act of increasing the altitude of an aircraft. The main climb phase is the increase of altitude following the takeoff to reach the cruise level. As a way to measure an aircraft's climb performance, it is common to set up a reference cruise altitude level, then, estimate the time needed to climb to the predetermined altitude at a constant airspeed. This climb performance measurement is called the Time To Climb (TTC).

**Balanced Field Length (BFL).** It refers to the shortest runway length at which a balanced field takeoff can be performed by an aircraft design while complying with safety regulations. A balanced field takeoff occurs when the required accelerate-stop distance is equal to the required takeoff distance. Accelerate-stop distance is the runway length required by an aircraft to accelerate to a specific speed, and then, in the event an engine fails, to stop safely on the runway. Thus, aircraft designers are sensitive to BFL as any changes in this require ensuring safety margins at takeoff are still respected.

**Design Configuration.** A set of 15 design variables are considered in the surrogate modeling based on experts' judgment on their influence on the above-mentioned quantities of interest. The design variables capture the wing design (i.e., aspect ratio, taper ratio, thickness ratio, and winglet span ratio), the fuselage geometrical structure, the engine thrust, and the mass of aircraft parts.

**Design of Experiment(DoE).** Our engineering collaborators configure the appropriate design of experiments (DOE) to collect the labeled data points. DoE assembles a set of tests that exercises the HF model across diverse design configurations to gather HF simulations. Indeed, the Latin Hypercube Sampling technique [38] is leveraged to efficiently sample from large, multivariable design spaces. LHS uniformly divides the range of each design variable into the same number of levels. It then systematically combines

independent samples of the levels of each factor to specify a variety of random data points in the design space. Running the HF model on these data points gives us HF simulations that map design factors to their corresponding responses, i.e., the measurable quantities of interest. Therefore, we obtain 2259 of training data points and 1960 of testing data points for each one of the study cases.

#### 4.1.2 Models.

**Surrogate.** Our base nonlinear regression model is a three-layer feedforward neural network(FNN) [5] that is trained using the Mean Squared Error (MSE) loss function with L2-norm regularization. Rectified linear units (ReLU) are used as hidden layer activation functions, and Adam is used as the optimization algorithm. In regards to the architecture, we followed the design principle of pyramidal neural structure [20], i.e., from low-dimensional to high-dimensional feature spaces/layers, as well as these dimensions are powers of 2 to achieve better performance on GPUs [56]. Regarding the hyperparameters tuning, we leverage the grid-search strategy and 5-fold cross validation to sample and try several possible settings. To outline their ranges, we denote  $linspace(a, b, n)$  to indicate the range of  $n$  equi-spaced values within  $[a, b]$  and  $logspace(c, d, base)$  to indicate the interval of  $base^c, \dots, base^d$ , where  $c < d$ . The FNN's width of layers are selected from  $logspace(5, 10, 2)$ , then, the optimizer's hyperparameters were tuned as follows: learning rate  $\eta \in s \cup 3 \times s$ , weight decay  $\lambda \in s \cup 5 \times s$ , where  $s = logspace(-4, -1, 10)$ . Batch size was tuned in  $logspace(3, 7, 2)$ , and epochs count in  $linspace(50, 500, 50)$ .

**OOD Detection.** Several ML classifiers (Logistic Regression [65], SVM [43], Random Forest [6], and Gradient Boosting [42]) and over-sampling techniques (SMOTE [11], BorderlineSMOTE [19], ADASYN [21], and SVMSMOTE [63]) have been evaluated on our OOD detection problems using 5-fold cross validation process. We have found that Gradient Boosting and SVMSMOTE are the best design choices. Then, we follow the same hyperparameters tuning method used for FNN, i.e., grid-search strategy and 5-folds cross validation. For the gradient boosting classifier, we examine the learning rate,  $\eta$  and the number of estimators,  $n$ , selected from the following ranges, respectively,  $s \cup 5 \times s$  and  $[100, 250, 500]$ , where  $s = logspace(-3, -1, 10)$  and  $[100, 250, 500]$ . In regards to SVMSMOTE oversampler, we consider  $k \in [5, 10, 15]$ , where  $k$  represents the count of nearest neighbors to used to construct synthetic samples, and  $r \in [0.25, 0.5, 0.75, 1.0]$ , where  $r$  indicates the desired ratio of the number of samples in the minority class over the number of samples in the majority class after resampling.

#### 4.1.3 Baseline.

**Surrogate.** We compare our FNNs with Gaussian Process(GP) regression that is already well researched for surrogate modeling [64], replacing expensive high-fidelity aircraft simulations. A GP [31] is a generalization of the Gaussian distribution to describe universal functions  $f(x)$ . The main ingredient of GP design is the selection of the covariance function  $k(x, x')$ , also called kernel. For our comparison with FNN, we are interested in approximating multiple quantities using the baseline. We chose the radial basis function (RBF) [8] that has been implemented to interpolate several aircraft design data [29, 46]. RBF [8] is an universal kernel function that can be used to fit any complex non-linear regression data. The

hyperparameters of the kernel are optimized during fitting by maximizing the log-marginal-likelihood (LML) on the training data. As the LML may have multiple local optima, we set up 10 to be the number of repetitive restarts of the optimizer in order to improve the convergence towards optimal results.

**OOD Detection.** A basic approach of hybrid GP [46] for surrogate aircraft design optimization mimics a rule of thumb in engineering analysis: if the new data is reasonably similar to the existing one, it can be assumed to be reliable. Given a design configuration to assess, one can compute the standard deviation,  $\sigma$ , of the differences between the output predicted by RBF,  $\hat{y}$ , and the actual outputs,  $y_1, y_2, \dots, y_n$ , of a set of  $n$  neighboring configurations. The standard deviation  $\sigma$  can be computed with differences  $\|\hat{y} - y_i\|$  with  $i \in [1, n]$ , as follows,  $\sigma = \text{std}(\|\hat{y} - y_1\|, \|\hat{y} - y_2\|, \|\hat{y} - y_n\| \dots)$ . Then, the obtained standard deviation is compared to a threshold,  $\sigma_t$ . If it is lower,  $\sigma < \sigma_t$ , the prediction of RBF surrogate is adopted; otherwise, the HF model must be requested instead. Regarding the threshold value  $t$ , a priori statistical analysis can be performed on the validation data to derive the threshold value that includes 95% of configurations. Using a 5-fold cross validation, we tune the number of neighbors considered in the computation of standard deviation within the range of  $linspace(4, 20, 4)$ , and we select the number that yields the highest correlation between the estimated standard deviation and the actual error.

**4.1.4 Evaluation Procedure.** To evaluate SmOOD's effectiveness, we conducted quantitative and qualitative analyses. In the following, we detail the metrics and the procedure used.

**Quantitative.** Due to the high cost of running HF simulations, we adopt a 10-fold cross-validation method for all experiments in order to have different splits for the training and validation datasets and quantify the target metrics by averaging their values over the 10 iterations. Besides, all the included estimated metrics like error rates and runtime, are computed as average values over 5 runs or more, in order to mitigate the effects of randomness inherent in statistical learning algorithms. Below, we introduce the different evaluation metrics that have been used in the empirical evaluations.

**NRMSE.** stands for normalized root mean square error, and it is scale-independent version of RMSE that allows comparison between models at different scales. Indeed, RMSE is the average deviation between predictions and actual outputs, measured on the same scale and with the same output unit, and can be formulated as follows:  $RMSE = \sqrt{\sum_{i=1}^n (y_i - \hat{y}_i)^2 / n}$ , where  $y_i$  is the  $i$ th observation of  $y$  and  $\hat{y}_i$  its corresponding prediction by the model.

In NRMSE, the expected model deviations are reported relative to the overall range of output values, and the formula becomes  $NRMSE = RMSE / (y_{\max} - y_{\min})$ .

**Confidence Interval (CI).** indicates the degree to which an estimate is expected to vary from the average, within a certain level of confidence. For example, a 95%CI of a model's error represents the upper and lower bounds within which the estimate of error will fall for 95 percent of the time (i.e., 95 samples out of 100 were taken). We use the bootstrapping method to produce accurate CIs because its non-parametric nature allows it to be used without making any prior assumptions about the estimate distribution.

**Precision.** represents the proportion of the positive identifications that were actually correct in a binary ML classification problem.

*Recall*. measures the proportion of actual positives that were correctly classified in a binary ML classification problem.

*Macro-Averages of Precision/Recall*. refer to the arithmetic averages of the precision and recall scores of individual classes.

*Precision-Recall (PR) curve*. is a useful evaluation plot for binary classifiers that return class membership probabilities, when there is moderate to large class imbalance [13]. It simply plots the precision versus recall obtained by a classifier using different probability thresholds.

*Area Under Precision-Recall (AUPR)*. approximates the area under the PR curve, ranging from 0.0 to 1.0.

*%Decr\_Err*. The decrease ratio of error we achieve by a certain improvement, can be formulated as follows:

$$\%Improv\_Err = \frac{pre-NRMSE - post-Err}{pre-NRMSE} \times 100 \quad (1)$$

*Speed up*. is a popular measure for the relative performance of two systems processing the same problem. In our case, we denote  $T_s$  be the surrogate compute time, and  $T_o$  the High-fidelity compute time. Then, the speedup due to surrogate modeling can be computed as follows,  $S_{pure} = T_o/T_s$ . Given our OOD detection strategy used in hybrid surrogate optimization, we obtain  $p$ , as a proportion of instances that can be predicted by the surrogate. Then, the remaining  $1 - p$  proportion of instances that require requests to HF model. This means that the speed up of the hybrid surrogate model can be formulated as follows,  $S_{hybrid} = T_o/(T_d + p_s + (1 - p) * T_o)$ , where  $T_d$  is the total computation time needed to predict whether each of the inputs is OOD or not.

**Qualitative.** In order to obtain the input of domain experts, we interviewed two senior aircraft engineers with more than 10 years of experience, who work with our industrial partner, Bombardier Aerospace. Both of them are proficient in the use of MDO in aircraft design. Furthermore, they are part of the team that developed the high-fidelity physics model; therefore, they are familiar with the QoIs associated with the case studies.

As a first step, we present the quantitative analysis to both aircraft engineers separately so that they can get detailed performance metrics, as well as all execution costs for all model development steps (training, tuning of hyperparameters, etc). Then, we ask them to provide a critique of the significance of the obtained performance metrics from an aircraft design standpoint. Specifically, we seek comparisons of FNN surrogates along with our co-designed OOD detection method, as opposed to conventional surrogate modeling setups with GP and uncertainty quantification. In addition, we request explanations on how SmOOD contributes to quality assurance and acceleration of data-driven aircraft design optimization. Afterwards, we compile and merge their opinions and comments into paragraphs. Next, we submit them for approval to ensure their claims are reflected in the written content, and that they both agree with the conclusions made. Final paragraphs are added to each related research question.

**4.1.5 Environment.** We use Pytorch [49], an established DL framework for modeling and training feedforward neural networks. We leverage GPy [18], a popular and maintained framework to design and train Gaussian processes. In terms of the hardware environment, we use CPU machines for all the experiments for fairness

comparisons. The HF models were running on a machine with Intel Xeon CPU E5-1630 v3 of 3.5Ghz and a 32 Gb of RAM. The FNNs and GPs were running on Standard Virtual Machines using 4 cores on Intel Xeon Platinum 8168 CPU of 2.70GHz and 8 Gb of RAM.

## 4.2 Experimental Results

In conducting the evaluation of the proposed approach, we studied the following research questions:

*RQ1. Is FNN a viable general-purpose approximator to model complex surrogate aircraft design performance models?*

**Motivation.** The objective is to evaluate the effectiveness of FNN as a universal approximator for surrogate aircraft design performance models in comparison to GP, which is the mainstream universal estimator for surrogate modeling. **Method.** We train and tune both FNN and GP surrogate models on all the three aircraft design performance factors. Then, we test the optimized models on a held-outs test dataset to assess their predictive performance. As performance metrics, we compute the NRMSE on validation datasets to gauge the fitness quality of both models and the NRMSE on test dataset to compare their generalizability capabilities. Besides, we infer the CI of each estimate of error with 99% of confidence level, as an accurate risk assessment of the possible range of the model's prediction errors, and especially, the upper bound that indicates how severe the error is expected to be. Last, we calculate the ratio of OOD inputs observed on validation and testing datasets according to our OOD labeling rule (Section 3.1), i.e., their actual prediction errors are higher than the maximum expected margin.

**Results.** Table 1 summarizes the obtained performance metrics for each pair of regression problem and surrogate model type. The results show that FNNs have reached lower validation and test estimation errors with even tighter confidence intervals than GP in all the studied surrogate modeling problems. Moreover, the ratio of model-dependent OOD inputs that have been experienced by GP is substantially higher than FNN. This demonstrates that GP's high average error and wide confidence interval are due to the prevalence of complex design configuration inputs on which GP fails and cannot provide a reliable assessment. It is on the time-to-climb (TTC) design problem that the gap between the two types of surrogate models is most pronounced. There are, indeed, high nonlinearities in TTC modeling, where two numerically-close inputs can result in distant takeoff behaviors. Because GP kernels are biased towards their predefined distributional priors and have a limited learning capacity in comparison to FNN, they do not generalize well to aircraft performance design problems with dispersed behaviors, such as the time to climb simulations, ranging from takeoff failures to accelerated climbing scenarios.

**Finding 1:** Feedforward neural network is a viable universal approximator, outperforming Gaussian Process, in building complex, highly-nonlinear surrogate aircraft performance models.

**Domain Expert Feedback.** Aircraft engineers find the comparison results enlightening as they demonstrate that data-driven surrogate modeling is still a viable solution even in cases of high nonlinearity input-output mappings, such as TTC models. Moreover, GP's results on TTC are expected, since it is the origin of this

**Table 1: Performance Metrics for different pairs of QoI and Surrogate Model.**

QoI	Model	Valid_Err	Test_Err	Valid_CI_Err@99%CL	Test_CI_Err@99%CL	%Valid_OODs	%Test_OODs
<b>MTOW</b>	GP	0.0415	0.0363	[0.0212,0.0596]	[0.014, 0.0547]	5.14%	4.39%
	FNN	0.0353	0.0319	[0.0163, 0.0523]	[0.0111, 0.0486]	1.59%	1.17%
<b>TTC</b>	GP	0.1707	0.1698	[0.1646, 0.1769]	[0.1632, 0.1762]	40.95%	41.33%
	FNN	0.0764	0.0803	[0.0662, 0.0865]	[0.0668, 0.0942]	5.0%	5.26%
<b>BFL</b>	GP	0.0454	0.0439	[0.0241, 0.0637]	[0.0168, 0.0664]	6.64%	5.97%
	FNN	0.0435	0.0361	[0.0198, 0.0637]	[0.0124, 0.0599]	1.42%	1.12%

initiative on reliable FNN surrogates. It was previously necessary for aircraft engineers to re-run the design optimization with the HF model in the objective function because the GP gives inaccurate TTC predictions and hinders the convergence of the optimizer, which does not find a feasible design.

*RQ2. Can local sensitivity profiles capture relevant information on FNN behaviors to accurately signal OODs across generated samples?*

**Motivation.** The aim is to determine how well pointwise local sensitivity profiles can discriminate between ID and OOD inputs compared with exploiting the uncertainty estimation of GP.

**Method.** For GP, the predictions,  $\hat{y}$  are made in a probabilistic way, i.e., described by mean,  $\mu_{\hat{y}}$ , and standard deviation,  $\sigma_{\hat{y}}$ , which is supposed to reflect the epistemic uncertainty, i.e., how certain the model is with respect to its prediction. We therefore statistically compare the distribution of  $\sigma_{\hat{y}}$  on ID versus OOD examples discovered during the validation. To do that, we compare their box plots and perform Mann-Whitney U nonparametric hypothesis tests to determine whether the  $\sigma_{\hat{y}}(\mathcal{D}_{id})$  and  $\sigma_{\hat{y}}(\mathcal{D}_{ood})$  were sampled from the same population. Regarding FNN, we compute local sensitivity profiles triggered by FNN on the validation data points, and then analyze the distribution of these profiles among the ID and OOD groups. To do that, we use Principal Components Analysis (PCA) to visualize the computed multivariate local sensitivity profiles into 2D graphs. Indeed, we reduce the dimensions into two orthogonal principal components that preserve the maximum amount of variance explained by the original multivariate points, as well as, we add the original variable vectors to show their correlations and directions w.r.t the principal components. Moreover, we colored the ID and OOD points with different colors to identify the differences between the two groups. In addition, we perform Mann-Whitney U tests for each of the sensitivity profile variables to statistically confirm if the ID and OOD groups are likely to yield values originating from two different distributions.

**Results.** Figure 2 shows the box plots of the GPs' standard deviations for both ID and OOD examples. As can be seen, the groups' boxes have very close sizes and almost overlap completely for all the regression problems. Additionally, hypothesis testing suggests no statistically significant differences between the uncertainty estimates predicted by GP on ID versus OOD. The median lines in the box plots indicate that the median of the standard out-of-distribution sample is lower than the standard identification sample, demonstrating how GP incorrectly produces a sense of overconfidence about out-of-distribution samples. Even worse, the comparison of the median lines within the boxes highlights that the median of  $\sigma_{\hat{y}}(\mathcal{D}_{ood})$  is lower than the median of  $\sigma_{\hat{y}}(\mathcal{D}_{id})$ ,

which shows how GP incorrectly tends to produce overconfident epistemic uncertainty when confronted with out-of-distribution samples. This is definitely an expected degradation in the reliability of uncertainty estimates, since the posterior distributions of predictions were calibrated using available samples, mostly ID, suggesting bias against distant examples. In contrast, Figure 3 present 2D plots of PCA analyses based on the sensitivity profiles of the regression models for the studied QoIs, MTOW, TTC, and BFL. As can be observed, the ID examples are grouped together and form a sort of cluster, whereas the OOD examples are quite dispersed and far from the centroid of the ID group. The hypothesis testing results confirm these observations, and there are statistically significant differences (with  $p\text{-value} < 1e - 6$ ) between the ID and OOD data groups in regards to all of the sensitivity profile's variables. Furthermore, the high dispersion and rarity of the OOD data points highlight that the identified scenarios represent corner-case and extreme behaviors with different patterns, rather than representing a novel distribution.

**Finding 2:** Pointwise local sensitivity profiles can accurately separate OOD and ID samples using FNN, whereas the reliability of GP's epistemic uncertainty declines on OOD samples.

**Domain Expert Feedback.** Aircraft engineering experts confirm they have experienced overconfidence of GP uncertainties before, especially when using objective functions in MDO that include the variance to better steer the search to high potential regions. In that case, over-confident variance is misleading to the optimizer. They perceive the added value of the FNN sensitivity profiling, which is able to capitalize on the smoothness of input-output mapping as it is an a priori system property and an implicit assumption in the HF model's differential equations.

*RQ3. How effective is SmOOD in the coverage of OOD inputs across generated samples?*

**Motivation.** The goal is to assess the performance of SmOOD when co-designed with FNN against the baseline strategy used with GP in terms of precision and coverage.

**Method.** We train and tune our gradient boosting classifier on the local sensitivity profiles yielded by FNNs during the validation for each regression problem. Then, we test the optimized classifier on the held-out testing dataset. Next, we draw PR curves along with their associated AUPR scores for each studied problem. Concerning the OOD detection baseline for GP introduced in Section 4.1.3, it provides only labels with no class scores, we compute the precision and recall scores for each class, separately, and their macro average for each pair of surrogate model and OOD detector.



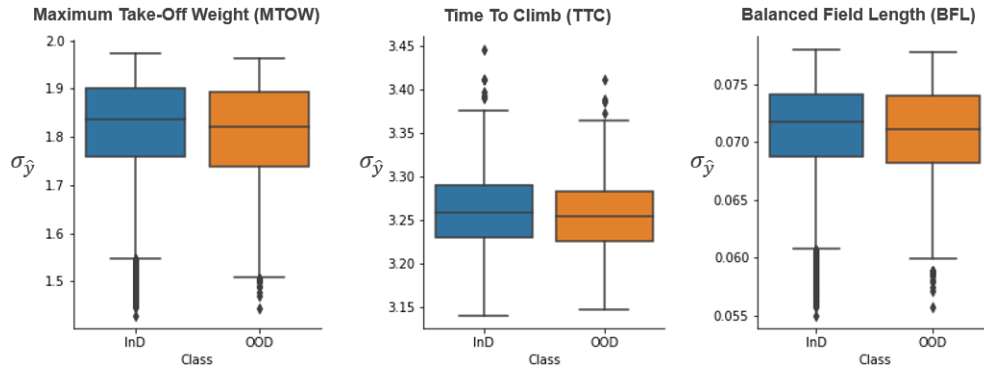


Figure 2: Comparison of GP’s standard deviation distribution between ID and OOD samples for each QoI

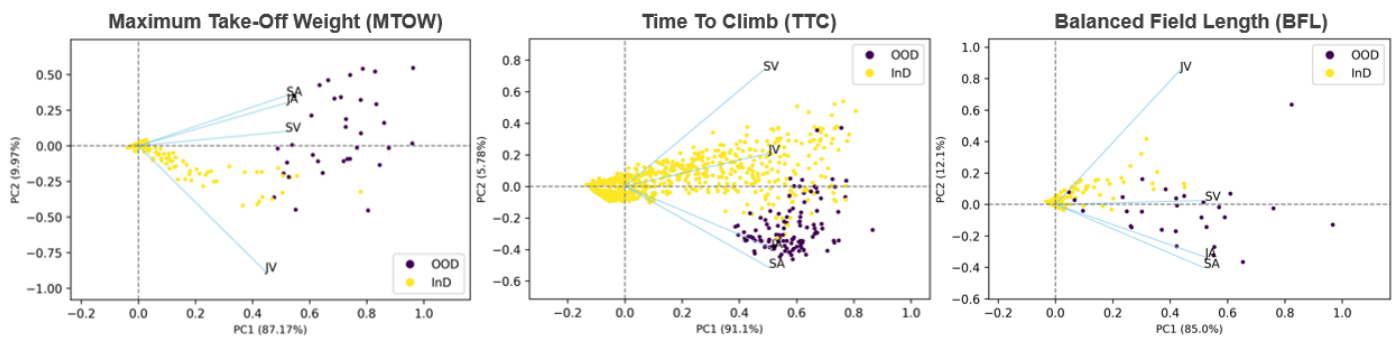


Figure 3: Graph of Pointwise Sensitivity Profiles for  $D_{id}$  and  $D_{ood}$  w.r.t the derived 2 Principal Components

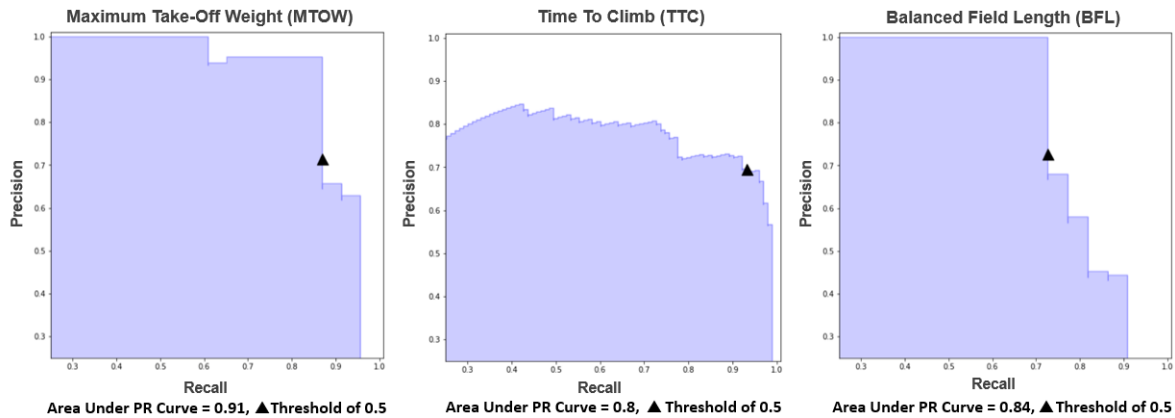


Figure 4: Precision-Recall Curves obtained by the optimized binary classifier for each QoI

**Results.** Figure 4 shows the precision-recall curves obtained by FNN for each aircraft performance problem and their associated AUPR scores. As can be seen, *SmOOD* was effective in fitting the collected local sensitivity profiles as evidenced by the relatively high AUPR scores. Although some levels of precision could be sacrificed to increase recall, we believe it is better to leave the default threshold at 0.5 since it yields a good trade-off between precision and recall, as shown by the marked points on the PR curves. Further pre-tuning of thresholds on validation dataset can

lead to overfit the validation dataset and inversely causes a degradation of the classifier’s performance. Besides, Table 2 summarizes all the classification scores: both of precision and recall for each class, as well as their macro-averages for each pair of surrogate model type and aircraft performance problem. Our method outperforms by far the baseline approach. This demonstrates the lack of generalizability expressed by the rigid smoothness apriori of the baseline OOD detection method, which could not hold for complex, highly-nonlinear aircraft design performance models. Our

**Table 2: Performance Comparison of SmOOD and Base for OOD Detection.**

Target	Detector	Model	Class	Precision	Recall
MTOW	Base	GP	ID	96%	96%
			OOD	<b>11%</b>	<b>10%</b>
			MA	53%	53%
	SmOOD	FNN	ID	100%	100%
			OOD	<b>72%</b>	<b>91%</b>
			MA	86%	95%
TTC	Base	GP	ID	57%	91%
			OOD	<b>14%</b>	<b>2%</b>
			MA	36%	47%
	SmOOD	FNN	ID	100%	98%
			OOD	<b>70%</b>	<b>93%</b>
			MA	85%	95%
BFL	Base	GP	ID	94%	96%
			OOD	<b>13%</b>	<b>10%</b>
			MA	54%	53%
	SmOOD	FNN	ID	100%	100%
			OOD	<b>76%</b>	<b>73%</b>
			MA	88%	86%

method, in contrast, retrieves pointwise local sensitivity profiles and fits a binary ML classifier to the high-fidelity simulations to serve as a calibrator and detects erroneous smoothness behavior triggered by the FNN during the inference on unseen test inputs.

**Finding 3:** *SmOOD* was effectively able to reveal at least 73% and up to 93% of the out-of-distribution examples on the different aircraft design performance models.

**Domain Expert Feedback.** Aircraft engineers emphasize the importance of developing such mitigation strategies against behavioral drifts of these black-box ML models to prevent their use under suspicious conditions. They point out that design optimization studies are actually driven by an end-to-end automated process, which results in one best configuration communicated to managers and related engineering teams. Thus, such 3-5% of unreliable predictions can enable the accelerated MDOs to reach feasible solutions at an early stage, however, the inefficiency of these solutions might remain hidden for quite a while.

*RQ4. What are the benefits of deploying SmOOD in hybrid surrogate optimization settings?*

**Motivation.** We aim to examine the effects of *SmOOD* on the prediction errors and computation time when it is combined with the FNN for hybrid surrogate aircraft design optimization.

**Method.** We compute the NRMSE and inference runtime values for all the surrogate variants: GP, FNN, HybridGP (GP + Baseline) and HybridFNN (FNN + *SmOOD*). Then, we derive their SpeedUp relative to the HF model, as well as, the Decrease rate of NRMSE achieved by the hybrid versions w.r.t. the pure surrogate counterparts. As a way to compare the surrogate model types, we calculated, separately, the computation durations required for all the design and evaluation steps (training, tuning, validation, and testing) for each type of model. Furthermore, we compute the same computation times for the design of OOD detection strategies to determine the overhead added when using an OOD detector to switch between

HF model and its surrogate counterpart in hybrid surrogate design optimization settings.

**Table 3: Improvement Evaluations for different pairs of QoI and Surrogate Model.**

QoI	Model	Test_ERR	%ERR_Decr	Runtime	SpeedUp
MTOW	HF	-	-	02h:57m:58s	1
	GP	0.0363	-	299.26ms	$3.57 \times 10^4$
	HybridGP	0.036	<b>0.83%</b>	07m:38s	$2.33 \times 10$
	FNN	0.0319	-	0.81ms	$1.32 \times 10^7$
TTC	HybridFNN	0.0169	<b>47.02%</b>	02m:38s	$6.76 \times 10$
	HF	-	-	02h:57m:58s	1
	GP	0.1698	-	289.47ms	$3.69 \times 10^4$
	HybridGP	0.1677	<b>1.24%</b>	10m:48s	$1.65 \times 10$
BFL	FNN	0.0803	-	15.51ms	$6.88 \times 10^3$
	HybridFNN	0.0441	<b>45.08%</b>	12m:31s	$1.42 \times 10$
	HF	-	-	02h:57m:58s	1
	GP	0.0439	-	290.07ms	$3.68 \times 10^4$
BFL	HybridGP	0.0435	<b>0.91%</b>	08m:32s	$2.08 \times 10$
	FNN	0.0361	-	4.58ms	$2.33 \times 10^6$
	HybridFNN	0.0317	<b>12.19%</b>	01m:54s	$1.42 \times 10$

**Table 4: Computation Times for different steps of Surrogate Model Design.**

QoI	Model	Train_Time	Tune_Time	Eval_Time	Test_Time
MTOW	GP	02m:50s	-	27m:28s	299.26ms
	FNN	00m:17s	05h:41m:39s	02m:53s	0.81ms
TTC	GP	04m:25s	-	42m:07s	289.47ms
	FNN	02m:10s	22h:49m:54s	21m:04s	15.51ms
BFL	GP	03m:11s	-	31m:12s	290.07ms
	FNN	00m:54s	20h:43m:09s	08m:46s	4.58ms

**Results.** According to Table 3, the inclusion of *SmOOD* with FNN surrogates contributed to significant decrease ratios in prediction error, from 12% to 47%, while the baseline leads to a maximum of 1.25% decrease rate on the GP's prediction errors. Hence, *SmOOD* was able to cover many problematic OOD inputs with high errors that lead to considerable reduction in FNN prediction errors. Besides, the speed up rates obtained by *SmOOD* plus FNN surrogates are mostly higher than those obtained by Baseline plus GP counterparts. Indeed, *SmOOD* was able to detect the underlying OODs more accurately with fewer false alarms than the baseline, which turns almost three hours of runtime to a few minutes. FNN surrogates have already better speed up in the inference time than their GP counterparts (see Table 3). Table 4 also reports lower values of training time, validation time and test time obtained by FNN surrogates compared to those yielded by GP counterparts. This can be explained by the FNN's deterministic mapping function that includes consecutive weighted sums and ReLU activations, and by its learning algorithm that applies loss gradients w.r.t. the parameters to iteratively update them. On the other hand, GP's mapping function is stochastic including multivariate posterior gaussian distributions, and its approximate Bayesian learning that updates our prior belief in our gaussian parameters in line with marginal log-likelihood estimates on the observed data. Nonetheless, Table 4 shows that only FNN surrogates require a time-consuming and hand-crafted hyperparameters tuning to select the structure of the

model (i.e., depth and width), configure the optimizer, adapt the regularizer strength, etc. All these choices are guided by trial-and-error processes on validation set. In contrast, the GP surrogate design is straightforward, and the kernel hyperparameters are determined systematically within the Bayesian posteriori optimization. The tuning of FNNs is common for any DL solution; so GPU-enabled parallelism and multi-machine distribution is straightforward when using modern DL frameworks.

**Table 5: Computation Times for different steps of OOD Detection Method Design.**

Method	Train_Time	Tune_Time	Valid_Time	Test_Time
<b>Baseline</b>	0.32s	178.99s	3.17s	271.08ms
<b>SmOOD</b>	0.43s	12.47s	4.22s	2.16ms

Table 5 demonstrates that *SmOOD* outperforms the baseline in terms of low design and evaluation workloads. The reason is the fast computation of predictions and derivatives using FNN surrogates to construct the local sensitivity profiles, and the ease with which the *SmOOD* inherent classifier and oversampler can be trained and tuned due to the low dimensionality of the precomputing sensitivity profiles. This is contrary to the baseline which requires repeated distance calculations between a requested design configuration and all the validation set to identify the nearest neighbors.

**Finding 4:** *SmOOD* plus FNN surrogate enables accelerated and accurate hybrid optimization, achieving, on average, 34.65% and 58.36× of decrease error rate and computation speed up rate.

**Domain Expert Feedback.** According to domain experts, as long as the tuning process is automated, the FNN is still viable given the achieved error rate regardless of the build time. Furthermore, they emphasize that these FNN surrogates are mapping aircraft design variables to a particular aircraft performance attribute. Hence, an optimized FNN surrogate can be involved in many design optimization studies, which compensates for its relatively-heavy creation procedure. Most changes applied to the HF simulations are adjustments to the design variables in order to accommodate new requirements recommended by marketing analysts. This may result in further training on FNN parameters rather than structural/hyperparameters tuning. Creating FNN surrogates from scratch can be only done in response to less-frequent, major updates to HF model inner functions and capabilities. Nevertheless, experts suggest taking full advantage of the hybrid mode and storing the detected OOD inputs to further fine-tune the FNN after every optimization cycle since the training is relatively fast.

## 5 CONCLUSION

In this research work, we demonstrate that feedforward neural networks outperforms Gaussian processes across several surrogate models of aircraft design performance, especially when the predicted quantity is highly nonlinear. These high-capacity black box models are prone to out-of-distribution issues, which restricts their direct use in aircraft design optimization. Due to the risk of overconfident uncertainties against these OOD samples, SmOOD

effectively spots them by identifying their suspicious local sensitivity behaviors that are far from the level of FNN smoothness observed across validation. Statistically, SmOOD reveals up to 93% of OOD samples, and its use as router in hybrid surrogate aircraft design optimization leads to 34.65% and 58.36× of error reduction and runtime speed up rates. Our empirical evaluation reinforces our prior belief in the a priori smoothness that exists over the HF simulations data. Indeed, we expect similar priors to exist in many other applications, on which our findings can be extrapolated, and our local sensitivity profiling can serve as a surrogate for model uncertainty, and as a discriminatory criterion to separate ID from OOD. The next step will be to develop an active version of SmOOD that stores the revealed OODs to fine-tune the FNN and expand further the ID regions, which results in fewer requests to the HF model and faster design assessments.

## ACKNOWLEDGMENT

This work is supported by the DEEL project CRDPJ 537462-18 funded by the National Science and Engineering Research Council of Canada (NSERC) and the Consortium for Research and Innovation in Aerospace in Québec (CRIAQ), together with its industrial partners Thales Canada inc, Bell Textron Canada Limited, CAE inc and Bombardier inc.

## REFERENCES

- [1] MYM Ahmed and Ning Qin. 2009. Surrogate-based aerodynamic design optimization: Use of surrogates in aerodynamic design optimization. In *International Conference on Aerospace Sciences and Aviation Technology*, Vol. 13. The Military Technical College, 1–26.
- [2] John David Anderson. 1999. *Aircraft performance and design*. Vol. 1. WCB/McGraw-Hill Boston.
- [3] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893* (2019).
- [4] Charles Audet and Michael Kokkolaras. 2016. Blackbox and derivative-free optimization: theory, algorithms and applications. , 2 pages.
- [5] George Bebis and Michael Georgiopoulos. 1994. Feed-forward neural networks. *IEEE Potentials* 13, 4 (1994), 27–31.
- [6] Gérard Biau and Erwan Scornet. 2016. A random forest guided tour. *Test* 25, 2 (2016), 197–227.
- [7] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. Weight uncertainty in neural network. In *International conference on machine learning*. PMLR, 1613–1622.
- [8] Martin D Buhmann. 2003. *Radial basis functions: theory and implementations*. Vol. 12. Cambridge university press.
- [9] Saikiran Bulusu, Bhavya Kaikhura, Bo Li, Pramod K Varshney, and Dawn Song. 2020. Anomalous example detection in deep learning: A survey. *IEEE Access* 8 (2020), 132330–132347.
- [10] Dalu Cao and Guang-Chen Bai. 2020. DNN-Based Surrogate Modeling-Based Feasible Performance Reliability Design Methodology for Aircraft Engine. *IEEE Access* 8 (2020), 229201–229218.
- [11] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16 (2002), 321–357.
- [12] Hyunsun Choi, Eric Jang, and Alexander A Alemi. 2018. Waic, but why? generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392* (2018).
- [13] Jesse Davis and Mark Goadrich. 2006. The relationship between Precision-Recall and ROC curves. In *Proceedings of the 23rd international conference on Machine learning*. 233–240.
- [14] Qiuwan Du, Tianyuan Liu, Like Yang, Liangliang Li, Di Zhang, and Yonghui Xie. 2022. Airfoil design and surrogate modeling for performance prediction based on deep learning method. *Physics of Fluids* 34, 1 (2022), 015111.
- [15] John C Duchi and Hongseok Namkoong. 2021. Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics* 49, 3 (2021), 1378–1406.
- [16] Yarin Gal, Jiri Hron, and Alex Kendall. 2017. Concrete dropout. *Advances in neural information processing systems* 30 (2017).

- [17] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [18] GPy. since 2012. GPy: A Gaussian process framework in python. <http://github.com/SheffieldML/GPy>.
- [19] Hui Han, Wen-Yuan Wang, and Bing-Huan Mao. 2005. Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning. In *International conference on intelligent computing*. Springer, 878–887.
- [20] Seyyed Hossein Hasanpour, Mohammad Rouhani, Mohsen Fayyaz, Mohammad Sabokrou, and Ehsan Adeli. 2018. Towards principled design of deep convolutional networks: Introducing simpnet. *arXiv preprint arXiv:1802.06205* (2018).
- [21] Haibo He, Yang Bai, Edwardo A Garcia, and Shutao Li. 2008. ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In *2008 IEEE international joint conference on neural networks (IEEE world congress on computational intelligence)*. IEEE, 1322–1328.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*. 1026–1034.
- [23] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. 2018. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606* (2018).
- [24] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. 1989. Multilayer feed-forward networks are universal approximators. *Neural networks* 2, 5 (1989), 359–366.
- [25] Emiliano Iuliano and Domenico Quagliarella. 2013. Proper orthogonal decomposition, surrogate modelling and evolutionary optimization in aerodynamic design. *Computers & Fluids* 84 (2013), 327–350.
- [26] Alex Kendall and Yarin Gal. 2017. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems* 30 (2017).
- [27] Durk P Kingma, Tim Salimans, and Max Welling. 2015. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems* 28 (2015).
- [28] Timoleon Kipouros, Massimiliano Molinari, William N Dawes, Geoffrey T Parks, Mark Savill, and Karl W Jenkins. 2007. An investigation of the potential for enhancing the computational turbomachinery design cycle using surrogate models and high performance parallelisation. In *Turbo Expo: Power for Land, Sea, and Air*, Vol. 47950. 1415–1424.
- [29] Spyridon G Kontogiannis, A Mark Savill, and Timoleon Kipouros. 2017. A Multi-Objective Multi-Fidelity framework for global optimization. In *58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. 0136.
- [30] György Kovács. 2019. An empirical comparison and evaluation of minority oversampling techniques on a large number of imbalanced datasets. *Applied Soft Computing* 83 (2019), 105662.
- [31] Daniel G Krige. 1951. A statistical approach to some basic mine valuation problems on the Witwatersrand. *Journal of the Southern African Institute of Mining and Metallurgy* 52, 6 (1951), 119–139.
- [32] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012).
- [33] Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. 2018. Accurate uncertainties for deep learning using calibrated regression. In *International conference on machine learning*. PMLR, 2796–2804.
- [34] Yuichi Kuya, Kenji Takeda, Xin Zhang, and Alexander IJ Forrester. 2011. Multifidelity surrogate modeling of experimental and computational aerodynamic data sets. *AIAA journal* 49, 2 (2011), 289–298.
- [35] Tian Li, Ahmad Beirami, Maziar Sanjabi, and Virginia Smith. 2020. Tilted empirical risk minimization. *arXiv preprint arXiv:2007.01162* (2020).
- [36] Dennis J Linse and Robert F Stengel. 1993. Identification of aerodynamic coefficients using computational neural networks. *Journal of Guidance, Control, and Dynamics* 16, 6 (1993), 1018–1025.
- [37] Jiashuo Liu, Zheyuan Hu, Peng Cui, Bo Li, and Zheyuan Shen. 2021. Heterogeneous risk minimization. In *International Conference on Machine Learning*. PMLR, 6804–6814.
- [38] Wei-Liem Loh. 1996. On Latin hypercube sampling. *The annals of statistics* 24, 5 (1996), 2058–2080.
- [39] Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017).
- [40] Lu Mi, Hao Wang, Yonglong Tian, Hao He, and Nir Shavit. 2019. Training-free uncertainty estimation for dense regression: Sensitivity as a surrogate. *arXiv preprint arXiv:1910.04858* (2019).
- [41] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. 2018. Do deep generative models know what they don't know? *arXiv preprint arXiv:1810.09136* (2018).
- [42] Alexey Natekin and Alois Knoll. 2013. Gradient boosting machines, a tutorial. *Frontiers in neurorobotics* 7 (2013), 21.
- [43] William S Noble. 2006. What is a support vector machine? *Nature biotechnology* 24, 12 (2006), 1565–1567.
- [44] Roman Novak, Yasaman Bahri, Daniel A Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. 2018. Sensitivity and generalization in neural networks: an empirical study. *arXiv preprint arXiv:1802.08760* (2018).
- [45] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. 2019. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems* 32 (2019).
- [46] Giampaolo Pagliuca, Timoleon Kipouros, and MA Savill. 2019. Surrogate modelling for wing planform multidisciplinary optimisation using model-based engineering. *International Journal of Aerospace Engineering* 2019 (2019).
- [47] Giampaolo Pagliuca and Sebastian Timme. 2017. Model reduction for flight dynamics simulations using computational fluid dynamics. *Aerospace Science and Technology* 69 (2017), 15–26.
- [48] Thierry Pardessus. 2004. Concurrent engineering development and practices for aircraft design at Airbus. In *Proceedings of the 24th ICAS Conf., Yokohama, Japan*, Vol. 1.
- [49] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*. 8026–8037.
- [50] J Slotnick PI, A Khodadoust PM, J Alonso, D Darmofal, W Gropp, E Lurie, and D Mavriplis. 2013. CFD vision 2030 study: a path to revolutionary computational aerosciences. (2013).
- [51] Stanislav Pidhorskyi, Ranya Almohsen, and Gianfranco Doretto. 2018. Generative probabilistic novelty detection with adversarial autoencoders. *Advances in neural information processing systems* 31 (2018).
- [52] Man Mohan Rai and Nateri K Madavan. 2000. Aerodynamic design using neural networks. *AIAA journal* 38, 1 (2000), 173–182.
- [53] Dushyanth Rajaram, Tejas G Puranik, Ashwin Renganathan, Woong Je Sung, Olivia J Pinon-Fischer, Dimitri N Mavris, and Arun Ramamurthy. 2020. Deep Gaussian process enabled surrogate models for aerodynamic flows. In *AIAA Scitech 2020 Forum*. 1640.
- [54] Mihaela Rosca, Theophane Weber, Arthur Gretton, and Shakir Mohamed. 2020. A case for new neural network smoothness constraints. (2020).
- [55] Mohammadreza Salehi, Hossein Mirzaei, Dan Hendrycks, Yixuan Li, Mohammad Hossein Rohban, and Mohammad Sabokrou. 2021. A Unified Survey on Anomaly, Novelty, Open-Set, and Out-of-Distribution Detection: Solutions and Future Challenges. *arXiv preprint arXiv:2110.14051* (2021).
- [56] Valerie Sarge, Michael Andersch, Lynsey Fabel, Paulius Micikevicius, and John Tran. 2019. *Tips for Optimizing GPU Performance Using Tensor Cores*. <https://developer.nvidia.com/blog/optimizing-gpu-performance-tensor-cores/>
- [57] Timothy W Simpson, Timothy M Maury, John J Korte, and Farrokh Mistree. 2001. Kriging models for global approximation in simulation-based multidisciplinary design optimization. *AIAA journal* 39, 12 (2001), 2233–2241.
- [58] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. 2017. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825* (2017).
- [59] András Sobester, Alexander Forrester, and Andy Keane. 2008. *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons.
- [60] Jaroslaw Sobieszcanski-Sobieski and Raphael T Haftka. 1997. Multidisciplinary aerospace design optimization: survey of recent developments. *Structural optimization* 14, 1 (1997), 1–23.
- [61] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [62] Gang Sun and Shuyue Wang. 2019. A review of the artificial neural network surrogate modeling in aerodynamic design. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering* 233, 16 (2019), 5863–5872.
- [63] Yuchun Tang, Yan-Qing Zhang, Nitesh V Chawla, and Sven Krasser. 2008. SVMs modeling for highly imbalanced classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 39, 1 (2008), 281–288.
- [64] Wendl. 2004. *Scattered Data Approximation (Cambridge Monographs on Applied and Computational Mathematics; 17)*. Cambridge University Press.
- [65] Raymond E Wright. 1995. Logistic regression. (1995).
- [66] Jing Kang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. 2021. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334* (2021).
- [67] Liang Zhao, KK Choi, and Ikjin Lee. 2011. Metamodeling method using dynamic kriging for design optimization. *AIAA journal* 49, 9 (2011), 2034–2046.