

Machine-learning-assisted continuous inverse design of multiring optical fibres for beam shaping

Marta Dziuba-Kozieł^{1*} 

¹ Faculty of Electrical Engineering and Computer Science, Lublin University of Technology,
 Nadbystrzycka 38A, 20-618 Lublin, Poland
 E-mail: m.dziuba@pollub.pl

ABSTRACT

This article presents a method for inverse design of multiring optical fibres using machine learning. It proposes fibre configurations expected to yield specified beam-shaping parameters and verifies selected off-grid configurations using independent OptiFiber simulations. The proposed method creates machine-learning surrogate models from numerical simulations performed in OptiFiber. It utilises differential evolution optimisation to predict and optimise the refractive index (RI) profiles of individual fibre layers for given beam parameters. The study reused a dataset comprising 721 valid three-layer optical-fibre configurations previously generated using OptiFiber. For each fibre configuration, the output light beam was simulated, and the full width at half maximum (FWHM), beam flatness and edge steepness were calculated. As the direct prediction of fibre configuration from beam parameters is ambiguous, a surrogate model was employed for subsequent use. The feasibility of using several regression models to create a surrogate model was analysed, including ExtraTrees, Random Forest, XGBoost, LightGBM, CatBoost, support vector regression, multi-layer perceptron, ridge regression and the k-nearest neighbours method. Ensemble tree-based models yielded the best predictive performance, with the ExtraTrees model providing the highest overall accuracy and a high degree of agreement between predicted and simulated beam parameters. An ablation study demonstrated that prediction accuracy can be improved by extending the feature set used to train the surrogate model to include descriptors of the differences between the RI values of individual optical fibre layers. The employed inverse design procedure enabled continuous estimation of the optical fibre configuration, reducing the need for exhaustive searches based on time-consuming simulations. The final inverse-design model was independently validated using 30 new OptiFiber simulations corresponding to continuous refractive-index configurations that were not present in the original simulation grid. For these off-grid cases, the ExtraTrees surrogate achieved R^2 values of 0.9982, 0.9922, and 0.9912 for FWHM, beam flatness, and edge steepness, respectively. The results demonstrate that the surrogate-assisted optimisation can generate previously unseen off-grid fibre configurations whose beam parameters are reproduced with high accuracy in independent OptiFiber simulations. The conclusions are restricted to interpolation within the investigated RI range of 1.447–1.455.

Keywords: machine learning, inverse design, optical fibres, beam shaping, surrogate modelling, refractive-index engineering, photonics, artificial intelligence

INTRODUCTION

Beam shaping in optical fibres is a subject of considerable interest due to its importance in laser systems [1,2], optical sensors [3,4], biomedical imaging [5], optical trapping [6], materials processing [7,8], and telecommunications [9]. In many practical applications, the shape of the light

beam influences the system's operating parameters. This requires control of the beam width, its flatness and the edge steepness [10,11]. Multiring cylindrical optical fibres can be used for this purpose, as the appropriate selection of refractive index (RI) values for the individual layers of the optical fibre allows a beam of the desired shape to be obtained [12]. However, the design of optical

fibre structures generating specified beam profiles remains a complex inverse problem due to the non-linear relationship between the RI distribution and the resulting shape of the light beam propagating through the optical fibre.

Traditional photonic design approaches typically rely on repeated numerical simulations combined with manual parameter tuning. Although simulation tools such as OptiFiber enable accurate modelling of optical propagation in complex fibre structures, exhaustive searches across multidimensional *RI* spaces are computationally expensive and become increasingly impractical for larger design problems. The inverse mapping from beam parameters to fibre configurations is also generally non-unique, meaning that multiple *RI* distributions may produce similar beam profiles. These limitations motivate the development of computationally efficient inverse-design methodologies capable of exploring large continuous parameter spaces [12].

Inverse design has become an important research direction in modern photonics. In integrated photonics and nanophotonics, optimisation-based inverse-design approaches are increasingly used for designing compact optical devices with complex functionality [13,14]. However, inverse design of optical fibres remains less explored, particularly for beam-shaping applications involving continuous RI optimisation. Existing fibre-design studies have focused mainly on modal engineering, dispersion optimisation, and few-mode propagation problems.

One of the most relevant studies related to the present work proposes a machine-learning-assisted design framework for inverted-gradient multiring optical fibres intended for flat-top beam generation. That work demonstrates that ensemble learning methods can successfully approximate the relationship between *RI* configurations and beam parameters such as flatness, edge steepness, and FWHM [12]. The authors proposed a random forest for predicting beam parameters based on the fibre configuration and demonstrated that this model can significantly accelerate beam shape optimisation compared with exhaustive simulation-based approaches. However, the study focused primarily on predictive performance and accuracy, while paying limited attention to error analysis, response-surface behaviour, and feature-ablation analysis. A similar direction was investigated in [15], where a machine-learning-aided inverse-design approach for flat-top beam fibres

using artificial neural networks was proposed. This method enabled the prediction of fibre configurations corresponding to desired beam parameters, demonstrating that neural network-based surrogates can effectively replace computationally intensive optical simulations during optimisation. However, the study considered a narrower set of beam descriptors and focused mainly on ANN-based modelling without broader comparison of different surrogate architectures.

Machine-learning-assisted inverse design has also been investigated for other types of optical fibres. In [16], the authors applied machine learning to the inverse design of ring-assisted few-mode fibres for weak-coupling optimisation, showing that surrogate-assisted optimisation can efficiently identify fibre structures with reduced intermodal coupling. Similarly, in [17] a tandem-neural-network framework for inverse design of hollow-core anti-resonant fibres was proposed. It demonstrates the applicability of deep-learning-based inverse modelling to more complex fibre geometries. This approach was extended in [18] where neural networks for inverse dispersion engineering to design photonic-crystal fibres were used.

The discussed studies confirm that machine learning is becoming an increasingly important tool for solving inverse photonic problems involving high-dimensional nonlinear parameter spaces.

Recent advances in artificial intelligence further demonstrate the effectiveness of deep learning and ensemble learning in modelling highly nonlinear systems. These methods are successfully applied in medicine [19], signal processing [20], engineering [21], and many other fields [22–24]. Some of these solutions can learn from small datasets, demonstrating the ability of modern AI architectures to learn complex multidimensional relationships from limited datasets, provided the feature set is appropriately adjusted. The stability of selected features strongly influences model robustness and interpretability in multidimensional data analysis [25]. In inverse design of optical fibres, the correct selection of characteristics is crucial for capturing the physically significant relationships between the fibre configuration and the shape of the light beam propagating through it. This is particularly important when, in addition to the refractive index values, derived descriptors are introduced to describe the relationships among them [12].

Despite recent progress in optical beam shaping, several limitations remain underexplored

in the literature. Most works focus primarily on forward-prediction accuracy, with limited investigation of surrogate reliability across the parameter space. Little attention is devoted to analysing how different feature representations affect inverse-design performance. Fibre inverse-design studies examine only a single surrogate architecture, making it difficult to assess trade-offs among predictive accuracy, smoothness of learned response surfaces, and optimisation stability. Independent simulation-based validation of continuously optimised, off-grid fibre configurations remains insufficiently addressed in optical-fibre inverse-design studies. In particular, low surrogate-model prediction errors on randomly selected test samples do not by themselves demonstrate that a continuously optimised configuration will reproduce the requested beam parameters in a new optical simulation.

Based on the identified research gap, the study addresses the following research questions:

- RQ1: Which surrogate-model architecture provides the most accurate prediction of FWHM, beam flatness, and edge steepness for the investigated multiring optical fibre?
- RQ2: How does the inclusion of derived RI-difference descriptors affect surrogate-model accuracy compared with using the original RI values alone?
- RQ3: Can a surrogate-assisted Differential Evolution procedure generate previously unseen off-grid RI configurations that reproduce the target beam parameters in independent OptiFiber simulations?

To answer these research questions, this study develops an integrated surrogate-assisted inverse-design workflow for multiring optical fibres. The workflow combines systematic benchmarking and hyperparameter optimisation of multiple surrogate architectures, feature-representation ablation, Differential Evolution-based continuous optimisation, response-surface analysis, and independent validation using new off-grid OptiFiber simulations.

The present study reuses the 721 valid OptiFiber simulations previously reported in [12]. Therefore, neither the original simulation dataset nor the definitions of the RI-based descriptors RI_1 – RI_3 and α_1 – α_3 are claimed as new contributions. In [12], these data were used primarily to train a Random Forest forward model, analyse feature importance, and identify favourable configurations within a discrete RI grid. In contrast,

the present work uses the dataset to compare and tune multiple surrogate architectures, assess alternative feature representations, and perform continuous surrogate-assisted inverse design. The resulting off-grid configurations are subsequently verified using 30 new OptiFiber simulations that were not used for model development.

The scientific novelty of this study lies in demonstrating, through independent optical simulations, that a systematically selected and tuned surrogate model can be coupled with Differential Evolution to generate previously unseen continuous refractive-index configurations that reproduce specified beam-shaping parameters. Unlike previous studies, the proposed framework jointly evaluates surrogate architecture, feature representation, response-surface behaviour, optimisation performance, and off-grid generalisation within a single validated inverse-design methodology. Thus, the novelty is methodological and validation-oriented rather than associated with the original simulation dataset or previously defined RI descriptors.

The main contributions of this work are:

- systematic tuning and comparison of multiple surrogate-model architectures using a common development and hold-out evaluation procedure,
- evaluation of alternative representations of the previously defined RI-based descriptors through a five-fold ablation study,
- continuous inverse design of off-grid RI configurations using Differential Evolution;
- explicit separation of hold-out testing, cross-validation stability analysis, reconstruction testing, and independent external validation,
- independent confirmation of continuous inverse-design solutions using 30 new OptiFiber simulations and local perturbation tests,
- comparative visualisation of response surfaces produced by different surrogate models.

The results demonstrate that machine-learning surrogates can effectively accelerate inverse photonic design while maintaining high predictive accuracy, enabling efficient exploration of multidimensional RI parameter spaces for beam-shaping applications. The rest of the paper is organised as follows: Section 2 describes the multiring optical fibre structure under investigation, the OptiFiber simulation procedure, the extracted beam parameters, the machine-learning surrogate models, and the proposed inverse-design framework. Section 3 presents the results,

including forward-prediction performance, a comparison of surrogate models, inverse-design validation, ablation analysis and visualisation of the learned response surfaces. Section 4 discusses the results, the limitations of the proposed approach, and potential future developments. Finally, Section 5 summarises the main conclusions of the study.

MATERIALS AND METHODS

Fibre configuration

The investigated structure was a multiring optical fibre composed of three concentric doped-glass layers with independently controlled refractive indices (RI). The structure also included pure silica regions with a refractive index of 1.447. Each ring layer was 2 μm thick, while the fibre core diameter was 4 μm . Additionally, the fibre was surrounded by a 10 μm thick pure-silica cladding with a refractive index of 1.447.

The fibre geometry was designed to enable beam shaping through refractive-index engineering of individual rings. In the investigated configuration, the refractive indices of the three layers, denoted RI_1 , RI_2 , and RI_3 , served as the primary design variables for the machine learning and inverse-design procedures.

A schematic representation of the investigated fibre structure is shown in Figure 1, and its parameters are listed in Table 1.

The refractive-index distribution within the fibre was modified by varying the RI values of the individual layers while preserving the geometrical dimensions of the structure. Such an approach enabled systematic investigation of the relationship between RI profile configuration and the resulting output beam characteristics.

Features used

The machine learning models were trained using descriptors derived from the refractive-index distribution of the investigated multiring optical fibre. We considered features identified as having a significant influence on the beam shape [12]. The input feature set included both the direct RI values of individual layers and additional physically motivated descriptors representing refractive-index contrasts between neighbouring regions:




Figure 1. Fibre scheme

$$\alpha_1 = RI_2 - RI_3 \quad (1)$$

$$\alpha_2 = RI_1 - RI_2 \quad (2)$$

$$\alpha_3 = (RI_1 - RI_3) \quad (3)$$

The introduced descriptors capture local and global RI gradients within the fibre structure and were intended to provide machine learning models with physically meaningful information on modal confinement and beam-propagation behaviour.

The performed analyses demonstrated that introducing RI -difference descriptors significantly improved predictive performance compared to using only the original RI values. The descriptors a_1 , a_2 and a_3 provided the best trade-off between model complexity and prediction accuracy.

The final surrogate models used for inverse design were therefore trained using the feature set composed of: RI_1 , RI_2 , RI_3 , a_1 , a_2 and a_3 . This representation enabled efficient learning of nonlinear relationships between fibre configuration and beam-shaping parameters.

Beam parameters

The optical beam distributions obtained from OptiFiber simulations were quantitatively characterised using three beam-shaping parameters:

- full width at half maximum (FWHM),
- beam flatness,
- edge steepness.

The parameters were calculated using the intensity profile extracted along the central

Table 1. Parameters of the investigated multiring optical fibre

Parameter	Value
Core diameter	4 μm
Ring thickness	2 μm
Cladding material and thickness	pure silica, 10 μm
RI range	1.447–1.455
Number of variable RI layers	3

diameter of the simulated beam distribution. The analysis was performed on the middle row of the intensity matrix generated by OptiFiber. FWHM is the width of a beam at half of its maximum height, where height represents light power. Edge steepness characterises the sharpness of beam edges and is calculated as a normalised difference between irradiation area at 10% of beam height ($A_{0.1}^i$) and 90% of beam height ($A_{0.9}^i$), and is calculated with the formula (ISO 13694:2018(en)):

$$s(z) = \frac{A_{0.1}^i - A_{0.9}^i}{A_{0.1}^i} \quad (4)$$

Beam flatness describes the uniformity of the beam intensity distribution within the central region of the beam. It is the average of all measured beam power samples greater than $1/e^2$. The method for calculating beam parameters is shown in Figure 2. Samples with values greater than $1/e^2$ are marked in orange.

The selected parameters jointly describe the most important beam-shaping characteristics of the investigated optical fibre structures and were

therefore used as target outputs for both surrogate modelling and inverse-design optimisation.

Data collection

The dataset used in the present study comprises 721 valid OptiFiber simulations previously generated and reported in [12]. The 30 off-grid OptiFiber simulations used for external validation were newly generated for the present study and were not part of the dataset reported in [12]. The same underlying simulation dataset was reused here to investigate research questions not addressed in the earlier publication, including systematic comparison of surrogate-model architectures, feature-ablation analysis, Differential Evolution-based inverse design, spatial analysis of prediction errors, and continuous off-grid optimisation.

All machine-learning experiments, train-test partitions, model comparisons, ablation analyses, optimisation procedures, and visualisations presented in this article were performed specifically for the present study.

The dataset used in this study was generated using OptiFiber simulations performed for different refractive-index configurations of the investigated multiring optical fibre. The simulations were conducted by systematically varying the RI_1 , RI_2 , and RI_3 while preserving all geometrical parameters of the structure. The refractive indices were sampled from the range: $RI \in [1.447, 1.455]$ using a discretisation step of 0.001. This resulted in a theoretical parameter space containing 729 possible configurations for the three-layer fibre.




Figure 2. Example beam power profile

Due to numerical convergence limitations of the OptiFiber software, valid simulation results were obtained for 721 fibre configurations. Only successfully simulated cases were included in the final dataset used for machine learning and inverse-design procedures. Each simulation produced two files: one containing the fibre's refractive-index configuration, and the other containing the simulated optical beam intensity distribution, represented as a 100×100 numerical matrix describing the spatial distribution of optical intensity at the fibre output. The beam characteristics were subsequently extracted from the central beam profile corresponding to the matrix diameter.

For each simulation, the following output beam parameters were calculated: FWHM, beam flatness, and edge steepness. Additionally, the set of input parameters was extended by a_1 , a_2 and a_3 values. The generated dataset was subsequently used for training and evaluating machine learning surrogate models as well as for inverse-design optimisation.

Machine learning surrogate models

Machine-learning surrogate models were used to approximate the nonlinear relationship between the refractive-index configuration of a multiring optical fibre and the resulting beam-shaping parameters. The models were trained to predict three output parameters simultaneously: *FWHM*, *flatness*, and *steepness* on the basis of input parameters: RI_1 , RI_2 , RI_3 , a_1 , a_2 and a_3 .

A comparative analysis of multiple regression models was performed in order to identify the most suitable surrogate architecture for inverse

photonic design. The investigated models and their main hyperparameters are listed in Table 2. They represent different machine learning paradigms, including:

- ensemble tree-based learning,
- boosting methods,
- kernel-based regression,
- neural networks,
- linear models,
- instance-based learning.

The dataset was randomly split into training and test sets at 80/20. In addition, k-fold cross-validation was applied during model evaluation to reduce dependence on a particular train-test partition and to obtain statistically reliable performance estimates.

The surrogate models were evaluated using the following metrics: coefficient of determination (R^2), mean absolute error (*MAE*), mean squared error (*MSE*), and root mean squared error (*RMSE*). Each metric was computed independently for each predicted beam parameter.

Tree-based ensemble methods were expected to perform particularly well due to their ability to model highly nonlinear relationships and interactions among *RI* descriptors. At the same time, smoother models such as SVR and MLP were also investigated for their potentially improved behaviour during continuous inverse optimisation.

The trained surrogate models were subsequently used as forward predictors within the inverse-design optimisation framework. Rather than running computationally expensive OptiFiber simulations during optimisation, the machine-learning models provided rapid approximations

Table 2. Machine-learning models and hyperparameters selected using RandomizedSearchCV

Model	Main hyperparameters
Random forest (RF)	676 trees; depth = unlimited; features = 0.5; minimum leaf size = 1; bootstrap enabled
ExtraTrees	301 trees; depth = unlimited; features = 0.75; minimum leaf size = 1; bootstrap disabled
XGBoost	773 estimators; learning rate = 0.0133; depth = 8; subsample = 0.8913
LightGBM	430 estimators; learning rate = 0.1123; 47 leaves; depth = 8
CatBoost	765 iterations; depth = 8; learning rate = 0.0174; L2 leaf regularisation = 0.6378
Gradient boost regr. (GBR)	534 estimators; learning rate = 0.0343; depth = 6; subsample = 0.7719
HistGradientBoosting regressor (HistGBR)	670 iterations; learning rate = 0.0105; 38 leaf nodes; minimum leaf size = 5
Support vector regr. (SVR)	C = 0.7069, gamma = 1.0, epsilon = 3.983×10^{-4}
Multilayer perceptron (MLP)	Hidden layers = (128, 64); ReLU; alpha= 3.588×10^{-5} ; learning rate = 0.002879
Ridge regression	alpha = 7.177
K-nearest neighbours (KNN)	n_neighbors = 6, weights = uniform, p = 1, leaf_size = 17

of beam parameters for arbitrary RI configurations across the design space under investigation.

This approach significantly reduced computational complexity and enabled efficient exploration of continuous RI values not explicitly included in the original simulation dataset.

Model selection, cross-validation, and final model fitting

The 721 valid simulation cases were randomly split into a development subset containing 80% of the data and an independent hold-out test subset containing 20%. The split was performed once, with $\text{random_state} = 42$. The same random seed was used in all stochastic procedures wherever supported to improve the reproducibility of the experiments. The hold-out test subset was not used during hyperparameter selection.

Hyperparameters were selected separately for each surrogate architecture using Randomized-SearchCV. For each model, 40 hyperparameter configurations were sampled. Model selection was performed only within the development subset using shuffled five-fold cross-validation with $\text{random_state} = 42$. The optimisation criterion was the mean normalised root-mean-square error across the three predicted beam parameters.

After hyperparameter selection, the best configuration of each model was fitted to the complete development subset and evaluated once on the untouched 20% hold-out test subset. The hold-out results were used for the final comparison of surrogate architectures.

To assess the stability of the selected Extra-Trees configuration, an additional five-fold cross-validation with shuffling was performed using the fixed hyperparameters selected during the preceding tuning procedure. In each fold, the estimator was cloned and fitted only to the corresponding training folds. This second analysis was used to report variation in the performance metrics across data partitions and was not treated as an additional independent model-selection result.

Finally, after ExtraTrees had been selected as the surrogate used for inverse design, a model with the same hyperparameters was fitted to all 721 original simulation cases. This full-dataset model was used only to generate continuous candidate configurations. Its off-grid predictions were subsequently assessed using 30 completely new OptiFiber simulations that had not been used for model fitting or hyperparameter selection.

All surrogate models were trained to predict FWHM, beam flatness, and edge steepness simultaneously. Models without native multi-output support were implemented using MultiOutputRegressor, and the evaluation metrics were calculated separately for each output parameter. For MLP, SVR, Ridge, and KNN, both input features and target variables were standardised using transformations fitted exclusively to the corresponding training data within each cross-validation fold. The predicted target values were subsequently transformed back to their original scales before the performance metrics were calculated.

Inverse design

The inverse design problem addressed in this study was to determine the refractive-index configuration of the multiring optical fibre that yields the desired beam-shaping parameters. The optimisation objective can be expressed as: $(FWHM, flatness, steepness) \rightarrow (RI_1, RI_2, RI_3)$. RI_1, RI_2, RI_3 values constrained to the interval $1.447 \leq RI \leq 1.455$.

Direct learning of the inverse mapping was not used because the relationship between fibre configuration and beam characteristics is generally non-unique, meaning that multiple RI configurations may yield similar optical beam profiles. Instead, the inverse-design framework was implemented using forward surrogate models and numerical optimisation.

The trained machine-learning surrogate models were used to rapidly predict beam parameters for arbitrary RI configurations across the investigated parameter space. During optimisation, the objective function minimised the normalised discrepancy between target beam parameters and surrogate-model predictions. The loss function was defined as:

$$L = \sum_{i=1}^N w_i \left(\frac{y_i^{pred} - y_i^{target}}{\Delta y_i} \right)^2 \quad (5)$$

where: y_i^{pred} is the predicted value of the i -th beam parameter obtained from the surrogate model, y_i^{target} is the desired target value of the i -th beam parameter specified for the inverse-design optimisation, Δy_i denotes the range of the i -th beam parameter in the training dataset, w_i is the weighting coefficient assigned to the i -th beam parameter, and N is the total number of beam parameters considered in the optimisation (in this study, $N = 3$: $FWHM$, $flatness$, $steepness$).

The loss function is therefore a weighted sum of squared normalised differences between the surrogate predictions and the target values, enabling prioritisation of certain parameters.

Parameter prioritisation allows the designer to assign greater weight to selected beam parameters while reducing the influence of others. This approach is essential because different applications may require optimising specific parameters over others. The inclusion of the normalisation term Δy_i compensates for the differences in numerical scale among *FWHM*, *flatness*, and *steepness*, ensuring balanced contributions of all parameters to the overall loss function.

Next, the optimisation process was performed using the Differential Evolution (DE) algorithm. This method was selected because it does not require gradient information, performs well for nonlinear optimisation problems, is relatively robust against local minima, and can efficiently explore continuous multidimensional parameter spaces [26, 27]. The complete inverse-design workflow is illustrated in Figure 3.

To evaluate optimisation consistency, inverse-design experiments were performed for target beam configurations selected from the original simulation dataset. The 12 reconstruction-validation targets were selected to represent low, medium, and high regions of each output parameter. Cases closest to the 0.1, 0.5, and 0.9 quantiles were initially selected, duplicate cases were removed, and the remaining positions were supplemented with configurations exhibiting high balanced multi-output scores or randomly selected unused cases using `random_state = 42`. Because the inverse mapping is non-unique, the optimised RI configurations were not required to reproduce the original RI values. Instead, their quality was assessed by comparing the target beam parameters with the values predicted by the surrogate model and, for the independent off-grid cases, with the values obtained from new OptiFiber simulations.

Differential Evolution used the best1bin strategy, a population-size multiplier of 12, corresponding to a nominal population of 36 individuals for the three-dimensional RI problem, a maximum of

300 iterations, mutation factors in the range 0.5–1.0, a recombination rate of 0.7, and a convergence tolerance of 10^{-8} . Final solution polishing was enabled. Equal weights were assigned to *FWHM*, *beam flatness*, and *edge steepness*. Ten independent runs were performed for the principal inverse-design problem. For each of the 12 reconstruction-validation targets, three independent optimisation runs were conducted, and the solution with the lowest objective-function value was retained.

Continuous RI optimisation

To refine the inverse design of the multiring optical fibre, continuous optimisation of the refractive index values was performed within the allowed design space. The goal of this process was to generate *RI* configurations that were not explicitly included in the original discrete OptiFiber simulation dataset, thereby enabling exploration of a continuous parameter space. DE efficiently searched the three-dimensional *RI* space $[RI_1, RI_2, RI_3] \in [1.447, 1.455]$, using the surrogate models to rapidly predict the resulting beam parameters for candidate configurations.

The optimisation loss function combined the normalised discrepancies between predicted and target values for *FWHM*, *flatness*, and *steepness*. Weighting coefficients were applied to prioritise specific beam parameters depending on the application requirements. Multiple optimisation runs with different random initialisation seeds were performed to assess solution stability and identify alternative *RI* configurations yielding similar beam characteristics. The resulting RI configurations were used to evaluate the surrogate model's ability to interpolate between the discrete RI levels of the original simulation grid. These optimised configurations serve as proposed fibre designs. This procedure enables inverse design at continuous RI values within the investigated range, while reducing the computational cost compared with exhaustive simulation-based searches.

The optimisation procedure enabled the generation of *RI* configurations not explicitly present in the original OptiFiber simulations.




Figure 3. Inverse design workflow

Independent off-grid OptiFiber validation

To independently validate the continuous inverse-design procedure, 30 additional *RI* configurations were simulated in OptiFiber after model development was complete. These configurations were not included in the original 721-case dataset and contained *RI* values between the discrete levels of the original grid.

The validation set included continuously optimised candidate designs, configurations covering different regions of the output-parameter space, and local perturbations around the selected inverse-design solution. For every new *RI* configuration, the final ExtraTrees surrogate predicted FWHM, beam flatness, and edge steepness. The same three quantities were then independently extracted from the corresponding new OptiFiber beam distribution using the same processing procedure as for the original dataset.

Agreement was quantified using R^2 , mean absolute error, mean squared error, root-mean-square error, maximum absolute error, and signed prediction error. These quantities represent prediction accuracy and were not interpreted as predictive uncertainty. Local consistency was additionally examined around the continuously optimised configuration: $RI_1 = 1.452067$, $RI_2 = 1.449008$, $RI_3 = 1.448929$

Each *RI* component was perturbed separately by ± 0.0001 and ± 0.0002 while the remaining components were kept constant. The surrogate and OptiFiber responses were then compared along the resulting one-dimensional local trajectories. Because FWHM and edge steepness were extracted from a discrete 100-point profile using threshold crossings, small physical changes could result in identical numerical values of these descriptors. The local experiment was therefore interpreted as a consistency check rather than a formal proof of mathematical continuity.

RESULTS

Forward prediction performance

Forward prediction was performed using the ExtraTrees model. The model's performance on the test set was evaluated using standard regression metrics: the coefficient of determination, mean absolute error, mean squared error and root mean squared error. Table 3 presents

the prediction accuracy for three beam parameters: FWHM, beam flatness and edge steepness. The values in Table 3 originate from a single untouched hold-out set. They are therefore different from the mean five-fold cross-validation values reported separately in Table 5.

The analysed surrogate model exhibits high predictive accuracy, particularly for FWHM and beam flatness. The slightly lower accuracy in predicting edge steepness stems from the more complex relationship between this parameter and the refractive index profile.

To visually assess the model's predictive accuracy, 30 off-grid configurations were used to generate parity plots comparing predicted and actual values for each parameter. The results are shown in Figure 4.

An analysis of the results presented in Figure 4 confirms that the ExtraTrees model accurately captures the non-linear relationship between the refractive index profile of a multiring optical fibre and the beam parameters. It therefore constitutes a reliable, simplified model that can be used for further optimisation via the inverse design method.

Comparison of surrogate models

A comprehensive comparison of surrogate models was conducted to assess their suitability for predicting beam-shaping parameters and to support the inverse design process. The comparison included selected machine-learning models representing various algorithmic paradigms, including ensemble tree-based methods, boosting algorithms, kernel-based regression and neural networks. The analysed models are compared in Table 4.

The results show that ensemble tree-based models, particularly ExtraTrees, achieve the highest predictive accuracy for beam flatness and FWHM and deliver stable, reliable results across all analysed beam parameters. CatBoost and XGBoost models produce slightly smoother response surfaces and have shorter prediction times per sample, which may be advantageous in inverse optimisation tasks requiring many computations. Ridge regression performs poorly, confirming that the relationship between the refractive index configuration and beam parameters is highly non-linear.

The independent parity plots presented in Figure 4 show strong agreement between ExtraTrees predictions and the corresponding results of the 30 new off-grid OptiFiber simulations.

Table 3. Performance of the tuned ExtraTrees surrogate on the independent 20% hold-out test set

Parameter	R^2	MAE	MSE	RMSE
FWHM	0.99	0.318	0.3425	0.59
Beam flatness	0.984	0.009	0.0001	0.01
Edge steepness	0.945	0.027	0.0029	0.05

Table 4. Independent hold-out test performance of the tuned surrogate models. Hyperparameters were selected using five-fold cross-validation on the 80% development subset

Model	Beam flatness			Edge steepness			FWHM			Training time per fold [s]	Pred. time / sample [ms]
	R^2	MAE	RMSE	R^2	MAE	RMSE	R^2	MAE	RMSE		
CatBoost	0.983	0.008	0.011	0.947	0.028	0.053	0.99	0.34	0.52	3.185 ± 0.482	0.05
ExtraTrees	0.984	0.009	0.011	0.945	0.027	0.054	0.99	0.318	0.585	0.288 ± 0.034	0.499
GBR	0.976	0.009	0.013	0.951	0.023	0.051	0.99	0.287	0.493	1.690 ± 0.055	0.078
HistGBR	0.979	0.009	0.012	0.946	0.029	0.053	0.99	0.313	0.549	22.016 ± 2.490	0.96
KNN	0.965	0.012	0.016	0.901	0.037	0.07	0.974	0.499	0.928	0.007 ± 0.001	0.025
LightGBM	0.977	0.01	0.013	0.92	0.037	0.063	0.983	0.465	0.751	1.486 ± 0.123	0.093
MLP	0.976	0.01	0.012	0.959	0.026	0.047	0.979	0.516	0.848	0.977 ± 0.261	0.019
RF	0.977	0.01	0.013	0.934	0.03	0.059	0.99	0.371	0.592	0.939 ± 0.088	1.036
Ridge	0.501	0.047	0.061	0.582	0.116	0.149	0.735	2.338	2.98	0.007 ± 0.001	0.014
SVR	0.98	0.009	0.012	0.893	0.041	0.075	0.972	0.479	0.968	0.108 ± 0.008	0.167
XGBoost	0.98	0.009	0.012	0.94	0.03	0.056	0.99	0.346	0.538	3.620 ± 0.313	0.091

Table 5. Five-fold cross-validation stability of the tuned ExtraTrees configuration

Parameter	R^2 mean ± SD	MAE mean ± SD	RMSE mean ± SD
FWHM	0.984 ± 0.008	0.343 ± 0.056	0.683 ± 0.16
Beam flatness	0.9801 ± 0.002	0.009 ± 0.0005	0.0125 ± 0.001
Edge steepness	0.9393 ± 0.022	0.025 ± 0.004	0.0553 ± 0.009




Figure 4. Parity plots for (a) beam flatness, (b) edge steepness and (c) FWHM

Considering its consistently high predictive performance across all three outputs, favourable computational cost, and stable cross-validation results, ExtraTrees was selected as the best overall balanced surrogate for the inverse-design procedure.

ExtraTrees model validation

The selected ExtraTrees surrogate was validated at two complementary levels. First, the stability of the tuned model configuration was assessed using shuffled five-fold cross-validation.

In each fold, the estimator with the previously selected hyperparameters was cloned and fitted only to the corresponding training subset. This analysis was used to quantify sensitivity to data partitioning and was kept separate from the independent 20% hold-out test reported in Table 3.

The five-fold results are summarised in Table 5. The mean R^2 values were 0.9840 for FWHM, 0.9801 for beam flatness, and 0.9393 for edge steepness, with corresponding RMSE values of 0.683, 0.0125, and 0.0553. The relatively small standard deviations indicate stable performance across folds, although edge steepness remained the most difficult parameter to predict.

Although cross-validation confirmed model stability, it did not directly demonstrate accurate prediction for continuous configurations located between the discrete RI levels of the original simulation grid. Therefore, the final ExtraTrees model was additionally evaluated using 30 new OptiFiber simulations that were not included in the original 721-case dataset and were not used during model fitting, hyperparameter selection, or cross-validation.

The independent validation set included off-grid configurations generated through Differential Evolution, cases covering different regions of the output space, and local perturbations around the primary inverse-design solution. The validation results are presented in Table 6. The model achieved R^2 values of 0.9982 for FWHM, 0.9922 for beam flatness, and 0.9912 for edge steepness. The corresponding normalised RMSE values were 1.76%, 3.27%, and 4.24%, respectively.

The mean signed error for FWHM was close to zero, indicating no substantial systematic bias. Beam flatness showed a slight average underestimation, whereas edge steepness showed a small tendency toward overestimation and the largest maximum absolute error. This is consistent with its lower cross-validation accuracy.

The primary continuously optimised configuration was: $RI_1 = 1.452067$, $RI_2 = 1.449008$, $RI_3 = 1.448929$.

For the target values FWHM = 21.0, flatness = 0.78, and edge steepness = 0.33, the surrogate predicted approximately 21.015, 0.780, and 0.330, respectively. The corresponding independent OptiFiber simulation produced approximately 21.010, 0.779, and 0.328. Thus, the principal off-grid solution generated through continuous optimisation was directly confirmed by a new optical simulation.

Local consistency was also examined by independently perturbing each RI component by ± 0.0001 and ± 0.0002 around the selected solution. Figure 5 compares the resulting beam-flatness values obtained from the surrogate and OptiFiber simulations.

The surrogate generally reproduced the local flatness trend, with the strongest agreement observed for RI_3 perturbations. Some discrepancies occurred for positive changes in RI_2 and RI_3 , indicating that the model captures the overall local response but not every small-scale variation.

The OptiFiber-derived FWHM and edge-steepness values remained unchanged for several

Table 6. Independent validation on 30 new off-grid OptiFiber simulations

Parameter	n	R^2	MAE	RMSE	Normalised RMSE	Maximum absolute error	Mean signed error
FWHM	30	0.9982	0.1744	0.2561	0.0176	0.8322	-0.0117
Beam flatness	30	0.9922	0.00721	0.01056	0.0327	0.03286	-0.00552
Edge steepness	30	0.9912	0.01176	0.02462	0.0424	0.09546	+0.01024




Figure 5. Local flatness response around the selected off-grid inverse-design solution

small perturbations because both descriptors were extracted from a discrete 100-point profile using threshold crossings. The local perturbation experiment was therefore interpreted as a consistency and robustness assessment rather than as a formal proof of mathematical continuity.

Overall, the hold-out evaluation, five-fold stability analysis, and independent validation on 30 new off-grid simulations demonstrate that the tuned ExtraTrees model provides accurate interpolation within the investigated RI range of 1.447–1.455. These results do not imply reliable extrapolation beyond this domain.

Ablation study

An ablation study was conducted to assess how different sets of input features affect the predictive performance of the ExtraTrees surrogate model. Because the descriptors α_1 , α_2 , and α_3 are derived from RI_1 , RI_2 , and RI_3 , the analysis should be interpreted as comparing alternative feature representations rather than entirely independent input variables.

Nine feature sets were evaluated using five-fold cross-validation. The configurations analysed included a baseline variant using only RI values, a variant based solely on α descriptors, the full RI – α set, and combinations in which selected RI values or α descriptors were removed. The results are summarised in Table 7.

The highest mean R^2 was obtained for the feature set comprising RI_1 , RI_2 , α_1 , α_2 , and α_3 , with $R^2 = 0.969$. The full six-feature set achieved a slightly lower R^2 of 0.9678 but the lowest mean MAE and RMSE, equal to 0.126 and 0.250, respectively. The small differences among the three highest-ranked representations indicate broadly comparable performance and suggest partial feature redundancy.

A comparison with the variant using only RI values showed that it achieved a lower mean R^2 of 0.9473, confirming that the α parameters improve prediction accuracy. However, the poorest results were obtained for the set containing only α values. This means that α values alone, representing the relationships between RI values, are insufficient and should be combined with the refractive index values of individual optical fibre layers.

The results also highlight the importance of the α_3 descriptor. Its removal reduced the model's performance to near that of the baseline variant based solely on RI values. Since α_3 captures the relationship between RI_1 and RI_3 , it likely provides key information about the overall gradient of the refractive index across the fibre radius.

In summary, the feature set comprising RI_1 , RI_2 , α_1 , α_2 , and α_3 achieved the highest mean R^2 . However, the full six-feature representation achieved the lowest mean MAE and RMSE. The small differences among the three highest-ranked variants indicate comparable overall performance and suggest partial redundancy among the RI -based descriptors.

Normalised prediction error

Normalised prediction-error statistics were calculated from out-of-fold predictions obtained using shuffled five-fold cross-validation. Each sample was therefore predicted by a model that had not been trained on that sample. To account for differences in the FWHM, beam flatness and edge steepness, and to enable comparison between them, the analysed parameters were normalised before estimation. Each beam parameter was normalised to the range [0,1] using its minimum and maximum values in the dataset. Predictions were then made using the ExtraTrees model. For each

Table 7. Ablation study results

Rank	Used features	R^2	MAE	RMSE
1	$RI_1, RI_2, \alpha_1, \alpha_2, \alpha_3$	0.969	0.137	0.253
2	$RI_2, RI_3, \alpha_1, \alpha_2, \alpha_3$	0.9686	0.131	0.258
3	$RI_1, RI_2, RI_3, \alpha_1, \alpha_2, \alpha_3$	0.9678	0.126	0.250
4	$RI_1, RI_2, RI_3, \alpha_1, \alpha_3$	0.9639	0.137	0.271
5	$RI_1, RI_2, RI_3, \alpha_2, \alpha_3$	0.9638	0.135	0.265
6	$RI_1, RI_3, \alpha_1, \alpha_2, \alpha_3$	0.9627	0.138	0.268
7	RI_1, RI_2, RI_3	0.9473	0.169	0.323
8	$RI_1, RI_2, RI_3, \alpha_1, \alpha_2$	0.9436	0.178	0.382
9	$\alpha_1, \alpha_2, \alpha_3$	0.8779	0.364	0.652

configuration, the absolute difference between the normalised predicted and actual values was calculated. The mean, median, and maximum normalised absolute errors were subsequently determined. For FWHM, they are: 0.026, 0.013 and 0.376, respectively. For beam flatness, these errors are 0.035, 0.026, and 0.245, respectively, and for edge steepness, they are 0.037, 0.009, and 0.518, respectively.

The results indicate that the ExtraTrees model achieved low mean normalised prediction errors for all three beam parameters. Edge steepness exhibited the largest maximum error, although further inspection showed that this value was associated with isolated cases. FWHM had the lowest mean normalised error, whereas edge steepness had the lowest median error. This suggests that predictions were generally accurate for most configurations, while occasional larger deviations occurred in selected regions of the design space. The spatial distribution of the normalised prediction errors is shown in Figure 6.

Areas characterised by higher normalised error values indicate lower reliability of the surrogate model's predictions. These typically correspond to ranges where the predicted beam parameters vary significantly (see Figure 8). These areas may require additional simulations to more accurately map changes in the predicted parameters by increasing the density of points at which the values of the analysed parameter are simulated. Increasing the simulation density in these ranges allows obtaining smooth response maps with a significantly smaller number of computationally expensive simulations. Furthermore, it is possible to identify areas where the variability of individual parameter values is low, thereby reducing the number of simulations in those areas. This approach is particularly useful for determining the minimum set of simulations required to train a surrogate model that ensures the highest possible prediction accuracy.

The error maps in Figure 6 also reveal artefacts resulting from the low resolution of the simulations. This observation is consistent with the main objective of this work: to develop a cost-effective method that enables approximate mapping of the relationship between fibre configuration and beam parameters using the smallest possible number of simulations. Configurations associated with greater prediction error should be verified through additional simulations or experimentally.

The analysis of normalised errors showed that the ExtraTrees surrogate model can successfully predict all three beam parameters, enabling scalable optimisation via inverse design.

Visualisation of learned response surfaces

To illustrate how the surrogate models work, example trained response surfaces – obtained using selected models – were visualised for all three beam-shaping parameters: FWHM, beam flatness, and edge steepness. These surfaces depict the refractive indices of individual optical fibre layers, enabling identification of their non-linear characteristics, areas of high variability, and potential local extrema. For each parameter, predictions were calculated on a grid of RI_1 and RI_2 values, whilst keeping RI_3 at a constant level corresponding to its median value. This approach allows the three-dimensional response to be projected onto a two-dimensional plane for visual analysis. Visualisations of the response surfaces for various surrogate models are shown in Figures 7–12. Figure 13 presents the response surfaces obtained from the simulations.

The visualised surfaces allow several key properties to be identified:

- The response surfaces exhibit strong non-linearity, manifesting in places as abrupt changes in parameter values in response to small changes in the refractive indices of the optical fibre layers.
- Areas with high gradients indicate ranges where beam parameters are particularly sensitive to changes in the refractive index. These areas are closely related to the maxima observed in the response-surface analysis.
- Ensemble tree models, such as ExtraTrees, generate constant-interval regions, leading to the occurrence of minor discontinuities. Boosting models, such as CatBoost and XGBoost, on the other hand, produce smoother surfaces that may be more suitable for visual interpretation and gradient-based optimisation.

The visualisations enable the identification of high-sensitivity regions, thereby facilitating the selection of starting points for numerical optimisation algorithms. They also show how changes in RI_1 and RI_2 affect individual beam parameters, thereby supporting the interpretation of the surrogate model in terms of physical properties. Areas characterised by discontinuities or




Figure 6. Normalised prediction-error maps for: (a) FWHM, (b) beam flatness, (c) edge steepness




Figure 7. Response surfaces for the ExtraTrees model




Figure 8. Response surfaces for the Random Forest model




Figure 9. Response surfaces for the CatBoost model




Figure 10. Response surfaces for the LightGBM model




Figure 11. Response surfaces for the SVR model




Figure 12. Response surfaces for the XGBoost model




Figure 13. Response surfaces obtained from OptiFiber simulations

high sensitivity indicate ranges where additional simulations may be necessary to ensure the surrogate model's accuracy.

FINAL REMARKS

The results show that surrogate machine-learning models can effectively approximate the non-linear relationship between refractive index configurations in multiring optical fibres and the resulting beam-shaping parameters. Ensemble tree models achieved the highest prediction accuracy among the methods analysed, confirming their suitability for modelling complex optical systems characterised by strongly non-linear relationships between input and output variables. The ExtraTrees model achieved the best overall performance, yielding high R^2 values for all three beam parameters analysed. Parity plots and normalised prediction error analysis confirmed that the surrogate model reproduces OptiFiber simulation results with high accuracy across most of the analysed RI space. However, the edge steepness remained more difficult to predict than the other parameters. This phenomenon was consistently observed in forward prediction and ablation analyses, suggesting that the steepness depends on more complex local interactions within the refractive index profile.

For beam flatness, the surrogate generally reproduced the local OptiFiber trend, with the strongest agreement observed for RI_3 perturbations. Some discrepancies were observed for positive RI_2 and RI_3 perturbations, indicating that the ExtraTrees surrogate does not reproduce every small-scale variation of the local response surface.

The OptiFiber-derived FWHM and edge-steepness values remained constant for several small perturbations. This behaviour results from the threshold-based extraction of these quantities from a discrete 100-point profile. A small variation in the intensity distribution does not change the descriptor until the corresponding threshold crossing moves to another sampling position. Therefore, these experiments do not demonstrate strict mathematical continuity of FWHM or edge steepness. They instead show that the selected off-grid design is locally stable with respect to small RI perturbations at the resolution of the employed descriptor-extraction procedure.

The comparative analysis also showed that prediction accuracy alone is insufficient for

evaluating surrogate models intended for the inverse design of multiring fibres. Although ExtraTrees achieved the highest R^2 values, the learned response surfaces exhibited locally discontinuous regions and piecewise-constant regions, characteristic of ensemble tree methods. Such behaviour may hinder optimisation in high-sensitivity RI regions. In contrast, boosting-based models, such as XGBoost and CatBoost, generated smoother response surfaces whilst maintaining competitive prediction accuracy. Therefore, these surrogate models may be more advantageous in terms of optimisation stability and physical interpretability, even if their regression metrics are slightly worse.

The superior performance of ExtraTrees can be attributed to its ability to model strongly nonlinear relationships and higher-order interactions among the RI variables and the derived α descriptors. The method combines a large number of highly randomised decision trees, which reduces variance and limits overfitting while preserving sensitivity to local changes in the input space. This is particularly advantageous for the present dataset, where the mapping from refractive-index configuration to beam parameters is nonlinear, contains interacting features, and may exhibit locally abrupt changes. In contrast, linear models cannot represent such relationships, while smoother models such as SVR and MLP may be less effective at capturing local irregularities with a relatively limited dataset.

However, this interpretation is based on the observed model behaviour and the structure of the problem; it should not be regarded as a formal proof of the superiority of ExtraTrees for other fibre geometries or datasets.

The ablation study provided additional insights into the physical representation of the input features. The results showed that combining selected absolute RI values with derived descriptors of RI differences between layers (α_1 – α_3) yields better predictive performance than using RI values alone or α parameter values alone. The α_3 parameter proved particularly significant, indicating that the difference in RI values between the inner and outer regions of the multiring structure strongly influences beam formation. At the same time, the full feature set, comprising all RI values and α descriptors, did not yield the best result, which confirms the partial redundancy of the features.

The validation of the inverse design confirmed that the proposed approach enables the

reconstruction of *RI* configurations matching the specified beam parameters with a relatively low reconstruction error. The optimisation experiments demonstrated that the surrogate-assisted approach can identify continuous *RI* values located between the discrete levels of the original simulation grid. The validity of this interpolation was independently assessed using 30 new OptiFiber simulations. This is particularly important because exhaustive simulation-based searches become computationally impractical as the dimensionality of the design space increases.

A significant practical advantage of the proposed methodology is a substantial reduction in computational costs. The numerical OptiFiber computation required approximately 0.3 s per configuration. However, the complete operator-assisted simulation workflow required approximately 1.5 min per case. The practical benefit of the surrogate, therefore, concerns not only numerical computation but also repeated model preparation, execution, result export, and operator involvement. Consequently, the machine learning-assisted approach to inverse design enables the rapid evaluation of a large number of candidate *RI* configurations during optimisation, leading to a significant acceleration of computations compared with approaches based on exhaustive simulations.

A direct exhaustive search over the original $9^3 = 729$ *RI* configurations would require approximately 219 s of pure OptiFiber computation at 0.3 s per case. However, when the complete operator-assisted workflow, including model preparation, execution, and result export, is considered, the estimated time increases to approximately 18.2 h at 1.5 min per configuration. In comparison, one Differential Evolution run with a nominal population of 36 individuals and 300 iterations requires at most approximately 10,836 surrogate evaluations. Ten independent runs, therefore, involve approximately 108,360 surrogate predictions. At sub-millisecond prediction times, this corresponds to roughly one minute of surrogate inference, excluding model training and hyperparameter tuning. The comparison is therefore intended as an approximate computational estimate rather than a hardware-independent benchmark.

External validation alters the interpretation of the inverse-design results. The low reconstruction errors obtained for targets selected from the original simulation set demonstrate optimisation consistency but do not independently establish off-grid generalisation. In contrast, the 30 new

OptiFiber simulations were not used during model fitting, hyperparameter selection, or optimisation-model assessment. Their high agreement with the surrogate predictions provides direct evidence that the method can interpolate between the discrete *RI* values of the original grid.

This conclusion is limited to the investigated three-layer structure and *RI* range of 1.447–1.455. The present results do not establish reliable extrapolation outside that interval or transferability to fibres with a different number of layers, layer thicknesses, wavelengths, or material compositions.

Despite the promising results, several limitations should be noted. The dataset comprised only 721 valid simulation cases, limiting the diversity of the learned *RI* space. Furthermore, the optical fibre structure under investigation consisted of only three layers, with the *RI* values varied. Whilst this is sufficient to demonstrate the effectiveness of the proposed methodology, more complex optical fibre structures may require larger datasets and more advanced surrogate model architectures.

The present results are consistent with previous studies demonstrating the usefulness of machine learning in optical-fibre inverse design [15–18]. Unlike direct inverse-regression approaches, the proposed method combines a forward surrogate with Differential Evolution, which is better suited to a potentially non-unique inverse relationship. Among the analysed models, ExtraTrees achieved the best overall predictive accuracy, although boosting and neural-network models produced smoother response surfaces. Direct comparison of error metrics with previous studies is not appropriate because the analysed fibre geometries, datasets, target parameters, and validation procedures differ. The main advantage of the present approach is the independent validation of continuously optimised off-grid configurations using 30 new OptiFiber simulations.

The proposed methodology is not limited in principle to a three-layer fibre structure. It can be extended to fibres with a larger number of layers, variable layer thicknesses, additional geometrical parameters, or different material compositions, provided that these quantities are included as input variables and that a sufficiently representative simulation dataset is generated. The main workflow, forward surrogate modelling followed by numerical optimisation, would remain unchanged. However, increasing the number of design variables would enlarge the dimensionality of the search space and would likely require

more simulations, stronger regularisation, more flexible surrogate architectures, and possibly active-learning or adaptive-sampling strategies. Therefore, the present study demonstrates the feasibility of the approach for a three-layer structure, while application to more complex geometries requires separate training, validation, and robustness assessment.

The present study evaluates numerical robustness with respect to small refractive-index deviations but does not provide experimental evidence of fibre manufacturability. Around the selected off-grid solution, each RI component was independently perturbed by ± 0.0001 and ± 0.0002 . The resulting OptiFiber simulations showed that the principal beam characteristics remained relatively stable at the resolution of the employed beam-parameter extraction procedure, although some local discrepancies were observed for beam flatness. These tests indicate local numerical tolerance to small RI deviations, but they should not be interpreted as a complete manufacturing-tolerance analysis.

The practical fabrication of the proposed three-layer RI profile would additionally depend on the attainable control of dopant concentration, radial layer dimensions, material composition, and process-induced RI variations. These technological factors were not modelled in the present study. Consequently, the optimised configurations should be regarded as numerically validated design candidates rather than fabrication-ready fibre specifications. Experimental fabrication and a systematic tolerance analysis covering simultaneous variations in RI values and geometrical parameters are required before practical applicability can be confirmed.

Future work will focus on extending the proposed approach to larger, multi-dimensional optical fibre structures, uncertainty-aware inverse design, and active learning strategies. In particular, iterative procedures in which the surrogate model automatically proposes new OptiFiber simulations in areas of high uncertainty could significantly improve the model's generalisation ability whilst reducing the number of costly numerical simulations. The author also plans to use deep learning-based solutions to map relationships in complex, multi-layer fibre optic configurations. The successful application of deep learning in other areas suggests its potential to tackle complex modelling tasks.

CONCLUSIONS

The results provide affirmative answers to the three research questions concerning surrogate-model selection, the value of RI-difference descriptors, and the feasibility of independently validated off-grid inverse design. The results show that surrogate machine-learning models can accurately approximate the non-linear relationship between refractive-index distributions and beam-shaping parameters. Among the analysed models, ensemble tree-based methods achieved the best predictive performance. In particular, the ExtraTrees model provided the highest overall prediction accuracy, with high R^2 values for FWHM, beam flatness and edge steepness.

The tuned ExtraTrees surrogate was independently validated using 30 new off-grid OptiFiber simulations. The resulting R^2 values were 0.9982 for FWHM, 0.9922 for beam flatness, and 0.9912 for edge steepness. The primary continuous configuration generated by Differential Evolution was also directly simulated and reproduced the requested beam parameters with small errors. These results provide strong evidence of accurate interpolation between the discrete RI levels of the original simulation grid within the investigated RI range.

Prediction error and within-ensemble dispersion were not interpreted as formal predictive uncertainty. The study provides empirical validation errors, while calibrated uncertainty quantification remains a topic for future work.

Local perturbation tests showed that the selected design is stable to small RI changes, although the discrete extraction of FWHM and edge steepness limits the interpretation of very small local variations. The proposed method should therefore be regarded as a validated surrogate-assisted interpolation framework rather than a general proof of accuracy throughout an unrestricted continuous design space.

The ablation study showed that combining selected absolute RI values with derived descriptors of the differences in refractive index between layers (α_1 – α_3) improves predictive accuracy compared to using either RI values alone or α descriptors alone. The α_3 descriptor proved to be particularly significant, suggesting that the difference in RI values between the inner and outer layers of the optical fibre strongly influences beam formation in the analysed multiring structure.

Inverse design experiments confirmed that the proposed approach enables reconstruction of optical fibre configurations that match specified light beam parameters, with relatively low reconstruction error. Furthermore, the optimisation process generated continuous RI configurations located between the discrete levels of the original simulation grid. Their agreement with 30 new OptiFiber simulations demonstrates that the method can support validated interpolation within the investigated RI range of 1.447–1.455. Compared with exhaustive simulation-based optimisation, the surrogate model-assisted approach significantly reduced computational cost and enabled rapid evaluation of a large number of candidate optical fibre configurations.

The methodology presented here demonstrates that machine learning can support the numerical inverse design of multiring optical fibres by enabling continuous exploration of refractive-index configurations within the investigated domain. The proposed configurations have been validated numerically in OptiFiber but have not yet been assessed through fibre fabrication or experimental beam measurements. Therefore, practical manufacturability and technological tolerances remain to be verified. Although the present study focused on a relatively simple, three-layer multiring fibre structure and a limited dataset, the proposed approach can be extended to more complex fibre geometries and additional optical parameters.

Acknowledgments

This research was supported by the Lublin University of Technology, grant number: FD-20/IT-3/059.

REFERENCES

1. Yan S, Off R, Yayak A bora, Wudy K, Aghajani-talesh A, Birg M, et al. Deep learning based phase retrieval with complex beam shapes for beam shape correction. *Optics Express* 2025;33:10806–34. <https://doi.org/10.1364/OE.547138>
2. Yan Q, Yingjie Z, Jing Z, Yanwei W, Mi Z, Chenxi S, et al. Research progress of laser beam shaping technology. *Laser & Optoelectronics Progress* 2024;61. <https://doi.org/10.3788/LOP231112>
3. Dziuba-Koziel M, Koziel G, Harasim D, Kisala P, Kochanowicz M. Method of automatic calibration

and measurement of the light polarisation plane rotation with tilted fibre Bragg gratings and discrete wavelet transform usage. *Advances in Science and Technology Research Journal* 2025;19:165–77. <https://doi.org/10.12913/22998624/194890>

4. Cieszczyk S, Skorupski K, Panas P, Wawrzyk M. Comparison of local mode shift and cut-off wavelength shift for demodulation of tilted fibre Bragg grating sensors. *Metrology and Measurement Systems* 2025;32. <https://doi.org/10.24425/mms.2025.154680>
5. Zhanglong F, Ruizhi L, Hongyi L, Dixiang S, Juncheng C. Progress in biomedical imaging based on terahertz quantum cascade lasers. *Chinese Journal of Lasers-Zhongguo Jiguang* 2020;47. <https://doi.org/10.3788/CJL202017.0207011>
6. Ueno K, Kishimoto T, Masui K, Hosokawa C. Optical trapping dynamics of nanoparticles by resonant optical response. In: Omatsu T, editor. *Optical Manipulation and Structured Materials Conference (OMC 2022)*, vol. 12479, 2022. <https://doi.org/10.1117/12.2659047>
7. Schroder SE, Yin W, Kristiansen M. Beam Shaping to Control Weld Pool Oscillations. In: Nandimpalli VK, Mohanty S, Jensen DJ, Defer MC, Pan Z, editors. *20th Nordic Laser Materials Processing Conference*, vol. 1332, 2025. <https://doi.org/10.1088/1757-899X/1332/1/012034>
8. Xiaodong H, Yuanlong L, Shaozhuang B, Ke M. Research progress of laser application in material removal. *Laser & Optoelectronics Progress* 2021;58. <https://doi.org/10.3788/LOP202158.0500008>
9. Su C, Liu D, Li Z, Tai Y, Huang Z, Bao W, et al. Ring-in-ring beam shaping based on a 3D nano-printed microlens on fiber tip. *2024 14th International Symposium on Communication Systems, Networks and Digital Signal Processing, CSNDSP 2024*, 2024, p. 351–3. <https://doi.org/10.1109/CSNDSP60683.2024.10636389>
10. Moon JB, Kim N, Han GW, Kim JW. Generation of a quasi-Bessel beam with a super-gaussian axial intensity profile. *Appl Opt* 2025;64:10180–5. <https://doi.org/10.1364/AO.580477>
11. Khoubyarian S, Mashko A, Cooper A. Correcting astigmatism and ellipticity in gaussian beams using a cylindrical lens pair with tunable focal lengths. *Opt Express* 2025;33:47292–320. <https://doi.org/10.1364/OE.570745>
12. Dziuba-Koziel M, Koziel G, Markiewicz J, Kochanowicz M, Dorosz D, Miluski P, et al. Machine learning-assisted design of inverted gradient multiring optical fibres for flat-top beam generation. *Optics Express* 2026;34:7159–73. <https://doi.org/10.1364/OE.587143>
13. Melati D, Xu D-X, Grinberg Y, Al-Digeil M. Inverse design of photonic integrated devices: Optimization and machine learning. *APL Photonics* 2025;10.

- <https://doi.org/10.1063/5.0283179>
14. Wang K, Tang Z, Wu Y, Li Y, Wu T, Liu D, et al. Robust inverse design of digital photonic devices for photonic integrated circuits. *Optics Express* 2025;33:130–43. <https://doi.org/10.1364/OE.544018>
15. Guo Y, Cheng Y, Jiang Y, Cao M, Tang M, Ren W, et al. Machine learning aided inverse design for flat-top beam fiber. *Optics Communications* 2022;524. <https://doi.org/10.1016/j.optcom.2022.128814>
16. He Z, Du J, Chen X, Shen W, Huang Y, Wang C, et al. Machine learning aided inverse design for few-mode fiber weak-coupling optimization. *Optics Express* 2020;28:21668–81. <https://doi.org/10.1364/OE.398157>
17. Meng F, Ding J, Zhao Y, Liu H, Su W, Yang L, et al. Artificial intelligence designer for optical Fibers: Inverse design of a Hollow-Core Anti-Resonant fiber based on a tandem neural network. *Results in Physics* 2023;46. <https://doi.org/10.1016/j.rinp.2023.106310>
18. Rodriguez-Guillen D, Diez A, Andres V M, Velazquez-Ibarra L. Inverse design of photonic crystal fibers for dispersion engineering using neural networks. *Optics Communications* 2025;587. <https://doi.org/10.1016/j.optcom.2025.131891>
19. Powroznik P, Skublewska-Paszkowska M, Nowomiejska K, Aristidou A, Panayides A, Rejdak R. Deep convolutional generative adversarial networks in retinitis pigmentosa disease images augmentation and detection. *Advances in Science and Technology Research Journal* 2025;19:321–40. <https://doi.org/10.12913/22998624/196179>
20. Chwaleba K, Wach W. A convolutional neural network-driven model with adaptive feature fusion for Polish national dance music recognition. *Advances in Science and Technology Research Journal* 2026;20:354–72. <https://doi.org/10.12913/22998624/210059>
21. Szczerska M, Cierpiak K, Talaska T. Machine learning assisted fiber optic photonic sensors for thermal monitoring of energy storage systems. In: Zhong H, Zhu R, Stranks SD, Wang J, editors. *Photonics for Energy V*, 2025, 13714. <https://doi.org/10.1117/12.3077663>
22. Korga ST. Rapid prototyping of optical lenses using the Isolation forest model for anomaly detection. *Advances in Science and Technology Research Journal* 2026;20. <https://doi.org/10.12913/22998624/214846>
23. Galka L. Optimized deep isolation forest. *Pattern Recognition Letters* 2025;197:88–94. <https://doi.org/10.1016/j.patrec.2025.07.014>
24. Matczuk M, Sulowska D, Gromaszek K. On the usability of isolation forest for 3D mesh analysis and watermarking. *Applied Sciences-Basel* 2025;15. <https://doi.org/10.3390/app152111364>
25. Lukaszuk T, Krawczuk J, Zyla K, Kesik J. Stability of feature selection in multi-omics data analysis. *Applied Sciences-Basel* 2024;14. <https://doi.org/10.3390/app142311103>
26. Piotrowski AP, Napiorkowski JJ, Piotrowska AE. Particle swarm optimization or differential evolution-a comparison. *Engineering Applications of Artificial Intelligence* 2023;121. <https://doi.org/10.1016/j.engappai.2023.106008>
27. Chen S-K, Wu G-H, Wu Y-H. Memory-based differential evolution algorithms with self-adaptive parameters for optimization problems. *Mathematics* 2025;13. <https://doi.org/10.3390/math13101647>