iScience

CellPress
OPEN ACCESS

Article

NEUTRON: Neural particle swarm optimization for material-aware inverse design of structural color

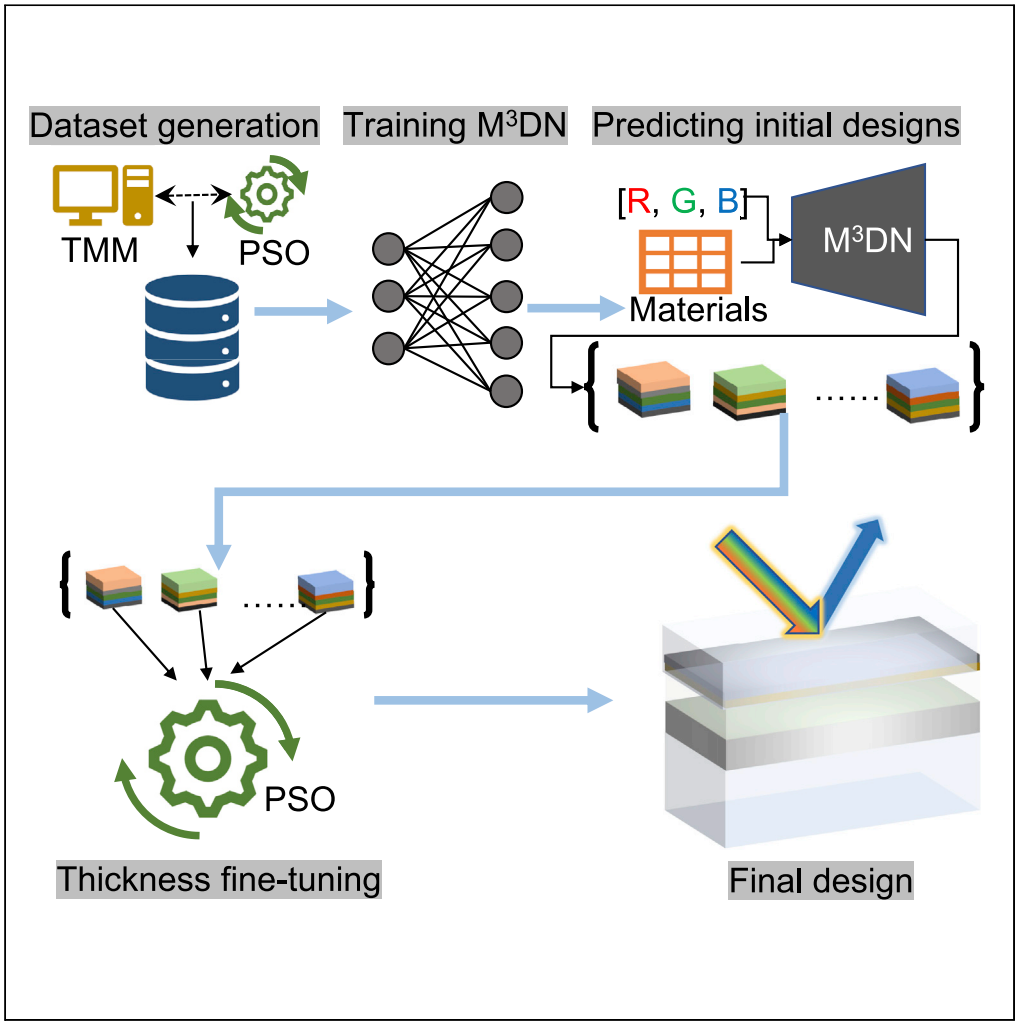

Haozhu Wang, L. Jay Guo

hzwang@umich.edu (H.W.)
guo@umich.edu (L.J.G.)

Highlights

NEUTRON combines machine learning and optimization for structural color designs

On two benchmark tasks, NEUTRON demonstrates high design accuracy and efficiency

Wang & Guo, iScience *25*, 104339
May 20, 2022
https://doi.org/10.1016/j.isci.2022.104339

iScience

CellPress
OPEN ACCESS

Article

NEUTRON: Neural particle swarm optimization for material-aware inverse design of structural color

Haozhu Wang[1,2,*] and L. Jay Guo[1,*]

SUMMARY

Designing optical structures for generating structural colors is challenging because of the complex relationship between the optical structures and the color perceived by human eyes. Machine learning-based approaches have been developed to expedite this design process. However, existing methods solely focus on structural parameters of the optical design, which could lead to suboptimal color generation because of the inability to optimize the selection of materials. To address this issue, an approach known as Neural Particle Swarm Optimization is proposed in this paper. The proposed method achieves high design accuracy and efficiency on two structural color design tasks; the first task is designing environment-friendly alternatives to chrome coatings, and the second task concerns reconstructing pictures with multilayer optical thin films. Several designs that could replace chrome coatings have been discovered; pictures with more than 200,000 pixels and thousands of unique colors can be accurately reconstructed in a few hours.

INTRODUCTION

Structural color refers to the color generated through the light interaction with patterned or layered optical structures. It is more stable than colors produced from chemical pigments and serves as an environment-friendly alternative. However, designing the structures for producing desired colors is challenging because of the complex relationship between the optical structures and their spectral properties. In addition, *color metamerism*, i.e., different spectra may correspond to the same color perceived by human eyes (Foster et al., 2006; Best, 2017), makes the relationship between the structures and the perceived color more complex because multiple different structures could have the same color appearance. Human experts often design optical structures based on the understanding of the physical properties of structures, including multilayer thin films (Wang et al., 2018; Yang et al., 2019), metasurfaces (Sun et al., 2017; Yang et al., 2020), and self-assembled colloidal particles (Kim et al., 2011; Liu et al., 2021), to name a few. Owing to the complex relationship between the structures and the generated color, the manual design process is often slow and could lead to suboptimal color production.

Recently, machine learning-based optical inverse design approaches have been developed to predict optical structures that can achieve user-specified color properties (Liu et al., 2018b; Ma et al., 2019; Unni et al., 2021; Yeung et al., 2021; Yao et al., 2019; Jiang et al., 2020). These inverse design methods often involve training a machine learning model such as deep neural networks (Gao et al., 2019) or support vector machines (Huang et al., 2019) on a curated dataset that contains a large number of data points mapping structural parameters to the corresponding color, e.g., represented by coordinates in CIE xyY (Gao et al., 2019) or LAB color space (Sajedian et al., 2019; Dai et al., 2021). Though previous methods have been demonstrated to be efficient in designing a wide range of colors, they often require materials constituting the optical structures to be fixed. Because the refractive index of materials affects their reflection and absorption properties, it could be challenging or even impossible to produce specific colors when the materials are not appropriately selected. Thus, the first step of screening appropriate materials for latter inverse design with machine learning models still requires extensive effort from human experts and is difficult for people without sufficient prior experience, which should be adequately addressed for the wide adoption of the developed machine learning models. For optical multilayer thin-film design, the recently reported reinforcement learning approach (Wang et al., 2021) addresses the material selection challenge by searching

[1]Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, USA

[2]Lead contact

*Correspondence: hzwang@umich.edu (H.W.), guo@umich.edu (L.J.G.)

https://doi.org/10.1016/j.isci.2022.104339

the material and thickness design space simultaneously. However, this method only designs a single color at a time because the reward function for training the reinforcement learning algorithm has to be defined for a specific color. It could be impractical when many colors need to be designed, e.g., reproducing all the colors found in a paint shop or designing a large array of reflective color pixels to reconstruct a colored picture, which would take an extremely long time using the reinforcement learning method reported in our previous work. To address these issues, we propose an inverse structural color design method that can efficiently predict both the materials and structural parameters in a synergistic manner.

The proposed method termed **Neu**ral Par**T**icle Swa**R**m **O**ptimizatio**N** (NEUTRON) is a hybrid approach that combines a **M**aterial-aware **M**ultitask **M**ixture **D**ensity **N**etwork (M^3DN) (Caruana, 1997; Bishop, 1994) and **P**article **S**warm **O**ptimization (PSO) (Kennedy and Eberhart, 1995). Instead of searching the material and thickness space simultaneously, NEUTRON first predicts the most suitable materials and provides a diverse set of initial guesses of the thicknesses in the form of probability distributions that could fulfill the target color and then applies particle swarm optimization to fine-tune the initial thickness designs. We demonstrate the effectiveness of the proposed approach on two optical multilayer thin film design tasks. The results show that our approach can lead to accurate color inverse designs efficiently. We believe that the proposed approach can be readily extended to many other optical design tasks where material selection and structural designs are both important.

BACKGROUND

Both numerical optimization and machine learning have been applied to the task of structural color design. Compared to machine learning, optimization-based methods for structural color design are often more accurate but suffer from low computational efficiency. Here, we briefly review the related works on numerical optimization and machine learning approaches for structural color design. For more comprehensive discussions on both approaches for broader nanophotonic design tasks, we refer readers to review articles (Campbell et al., 2019; So et al., 2020; Wiecha et al., 2021; Ma et al., 2021; Jiang et al., 2021).

Various optimization methods including PSO (Kennedy and Eberhart, 1995) and genetic algorithms (GAs) (Whitley, 1994) have been applied to design optical structures for structural color applications. While applying optimization methods to design optical structures for structural color, the designs are iteratively updated through the feedback from the electromagnetic simulations. Owing to the active feedback loop, optimization methods can often lead to highly accurate results but may suffer from low computational efficiency because of the large number of EM simulations (Schneider et al., 2019). Shokooh-Saremi and Magnusson applied both PSO and GA to the design of grating filters with a sharp reflection peak around 550 nm (Shokooh-Saremi and Magnusson, 2007). Both optimization methods took a few hours to converge for a single design target. Yang et al. applied PSO to design 2D grating structures for reflective structural color (Yang et al., 2013). The reflective color maintains CIEDE2000 less than eight for light incidence angles up to 45°. Rabady and Ababneh optimized the layer thicknesses of optical multilayer thin films with PSO (Rabady and Ababneh, 2014), where the alternating materials with high and low refractive index are fixed. Recently, Liu et al. applied GA to optimize the geometry parameters of metasurfaces with multiple meta-atoms in a single cell to produce highly pure primary red, blue, and green colors. Compared to metasurfaces without multiple meta-atoms, the achievable area in the CIE diagram is expanded by 106% (Liu et al., 2020). All aforementioned works only apply optimization methods to a few design targets because of the low computational efficiency.

Machine learning methods, especially deep learning, have been recently applied to the task of structural color designs. Unlike numerical optimizations, machine learning methods learn models from the data points that can later be used to efficiently predict designs corresponding to user-specified color targets. However, machine learning could lead to lower design performance than optimizations because of the lack of active evaluations with EM simulations. Tandem networks combine a forward network and an inverse network to directly predict the designs corresponding to a color target (Liu et al., 2018a). The forward network predicted optical properties given the design, whereas the inverse network performs the opposite. Gao et al. optimized the reflective color of silicon metasurface (Gao et al., 2019) and Dai et al. (2021) designed three layer Fabry-Perot color filters with tandem networks. However, both works can only output a single design given a user-specified color input, whereas our proposed method can output a set of potential solutions. The capability of outputting a set of solutions allows researchers to select designs that are amenable to the fabrication process (Ma et al., 2022). Mixture density networks (MDNs)

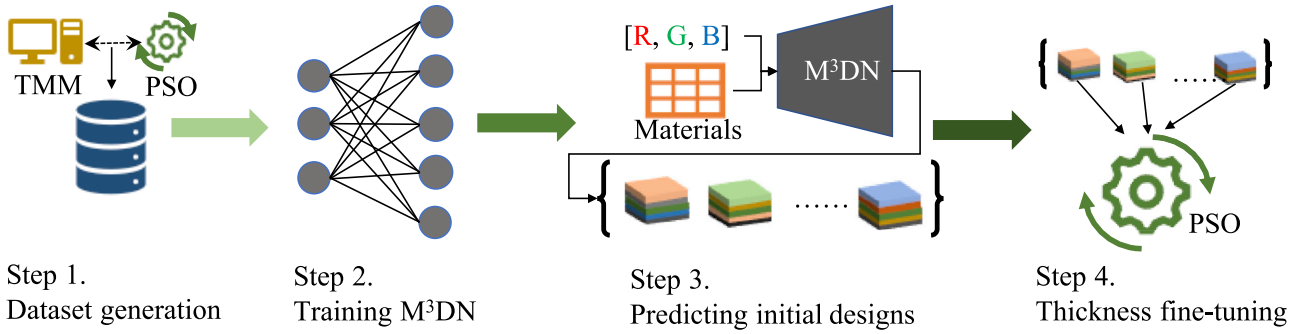

Figure 1. Pipeline of NEUTRON

We used a dataset generated through transfer-matrix method to train a M^3DN. After the training, M^3DN can output a set of designs based on different candidate materials. The initial designs' thicknesses are further fine-tuned through particle swarm optimization.

model the potential designs as a mixture of Gaussian distributions and are able to output a set of designs. Most closely related to our work, Unni et al. applied MDNs to design multilayer thin films with high reflection (Unni et al., 2021). Unlike our proposed approach, their method fixes the material selections for the alternating high refractive index and low refractive index layers, a well-established structure for optical filter design. Moreover, they select the best design from the set of MDN's outputted designs by comparing the optical properties predicted by a forward network. Instead of relying on predicting the optical properties with the forward model, our method uses PSO to fine-tune the thickness predictions by the MDN through the TMM simulations' feedback. Thus, our method does not suffer from the potential prediction errors from the forward model. Owing to the ability to search the optimal materials and accurately fine-tune the designed thicknesses with PSO, our proposed method could achieve a better design accuracy, with improved efficiency for a large number of color designs.

RESULTS

Here, we first provide an overview of the proposed method. Then, details on how the training dataset is generated are discussed. We highlight the exceptional performance of the proposed approach with results on two design tasks, including 1) environment-friendly chrome coating replacement design and 2) picture reconstructions. Finally, we conduct analysis to understand the effects of major components of the proposed method.

Overview of the proposed method

The proposed method consists of four steps for structural color designs (Figure 1). In the first step, we collect training data with a wide range of colors. Details of the dataset collection are described in the next section. Then, a novel neural network model called material-aware multitask mixture density network (M^3DN) is trained on the collected dataset. Next, for a given color design target, M^3DN outputs a set of potential designs with different material combinations and layer thicknesses. Finally, the set of initial designs are fine-tuned with PSO to obtain the final designs with optimized material selections and layer thicknesses.

Dataset generation

Previous research shows that five-layer optical thin films with two absorbing layers sandwiched by two dielectric layers and a bottom metal reflecting layer (Figure 2A) can achieve high color purity and brightness (Yang et al., 2019), where the layers can be easily deposited by physical vapor evaporation. Owing to the high performance and feasibility for large-scale fabrications of such structures, we synthesize a dataset with diverse designs based on the same five-layer structural template by varying both the material and thickness of each layer. All designs are based on randomly sampled materials from ten candidate metal materials Au, Ag, Al, Cu, Cr, Ge, Ni, Ti, W, Zn, and ten dielectric materials Al_2O_3, Fe_2O_3, HfO_2, MgF_2, SiO_2, Ta_2O_5, TiO_2, ZnO, ZnS, and ZnSe. Both absorber layers and the bottom reflective layers are composed of metals, whereas the other two layers are based on dielectric materials. Including a wide range of candidate materials with different refractive indices makes it possible to search for the most suitable materials combinations for specific color targets.

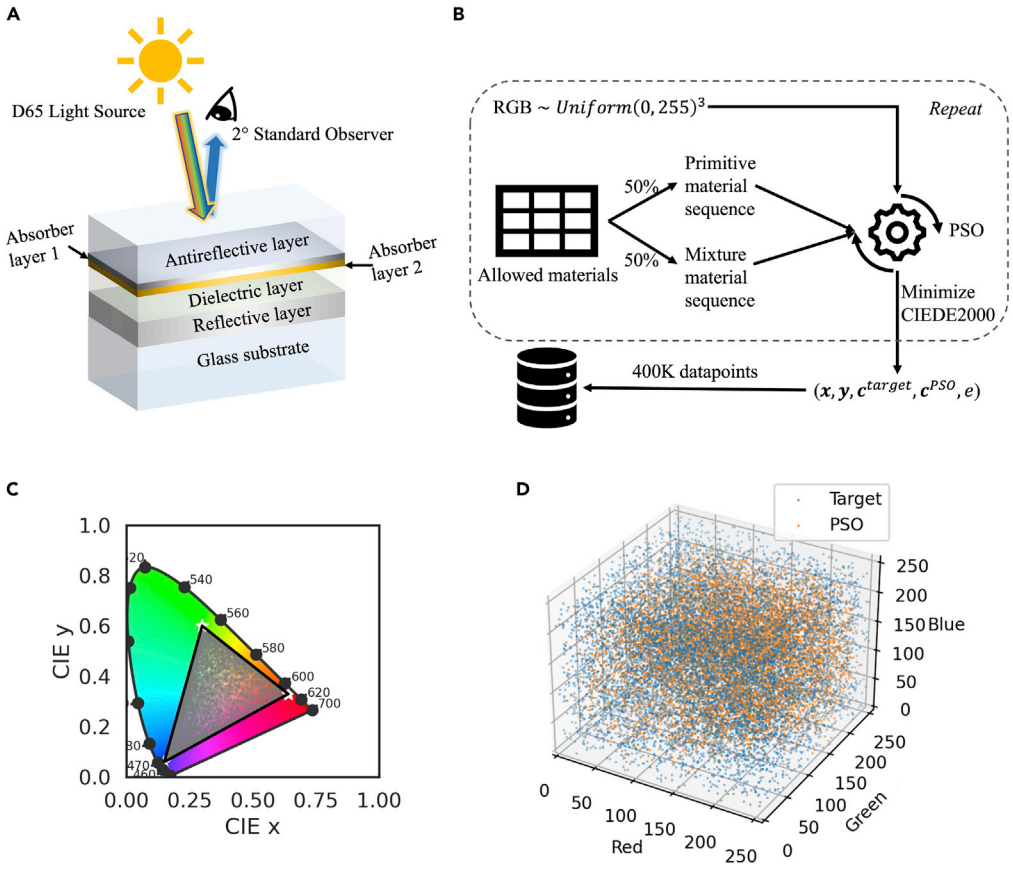

Figure 2. Five-layer optical structure and the data distribution

(A) A five-layer optical layer thin-film structure for generating a wide range of reflective colors covering the sRGB color gamut. Both the materials and their thicknesses are designed to obtain the target RGB color.

(B) The dataset generation process. 400,000 data points are obtained through PSO for randomly sampled RGB color target and material combinations.

(C) The color distribution of the validation set is visualized in the CIE 1931 xy space. The data points achieve good coverage of the entire sRGB color gamut spanned by the standard Red, Blue, and Green colors.

(D) Randomly sampled RGB color target and the obtained RGB color through PSO.

While sampling the materials for the two absorbing layers, we introduce a constraint that the two adjacent absorber layers must be composed of different materials; otherwise, it is simply a single absorber material with a larger thickness, which cannot satisfy the design requirement based on previous findings (Yang et al., 2019). Thus, the total number of unique material combinations of the five-layer stack is $10 \times 10 \times 9 \times 10 \times 10 = 9 \times 10^4$. Because sRGB is widely used in display industry design and production, we aim to provide the best coverage over the sRGB color gamut through our model. Most previous studies randomly sample the layer thicknesses from a prefixed range to generate data points with different designs and color properties. However, this approach often leads to nonuniform coverage in the color space, i.e., the density of data points in a certain color region is higher than other regions in the color space. Because nonuniform data coverage could lead to an undesirable skewed inverse design performance where the inverse design model is more accurate for the color region with more data points, we propose to directly sample from the color target space uniformly to avoid this problem. To this end, instead of randomly sampling the thicknesses of the layers as in most previous studies, we randomly sample the sRGB color target from the 3D sRGB space with a value range [0, 255] for each dimension. In Figures 2C and 2D, we show that this random color target sampling ensures uniform coverage of the color space.

Then, for each randomly sampled material combination and color target pair, we use PSO to optimize the thickness of the top four layers to minimize the color difference between the color of the designed structure and the color target while fixing the bottom reflective layer at 100 nm. PSO is a widely used global

optimization algorithm for optical designs (Shokooh-Saremi and Magnusson, 2007; Yang et al., 2013; Rabady and Ababneh, 2014). When solving a minimization problem, PSO maintains a group of particles (i.e., solutions) that individually explore the solution space and communicate with each other to share information about the explored solution space. All particles' positions are iteratively updated based on the best solution each particle has found and the best solution found by the entire group until a convergence criterion is met. In our five-layer optical thin film design task, each particle's "position" is a 4D vector corresponding to the top four layers' thickness. The final position all particles converge to is the final thickness design obtained through PSO. More details on PSO can be found in the STAR Methods section. Throughout the optimization process, the reflection spectrum of multilayer designs is computed with transfer-matrix method (Byrnes, 2016) for the wavelength range [400, 700] nm with a step size of 10 nm. Thus, each material's complex refractive index is described by a 62-dimensional vector for the entire 31 wavelength points. We use the Python package Colour (Mansecal et al., 2021) for the conversion of reflection spectrum to sRGB and Lab color coordinates. In addition, we use the industry-standard CIEDE2000 metric based on Lab coordinates to measure the color difference (Sharma et al., 2005) between the target color and the color obtained through the designs. After particle swarm optimization converges, if CIEDE2000 for the given material combination and the color target is lower than or equal to 2, we consider that the specific material combination allows accurate generation of the target color because CIEDE2000 lower than or equal to two is almost imperceivable by untrained eyes (ViewSonic, 2021), and we assign a label $e = 1$; otherwise the label is $e = 0$. Overall, we observe that the average success ratio of obtaining the accurate color following the random data generation process is about 25%. Both classes of data points ($e = 0$ and $e = 1$) are included in the final dataset. The Lab color coordinates and the sRGB coordinates for the same design have a one-to-one correspondence, and we only use Lab color coordinates for computing CIEDE2000. In contrast, sRGB coordinates are used explicitly as the color target. If users need to design specifically for a given Lab or CIE xyY target, sRGB coordinates can be uniquely obtained from the provided color coordinates defined on other color spaces.

In terms of the thickness of each layer, we set the thickness range for both dielectric layers from the range [5, 250] nm and the thickness of the absorbing layers to be in the range [5, 15] nm during the particle swarm optimization process. Again, the thickness of the bottom metal layer is fixed to be 100 nm, because it is used as a reflector whose thickness has a negligible effect on the reflective color as long, as the layer is thick enough to reflect the light completely. In addition, we constrain the thicknesses of all layers to be integers to allow fabrications. Thus, the size of the design space after considering variations in materials combinations and thickness designs is $9 \times 10^4 \times 11^2 \times 246^2 = 6.6 \times 10^{11}$.

Because the small number of primitive materials limits variations in the refractive index data, we augment the dataset by randomly mixing two dielectrics or two metals to form synthetic dielectric or metal composite materials to broaden the training dataset distribution through convex combinations of their complex refractive indices. Inspired by our previous finding that the mixture of two thin material layers can be considered as a linear combination of the two materials and helpful to improve color purity Yang et al. (2019), we obtain the complex refractive index of synthetic mixture materials through the following formula:

$$\boldsymbol{n}_i = \beta_i \bullet \boldsymbol{n}_{\mathcal{M}_{1,i}} + (1 - \beta_i) \bullet \boldsymbol{n}_{\mathcal{M}_{2,i}}, \qquad \text{(Equation 1)}$$

where $\mathcal{M}_{1,i}$ and $\mathcal{M}_{2,i}$ are the two sampled materials for ith layer, $\beta_i \sim (0, 1)$ is the mixing factor for the ith layer, $\boldsymbol{n}_{\mathcal{M}_{1,i}}, \boldsymbol{n}_{\mathcal{M}_{2,i}}$, and $\boldsymbol{n}_i$ are the complex refractive indices for randomly selected two primitive materials and the composite material synthesized from them for the ith layer. The composite material refractive index is synthetic and does not require to be experimentally achievable. It only serves as additional data to improve the data variation so that a machine learning model can learn more efficiently from the data. Similar data augmentation strategies have been widely adopted in other fields (Zhang et al., 2018). Following the described data generation procedure, we generated a total of 400,000 data points and about half of the samples are based on primitive (i.e., actual) materials, whereas the others are based on synthetic materials via linear superposition. The entire dataset is randomly split into 380,000/10,000/10,000 data points for training/validation/testing. Importantly, both validation and test sets are based on primitive materials only, because we are interested in the model's design accuracy for designs based on attainable actual materials, and mixture materials are introduced only in the training set for the data augmentation purpose. Unless stated otherwise, all results reported in the paper are based on the test sets. Each data point $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{c}^{target}, \boldsymbol{c}^{PSO}, e)$ is a tuple comprised of the refractive index data for all five layers concatenated as a single vector $\boldsymbol{x}$ with 310 dimensions (each material's complex refractive data is a 62-dimensional vector), the thickness of top four layers $\boldsymbol{y}$ obtained with PSO, the target color sRGB color

$\boldsymbol{c}^{target}$, the obtained color through PSO $\boldsymbol{c}^{PSO}$, and the binary label $e \in [0, 1]$, indicating whether the CIEDE2000 between $\boldsymbol{c}^{target}$ and $\boldsymbol{c}^{PSO}$ is lower than or equal to 2. We transform continuous variables $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{c}^{target}, \boldsymbol{c}^{PSO})$ to the range [-1, 1]. A pictorial illustration of the full data generation pipeline is included in Figure 2B.

Chrome coating replacement design

In the first task, we aim to demonstrate NEUTRON's performance for designing a single target color. Specifically, we choose a highly practical application where we design an environmental-friendly five-layer optical thin film stack with an appearance similar to chrome, which could potentially replace the traditional chrome coatings produced by a highly toxic electroplating process that poses great dangers to both workers and the environment. This is an important task because chrome coatings are widely used in the automotive industry for aesthetic purposes because of their pleasing appearance. Traditionally, the chrome layer is electroplated to the surface of the object by submerging it in a chemical electrolyte solution that contains hexavalent chrome (Cr (6)). However, Cr (VI) is a strong human carcinogen and has been found to greatly increase the risk of lung cancer, nasal cancer, and sinus cancer OSHA (2013). Cr (VI) can also cause severe nasal septum ulcerations and perforations, gastritis, and gastrointestinal ulcers Lieberman et al. (1941); Lindberg and Hedenstierna (1983). Workers can be exposed to Cr (VI) at the workplace from the mist generated during plating. The electroplating process can also lead to air pollution through the emissions of toxic air containing cadmium and cyanide, which could impact the nervous system, hearts, and lungs of millions of people EPA (2005). Thus, it is highly desirable to develop alternative solutions that can produce the chrome appearance but do not require the dangerous chrome plating process.

In the material screening phase, we exclude Cr from the available metals and predicted the likelihood $p_{\Delta E_{00} \leq 2}(\boldsymbol{c}^{target}|\boldsymbol{x})$ for all $10 \times 9 \times 8 \times 10 \times 10 = 720,000$ possible material combinations. Then, we sample initial thickness designs from the M^3DN and fine-tune the designs with PSO for the top 100 combinations with the largest likelihoods. The emulated colors of the five best designs are shown in Figure 3. We also provide the detailed designs and CIEDE2000 of all five designs in Table 1. PSO significantly improved the initial designs predicted by the M^3DN in terms of the CIEDE2000. All five final designs obtained CIEDE2000 lower than two and are highly promising for replacing the traditional Cr plating process through the environmental-friendly thermal evaporation process.

Picture reconstruction

Next, we apply NEUTRON to design optical thin film structures for a large set of color targets. We choose the picture reconstruction task (Gao et al., 2019; Dai et al., 2021), because it has been previously studied in machine learning-based structural color inverse design works. Because most high-resolution pictures have a large number of pixels with unique sRGB values, which could take a long time to design when the PSO fine-tuning step is involved, we quantize the pixel values and find the quantization step size of 10 to reduce the number of unique pixels of a full-color picture by more than ~ 50 X while maintaining the quality of the original pictures with more than 200,000 pixels. Details of picture quantization can be found in the STAR Methods section. The detailed design pipeline is provided in Figure 4A, where the top material combination is selected based on the average $p_{\Delta E_{00} \leq 2}$ computed over the entire picture to be reconstructed. Note that we select the same material combination of all pixels across the entire picture to make it possible for fabrication through grayscale lithography Wang et al. (2018). We plot average $p_{\Delta E_{00} \leq 2}$ distributions for all allowed material combinations in Figure 4B, which show that only a few material combinations can lead to high average success design rates and also demonstrate the importance of the material classification network. Again, in Figure 5C, we show that PSO significantly improved the average design accuracy upon the initial designs based on the trained MDN.

As another demonstration, we applied our method to pictures obtained through neural style transfer Gatys et al. (2016), which is a computer vision method that can transform images to possess an appearance similar to a style source image. As shown in Figure 6, we first apply the neural style transfer method to transfer real photos into stylized pictures, and then we apply NEUTRON to obtain designs for reconstructing the transferred pictures.

Results in Figures 4 and 5 demonstrate that NEUTRON allows accurate picture reconstructions with almost unnoticeable errors. In addition, the entire design process can be accomplished within a few hours for high-resolution pictures (Table 2), which makes NEUTRON a highly practical algorithm for real applications.

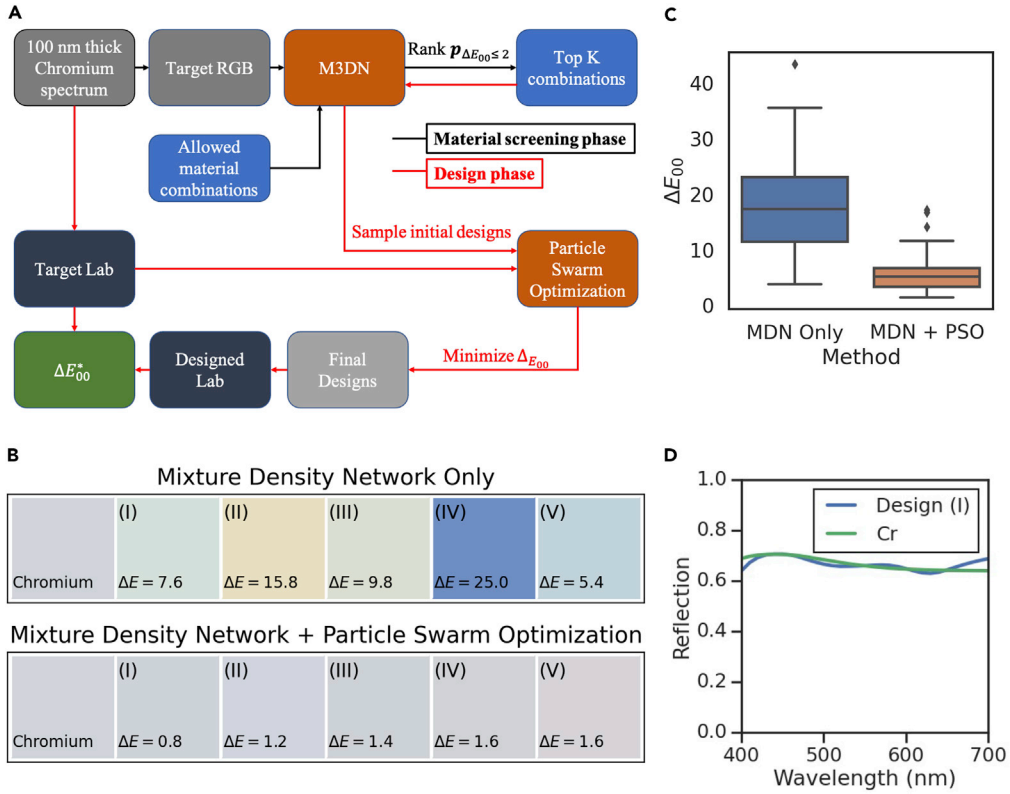

Figure 3. Single color design pipeline and chrome color design results

(A) The top *K* material combinations obtaining the target color is obtained through ranking the predicted probability vector $p_{\Delta E_{00}}$ for all allowed material combinations. Then a set of initial designs are sampled from the M^{3}DN as the initial solutions for particle swarm optimization to fine-tune the thickness of each layer.

(B) Top five designs based on the final CIEDE2000 values. The particle swarm optimization step leads to significant performance improvement compared to the designs sampled directly from the mixture density network.

(C) CIEDE2000 before and after the particle swarm optimization fine-tuning for the top 100 material combinations.

(D) The reflection spectrum of the best design and a 100 nm thick Cr film.

Although PSO also achieves good performance in terms of ΔE_{00}, the PSO results reported in Table 2 are based on the material combination predicted by the classification model of the trained M^{3}DN and does not take the amount of time required for screening the materials when such a model is not available. In

Table 1. Chrome color designs based on M^{3}DN only and the fine-tuned design by PSO (bold)

Design	Antireflective Layer (nm)	Absorber Layer I (nm)	Absorber Layer II (nm)	Dielectric Layer (nm)	Reflective Layer (nm)	ΔE_{00}
I	HfO_2	W	Ge	SiO_2	Zn	
	120/**105**	13/**15**	15/**10**	122/**83**	100/**100**	7.6/**0.8**
II	HfO_2	Zn	Ge	Al_2O_3	Ni	
	12/**7**	14/**14**	8/**13**	160/**98**	100/**100**	15.8/**1.2**
III	SiO_2	Ag	Ge	SiO_2	Au	
	273/**250**	14/**14**	15/**15**	147/**130**	100/**100**	9.8/**1.4**
IV	HfO_2	Zn	Ge	Al_2O_3	Au	
	61/**5**	14/**14**	15/**15**	202/**207**	100/**100**	25.0/**1.6**
V	HfO_2	Zn	Ge	MgF_2	Au	
	97/**93**	15/**15**	12/**11**	249/**218**	100/**100**	5.4/**1.6**

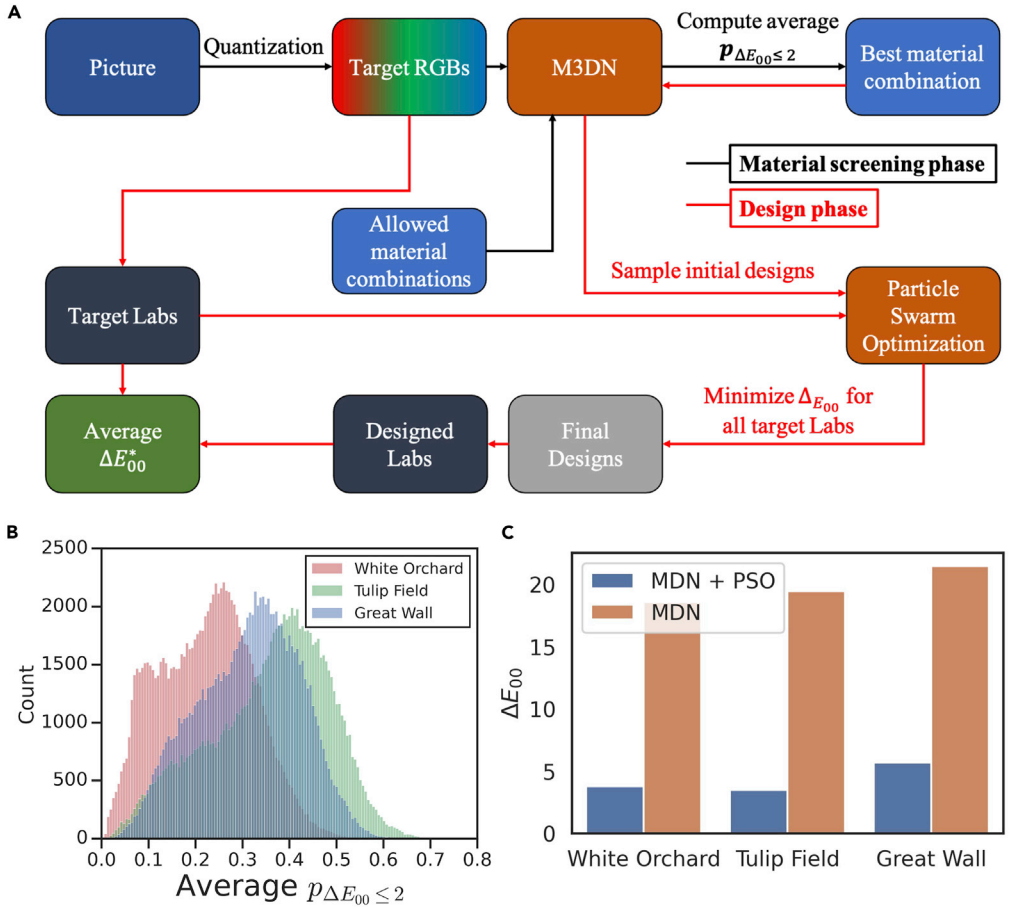

Figure 4. Picture reconstruction design pipeline

(A) Similar to the design pipeline of a single color, the material screening is the first phase of the design process for identifying the best material for each layer to reconstruct the entire picture. Unlike the single-color case, the average probability is computed for all target colors through which the material combination that gives the largest probability is selected.

(B) Average $p_{\Delta E_{00}\leq 2}$ distribution for all 100,000 allowed material combinations.

(C) Average CIEDE2000 computed over entire images before and after particle swarm optimization fine-tuning.

addition, when initializing the PSO with the initial designs predicted by the M^{3}DN, the design accuracy ΔE_{00} can be further improved upon using randomly-initialized PSO.

All results reported in this section are based on optimizing the top four layers' thickness to obtain designs across the entire picture. Though the same material combination is selected for the entire picture with the feasibility of fabrication in mind, it would still be highly challenging in practice. Thus, we also explored the feasibility of picture reconstruction when only varying the bottom dielectric layer while keeping all other layers' thickness fixed. Compared to adjusting all four layers' thickness, adjusting a single layer's thickness leads to a deteriorated average CIEDE2000 of 13.36 (see Figure S1) compared to 3.48 when tuning the top four layers for the Tulip Field painting. When fabrication-friendly designs are preferred over design accuracy, varying a single layer could be applied for the picture reconstruction application.

Understanding algorithmic design choices of NEUTRON

NEUTRON involves multiple hyperparameters that affect its performance (Table 3). We conducted a series of experiments to understand the effects of three important algorithmic design choices for the proposed algorithm NEUTRON, including the number of mixture components, multitask learning, and dataset augmentation with mixture materials.

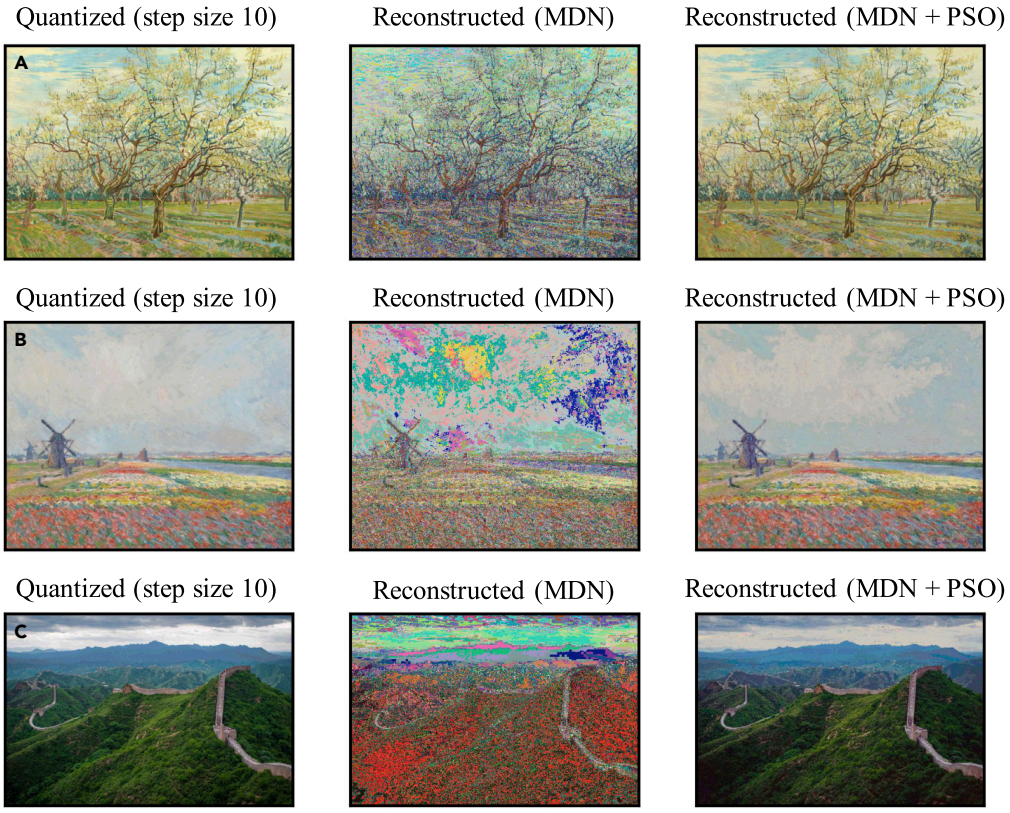

Figure 5. Picture reconstruction results

The first column is the target picture. The second column is the reconstruction based on NEUTRON. The last column is without the particle swarm fine-tuning.

(A) *The White Orchard* by Vincent van Gogh.

(B) *The Tulip Field* by Vincent van Gogh.

(C) A photo of The Great Wall of China taken by photographer Severin Stalder.

(A and B) are reproduced with the permission from the Van Gogh Museum, Amsterdam (Vincent van Gogh Foundation).

(C) is reproduced with the permission from Wikimedia Commons.

Effect of number of mixture components

The optimal number of mixture components m is dependent on the one-to-many mapping between the input $(\boldsymbol{x}, \boldsymbol{c})$ and the potential designs $\{\boldsymbol{y}\}$. A large m is required when many designs could lead to the similar color target. Here, we vary the number of mixtures from 200 to 51,200. As shown in Figure 7, $m = 51,200$ leads to the best performance, whereas reducing m causes the validation CIEDE2000 to increase significantly.

Effect of multitask learning

We vary the weighting parameter α to investigate if positive transfer exists between the material classification task and the design prediction task. Results in Figures 7C and 7D show that $\alpha = 5$ leads to the best CIEDE2000 measured on the validation set while the classification AUC is insensitive to nonzero α values. This indicates that positive knowledge transfer exists between the material classification task and the thickness prediction task. Thus, incorporating the material classification task not only allows users to select the best material combinations for downstream tasks but also improves the thickness inverse design accuracy because of the benefit of multitask learning.

Effect of dataset augmentation

Introducing data points simulated from a mixture of materials increases the variation of the input refractive data $\boldsymbol{x}$, which could allow the trained M^{3}DN to generalize better. To test this hypothesis, we first train M^{3}DN on a dataset with 100,000 data points simulated entirely based on primitive materials. Then, we add

Figure 6. Picture reconstruction results based on neural style transferred photos

(A) Flower.

(B) Street. Both original photos in (A and B) were taken by the authors.

(C) The source style painting *Vase with Gladioli and Chinese Asters* for generating the style transferred photos is reproduced with the permission from the Van Gogh Museum, Amsterdam (Vincent van Gogh Foundation).

additional 50,000 data points that were either simulated from primitive materials or mixture materials and train two networks separately. By comparing the validation ΔE_{00}, we observe that adding 50,000 additional data points from mixture materials leads to greater performance improvement than adding 50,000 data points simulated from primitive materials.

DISCUSSION

In both the chrome color design and the picture reconstruction tasks, NEUTRON achieves exceptional design accuracy efficiently through combining a material classification network, a mixture density network for thickness predictions and the final PSO fine-tuning. Unlike previous methods that either assume fixed materials (Gao et al., 2019; Dai et al., 2021) or search the materials and thickness designs simultaneously Wang et al. (2021), our method split the materials and thickness design into two stages to enable an approach that can search for the best materials and design the thickness accordingly in an efficient manner. The reason why splitting the material and thickness design is helpful can be understood by the fact that the entire design space is much larger than the material design space or the thickness design space alone. In

Table 2. The best material combinations and ΔE_{00} with M^{3}DN only, PSO only, and NEUTRON for reconstructing pictures in Figures 4 and 5

Picture	Material Combination					MDN	PSO	NEUTRON	Time (s)
White Orchard	Al_2O_3	Ge	Al	Al_2O_3	Zn	18.55	5.16	**3.76**	9,917
Tulip Field	Ta_2O_5	Au	Ti	Al_2O_3	Au	19.45	4.34	**3.48**	9,246
Great Wall	HfO_2	Au	Ti	SiO_2	Au	21.46	6.45	**5.69**	6,835
Flower	Ta_2O_5	Au	Ti	Al_2O_3	Au	19.37	5.00	**3.93**	11,249
Street	ZnO	Ti	Au	MgF_2	Au	21.60	4.45	**3.28**	7,349

The results reported for the last two pictures in italic font are for the style transferred photos. Bold values indicate the lowest errors (i.e., best performance among MDN, PSO, and NEUTRON).

Table 3. Hyperparemeter search values for training the M^{3}DN model

Hyperparameters	Values
Learning rate	loguniform(10^{-5},10^{-3})
Hidden dimension of FC layers	[128, 256, 512, 1024]
Number of mixtures	[200, 400, 800, 1600, 3200, 6400, 12800, 25600, 51200]
Weight decay factor	loguniform(10^{-6},10^{-4})
α	[1, 2, 5, 10]

our five-layer optical thin film design task, the material design space size is 9×10^4, and the thickness design space size is 7.3×10^6, which lead to a huge full design space with a size of $(9 \times 7.3) \times 10^{10} = 6.3 \times 10^{11}$ when considering the material design space and the thickness design space simultaneously. With the two-step process, we first narrow down the promising material combinations with the material classification model and only optimize the thickness designs for the selected small set of materials. Thus, we only need to consider a design space that is one or multiple times of the thickness design space, which is a $\sim$ 10,000 X complexity reduction compared to searching both materials and thickness in one step.

We combine M^{3}DN and PSO by initializing the particle positions with designs sampled from the M^{3}DN. This process is easy to implement and highly effective in obtaining optimal designs because of the probabilistic nature of the mixture density network. On both design tasks, we show that PSO significantly improved the initial solutions by the M^{3}DN. This result is not surprising, because PSO involves iterative updates of the design parameters based on the feedback from optical simulations. In addition, the exceptional design accuracy after fine-tuning the initial M^{3}DN designs with PSO indicates that probabilistic

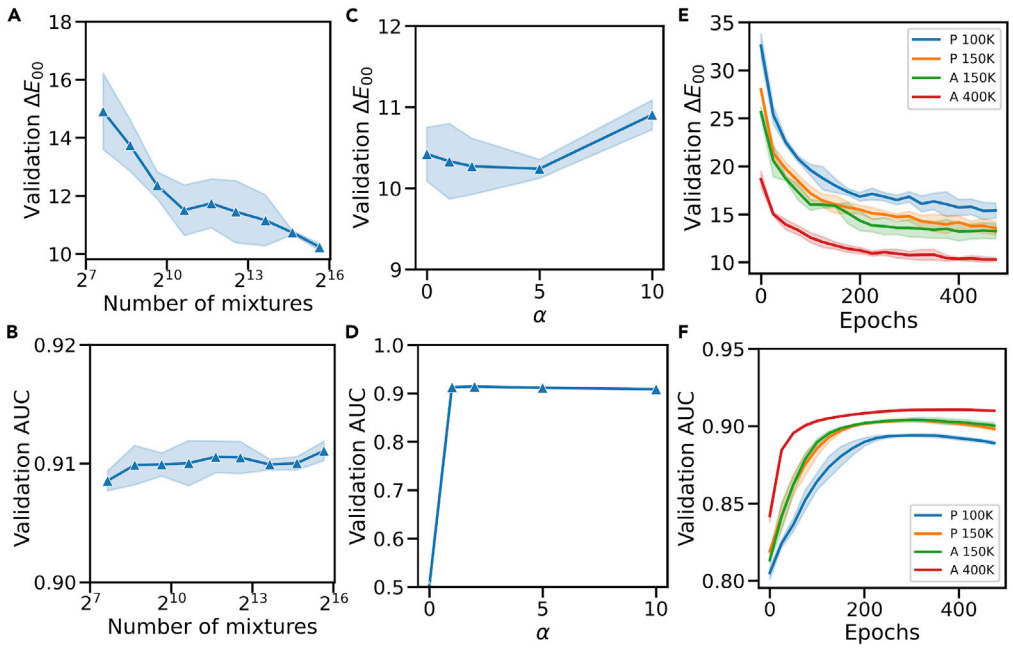

Figure 7. Effects of important algorithmic design choices

All results are based on three runs initialized with different random seeds. The shaded area corresponds to one SD.

(A and B) Varying the number of mixtures. The network with 51,200 mixture components achieves the best validation performance among the nine experiments.

(C and D) Effect of weighting parameter for the classification loss.

(E and F) Effect of dataset augmentation with mixed materials. By adding 50,000 data points from mixed materials (A 150K), the trained network achieves a more significant improvement of CIEDE2000 over the dataset with 100,000 data points simulated from primitive materials (P 100K), compared to adding 50,000 additional data points simulated from only primitive materials (P 150K). Further increasing the dataset size to 400,000 data points with 50% of the data simulated from the mixed materials leads to further performance improvement of both validation CIEDE2000 and the validation AUC.

machine learning models that can output diverse predictions are highly compatible with PSO, which requires an initial population of designs to begin with.

Our experiment on the number of mixtures suggests that a strong degree of one-to-many mapping exists (Figures 7A and 7B) for the problem of inverse designing the structural color with a five-layer optical stack, thus a large number of mixture components of 51,200 is required to obtain an accurate inverse prediction of the designs. Owing to the limit of GPU memory, we cannot further increase the number of mixture components, but we expect the design accuracy to improve as the number is increased further.

Another interesting finding on multitask learning suggests that future optical inverse design models should leverage the task relationships by learning the main and auxiliary tasks simultaneously. As shown in the experiment for varying the classification loss weight α, we observe optimal inverse design performance at $\alpha = 5$. In optical design problems, many tasks are related, such as the electric field distribution prediction, far-field intensity prediction, spectrum prediction, and color prediction, etc. Learning multiple tasks simultaneously could allow researchers to obtain high inverse design accuracy with fewer data points.

Moreover, our results show that data augmentation that increases the refractive index data variation by mixing the primitive materials is an efficient strategy to improve the inverse design accuracy. We believe such data augmentation should be incorporated for training material-aware inverse design models when possible.

NEUTRON combines a classification model, a probabilistic regression model, and particle swarm optimization for fine-tuning the designs predicted by the M^3DN. Compared to using ML or the optimization approach alone, NEUTRON brings the best of both worlds to allow efficient and accurate inverse designs for large-scale design tasks.

In sum, we developed a material-aware inverse design algorithm that combines a multitask mixture density network and particle swarm optimization to select the optimal material combination and optimize structural parameters. On two practical tasks for producing reflective structural color with a five-layer optical stack, NEUTRON demonstrates the proposed algorithm's exceptional design accuracy and efficiency. Though we develop NEUTRON for structural color design based on multilayer optical thin films, it can be adopted for a wide range of other optical design tasks that require optimizing the material selections and structural parameters.

Limitations of the study

The training dataset for NEUTRON consists of PSO-optimized designs. To achieve accurate design predictions, we collected 400K designs as the entire dataset. However, when applying NEUTRON to optical structures that require a longer simulation time, collecting 400K PSO-optimized designs might be infeasible. Thus, future study on improving the sample efficiency of NEUTRON is required.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
  - Lead contact
  - Materials availability
  - Data and code availability
- METHOD DETAILS
  - Material-aware multitask mixture density networks
  - Picture quantization details
  - Model implementation details

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.isci.2022.104339.

ACKNOWLEDGMENTS

The authors acknowledge the support by a data science supplement of NSF SNM 092380 program. H.W. would like to thank Taigao Ma and Prof. Cheng Zhang for helpful discussions.

AUTHOR CONTRIBUTIONS

Conceptualization, H.W. and L.J.G.; Methodology, H.W. and L.J.G.; Investigation, H.W. and L.J.G.; Writing – Original Draft, H.W. and L.J.G.; Writing – Review & Editing, H.W. and L.J.G.; Funding Acquisition, H.W. and L.J.G.; Resources, L.J.G.; Supervision, L.J.G.

DECLARATION OF INTERESTS

The authors declare that this publication is the subject of a related patent application.

Received: December 5, 2021
Revised: April 10, 2022
Accepted: April 26, 2022
Published: May 20, 2022

REFERENCES

Best, J. (2017). Colour Design: Theories and Applications (Woodhead Publishing).

Bishop, C.M. (1994). Mixture density networks. In Neural Computing Research Group Report (Aston University), pp. 1–25.

Byrnes, S.J. (2016). Multilayer optical calculations. Preprint at arXiv. arXiv:1603.02720.

Campbell, S.D., Sell, D., Jenkins, R.P., Whiting, E.B., Fan, J.A., and Werner, D.H. (2019). Review of numerical optimization techniques for meta-device design [Invited]. Opt. Mater. Express *9*, 1842–1863. https://doi.org/10.1364/ome.9.001842.

Caruana, R. (1997). Multitask learning. Machine Learn. *28*, 41–75. https://doi.org/10.1023/a:1007379606734.

Dai, P., Wang, Y., Hu, Y., de Groot, C.H., Muskens, O., Duan, H., and Huang, R. (2021). Accurate inverse design of fabry–perot-cavity-based color filters far beyond srgb via a bidirectional artificial neural network. Photon. Res. *9*, B236–B246. https://doi.org/10.1364/prj.415141.

EPA (2005). Reducing air pollution from: electroplating operations. https://www.epa.gov/sites/default/files/2017-06/documents/electroplating_comm_info.pdf.

Foster, D.H., Amano, K., Nascimento, S.M.C., and Foster, M.J. (2006). Frequency of metamerism in natural scenes. JOSA A *23*, 2359–2372. https://doi.org/10.1364/josaa.23.002359.

Gao, L., Li, X., Liu, D., Wang, L., and Yu, Z. (2019). A bidirectional deep neural network for accurate silicon color design. Adv. Mater. *31*, 1905467. https://doi.org/10.1002/adma.201905467.

Gatys, L.A., Ecker, A.S., and Bethge, M. (2016). Image style transfer using convolutional neural networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2414–2423.

Huang, Z., Liu, X., and Zang, J. (2019). The inverse design of structural color using machine learning. Nanoscale *11*, 21748–21758. https://doi.org/10.1039/c9nr06127d.

Jiang, J., Chen, M., and Fan, J.A. (2020). Deep neural networks for the evaluation and design of photonic devices. Nat. Rev. Mater. *6*, 679–700. https://doi.org/10.1038/s41578-020-00260-1.

Jiang, J., Chen, M., and Fan, J.A. (2021). Deep neural networks for the evaluation and design of photonic devices. Nat. Rev. Mater. *6*, 679–700. https://doi.org/10.1038/s41578-020-00260-1.

Kennedy, J., and Eberhart, R. (1995). Particle swarm optimization. In Proceedings of ICNN'95-international conference on neural networks (IEEE), pp. 1942–1948.

Kim, S.H., Lee, S.Y., Yang, S.M., and Yi, G.R. (2011). Self-assembled colloidal structures for photonics. NPG Asia Mater. *3*, 25–33. https://doi.org/10.1038/asiamat.2010.192.

Li, C., and Lee, G.H. (2019). Generating multiple hypotheses for 3d human pose estimation with mixture density network. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9887–9895.

Lieberman, H., et al. (1941). Chrome ulcerations of the nose and throat. New Engl. J. Med. *225*, 132–133. https://doi.org/10.1056/nejm194107242250402.

Lindberg, E., and Hedenstierna, G. (1983). Chrome plating: symptoms, findings in the upper airways, and effects on lung function. Arch. Environ. Health Int. J. *38*, 367–374. https://doi.org/10.1080/00039896.1983.10545822.

Liu, C., Maier, S.A., and Li, G. (2020). Genetic-algorithm-aided meta-atom multiplication for improved absorption and coloration in nanophotonics. ACS Photon. *7*, 1716–1722. https://doi.org/10.1021/acsphotonics.0c00266.

Liu, D., Tan, Y., Khoram, E., and Yu, Z. (2018a). Training deep neural networks for the inverse design of nanophotonic structures. Acs Photon. *5*, 1365–1369. https://doi.org/10.1021/acsphotonics.7b01377.

Liu, T., VanSaders, B., Keating, J.T., Glotzer, S.C., and Solomon, M.J. (2021). Effect of particles of irregular size on the microstructure and structural color of self-assembled colloidal crystals. Langmuir *37*, 13300–13308. https://doi.org/10.1021/acs.langmuir.1c01898.

Liu, Z., Zhu, D., Rodrigues, S.P., Lee, K.T., and Cai, W. (2018b). Generative model for the inverse design of metasurfaces. Nano Lett. *18*, 6570–6576. https://doi.org/10.1021/acs.nanolett.8b03171.

Ma, T., Tobah, M., Wang, H., and Guo, L.J. (2022). Benchmarking deep learning-based models on nanophotonic inverse design problems. Opto-Electronic Sci. *1*, 210012. https://doi.org/10.29026/oes.2022.210012.

Ma, W., Cheng, F., Xu, Y., Wen, Q., and Liu, Y. (2019). Probabilistic representation and inverse design of metamaterials based on a deep generative model with semi-supervised learning strategy. Adv. Mater. *31*, 1901111. https://doi.org/10.1002/adma.201901111.

Ma, W., Liu, Z., Kudyshev, Z.A., Boltasseva, A., Cai, W., and Liu, Y. (2021). Deep learning for the design of photonic structures. Nat. Photon. *15*, 77–90. https://doi.org/10.1038/s41566-020-0685-y.

Mansecal, T., Mauderer, M., and Parsons, M. (2021). Colour. https://github.com/colour-science/colour.

OSHA (2013). Controlling hexavalent chromium exposures during electroplating. https://www.osha.gov/sites/default/files/publications/OSHA_FS-3648_Electroplating.pdf.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic Differentiation in Pytorch.

Rabady, R.I., and Ababneh, A. (2014). Global optimal design of optical multilayer thin-film filters using particle swarm optimization. Optik *125*, 548–553. https://doi.org/10.1016/j.ijleo.2013.07.028.

Sajedian, I., Badloe, T., and Rho, J. (2019). Optimisation of colour generation from dielectric nanostructures using reinforcement learning. Opt. express *27*, 5874–5883. https://doi.org/10.1364/oe.27.005874.

Schneider, P.I., Garcia Santiago, X., Soltwisch, V., Hammerschmidt, M., Burger, S., and Rockstuhl, C. (2019). Benchmarking five global optimization approaches for nano-optical shape optimization and parameter reconstruction. ACS Photon. *6*, 2726–2733. https://doi.org/10.1021/acsphotonics.9b00706.

Sharma, G., Wu, W., and Dalal, E.N.; Endorsed by Inter-Society Color Council, The Colour Group (Great Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de la Couleur (2005). The CIEDE2000 color-difference formula: implementation notes, supplementary test data, and mathematical observations. Color Res. Appl. *30*, 21–30.

Shokooh-Saremi, M., and Magnusson, R. (2007). Particle swarm optimization and its application to the design of diffraction grating filters. Opt. Lett. *32*, 894–896. https://doi.org/10.1364/ol.32.000894.

So, S., Badloe, T., Noh, J., Bravo-Abad, J., and Rho, J. (2020). Deep learning enabled inverse design in nanophotonics. Nanophotonics *9*, 1041–1057. https://doi.org/10.1515/nanoph-2019-0474.

Sun, S., Zhou, Z., Zhang, C., Gao, Y., Duan, Z., Xiao, S., and Song, Q. (2017). All-dielectric full-color printing with TiO2 metasurfaces. ACS nano *11*, 4445–4452. https://doi.org/10.1021/acsnano.7b00415.

Unni, R., Yao, K., Han, X., Zhou, M., and Zheng, Y. (2021). A mixture-density-based tandem optimization network for on-demand inverse design of thin-film high reflectors. Nanophotonics *10*, 4057–4065. https://doi.org/10.1515/nanoph-2021-0392.

ViewSonic (2021). DeltaE ≤2 color accuracy. https://color.viewsonic.com/explore/content/DeltaE<=2Color-Accuracy_2.html.

Wang, H., Zheng, Z., Ji, C., and Jay Guo, L. (2021). Automated multi-layer optical design via deep reinforcement learning. Machine Learn. Sci. Technol. *2*, 025013. https://doi.org/10.1088/2632-2153/abc327.

Wang, Y., Zheng, M., Ruan, Q., Zhou, Y., Chen, Y., Dai, P., Yang, Z., Lin, Z., Long, Y., Li, Y., et al. (2018). Stepwise-nanocavity-assisted transmissive color filter array microprints. Research *2018*, 1–10.

Whitley, D. (1994). A genetic algorithm tutorial. Stat. Comput. *4*, 65–85. https://doi.org/10.1007/bf00175354.

Wiecha, P.R., Arbouet, A., Girard, C., and Muskens, O.L. (2021). Deep learning in nano-photonics: inverse design and beyond. Photon. Res. *9*, B182–B200. https://doi.org/10.1364/prj.415960.

Yang, C., Hong, L., Shen, W., Zhang, Y., Liu, X., and Zhen, H. (2013). Design of reflective color filters with high angular tolerance by particle swarm optimization method. Opt. express *21*, 9315–9323. https://doi.org/10.1364/oe.21.009315.

Yang, W., Xiao, S., Song, Q., Liu, Y., Wu, Y., Wang, S., Yu, J., Han, J., and Tsai, D.P. (2020). All-dielectric metasurface for high-performance structural color. Nat. Commun. *11*, 1864–1868. https://doi.org/10.1038/s41467-020-15773-0.

Yang, Z., Ji, C., Liu, D., and Guo, L.J. (2019). Enhancing the purity of reflective structural colors with ultrathin bilayer media as effective ideal absorbers. Adv. Opt. Mater. *7*, 1900739. https://doi.org/10.1002/adom.201900739.

Yao, K., Unni, R., and Zheng, Y. (2019). Intelligent nanophotonics: merging photonics and artificial intelligence at the nanoscale. Nanophotonics *8*, 339–366. https://doi.org/10.1515/nanoph-2018-0183.

Yeung, C., Tsai, R., Pham, B., King, B., Kawagoe, Y., Ho, D., Liang, J., Knight, M.W., and Raman, A.P. (2021). Global inverse design across multiple photonic structure classes using generative deep learning. Adv. Opt. Mater. *9*, 2100548. https://doi.org/10.1002/adom.202100548.

Zhang, H., Cisse, M., Dauphin, Y.N., and Lopez-Paz, D. (2018). mixup: beyond empirical risk minimization. In International Conference on Learning Representations. https://openreview.net/forum?id=r1Ddp1-Rb.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Color datasets	This paper	Mendely Data:https://data.mendeley.com/datasets/54d84d88p8/1
Software and algorithms		
NEUTRON code base	This paper	https://github.com/hammer-wang/NEUTRON
TMM Python package	GitHub	https://github.com/sbyrnes321/tmm
Color Python package	GitHub	https://github.com/colour-science/colour

RESOURCE AVAILABILITY

Lead contact

Further information and requests should be directed to either Haozhu Wang (hzwang@umich.edu) or L. Jay Guo (guo@umich.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- All data reported in this paper can be accessed from Mendeley Data Repository at https://data.mendeley.com/datasets/54d84d88p8/1.
- All original code is publicly available at https://github.com/hammer-wang/NEUTRON.
- Any additional information required to implement the methods and analyze the results is available from the lead contact upon request.

METHOD DETAILS

Material-aware multitask mixture density networks

Mixture density networks (MDNs) have been previously applied for inverse problems extensively (Bishop, 1994; Li and Lee, 2019; Unni et al., 2021). Instead of mapping the input to a single output, which is done by deterministic neural network models, an MDN maps the input to the probability density function of a multivariate Gaussian Mixture with m isotropic Gaussian mixture components. Because each Gaussian mixture component can learn a different mean and standard deviation value, MDNs are able to learn a one-to-many mapping, which is critical to solving the often ill-posed inverse design problems with non-unique solutions. The probability density function given by a Gaussian mixture for a pair of refractive data input, thickness design, and corresponding color $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{c})$ can be written as :

$$p(\boldsymbol{y}|\boldsymbol{x},\ \boldsymbol{c}) = \sum_{j=0}^{m-1} \boldsymbol{\pi}(\boldsymbol{x},\boldsymbol{c})_j p_j\left(\boldsymbol{y}|\boldsymbol{x},\ \boldsymbol{c}\right) = \sum_{j=0}^{m-1} \boldsymbol{\pi}_j \frac{1}{(2\pi)^{\frac{D}{2}}\prod_{d=0}^{D-1}\boldsymbol{\sigma}_{j,d}(\boldsymbol{x},\boldsymbol{c})} e^{-\left(\sum_{d=0}^{D-1}\frac{\left(y_d-\boldsymbol{\mu}_{j,\ d}(\boldsymbol{x},\boldsymbol{c})\right)^2}{2\boldsymbol{\sigma}^2_{j,\ d}(\boldsymbol{x},\boldsymbol{c})}\right)},$$

(Equation 2)

where $\boldsymbol{\pi}(\boldsymbol{x},\boldsymbol{c})$, $\boldsymbol{\mu}(\boldsymbol{x},\boldsymbol{c})$, $\boldsymbol{\sigma}(\boldsymbol{x},\boldsymbol{c})$ are the mixing weights, mean, and standard deviation outputted by the MDN with the input refractive data $\boldsymbol{x}$ and the color target $\boldsymbol{c}$. D is the dimension of the design parameter vector $\boldsymbol{y}$, which equals 4 in our design problem. Note that the mixing weights sum up to one (i.e., $\sum_{j=0}^{m-1}\boldsymbol{\pi}_j = 1$) so that the mixed function $p(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{c})$ is a proper probability density function.

In addition to predicting the thickness, we also predict whether a given refractive index data $\boldsymbol{x}$ can lead to accurate designs for a target color $\boldsymbol{c}$ with CIEDE2000 ≤ 2 (i.e., $\Delta E_{00} \leq 2$) after the thickness of each layer has

been optimized. To this end, we include a sub-network to form a material-aware multitask mixture density network (M^3DN) that is composed of a mixture density network and a classification network (see Figure S2A).

To train the M^3DN, we minimize the multitask loss function:

$$\mathcal{L}_{M^3DN} = \sum_{i=0}^{N-1} -\log p\left(\boldsymbol{y}_i|\boldsymbol{x}_i,\ \boldsymbol{c}_i^{PSO}\right) + \alpha \cdot \mathrm{BCE}\left(e_i,\ p_{\Delta E_{00}\le 2}\left(\boldsymbol{x}_i,\ \boldsymbol{c}_i^{target}\right)\right) \quad \text{(Equation 3)}$$

where the first term on the right hand side is the negative log likelihood for training the M^3DN, and the second term is the binary cross entropy loss for training the classifier, i.e.,

$$\begin{aligned}\mathrm{BCE}\left(e_i,\ p_{\Delta E_{00}\le 2}\left(\boldsymbol{x}_i,\ \boldsymbol{c}_i^{target}\right)\right) = &-e_i \cdot \log p_{\Delta E_{00}\le 2}\left(\boldsymbol{x}_i,\ \boldsymbol{c}_i^{target}\right)\\ &-(1-e_i)\cdot\log\left(1-p_{\Delta E_{00}\le 2}\left(\boldsymbol{x}_i,\ \boldsymbol{c}_i^{target}\right)\right).\end{aligned}$$

The hyperparameter α in the M^3DN loss controls the knowledge sharing among the material screening task and the thickness prediction task. Note that, $\boldsymbol{c}_i^{target}$ is used as input for training the M^3DN while $\boldsymbol{c}_i^{PSO}$ is used as input for training the classifier because the thickness $\boldsymbol{y}_i$ corresponds to $\boldsymbol{c}_i^{PSO}$, which could differ slightly from the $\boldsymbol{c}_i^{target}$ in the generated dataset because PSO may only be able to find solutions with a color coordinate slightly different from the target color, i.e., $\boldsymbol{c}_i^{PSO} = \boldsymbol{c}_i^{target} + \boldsymbol{\delta}_i$.

Through an extensive hyperparameter search based on the CIEDE2000 measured on the validation set (details can be found in the STAR Methods section), we found a seven-layer neural network with four separate output heads to give the best inverse design accuracy (see Figure S2A). The number of of Gaussian mixture components is 51,200. We visualize the learned probability density function and the mixing weights (see Figure S3). Note that we use the *softplus* activation function for the $\boldsymbol{\sigma}$ output head to ensure the standard deviation prediction is always greater than 0, and *softmax* activation function is used for the mixing weight $\boldsymbol{\pi}$ output so that all mixing weights sum up to 1. We use *ELU* activations for all other layers except the last layer in the base network and the output layer for μ, where *tanh* is used to ensure their outputs are in the range [-1, 1].

The advantage of training a single multitask network for both classification and thickness prediction tasks is two-folded: 1) the classification module allows users to screen among possible material combinations to select those that could lead to an accurate generation of the target color; 2) improved sample efficiency through sharing knowledge among the related classification and thickness prediction tasks.

The test results for the classification AUC (area under the receiver operating curve, or AUROC) and the inverse design accuracy in terms of CIEDE2000 for both the target color and the PSO obtained color are visualized (see Figures S2B and S2C). The material classification network achieves excellent performance with an AUC of 0.91. Both average test CIEDE2000 values are slightly above ten and can be further improved by finetuning with the PSO.

We train the M^3DN with PyTorch (Paszke et al., 2017) on NVIDIA RTX 3090 GPUs with 24 GB of internal memory. Details on implementation and hyperparameter tuning can be found in the STAR Methods section. Both our code and data are publicly available (links included in the key resource table).

Based on the downstream task, the material screening module can be used in various ways with the predicted probabilities $p_{\Delta E_{00}\le 2}$. On one hand, when the goal is to select the best material corresponding to a single color (e.g. in the first example below searching for chrome color replacement), the top K material combinations with largest $p_{\Delta E_{00}\le 2}$ among all possible combinations can be chosen for further examinations of their performance. On the other hand, when researchers have a set of target colors to produce (e.g. in the second example of reproducing a color picture), the average $p_{\Delta E_{00}\le 2}$ over the set of target colors can be computed for ranking and selecting the material combinations.

Given the selected materials, the thickness designs are predicted based on the color target and the material index data. However, since the output of the mixture density network is probabilistic, it is possible that the thickness output by the trained mixture density network does not correspond to the optimal value. To address this issue, we further fine-tune the designs with PSO. Combining the dataset generation process, M^3DN, and the PSO finetuning process completes the proposed NEUTRON method for structural color design.

In all of our experiments, for each material combination, we randomly sampled 32 thickness designs from the MDN to be used as the initial positions of the particle swarm optimization. Then, we optimize until convergence with a tolerance level of 1×10^{-5}. In each optimization iteration, the designs are evaluated through the transfer matrix method (Byrnes, 2016) and compared with the target spectrum. Compared to randomly initializing the positions of the particles, starting with designs sampled from the M^{3}DN allows PSO to find better solutions. More importantly, conventional PSO requires the materials to be provided by the user while NEUTRON predicts the best materials to be used with the PSO. Since the manual material screening process could be time-consuming, using NEUTRON for structural color design can significantly speed up the entire design process by directly predicting the best materials.

Picture quantization details

High-resolution pictures generally have more than $\sim 10^5$ unique pixel RGB values, which makes it impossible to directly reconstruct when iterative optimizations are is involved. In our work, the particle swarm optimization process for finetuning a single pixel requires $\sim$ 10 seconds, and reconstructing the entire picture takes $\sim 10^6$ seconds. To speed up the reconstruction, we apply a quantization approach to group similar pixels values to the nearest integer values that can be divided by a quantization step size s. We explored quantized step size $s \in [5, 10, 20, 30]$ and found the step size s = 10 led to the best tradeoff between the compression ratio of unique pixels values and the quality of the quantized picture. We compare the pictures before and after quantization with a step size 10 (see Figure S4). With a negligible difference in picture appearance, we achieved more than **40X** compression ratio of unique pixel values for all pictures.

Model implementation details

For training the M^{3}DN, we randomly search hyperparameters from the range listed in Table 3. The model with the best CIEDE2000 on the validation set has a learning rate of 0.000176, hidden dimension 128, the number of mixtures 51,200, weight decay of 0.0000117, and $\alpha = 5$. We train the model on the training set for 500 epochs and store the model with the model checkpoint with the best validation CIEDE2000. Our reported results in the main text are all based on the best model unless stated otherwise.

In particle swarm optimization, each particle's position at time t is updated with a velocity that is determined by both the best solution found by the particle itself and the best solution found by the entire group. In the optical multilayer thin-film design task, the position vector $\boldsymbol{x}$ is the thickness designs for all layers. The velocity $\boldsymbol{v}$ is the update to each layer's thickness between two optimization steps. The velocity is calculated as:

$$\mathbf{v}_i^{t+1} = w \cdot \mathbf{v}_i^t + c_1 \cdot r_1 \cdot \left(\mathbf{x}_{i,best}^t - \mathbf{x}_i^t\right) + c_2 \cdot r_2 \cdot \left(\mathbf{x}_{best}^t - \mathbf{x}_i^t\right), \quad \text{(Equation 4)}$$

where $\mathbf{x}_{i,best}^t$ is the best solution found by the ith particle by time step t, while $\mathbf{x}_{best}^t$ is the best solution found by the entire group of particles by time step t. r_1, r_2 are two random numbers in the range [0, 1]. w is the inertia coefficient that controls how much the previous velocity is maintained, c_1 is the cognitive coefficient that controls how much the best solution found by the particle itself contributes to the velocity, and c_2 is the social coefficient that determines how the best solution found by the entire group alters the velocity. In our implementation, we set the $w = 0.9, c_1 = 0.5, c_2 = 0.3$ based on the results from a manual search. After the velocity is obtained at the $(t+1)$ th step, we update the particle's position by:

$$\boldsymbol{x}_i^{t+1} = \boldsymbol{x}_i^t + \boldsymbol{v}_i^{t+1}. \quad \text{(Equation 5)}$$