
GENERATIVE INVERSE DESIGN OF METAMATERIALS WITH FUNCTIONAL RESPONSES BY INTERPRETABLE LEARNING

Wei (Wayne) Chen

J. Mike Walker '66 Department of Mechanical Engineering
Texas A&M University
College Station, TX 77843
w.chen@tamu.edu

Rachel Sun

Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139
rmsun@mit.edu

Doksoo Lee

Department of Mechanical Engineering
Northwestern University
Evanston, IL 60208
dslee@northwestern.edu

Carlos M. Portela

Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139
cportela@mit.edu

Wei Chen

Department of Mechanical Engineering
Northwestern University
Evanston, IL 60208
weichen@northwestern.edu

January 2, 2024

ABSTRACT

Metamaterials with functional responses, such as wave-based responses or deformation-induced property variation under external stimuli, can exhibit varying properties or functionalities under different conditions. Herein, we aim at rapid inverse design of these metamaterials to meet target qualitative functional behaviors. This inverse problem is challenging due to its intractability and the existence of non-unique solutions. Past works mainly focus on deep-learning-based methods that are data-demanding, require time-consuming training and hyperparameter tuning, and are non-interpretable. To overcome these limitations, we propose the Random-forest-based Interpretable Generative Inverse Design (RIGID), a single-shot inverse design method to achieve the fast generation of metamaterial designs with on-demand functional behaviors. Unlike most existing methods, by exploiting the interpretability of the random forest, we eliminate the need to train an inverse model mapping responses to designs. Based on the likelihood of target satisfaction derived from the trained forward model, one can sample design solutions using Markov chain Monte Carlo methods. The RIGID method therefore functions as a generative model that captures the conditional distribution of satisfying solutions given a design target. We demonstrate the effectiveness and efficiency of RIGID on both acoustic and optical metamaterial design problems where only small datasets (less than 250 training samples) are available. Synthetic design problems are created to further illustrate and validate the mechanism of likelihood estimation in RIGID. This work offers a new perspective on solving on-demand inverse design problems, showcasing the potential for incorporating interpretable machine learning into generative design and eliminating its large data requirement.

1 Introduction

Metamaterials with functional responses are engineered materials that exhibit varying properties or behaviors under different conditions. One example is metamaterials whose electromagnetic, acoustic, or elastic wave propagation behaviors change with wavelengths or frequencies [1]. Another example is metamaterials that exhibit changing properties or functionalities due to deformation in response to external stimuli like temperature [2] and magnetic fields [3]. Tailoring functional responses of these metamaterials is of interest to applications such as sound and vibration control, analog computing, medical imaging, sensing, communication, and soft robotics. In many use cases, rather than precisely controlling the complete functional responses, we only care about *qualitative* behaviors under certain conditions. For example, acoustic metamaterials were usually designed to have bandgaps at specified frequencies to achieve functionalities like wave-guiding [4, 5], focusing [6, 7], and vibration mitigation [8, 9, 10]. However, it is unnecessary and computationally expensive to design for the whole dispersion relation [11, 12, 13, 14, 15]. Similarly, we may design optical metamaterials to qualitatively manipulate optical properties (e.g., high or low absorption/reflection/transmission) under certain wavelengths, without requiring the entire spectral response to match an exact target [16, 17].

Identifying metamaterial designs from a given target forms an inverse design problem. Unlike many forward problems where one can obtain solutions (e.g., spectral responses or material properties under external stimuli) by modeling the physics or conducting experiments, inverse design problems are usually intractable. Traditionally, these problems are solved by iterative optimization (i.e., minimizing the difference between the actual quantity of interest and the target) [11, 12, 14]. This, however, requires repeatedly updating the design solution and solving forward problems. When the design target changes, one needs to rerun the entire optimization process. Thus, inverse design by iterative optimization becomes impractical if solving the forward problem (by simulations or experiments) is time-consuming or if the design target needs to change frequently. To accelerate the optimization approach, prior works replaced simulations or experiments with machine learning models [18, 19]. However, the efficiency and quality of final solutions are highly dependent on both the machine learning model and the optimization algorithm. On the other hand, a single run of optimization usually only returns one final solution, although multiple designs might satisfy a given target (i.e., the *non-uniqueness* of solutions). For example, multiple acoustic metamaterial designs may have bandgaps within the same target frequency range. This non-uniqueness nature of inverse design problems was also shown for optical metasurfaces [20, 21, 22]. The optimization approach eliminates the opportunity to explore diverse alternative solutions.

To avoid iterative optimization and enable fast on-demand inverse design, prior research attempted to realize single-shot (iteration-free) inverse design using machine learning. There are three mainstream models (their schematic diagrams are shown in *Appendix, Fig. 7*). The most straightforward approach is to learn a *direct inverse mapping* from the response to design variables. Neural networks are the most commonly used machine learning model for this purpose, due to their high flexibility in approximating arbitrary nonlinear input-output relationships [23, 13]. Despite the simplicity of the direct inverse mapping, its underlying assumption of the response-design mapping being one-to-one does not hold in many cases due to the non-uniqueness of solutions, as mentioned earlier. Such non-uniqueness will cause conflicting training instances where the same input (response) is associated with distinct outputs (designs), which will destabilize the convergence during neural network training [20, 24]. To avoid this issue, past work proposed the *Tandem Neural Network (T-NN)* that cascades an inverse-design network with a forward-modeling network [20, 25, 26, 27, 28]. Its training is split into two steps: (1) pretraining the forward-modeling network to approximate the design-response mapping and (2) training the cascaded network by freezing the weights of the pretrained forward-modeling network. There is no loss function that forces designs at the intermediate layer to match data (which contains conflicting instances), hence the training convergence issue is avoided. Nonetheless, the original T-NNs still learn a one-to-one response-design mapping and cannot account for the non-uniqueness of design solutions. To fundamentally solve this problem, one needs to learn a one-to-many mapping. Bastek et al. [28] integrated stochastic sampling into the inverse-design network to allow the generation of multiple feasible solutions. A large body of recent works achieved the goal of learning one-to-many mapping by using *conditional generative models*, typically conditional generative adversarial networks (cGANs) [29, 30, 31, 32, 22], conditional variational autoencoders (cVAEs) [21], and conditional diffusion models [33]. These models can generate multiple designs given a target response by learning the distribution of designs conditioned on the response. Different generative models have distinct ways of learning conditional distributions. In general, this is realized by training neural networks to transform responses and random noise (or latent variables) into designs, so that the trained network can generate a non-deterministic design solution from a given target response and randomly sampled noise, which is equivalent to sampling from a conditional distribution.

Although conditional generative models have demonstrated success in solving inverse design problems, they still have issues, such as high data demand, exhaustive hyperparameter tuning, slow training, and low interpretability, especially compared to traditional machine learning models like decision trees and random forests. On the other hand, Elzouka et al. [34] proposed to use the decision tree as a more interpretable model to solve both the forward prediction and

inverse design problem. After training a decision tree for forward prediction, one can identify explicit design rules (i.e., feasible regions in the design space) by tracing from target leaf nodes to the root node. This approach also captures the one-to-many mapping nature of inverse design problems since it gives feasible design variable ranges rather than a single solution. However, there remain some limitations. Firstly, for solutions identified by the design rules, the method does not differentiate their likelihood of target satisfaction. Yet in reality, solutions always have different likelihoods due to the uncertainty of model estimation. Secondly, the method has to train two models: a random forest was trained first to ensure model accuracy and robustness, and then a large decision tree was trained to emulate the performance of the random forest and provide design rules. This is due to the challenge of deriving explicit design rules from an ensemble model like the random forest. Finally, the method was demonstrated on a problem with more than 10^4 training data, while the effectiveness on smaller datasets (i.e., data with orders of magnitude smaller sample sizes) was not studied.

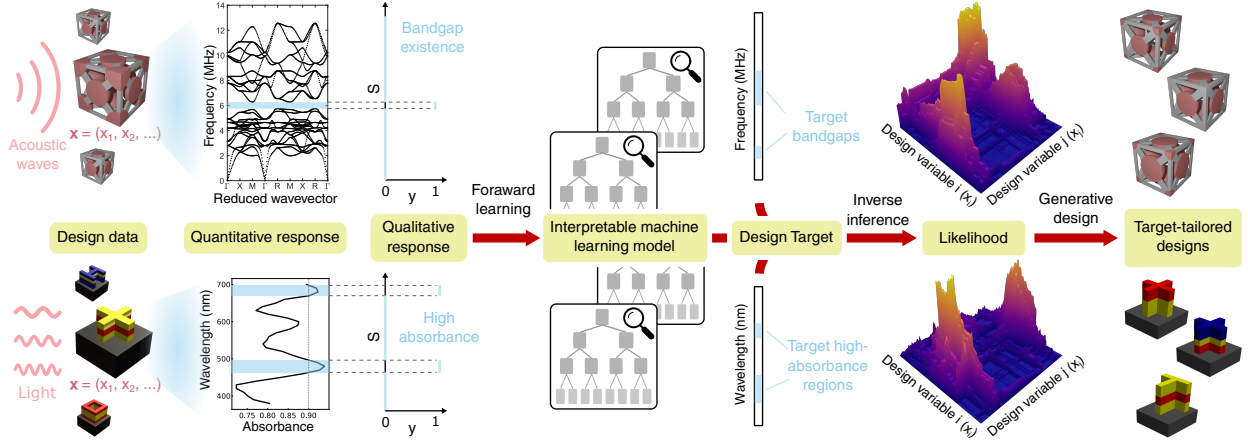


Figure 1: Schematic diagram of the RIGID method. We first train a random forest on a design-response dataset to learn the forward design-response relation — predicting qualitative responses (e.g., bandgap existence at any given wave frequency) of designs. Then given a design target, we can infer the likelihood of any design satisfying the target by probing into the trained random forest. New designs with tailored responses can be generated by sampling the design space based on the likelihood estimation.

This work aims to address the aforementioned problems by proposing a method called *Random-forest-based Interpretable Generative Inverse Design (RIGID)*. Figure 1 shows an overview of this method. Specifically, we first train a forward prediction random forest. Then given a design target, we can probe the trained random forest to infer the likelihood of any design satisfying the target. To generate new designs tailored to the target, we can sample from the design space according to the likelihood. Compared to the most widely studied neural-network-based methods, RIGID has a much lower cost in training and hyperparameter tuning, and works more robustly on small-size datasets (as random forests are less prone to overfitting). Similar to deep generative models, it can generate a desired number of solutions, allowing one to explore alternative solutions that might have desired properties or functionalities beyond the ones considered as the target. The explicit likelihood estimation also offers an interpretable characterization of a design’s target satisfaction probability and allows an exploitation-exploration trade-off when selecting generated designs. We validate the RIGID method on two metamaterial design examples — an acoustic metamaterial design example, where the target is to generate metamaterials with specific bandgaps, and an optical metasurface design example, where the target is to generate metasurfaces with high absorbance at specified wavelengths.

Our contributions are three-fold. First, we propose a single-shot inverse design method that is fast, generative, interpretable, and small-data-compatible. Secondly, we demonstrate the effectiveness of the proposed method on acoustic and optical metamaterial design examples, and propose both qualitative and quantitative ways of assessing our method. Finally, we create two synthetic test cases for fast examination and validation of model performance. These test cases can be used for future benchmarking studies of related methods.

2 Methods

The functional response of metamaterials can be modeled as $y = f(\mathbf{x}, s)$, where \mathbf{x} denotes metamaterial design variables (e.g., materials and geometry parameters), s is an *auxiliary variable* representing the independent variable (or the “ x -axis”) of the response (e.g., the frequency/wavelength or the external stimuli such as temperature), and y indicates

the value of the response associated with our design target. In this paper, we assume $y \in \{0, 1\}$ since we only focus on qualitative behaviors at specified frequencies (e.g., for an acoustic metamaterial or an optical metamaterial design, whether a bandgap exists or whether the energy absorbance is higher than a threshold within a range of frequencies). We leave the more challenging problem of tailoring quantitative behaviors as future work.

We use a random forest to approximate the function f . A random forest is an ensemble learning method that combines the predictions of multiple decision trees to improve accuracy and reduce overfitting [35]. The trained random forest serves as a *forward prediction* model that predicts the outcome y given design variables \mathbf{x} and the auxiliary variable s . Compared to the widely used neural networks, the random forest as a forward prediction model offers (1) significantly faster training, (2) less hyperparameters to tune, (3) less susceptible to overfitting on small data, and (4) interpretability (i.e., the decision-making of each tree in the random forest is transparent). More importantly, this interpretability also allows us to realize *inverse design* without training a separate inverse model.

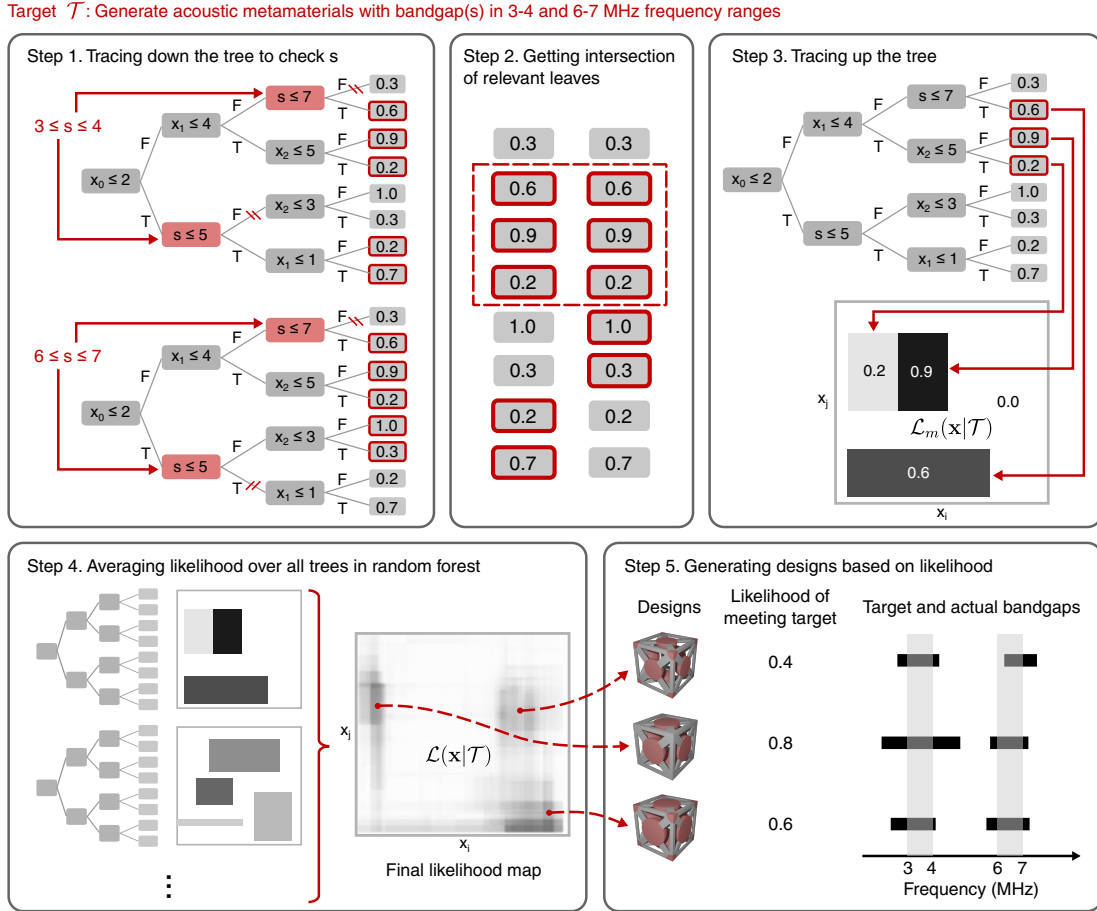


Figure 2: The inverse design pipeline of the proposed RIGID method (using the inverse design of acoustic metamaterials as an example). Given design parameters \mathbf{x} and the auxiliary variable s (e.g., wave frequency), a trained random forest predicts the probability of the qualitative response y (e.g., bandgap existence). Each tree in the random forest splits the joint space of \mathbf{x} and s into regions, each associated with a specific prediction (shown on leaf nodes). The splitting criteria are encoded in tree nodes. “T” means meeting a criterion and “F” means not meeting it. RIGID first identifies leaf nodes that are relevant to the considered range of auxiliary variable s by checking splitting criteria related to s and pruning tree branches that are irrelevant (Step 1). If the considered range of s has multiple parts, we repeat this step for each part and take the intersection of relevant leaves (Step 2). Each relevant leaf node corresponds to a decision path indicating a region in the design space, as well as a predicted probability of target satisfaction, which is a score we assign to the corresponding design space region (Step 3). With multiple trees in a random forest, we can average the scores predicted by each tree and use the average score as our likelihood estimation (Step 4). We can then sample from the design space based on the likelihood distribution to generate new designs tailored to the target (Step 5).

Figure 2 shows how, by probing the trained random forest, one can estimate a likelihood distribution for target satisfaction of solutions over the entire design space and sample (generate) new designs based on this likelihood distribution. Since we target qualitative (binary) behaviors at specified s (e.g., a bandgap in 3-4 MHz frequency or high absorption at a wavelength of 400-500 nm), we first identify the leaf nodes (on each decision tree in the random forest) that are relevant to the s in the target (Fig. 2, Step 1). We do this by tracing down each tree, checking only the nodes that use s as the splitting feature, and pruning the branches that are irrelevant to the s in the target. For example, as shown in Fig. 2, there are two tree nodes using s as the splitting feature, with splitting criteria at $s \leq 5$ and $s \leq 7$. Given the target frequency range of $3 \leq s \leq 4$, we can remove the right branches of both nodes as these branches are only relevant to $s > 5$ and $s > 7$, respectively, which conflicts with the target range of $3 \leq s \leq 4$. After pruning these branches, we end up with a set of leaves relevant to the target (highlighted in Fig. 2, Step 1). When we have a combined target (e.g., bandgaps in both 3-4 MHz and 6-7 MHz, as shown in Fig. 2), we need to get the intersection of all the sets of relevant leaves and use that as the final set of relevant leaves (highlighted in Fig. 2, Step 2). Note that a combined target includes cases where there are multiple nonadjacent target ranges (e.g., 3-4 MHz and 6-7 MHz) or when a target range is split by a tree node (e.g., a target range of 4-6 MHz can be split by the node “ $s \leq 5$ ”, thus we need to consider it as the combination of two target ranges — 4-5 MHz and 5-6 MHz). A more detailed discussion of this step is in *Appendix, Sec. B*.

The next step is to trace up the tree from the N relevant leaves, obtained by Step 2, to the root node (Fig. 2, Step 3). This will result in N decision paths, along which are nodes indicating splitting criteria for design variables \mathbf{x} . Thus, each decision path represents a set of design variable ranges, or in other words, a region in the design space. We assign each region a score equal to the predicted probability at each corresponding leaf. This probability is learned from the training data and equals the proportion of positive data in a leaf. It indicates the tree’s belief in the probability of a design \mathbf{x} satisfying the target \mathcal{T} if the design falls in the design space region corresponding to the leaf. Therefore, with a single decision tree i , we already have the map of likelihood $\mathcal{L}_m(\mathbf{x}|\mathcal{T}) = \mathbb{P}_m(\mathcal{T}|\mathbf{x})$ for target satisfaction: each of the N regions has a uniformly distributed likelihood equal to the predicted probability at the corresponding leaf, and the rest of the design space has a likelihood of 0 (Fig. 2, Step 3).

Since a single decision tree usually lacks accuracy, robustness, and a way to quantify estimation uncertainty, we still want to take advantage of the random forest as an ensemble model for inverse design. We use Steps 1-3 to derive the likelihood distribution for each of the M trees in the random forest, and simply use the average of these M likelihood distributions as the final likelihood for target satisfaction, $\mathcal{L}(\mathbf{x}|\mathcal{T}) = \sum_m^M \mathcal{L}_m(\mathbf{x}|\mathcal{T})/M$, which is a more complex and smooth function (Fig. 2, Step 4). If more trees believe a design \mathbf{x} has a higher likelihood of satisfying the target, then the design will have a higher likelihood $\mathcal{L}(\mathbf{x}|\mathcal{T})$. Finally, to generate new designs, we can sample from $\mathcal{L}(\mathbf{x}|\mathcal{T})$ using Markov chain Monte Carlo (MCMC) methods such as Metropolis-Hastings [36] (Fig. 2, Step 5).

Compared to prior works, RIGID provides the following unique benefits:

1. It is effective on small data problems as the random forest is less susceptible to overfitting.
2. The training is fast (in seconds of wall time) and does not require computationally-demanding hyperparameter tuning. Once the training is done, no further training or iterative optimization is required to generate designs for different targets.
3. The model is interpretable as one can easily probe the trained model to understand its reasoning behind any decision-making.
4. It estimates the explicit likelihood of target satisfaction for every possible solution in the design space. Given a design target of specific functional behavior, we can generate an unlimited number of solutions based on the likelihood, allowing us to explore alternative solutions that might have desired properties or functionalities beyond the ones considered as the target.
5. When generating design solutions, one can use a single parameter — the sampling threshold — to easily tune the trade-off between exploitation (i.e., generated designs have higher chances of satisfying the target) and exploration (i.e., generated designs cover a broader area of the design space), as we will demonstrate in Results.

3 Results

We demonstrate our RIGID method on an acoustic metamaterial design problem, an optical metasurface design problem, and two synthetic design problems. Based on a recent review article by Lee et al. [24] and other related works (e.g., [28]), existing single-shot inverse design methods were demonstrated on training data size ranging from 10^3 to 10^6 in scale. Here we show that our method can work with much smaller-scale datasets (less than 250 training samples).

For all the test problems, we used the same random forest hyperparameter settings and did not perform hyperparameter tuning. Specifically, each random forest contains 1,000 trees with a minimum of 2 samples required to split an internal node and a minimum of 1 sample required to be at a leaf node. Gini impurity [37] was used as the splitting criterion at tree nodes. The train-test split ratio was 4:1. Since the positive/negative training data can be highly imbalanced (e.g., the frequency ranges with bandgaps are much narrower than those without), we used the Synthetic Minority Over-sampling TEchnique (SMOTE) [38] to over-sample the positive class. For all the case studies, the random forest training took less than 10 seconds on an Intel Core i5-9300H CPU 2.4GHz and 8GB memory.

After training, we generate new designs by sampling from the resulting likelihood distribution using Metropolis-Hastings. In practice, Metropolis-Hastings can generate identical samples, which provides no benefits for design purposes. Thus in this work, we reject the designs identical to previous ones during sampling.

3.1 Applying RIGID to Design Acoustic Metamaterials with Target Bandgaps

Here we consider acoustic metamaterials that can control elastic wave propagation at ultrasound (MHz) frequencies. Varying the microscale geometries of acoustic metamaterials changes the dynamic properties of a material, such as bandgaps [10] (i.e., forbidden frequency ranges of a material) and wave propagation direction [4]. These materials promise applications in waveguides [4, 5], lenses [6, 7], and vibration mitigation [9]. We present the braced cubic design framework (Fig. 3A-B) as a method to tune the size and location of bandgaps (Fig. 3C). In particular, spherical micro-inertia are added to the center and corner of a braced cubic unit cell with strut radius r_{strut} . Micro-inertia placed at the center of the brace has radius r_{center} while micro-inertia placed at the corner of the cubic unit cell has radius r_{corner} .

We randomly created 284 sets of geometric parameters $\mathbf{x} = (r_{\text{strut}}, r_{\text{center}}, r_{\text{corner}})$ with $4 \leq r_{\text{strut}} \leq 6.41$, $0 \leq r_{\text{center}} \leq 20$, and $0 \leq r_{\text{corner}} \leq 20$ (unit: μm). The unit cell size was set at $a = 60 \mu\text{m}$. For each of these designs, we performed Bloch-wave analysis to compute its acoustic dispersion relation. Bandgap location and width were extracted for each design based on its dispersion relation.

Out of the 284 sets of design variables and bandgap data, we used 227 samples as training data. We first discretized the entire frequency range into 100 intervals, and trained a random forest to predict bandgap existence $y \in \{0, 1\}$ at a specific interval s for a given design \mathbf{x} . The trained model has a test F1 score of 0.82. The resulting confusion matrix on test data is shown in *Appendix, Tab. 1*.

To test the inverse design capability of RIGID, we randomly created 10 targets, each containing 1-2 frequency ranges in which bandgap(s) should exist. We generated 30 designs for each target by sampling from the resulting likelihood distribution over the design space¹. Bandgaps were identified from dispersion relations computed using Bloch-wave analysis. Figure 3D shows the kernel density estimation (KDE) for the likelihood of the 300 generated designs, conditioned on their target satisfaction. We use \mathcal{D} and $\mathcal{D}_{\text{feas}}$ to represent the complete set of generated designs and the set of generated designs that actually satisfy the target, respectively. Then $\mathcal{D} \setminus \mathcal{D}_{\text{feas}}$ denotes the set of generated designs that cannot fulfill the target in reality. In an ideal scenario, all solutions in \mathcal{D} would satisfy the target, which means $\mathcal{D} = \mathcal{D}_{\text{feas}}$, and their density profiles should coincide. However, this ideal scenario is not possible due to limited model accuracy. Conveniently, the estimation of target satisfaction likelihood offers us an indicator of what solution is more likely to violate the target. For a reasonable model, most designs in $\mathcal{D} \setminus \mathcal{D}_{\text{feas}}$ should have low estimated likelihood values. Consequently, the density of $\mathcal{D}_{\text{feas}}$'s likelihood is a result of shifting some of \mathcal{D} 's density from left (low likelihood) to right (high likelihood). This expectation aligns with the observation in Fig. 3D.

When sampling new designs or selecting solutions from generated designs, we can put a *sampling threshold* $\tau \in (0, 1)$ on the likelihood values to filter out "less promising" solutions. To further examine model behavior and quantify how τ affects the inverse design outcome, we define the following metrics:

$$\begin{aligned} \text{Selection Rate} &= \frac{|\mathcal{D}_{\phi \geq \tau}|}{|\mathcal{D}|}, \\ \text{Satisfaction Rate} &= \frac{|\mathcal{D}_{\phi \geq \tau} \cap \mathcal{D}_{\text{feas}}|}{|\mathcal{D}_{\phi \geq \tau}|}, \\ \text{Average Score} &= \frac{1}{|\mathcal{D}_{\phi \geq \tau}|} \sum_{i=1}^{|\mathcal{D}_{\phi \geq \tau}|} q_i, \end{aligned} \tag{1}$$

where $\mathcal{D}_{\phi \geq \tau}$ is the set of generated designs with the likelihood of at least τ (i.e., the selected designs) and q_i denotes the percentage overlap between the target and the actual behavior of selected designs. The satisfaction rate evaluates how

¹Note that it is possible for the likelihood to be zero everywhere in the design space when the model believes the target is unachievable. We ignore these cases as it is meaningless and impossible to sample designs from such likelihood distribution.

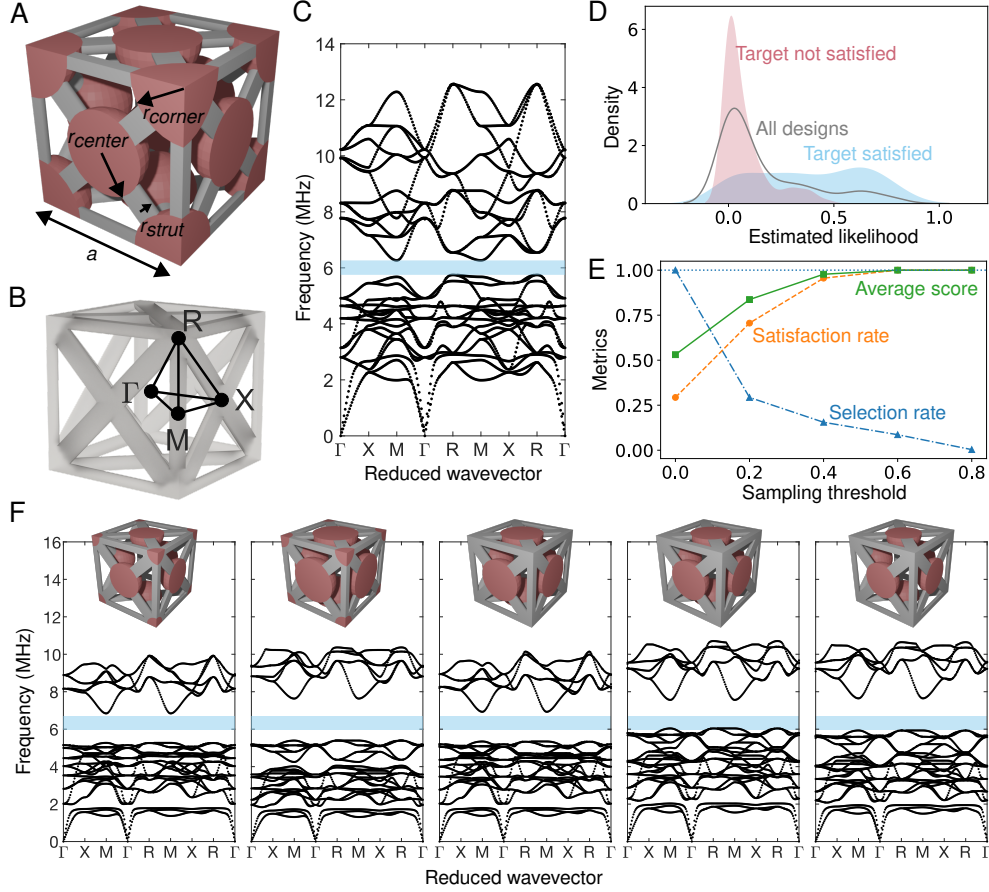


Figure 3: Acoustic metamaterial design problem configuration and results. (A) Design variables of center and corner mass radii (r_{center} and r_{corner}) and strut radius (r_{strut}). (B) High symmetry points of the cubic irreducible Brillouin zone. (C) A sample dispersion relation and bandgap (marked by the highlighted zone). The design objective is to generate new acoustic metamaterial designs with target bandgaps. (D) KDE of the estimated likelihood for generated designs. (E) Validation metrics under varying sampling thresholds. (F) Geometries and corresponding dispersion relations of five generated designs with the highest likelihood of satisfying a specified target bandgap (marked as highlighted frequency regions). Here only the fourth design fails to meet a small portion (at around 6 MHz) of the target bandgap, whereas the others meet it. Generated designs for other targets are shown in *Appendix, Figs. 8-10*.

many selected designs satisfy the target based on a binary criterion (i.e., whether or not a design satisfies the complete target), whereas the average score provides a soft measure where partial satisfaction is also counted. The average score is lower-bounded by the satisfaction rate.

As shown in Fig. 3E, the selection rate decreases when τ increases since more solutions are filtered out. On the other hand, both the satisfaction rate and the average score increase with τ , which indicates a high correlation between the estimated likelihood of a solution and its probability of actually achieving the target. As τ reaches 0.6, the satisfaction rate and the average score reach 1, indicating that all generated designs satisfy their targets. When sampling or selecting new solutions, we can use the sample threshold τ to tune the trade-off between exploitation and exploration — a low τ favors exploration as sampled solutions will cover a larger area of the design space, while a high τ favors exploitation as sampled solutions will have a higher chance of satisfying the target.

Figure 3F visualizes the geometries and dispersion relations of designs generated based on a randomly created bandgap target. Only the top five designs with the highest likelihood values are shown. In this example, our method generates geometrically different designs that have a high probability of achieving target bandgaps, each yielding a slightly different dispersion relation. This is promising in design applications requiring other material properties, such as dynamic wave velocity or quasi-static stiffness, in which the user can select from a menu of designs with the same

target bandgap but other varying properties. Generated designs based on the other nine bandgap targets can be found in *Appendix, Figs. 8-10*.

3.2 Applying RIGID to Design Optical Metasurfaces with Target High-Absorbance Wavelengths

Optical metasurfaces are artificially engineered systems that can support exotic light propagation building on subwavelength inclusions [39, 40, 41, 42, 43, 44]. Among a diverse array of devices, metamaterial absorbers [45, 46, 47, 48, 49, 50, 51] have been intensely studied for medical imaging, sensing, and wireless communications.

In this case study, we consider four types of cross-sections ($c \in \{1, 2, 3, 4\}$) chosen from the literature (Fig. 4B). It is assumed that a 3D geometric instance is composed of a stack of three layers of prismatic unit cells, each of which is vertically extruded and stacked (Fig. 4A). The geometries constructed in this way can be regarded as an instantiation of multilayered metasurfaces [52, 53, 54, 55, 56], which offer richer design freedom than the single-layer counterpart. The height of each layer ($h_l, l = 1, 2, 3$) is allowed to continuously vary between 50 and 150 nm. Herein we do not consider parametric variations of a given type of unit cell cross-section; yet those can be trivially incorporated in the proposed design framework if necessary.

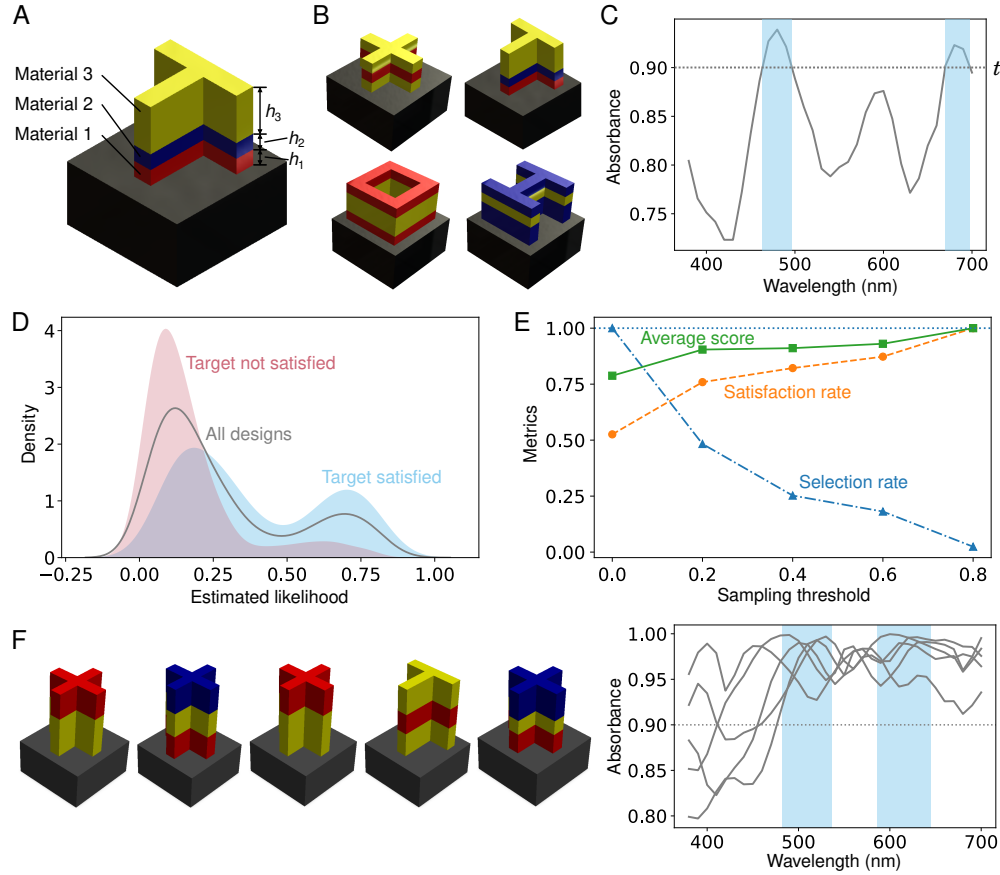


Figure 4: Optical metasurface design problem configuration and results. (A-B) Design variables (materials, layer heights, and cross-section geometry types). (C) A sample absorbance spectrum and the wavelength intervals (highlighted wavelength regions) corresponding to absorbance above the threshold t . The design objective is to generate new optical metasurface designs that exhibit higher absorbance than a threshold t at the user-defined wavelength interval(s). (D) KDE of the estimated likelihood for generated designs. (E) Validation metrics under varying sampling thresholds. (F) Designs (geometries and material selections) and corresponding absorbance spectra of five generated metasurfaces with the highest likelihood of satisfying specified target high-absorbance regions (marked as highlighted wavelength regions). Here all five designs satisfy the target. Generated designs for other targets are shown in *Appendix, Figs. 11-12*.

We also design the material of each layer ($m_l, l = 1, 2, 3$). Three dielectric materials of interest, each of which is assigned to a different color in Fig. 4A, are Ti (red), Si (blue), and Ge (yellow). In general, a dielectric material is

characterized through a complex refractive index $\tilde{n} \in \mathbb{C}$ defined as $\tilde{n} = n + jk$, where $j = \sqrt{-1}$ is the imaginary unit, $n \in \mathbb{R}$ involves the speed at which the light propagates through the material, and $k \in \mathbb{R}$ is the extinction coefficient that dictates the energy loss due to the material. Within the frequency regime of interest, those exhibit nonlinear dispersion; both the real and imaginary terms in general are a non-analytic function of excitation wavelength s , i.e., $n(s)$ and $k(s)$. In addition, the impact of the same material choice on the spectral response $A(s)$ varies depending on the layer location at which the material is placed. Thus the highlight of this case study is the combinatorial nature of design, whose spectral responses are affected by the joint contributions of geometry and material.

Based on the above configuration, we randomly sampled 258 sets of design variables $\mathbf{x} = (c, h_1, h_2, h_3, m_1, m_2, m_3)$ and computed their corresponding absorbance spectra using wave analysis. We set $t = 0.9$ as the absorbance threshold, so that “high” absorbance means the absorbance $A(s)$ is no less than 0.9. We trained a random forest on 206 training data (i.e., 80% of the 258 designs and corresponding absorbance spectra) to predict whether “high” absorbance is presented (i.e., the binary indicator $y = 1$) at a wavelength s for a given design \mathbf{x} . The trained random forest has a test F1 score of 0.83. The confusion matrix on test data is shown in *Appendix, Tab. 2*.

Note that this problem involves inverse design with both continuous and categorical variables, which common optimization and generative modeling-based inverse design cannot handle well without special treatment [57, 58, 59]. On the other hand, our random forest-based method can naturally address such mixed-variable problems without any issues.

Similar to the acoustic metamaterial design problem, we use 10 randomly created targets to evaluate the inverse design performance of RIGID, except that here a target is represented as the wavelength range(s) within which absorbance should be at least 0.9. We generated 100 designs for each target by sampling from the estimated likelihood distribution. Among the 1,000 generated solutions, we successfully conducted wave analysis for 911 designs and obtained their absorbance spectra. Figure 4D shows the KDE for the likelihood of these 911 designs, conditioned on their target satisfaction. The densities share similar behavior as in the acoustic problem (Fig. 3D) —unsatisfied/infeasible designs $\mathcal{D} \setminus \mathcal{D}_{\text{feas}}$ are concentrated at low likelihood regions, which causes the likelihood density of satisfied/feasible designs $\mathcal{D}_{\text{feas}}$ to be a result of shifting some of \mathcal{D} ’s density from left (low likelihood) to right (high likelihood). The sampling threshold and metrics relation shown in Fig. 4E also follow the same trend as in the acoustic problem (Fig. 3E), which again demonstrates a strong positive correlation between the estimated likelihood and the probability of generated designs actually achieving their targets.

Figure 4F shows generated optical metasurface designs with the top five likelihood estimations for a randomly created target. While the materials, cross-section geometries, and layer heights of generated designs can be different, all the designs satisfy the target (Fig. 4F, right panel). To further enhance the diversity of final solutions, we can use sampling strategies such as the one proposed in Ref. [60, 61] to identify a subset (of generated solutions) that simultaneously exhibits high likelihood and high diversity. Generated designs based on the other nine targets can be found in *Appendix, Figs. 11-12*.

3.3 Synthetic Design Problems for Rapid Validation and Visualization

While the above metamaterial design problems represent practical use cases, the validation study is time-consuming due to the expensive computation of metamaterials’ responses such as dispersion relations and absorbance spectra. To allow fast validation of the proposed method and easier inspection of the estimated likelihood in the design space, we create two synthetic case studies. Both problems have 2-dimensional “design spaces” that allow easy visualization.

SqExp Problem To construct the first synthetic problem, we used a *squared exponential function* with tunable parameters a and b to mimic the quantitative functional response of metamaterials. The qualitative response (e.g., “high” or “low” energy absorption at a wavelength) is defined as:

$$I(a, b; s) = \begin{cases} 0, & \text{if } z = \exp\left(-\left(\frac{s-a}{0.3b+0.1}\right)^2\right) < t, \\ 1, & \text{otherwise} \end{cases}, \quad (2)$$

where z represents quantitative response and t is a threshold that converts z into a qualitative response $I(a, b; s)$. Specifically, $I(a, b; s) = 1$ can mean the existence of a bandgap or high absorbance at a frequency s . Then $\{s | I(a, b; s) = 1\}$ represents a range of s that mimics our design targets, such as the bandgap or the frequency range of high absorbance. By varying a and b , we can produce different synthetic responses and ranges. Therefore, we can use a and b as synthetic design variables. There is a clear relation between these design variables and the range that Eq. 2 creates — a and b control the center location and the width of the range, respectively.

In this design problem, we sampled 100 sets of a and b uniformly at random. We set t as 0.9. Based on Eq. 2, we obtained the corresponding responses (Fig. 5A). These sets of a , b , and responses constitute a dataset for training and testing our model.




Figure 5: Synthetic data creation for (A) the SqExp problem and (B) the SupSin problem. For each problem, the left panel shows 100 functions with randomly sampled parameters a and b . We treat a and b as synthetic design variables, and the corresponding functions as quantitative responses (e.g., absorbance spectra of optical metasurfaces). The right panel shows qualitative responses (e.g., high-absorbance wavelength ranges or bandgaps) are simulated by synthetic ranges, derived by thresholding the 100 functions (Equations 2 and 3 with threshold $t=0.9$).

SupSin Problem Another synthetic design problem was constructed by replacing the squared exponential function in the SqExp problem with a *superposed sine function*. Given synthetic design variables a and b , we can produce qualitative responses using the following equation:

$$I(a, b; s) = \begin{cases} 0, & \text{if } z = \sin(2\pi(s + a)) + \sin(3\pi(s + b)) < t \\ 1, & \text{otherwise} \end{cases} \quad (3)$$

Same as in the SqExp problem, we set $t = 0.9$ and created a dataset with 100 sets of synthetic design variables and corresponding ranges derived from synthetic responses (Fig. 5B). Unlike the squared exponential function, the superposed sine function can be multimodal, which means it can result in multiple synthetic ranges to mimic, for example, multiple bandgaps. The bandgap locations are controlled by a and b .

For each synthetic example, we split the data into 80 training data and 20 test data, and trained a random forest, with the same hyperparameter settings as the other problems, to predict the binary indicators $I(a, b; s)$. The F1 scores are 0.85 and 0.86 for the SqExp and the SupSin problems, respectively. The resulting confusion matrices are shown in *Appendix, Tables 3-4*. We evaluated the inverse design performance with the trained models.

Due to the fast evaluation of Equations 2 and 3, we can exhaust all the possible solutions in the design space to obtain the ground-truth feasible region(s) for a target. Figure 6A shows the estimated likelihood values and the ground-truth feasible regions under randomly created targets. In general, high-likelihood regions match actual feasible regions well, which further demonstrates the effectiveness of RIGID. We can also observe that feasible regions in the SqExp and the SupSin problems follow distinct patterns. In the SqExp problem, a and b control the center location and the width of the output range, respectively. Therefore, the position of the feasible region along the a -axis moves with the location of the target range, while the feasible region gradually shrinks as b decreases since the decrease of b (i.e., output range width) causes the choice of a (i.e., output range center location) to be more restricted to fit the target range. In the SupSin problem, there might be multiple bandgaps appearing at the peaks of the superposed sine function in Eq. 3. Design

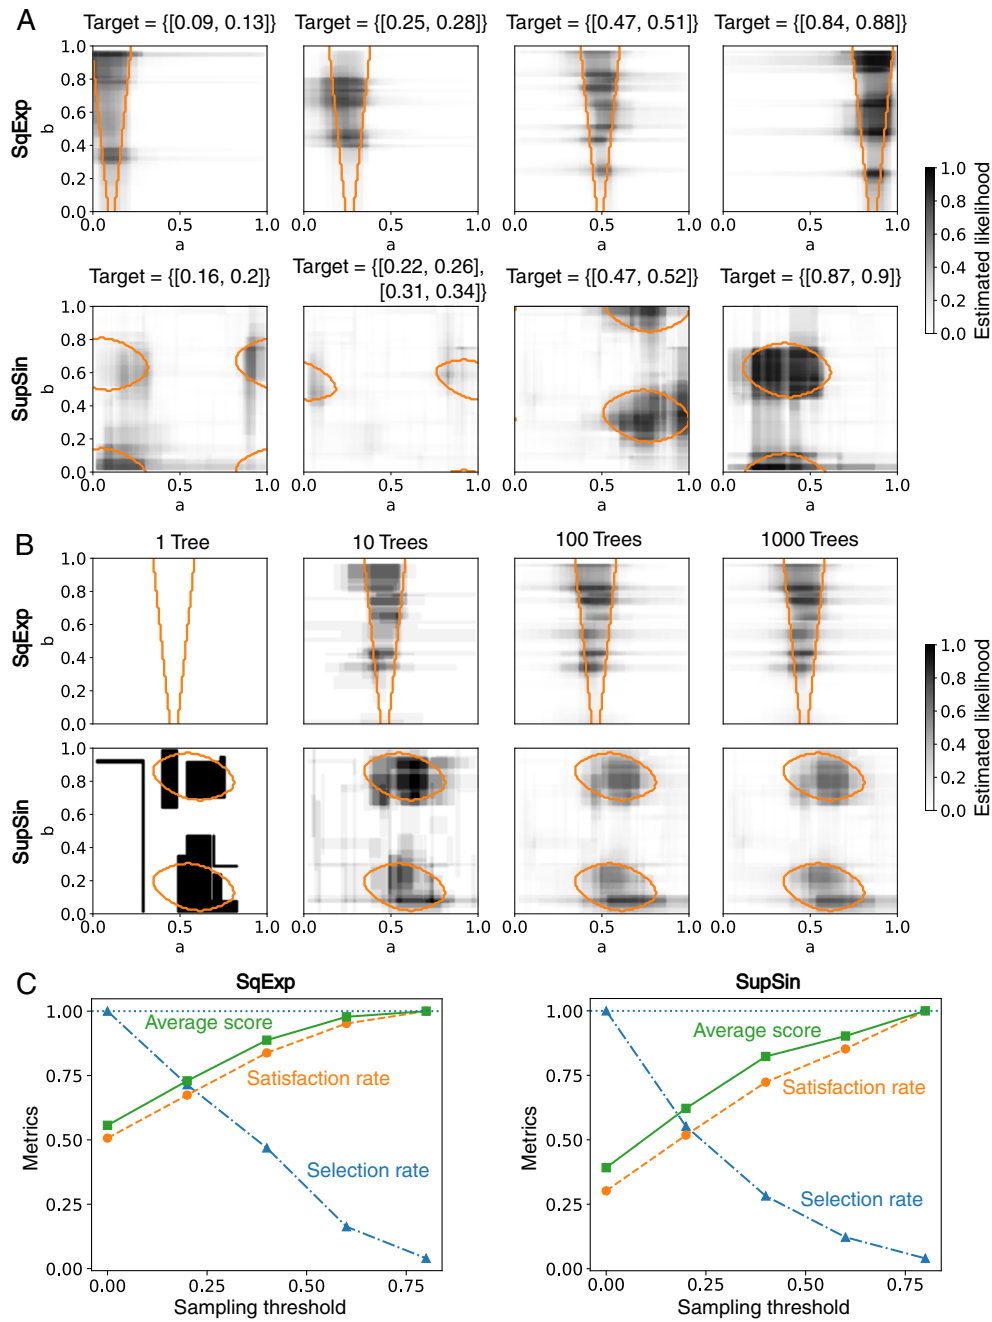


Figure 6: Visualization of estimated likelihood and validation metrics for synthetic problems. (A) Likelihood function values for randomly created design targets. Orange lines show boundaries of actual feasible regions associated with the targets. (B) Likelihood function values estimated by random forests with varying numbers of decision trees. The design target is set as $\{[0.45, 0.48]\}$ for the SqExp problem and $\{[0.63, 0.68], [0.69, 0.71]\}$ for the SupSin problem. (C) Validation metrics for inverse design generation.

variables a and b control bandgap locations by translating each sine function. Due to the sine function’s periodicity, we can obtain multiple feasible regions along both a - and b -axes. Figure 6A shows that the likelihood estimation by RIGID successfully captured the above-mentioned patterns of feasible regions.

Figure 6B demonstrates how the estimated likelihood varies when increasing the number of trees in a random forest. With a single decision tree, the estimated likelihood function is almost a binary function and highly inaccurate. The likelihood in the SqExp case is even zero everywhere, which makes it impossible to sample designs based on the likelihood. As the number of trees increases, the likelihood function becomes smoother and eventually converges.

Besides these qualitative visual inspections, we also calculated the metrics proposed in Eq. 1, as shown in Fig. 6C. For each of the two synthetic problems, these metrics were computed on 500 designs generated by giving five random target ranges. Again, the satisfaction rate and the average score increase with the sampling threshold, indicating a strong correlation between the sampling threshold and the probability of generated designs actually achieving their targets. In both problems, all the selected designs satisfy their targets (i.e., the satisfaction rates and average scores reach 1) when the sampling threshold reaches 0.8.

4 Summary and Outlook

We proposed RIGID, a single-shot inverse design method that generates metamaterials to satisfy qualitative behaviors of functional responses. Such qualitative behaviors are important design targets in many applications such as tailoring bandgaps of acoustic metamaterials for wave-guiding, focusing, and vibration mitigation, or tailoring the absorption level of optical metasurfaces at certain wavelengths for medical sensing, imaging, and communication applications. Unlike most existing inverse design methods that require training an inverse model to map targets to designs, the RIGID method takes advantage of the random forest’s interpretability and derives the likelihood of target satisfaction by probing the trained forward model. Incorporated with MCMC, one can sample a desired number of new designs based on the estimated likelihood. Therefore, RIGID functions as a generative model that can capture the conditional distribution of satisfying designs given a target, or in other words, the one-to-many mapping from the target to satisfying designs. Using both real-world and synthetic design problems, we demonstrated that RIGID is efficient and effective on datasets with training sample sizes smaller than 250. Thus, RIGID is particularly useful when data collection is expensive, as in many cases where high-fidelity simulation or experimental data are needed. We used both qualitative and quantitative approaches to validate the proposed method. The quantitative results revealed a strong correlation between the estimated likelihood of a solution and its probability of actually achieving the target, which demonstrated the effectiveness of the likelihood estimation. Due to the fast evaluation of output responses and the transparency of ground-truth solutions, the proposed synthetic problems can be used for future benchmarking studies of metamaterial design problems.

While we address qualitative design targets in this study, the idea of using random forest-based models for inverse design has the potential to generalize to quantitative targets. Such problems can be, for example, generating optical metasurface designs with specific optical spectra [62, 21], generating functional materials with target nonlinear constitutive relations [33, 63], or generating programmable metamaterials with prescribed functional responses [64, 65]. It is also straightforward to adjust the target to achieve multifunctionality (e.g., negative/positive Poisson’s ratio under low/high compression rate [66]).

Although this study only demonstrates the RIGID method on parametric design (i.e., designs are represented by geometric and/or material parameters), the method also applies to shape or topological design problems where the shape or topology of designs can vary without being restricted to a limited number of geometric parameters [62, 29, 21, 31, 67]. In those cases, as valid designs only lie on a lower-dimensional manifold of the design space, the likelihood of target satisfaction will be zero almost everywhere in the original design space. Thus before applying RIGID, we need to obtain a latent representation that compactly captures the manifold of valid designs [68, 69], and use the latent representation as design variables for inverse design.

Acknowledgement

This work was supported by the startup funds from the J. Mike Walker ’66 Department of Mechanical Engineering at Texas A&M University, the National Science Foundation (NSF) BRITE Fellow program (CMMI 2227641), the NSF CSSI program (OAC 1835782), the Kansas City National Security Campus (PDRD #705288), and NSF CAREER Award (CMMI-2142460). R.S. acknowledges financial support from the NSF Graduate Research Fellowship Program.

References

- [1] Shan Jiang, Xuejun Liu, Jianpeng Liu, Dong Ye, Yongqing Duan, Kan Li, Zhouping Yin, and YongAn Huang. Flexible metamaterial electronics. *Advanced Materials*, 34(52):2200070, 2022.
- [2] J William Boley, Wim M Van Rees, Charles Lissandrello, Mark N Horenstein, Ryan L Truby, Arda Kotikian, Jennifer A Lewis, and L Mahadevan. Shape-shifting structured lattices via multimaterial 4d printing. *Proceedings of the National Academy of Sciences*, 116(42):20856–20862, 2019.
- [3] Chunping Ma, Yilong Chang, Shuai Wu, and Ruike Renee Zhao. Deep learning-accelerated designs of tunable magneto-mechanical metamaterials. *ACS Applied Materials & Interfaces*, 14(29):33892–33902, 2022.
- [4] Filippo Casadei, Tommaso Delpero, Andrea Bergamini, Paolo Ermanni, and Massimo Ruzzene. Piezoelectric resonator arrays for tunable acoustic waveguides and metamaterials. *Journal of Applied Physics*, 112(6):064902, 09 2012.
- [5] Guang-Sheng Liu, Yang Zhou, Ming-Hao Liu, Ying Yuan, Xin-Ye Zou, and Jian-Chun Cheng. Acoustic waveguide with virtual soft boundary based on metamaterials. *Scientific Reports*, 10(1):981, 2020.
- [6] Gunho Kim, Carlos M Portela, Paolo Celli, Antonio Palermo, and Chiara Daraio. Poroelastic microlattices for underwater wave focusing. *Extreme Mechanics Letters*, 49:101499, 2021.
- [7] Yangbo Xie, Yangyang Fu, Zhetao Jia, Junfei Li, Chen Shen, Yadong Xu, Huanyang Chen, and Steven A Cummer. Acoustic imaging with metamaterial luneburg lenses. *Scientific reports*, 8(1):16188, 2018.
- [8] Kathryn H Matlack, Anton Bauhofer, Sebastian Krödel, Antonio Palermo, and Chiara Daraio. Composite 3d-printed metastructures for low-frequency and broadband vibration absorption. *Proceedings of the National Academy of Sciences*, 113(30):8386–8390, 2016.
- [9] Sebastian Krödel, Nicolas Thomé, and Chiara Daraio. Wide band-gap seismic metastructures. *Extreme Mechanics Letters*, 4:111–117, 2015.
- [10] Alireza Bayat and Stavros Gaitanaros. Wave directionality in three-dimensional periodic lattices. *Journal of Applied Mechanics*, 85(1):011004, 2018.
- [11] Henrik Ronellenfitsch, Norbert Stoop, Josephine Yu, Aden Forrow, and Jörn Dunkel. Inverse design of discrete mechanical metamaterials. *Physical Review Materials*, 3(9):095201, 2019.
- [12] Heedong Goh and Loukas F Kallivokas. Inverse band gap design of elastic metamaterials for p and sv wave control. *Computer Methods in Applied Mechanics and Engineering*, 370:113263, 2020.
- [13] Xiang Li, Shaowu Ning, Zhanli Liu, Ziming Yan, Chengcheng Luo, and Zhuo Zhuang. Designing phononic crystal with anticipated band gap through a deep learning based data-driven method. *Computer Methods in Applied Mechanics and Engineering*, 361:112737, 2020.
- [14] Xiaopeng Zhang, Jian Xing, Pai Liu, Yangjun Luo, and Zhan Kang. Realization of full and directional band gap design by non-gradient topology optimization in acoustic metamaterials. *Extreme Mechanics Letters*, 42:101126, 2021.
- [15] Zhi Chen, Alexander Ogren, Chiara Daraio, L Catherine Brinson, and Cynthia Rudin. How to see hidden patterns in metamaterials with interpretable machine learning. *Extreme Mechanics Letters*, 57:101895, 2022.
- [16] Robin Vismara, Nils Odebo Länk, Ruggero Verre, Mikael Käll, Olindo Isabella, and Miro Zeman. Solar harvesting based on perfect absorbing all-dielectric nanoresonators on a mirror. *Optics Express*, 27(16):A967–A980, 2019.
- [17] Yongzhi Cheng, Fu Chen, and Hui Luo. Plasmonic chiral metasurface absorber based on bilayer fourfold twisted semicircle nanostructure at optical frequency. *Nanoscale Research Letters*, 16:1–9, 2021.
- [18] Sandeep Inampudi and Hossein Mosallaei. Neural network based design of metagratings. *Applied Physics Letters*, 112(24), 2018.
- [19] Maksym V Zhelyeznyakov, Steve Brunton, and Arka Majumdar. Deep learning to accelerate scatterer-to-field mapping for inverse design of dielectric metasurfaces. *ACS Photonics*, 8(2):481–488, 2021.
- [20] Dianjing Liu, Yixuan Tan, Erfan Khoram, and Zongfu Yu. Training deep neural networks for the inverse design of nanophotonic structures. *Acs Photonics*, 5(4):1365–1369, 2018.
- [21] Wei Ma, Feng Cheng, Yihao Xu, Qinlong Wen, and Yongmin Liu. Probabilistic representation and inverse design of metamaterials based on a deep generative model with semi-supervised learning strategy. *Advanced Materials*, 31(35):1901111, 2019.

- [22] Sensong An, Bowen Zheng, Hong Tang, Mikhail Y. Shalaginov, Li Zhou, Hang Li, Myungkoo Kang, Kathleen A. Richardson, Tian Gu, Juejun Hu, Clayton Fowler, and Hualiang Zhang. Multifunctional Metasurface Design with a Generative Adversarial Network. *Advanced Optical Materials*, 9(5):1–10, 2021.
- [23] Itzik Malkiel, Michael Mrejen, Achiya Nagler, Uri Arieli, Lior Wolf, and Haim Suchowski. Plasmonic nanostructure design and characterization via Deep Learning. *Light: Science and Applications*, 7(1), 2018.
- [24] Doksoo Lee, Wei Wayne Chen, Liwei Wang, Yu-Chin Chan, and Wei Chen. Data-driven design for metamaterials and multiscale systems: A review. *arXiv preprint arXiv:2307.05506*, 2023.
- [25] Sensong An, Clayton Fowler, Bowen Zheng, Mikhail Y Shalaginov, Hong Tang, Hang Li, Li Zhou, Jun Ding, Anuradha Murthy Agarwal, Clara Rivero-Baleine, et al. A deep learning approach for objective-driven all-dielectric metasurface design. *ACS Photonics*, 6(12):3196–3207, 2019.
- [26] Siddhant Kumar, Stephanie Tan, Li Zheng, and Dennis M Kochmann. Inverse-designed spinodoid metamaterials. *npj Computational Materials*, 6(1):1–10, 2020.
- [27] Christopher Yeung, Ju Ming Tsai, Brian King, Benjamin Pham, David Ho, Julia Liang, Mark W. Knight, and Aaswath P. Raman. Multiplexed supercell metasurface design and optimization with tandem residual networks. *Nanophotonics*, 10(3):1133–1143, 1 2021.
- [28] Jan-Hendrik Bastek, Siddhant Kumar, Bastian Telgen, Raphaël N Glaesener, and Dennis M Kochmann. Inverting the structure–property map of truss metamaterials by deep learning. *Proceedings of the National Academy of Sciences*, 119(1):e2111505119, 2022.
- [29] Jiaqi Jiang, David Sell, Stephan Hoyer, Jason Hickey, Jianji Yang, and Jonathan A Fan. Free-form diffractive metagrating design based on generative adversarial networks. *ACS nano*, 13(8):8872–8878, 2019.
- [30] Sunae So and Junsuk Rho. Designing nanophotonic structures using conditional deep convolutional generative adversarial networks. *Nanophotonics*, 8(7):1255–1261, 2019.
- [31] Fufang Wen, Jiaqi Jiang, and Jonathan A Fan. Robust freeform metasurface design based on progressively growing generative networks. *ACS Photonics*, 7(8):2098–2104, 2020.
- [32] Caglar Gurbuz, Felix Kronowetter, Christoph Dietz, Martin Eser, Jonas Schmid, and Steffen Marburg. Generative adversarial networks for the design of acoustic metamaterials. *The Journal of the Acoustical Society of America*, 149(2):1162–1174, 2021.
- [33] Andrew J Lew and Markus J Buehler. Single-shot forward and inverse hierarchical architected materials design for nonlinear mechanical properties using an attention-diffusion model. *Materials Today*, 64:10–20, 2023.
- [34] Mahmoud Elzouka, Charles Yang, Adrian Albert, Ravi S Prasher, and Sean D Lubner. Interpretable forward and inverse design of particle spectral emissivity using common machine-learning models. *Cell Reports Physical Science*, 1(12), 2020.
- [35] Leo Breiman. Random forests. *Machine learning*, 45:5–32, 2001.
- [36] WK Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, pages 97–109, 1970.
- [37] Leo Breiman, Jerome Friedman, Charles J Stone, and R A Olshen. *Classification and Regression Trees*. Chapman and Hall/CRC, 1984.
- [38] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- [39] John Brian Pendry. Negative refraction makes a perfect lens. *Physical review letters*, 85(18):3966, 2000.
- [40] Nanfang Yu, Patrice Genevet, Mikhail A Kats, Francesco Aieta, Jean-Philippe Tetienne, Federico Capasso, and Zeno Gaburro. Light propagation with phase discontinuities: generalized laws of reflection and refraction. *science*, 334(6054):333–337, 2011.
- [41] Alexander V Kildishev, Alexandra Boltasseva, and Vladimir M Shalaev. Planar photonics with metasurfaces. *Science*, 339(6125):1232009, 2013.
- [42] Tie Jun Cui, Mei Qing Qi, Xiang Wan, Jie Zhao, and Qiang Cheng. Coding metamaterials, digital metamaterials and programmable metamaterials. *Light: science & applications*, 3(10):e218–e218, 2014.
- [43] Syed S Bukhari, J Yiannis Vardaxoglou, and William Whittow. A metasurfaces review: Definitions and applications. *Applied Sciences*, 9(13):2727, 2019.
- [44] Jie Hu, Sankhyabrata Bandyopadhyay, Yu-hui Liu, and Li-yang Shao. A review on metasurface: from principle to smart metadevices. *Frontiers in Physics*, 8:586087, 2021.

- [45] N Iê Landy, S Sajuyigbe, Jack J Mock, David R Smith, and Willie J Padilla. Perfect metamaterial absorber. *Physical review letters*, 100(20):207402, 2008.
- [46] Xianliang Liu, Tatiana Starr, Anthony F Starr, and Willie J Padilla. Infrared spatial and frequency selective metamaterial with near-unity absorbance. *Physical review letters*, 104(20):207403, 2010.
- [47] Jiaming Hao, Jing Wang, Xianliang Liu, Willie J Padilla, Lei Zhou, and Min Qiu. High performance optical absorber based on a plasmonic metamaterial. *Applied Physics Letters*, 96(25), 2010.
- [48] Claire M Watts, Xianliang Liu, and Willie J Padilla. Metamaterial electromagnetic wave absorbers. *Advanced materials*, 24(23):OP98–OP120, 2012.
- [49] Yanxia Cui, Yingran He, Yi Jin, Fei Ding, Liu Yang, Yuqian Ye, Shoumin Zhong, Yinyue Lin, and Sailing He. Plasmonic and metamaterial structures as electromagnetic absorbers. *Laser & Photonics Reviews*, 8(4):495–520, 2014.
- [50] Xinyu Liu, Kebin Fan, Ilya V Shadrivov, and Willie J Padilla. Experimental realization of a terahertz all-dielectric metasurface absorber. *Optics express*, 25(1):191–201, 2017.
- [51] Gun-Yeal Lee, Gwanho Yoon, Seung-Yeol Lee, Hansik Yun, Jaebum Cho, Kyookeun Lee, Hwi Kim, Junsuk Rho, and ByoungHo Lee. Complete amplitude and phase control of light using broadband holographic metasurfaces. *Nanoscale*, 10(9):4237–4245, 2018.
- [52] You Zhou, Ivan I Kravchenko, Hao Wang, J Ryan Nolen, Gong Gu, and Jason Valentine. Multilayer noninteracting dielectric metasurfaces for multiwavelength metaoptics. *Nano letters*, 18(12):7529–7537, 2018.
- [53] You Li, Qunsheng Cao, and Yi Wang. A wideband multifunctional multilayer switchable linear polarization metasurface. *IEEE Antennas and Wireless Propagation Letters*, 17(7):1314–1318, 2018.
- [54] Giuseppe Marino, Davide Rocco, Carlo Gigli, Grégoire Beaudoin, Konstantinos Pantzas, Stéphan Suffit, Pascal Filloux, Isabelle Sagnes, Giuseppe Leo, and Costantino De Angelis. Harmonic generation with multi-layer dielectric metasurfaces. *Nanophotonics*, 10(7):1837–1843, 2021.
- [55] Stephanie C Malek, Adam C Overvig, Andrea Alù, and Nanfang Yu. Multifunctional resonant wavefront-shaping meta-optics based on multilayer and multi-perturbation nonlocal metasurfaces. *Light: Science & Applications*, 11(1):246, 2022.
- [56] Yuanjian Zhang, Yingting Yi, Wenxin Li, Shiri Liang, Jing Ma, Shubo Cheng, Wenxing Yang, and Yougen Yi. High absorptivity and ultra-wideband solar absorber based on ti-al₂o₃ cross elliptical disk arrays. *Coatings*, 13(3):531, 2023.
- [57] Yichi Zhang, Daniel W Apley, and Wei Chen. Bayesian optimization for materials design with mixed quantitative and qualitative variables. *Scientific reports*, 10(1):4924, 2020.
- [58] Chao Ma, Sebastian Tschiatschek, Richard Turner, José Miguel Hernández-Lobato, and Cheng Zhang. Vaem: a deep generative model for heterogeneous mixed type data. *Advances in Neural Information Processing Systems*, 33:11237–11247, 2020.
- [59] Lei Xu and Kalyan Veeramachaneni. Synthesizing tabular data using generative adversarial networks. *arXiv preprint arXiv:1811.11264*, 2018.
- [60] Yu-Chin Chan, Faez Ahmed, Liwei Wang, and Wei Chen. Metaset: Exploring shape and property spaces for data-driven metamaterials design. *Journal of Mechanical Design*, 143(3):031707, 2021.
- [61] Doksoo Lee, Yu-Chin Chan, Wei Chen, Liwei Wang, Anton van Beek, and Wei Chen. T-metaset: Task-aware acquisition of metamaterial datasets through diversity-based active learning. *Journal of Mechanical Design*, 145(3):031704, 2023.
- [62] Zhaocheng Liu, Dayu Zhu, Sean P Rodrigues, Kyu-Tae Lee, and Wenshan Cai. Generative model for the inverse design of metasurfaces. *Nano letters*, 18(10):6570–6576, 2018.
- [63] Nikolaos N Vlassis and WaiChing Sun. Denoising diffusion algorithm for inverse design of microstructures with fine-tuned nonlinear material properties. *Computer Methods in Applied Mechanics and Engineering*, 413:116126, 2023.
- [64] Weichen Li, Fengwen Wang, Ole Sigmund, and Xiaojia Shelly Zhang. Digital synthesis of free-form multimaterial structures for realization of arbitrary programmed mechanical responses. *Proceedings of the National Academy of Sciences*, 119(10):e2120563119, 2022.
- [65] Xin Lin, Fei Pan, Yong Ma, Yuling Wei, Kang Yang, Zihong Wu, Juan Guan, Bin Ding, Bin Liu, Jinwu Xiang, et al. Mechanical fourier transform for programmable metamaterials. *Proceedings of the National Academy of Sciences*, 120(37):e2305380120, 2023.

- [66] Aleksi Bossart, David MJ Dykstra, Jop Van der Laan, and Corentin Coulais. Oligomodal metamaterials with multifunctional mechanics. *Proceedings of the National Academy of Sciences*, 118(21):e2018610118, 2021.
- [67] Sensong An, Bowen Zheng, Hong Tang, Mikhail Y Shalaginov, Li Zhou, Hang Li, Myungkoo Kang, Kathleen A Richardson, Tian Gu, Juejun Hu, et al. Multifunctional metasurface design with a generative adversarial network. *Advanced Optical Materials*, 9(5):2001433, 2021.
- [68] Liwei Wang, Yu-Chin Chan, Faez Ahmed, Zhao Liu, Ping Zhu, and Wei Chen. Deep generative modeling for mechanistic-based learning and design of metamaterial systems. *Computer Methods in Applied Mechanics and Engineering*, 372:113377, 2020.
- [69] Wei Chen, Doksoo Lee, Oluwaseyi Balogun, and Wei Chen. Gan-duf: Hierarchical deep generative models for design under free-form geometric uncertainty. *Journal of Mechanical Design*, 145(1):011703, 2023.
- [70] COMSOL AB. Comsol multiphysics (version 5.6). <https://www.comsol.com>, 2020. Stockholm, Sweden.
- [71] David E Aspnes and AA Studna. Dielectric functions and optical parameters of si, ge, gap, gaas, gasb, inp, inas, and insb from 1.5 to 6.0 ev. *Physical review B*, 27(2):985, 1983.

Appendices

A Related Works on Single-Shot Inverse Design

Figure 7 shows the schematic diagrams of three mainstream machine learning models for single-shot inverse design of metamaterials.

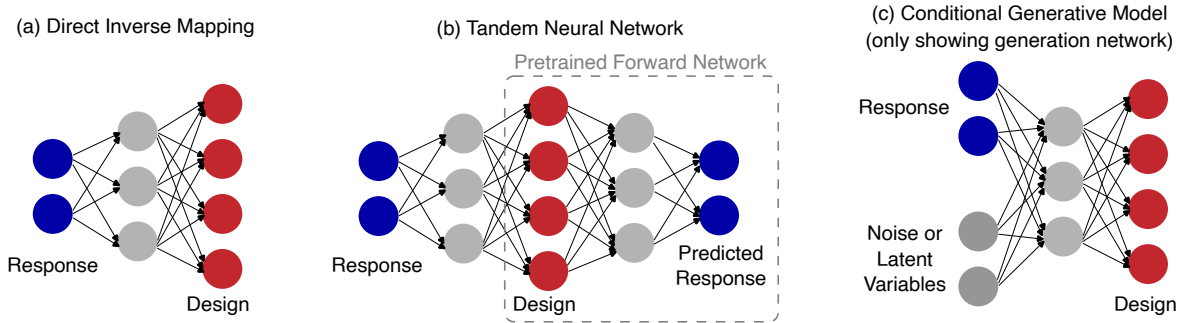


Figure 7: Three major machine learning models for single-shot inverse design. Note that these are only conceptual plots and the networks can be more complicated in reality. Also, due to the diversity and complexity of conditional generative model architectures, we only show the network for design generation based on commonly used models (e.g., generators in cGANs or decoders in cVAEs), while omitting networks that are necessary for training (e.g., discriminators in cGANs or encoders in cVAEs).

B Explanation of Step 2 in RIGID

The purpose of Step 2 is to obtain the intersection of relevant design space regions for all the ranges of s in a target. We are approximating this goal by simply obtaining the intersection of relevant leaves. However, some non-intersecting leaves may still have intersecting design space regions. When assigning the probability to the intersecting region of two non-intersecting leaves A and B , we need to consider the predicted probabilities at both leaves (P_A and P_B). Specifically, the assigned probability at this intersecting region should be $P_A P_B$, which can be small. Therefore, we adopt the simplification of only considering the intersection of relevant leaves and ignoring the intersecting regions associated with non-intersecting leaves. The results also demonstrate that this is a reasonable approximation.

C Computational Methods for Data Collection and Validation

C.1 Computation of Acoustic Dispersion Relations

We performed Bloch-wave analysis in COMSOL Multiphysics to compute the dispersion relations of acoustic metamaterials. Poisson’s ratio of 0.49, Young’s modulus of 2.7 GPa, and density of 1170 kg/m³ were set as material properties with $\sim 1.5 \times 10^4$ mesh elements per unit cell. We used Floquet-Bloch periodic boundary conditions to obtain the first 60 eigenfrequencies along all symmetry domains of the cubic irreducible Brillouin zone (Fig. 3B) for all lattices, thus generating a dispersion relation.

C.2 Computation of Optical Absorbance Spectra

We computed the absorbance spectra for optical metasurfaces using wave analysis inspired by Zhang et al. [56]. The RF Module of COMSOL Multiphysics® [70] was used to evaluate the spectral response of concern, which is the energy absorbance $A(s)$ in the visible regime (380-700 nm). An absorbance spectrum is computed with respect to 33 wavelength components s_k that are uniformly discretized over the specified range. An incident plane wave is assumed to be given from the port, located at the top face of the analysis domain. We set the periodicity of the analysis domain as 400 nm. The periodic boundary condition on electromagnetic fields is imposed on the lateral faces of the analysis domain. A substrate made of SiO₂ is placed right below a given unit cell instance (the black layers in Fig. 4A-B). With full electric fields computed through the wave analysis, the energy absorbance at a single wavelength s is quantified as $A(s) = 1 - |S_{11}(s)|^2$, where S_{ij} is the component of the S -parameter matrix that specifies energy transfer between ports. We use the data presented in Ref. [71] to set the material dispersion of the dielectric.

D Test performance of forward prediction

Test data for all the design problems are designs \mathbf{x} (that the random forests have never seen during training) and their corresponding qualitative behaviors $y \in \{0, 1\}$ based on functional responses.

In the acoustic metamaterial design problem, we have 57 test designs, yielding 5,700 test points as the entire frequency range of dispersion relations is discretized into 100 intervals for each design. In the optical metasurface design problem, we have 52 test designs, with each functional response discretized into 33 points, which yields 1,716 test points in total. In each of the synthetic design problems, we have 20 synthetic test designs, with each synthetic response discretized into 100 points. This results in 2,000 test data points in total. Confusion matrices showing test performances are in Tables 1-4.

Table 1: Confusion matrix on test data in the acoustic metamaterial design problem. The test F1 score is 0.82.

		Predicted	
		Yes	No
Actual	Yes	393	131
	No	46	5130

Table 2: Confusion matrix on test data in the optical metasurface design problem. The test F1 score is 0.83.

		Predicted	
		Yes	No
Actual	Yes	776	157
	No	151	632

E Generated Designs and Corresponding Responses

Figures 8-12 show the rest of generated designs and their corresponding responses in addition to those in Figures 3F and 4F.

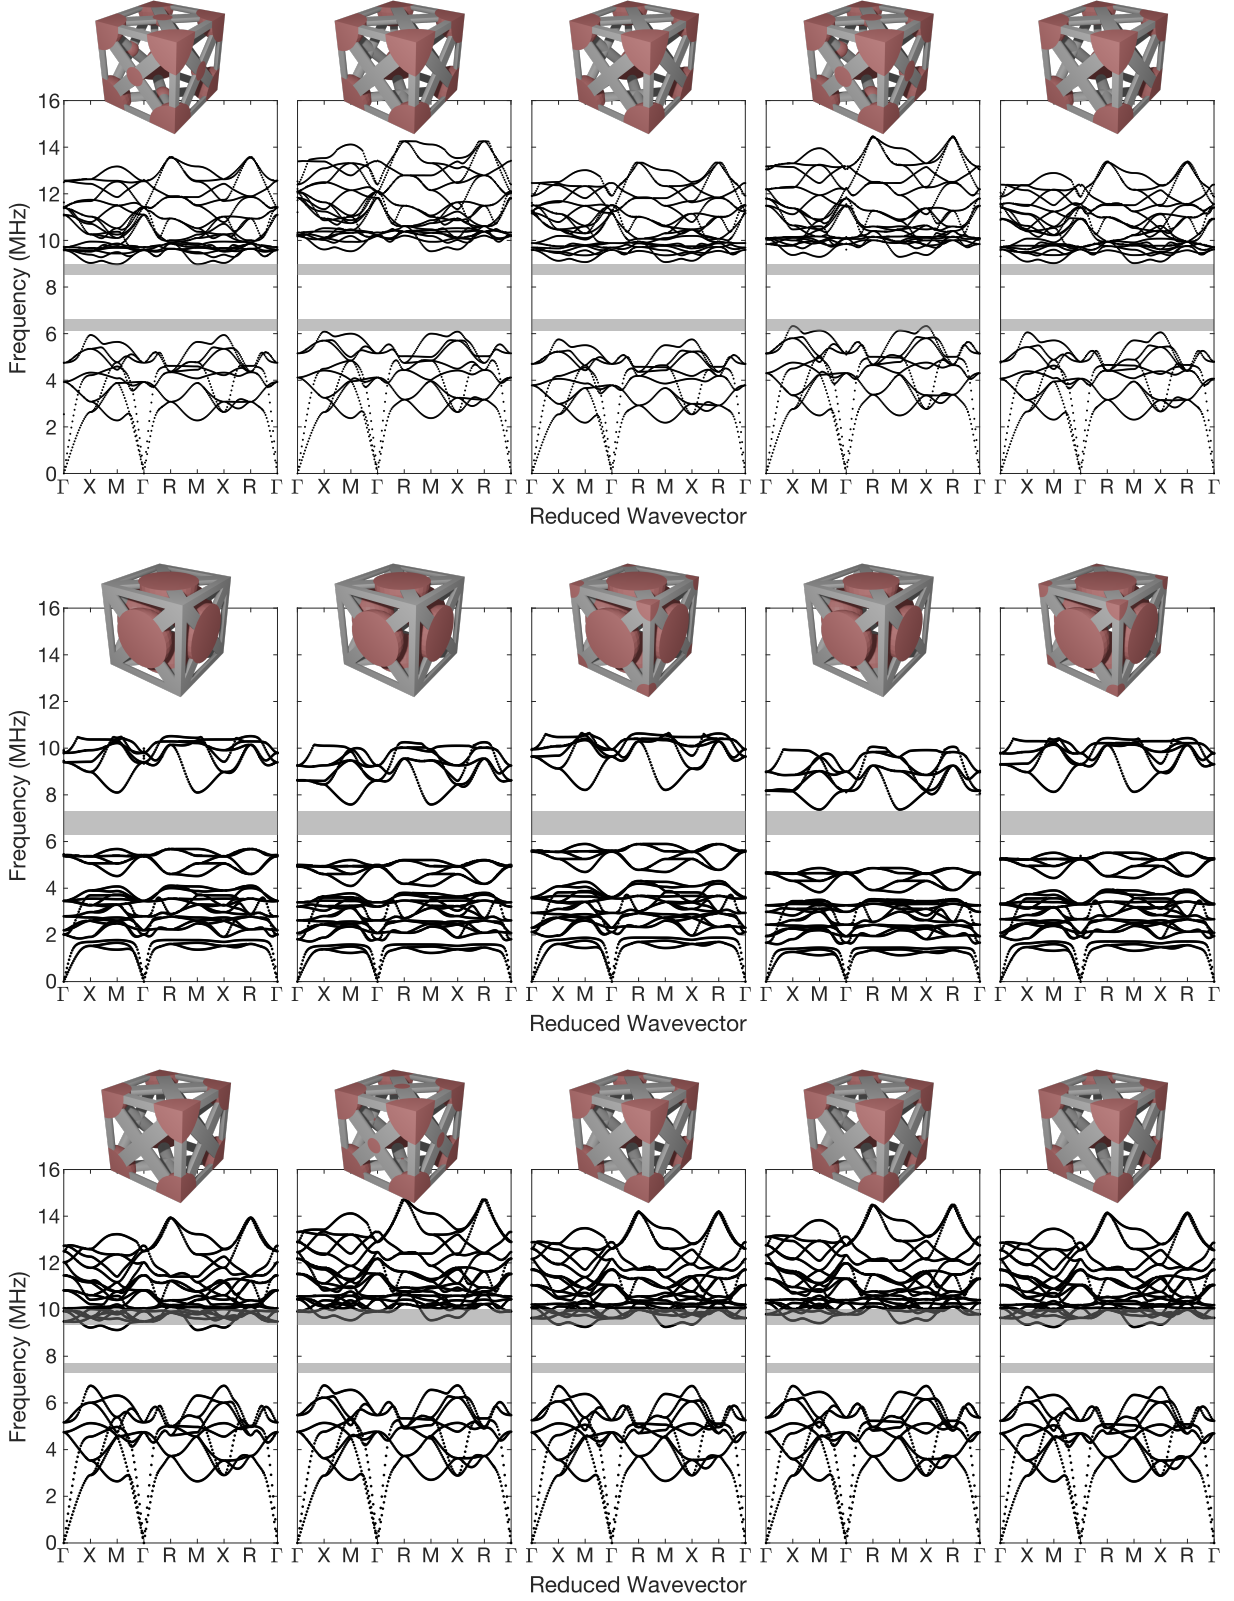


Figure 8: Geometries and corresponding dispersion relations of generated acoustic metamaterial designs for bandgap targets II-IV (marked as gray-shaded frequency regions). Each row shows five designs with the highest likelihood of satisfying a target.

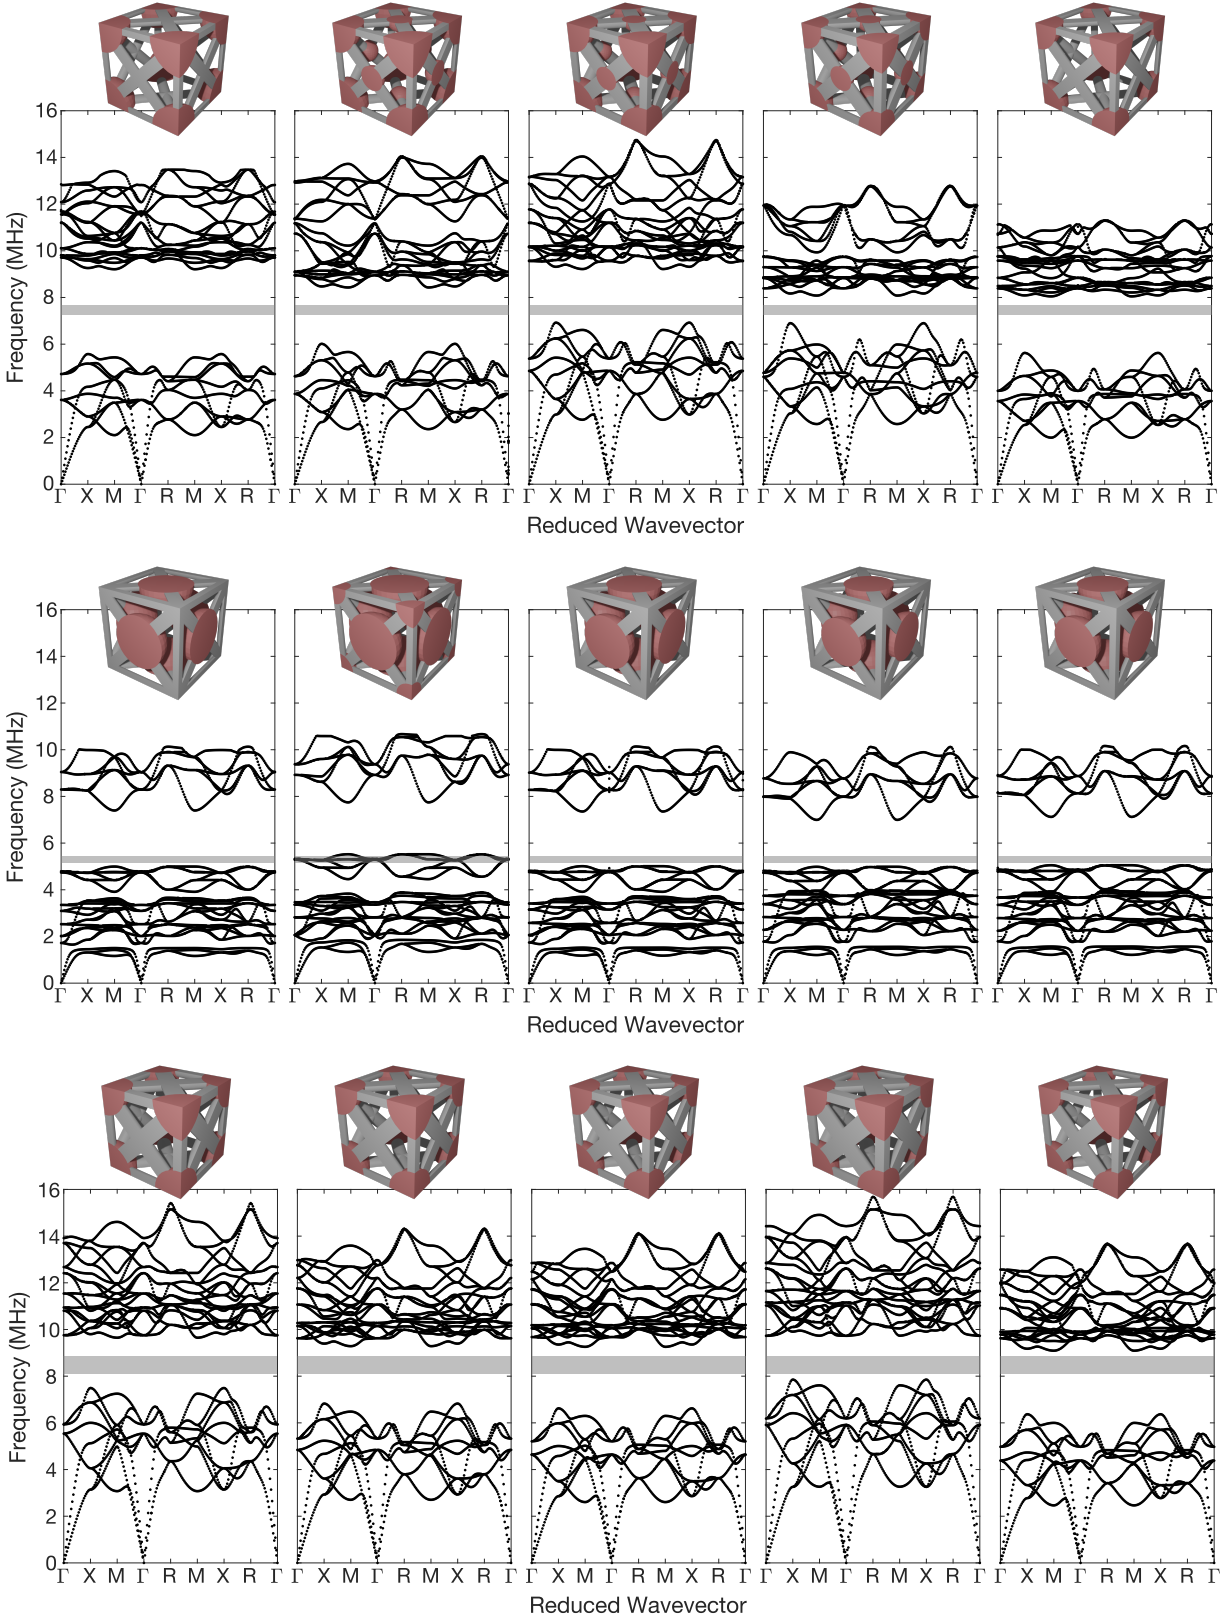


Figure 9: Geometries and corresponding dispersion relations of generated acoustic metamaterial designs for bandgap targets V-VII (marked as gray-shaded frequency regions). Each row shows five designs with the highest likelihood of satisfying a target.

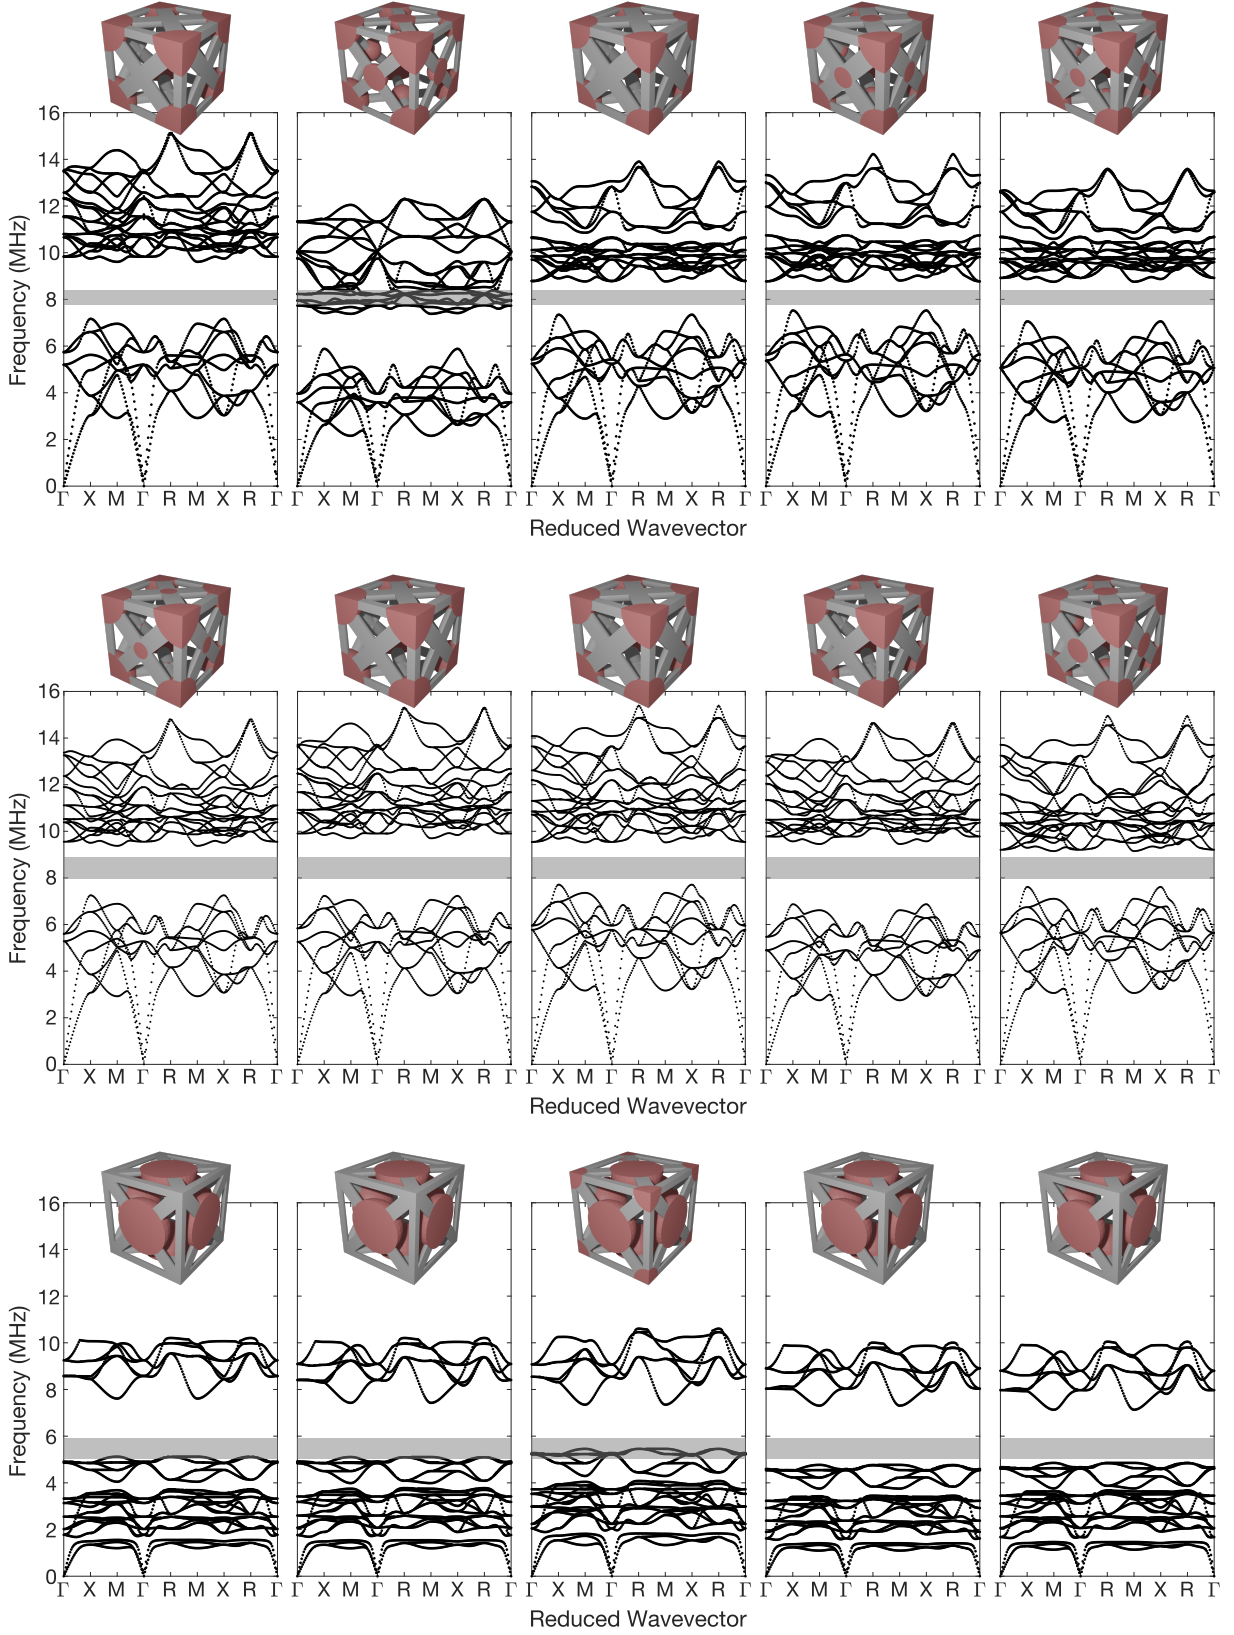


Figure 10: Geometries and corresponding dispersion relations of generated acoustic metamaterial designs for bandgap targets VIII-X (marked as gray-shaded frequency regions). Each row shows five designs with the highest likelihood of satisfying a target.

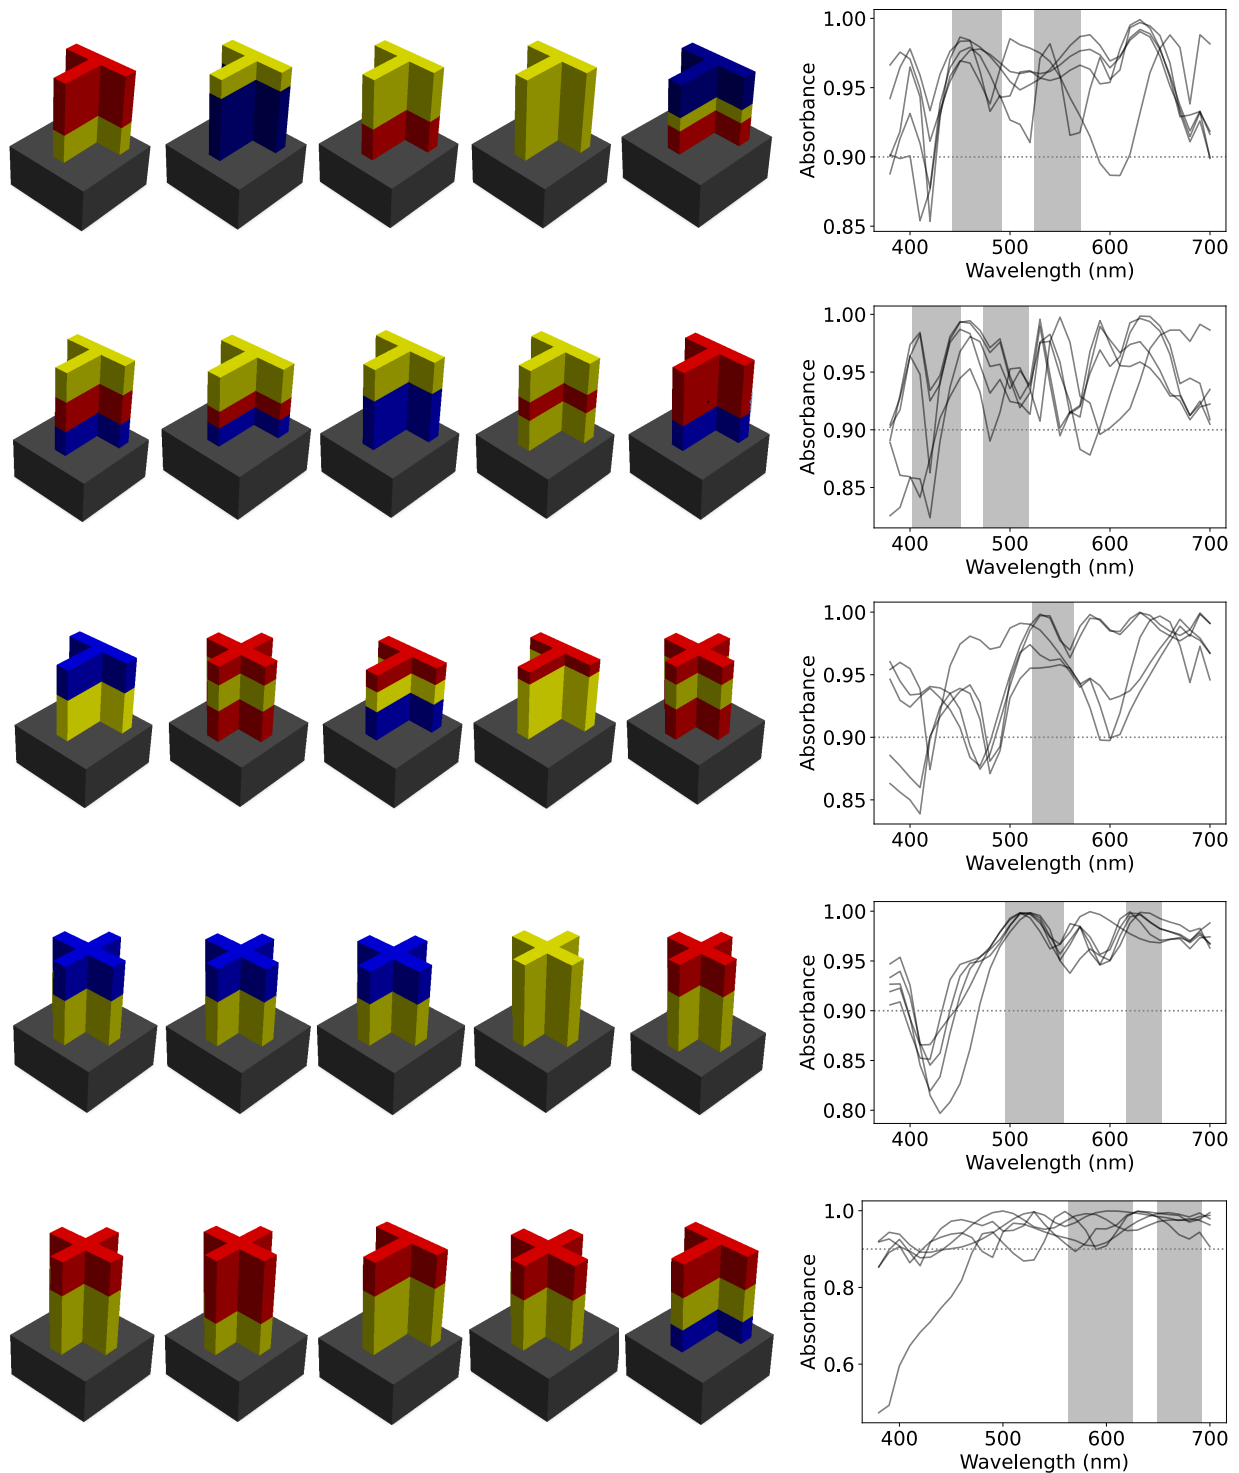


Figure 11: Designs and corresponding absorbance spectra of generated optical metasurface designs for design targets II-VI (marked as gray-shaded wavelength regions). Each row shows five designs with the highest likelihood of satisfying a target (left panel) and their absorbance spectra (right panel).

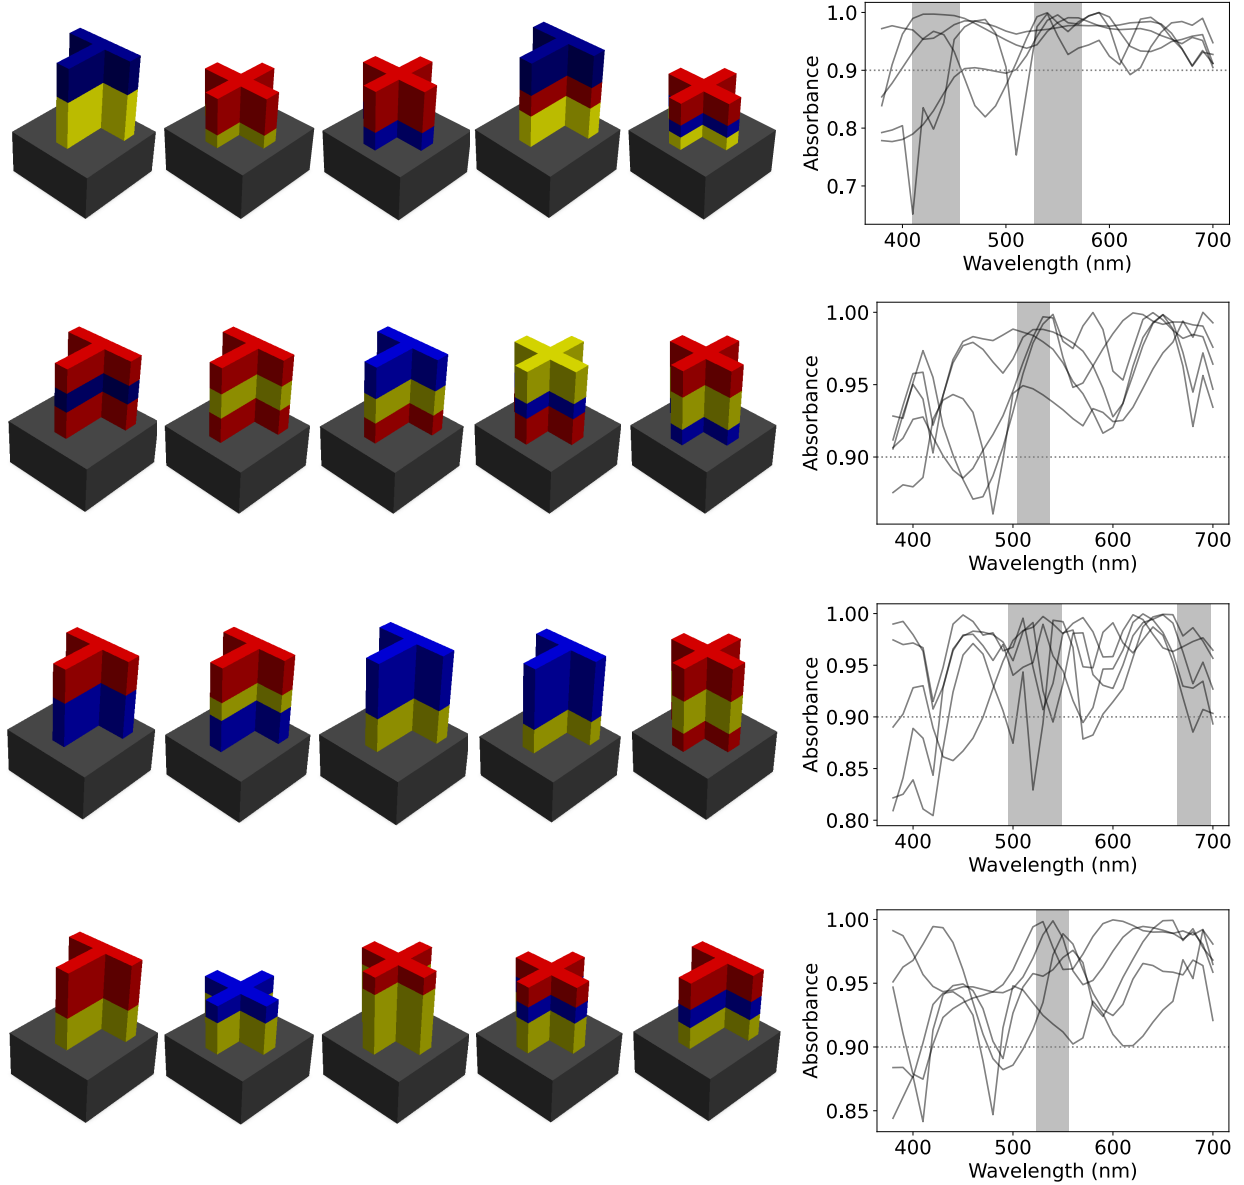


Figure 12: Designs and corresponding absorbance spectra of generated optical metasurface designs for design targets VII-X (marked as gray-shaded wavelength regions). Each row shows five designs with the highest likelihood of satisfying a target (left panel) and their absorbance spectra (right panel).

Table 3: Confusion matrix on test data in the SqExp problem. The test F1 score is 0.85.

		Predicted	
		Yes	No
Actual	Yes	229	68
	No	13	1690

Table 4: Confusion matrix on test data in the SupSin problem. The test F1 score is 0.86.

		Predicted	
		Yes	No
Actual	Yes	355	81
	No	37	1527