

Computational Design of Nanostructural Color for Additive Manufacturing – Supplemental Material

THOMAS AUZINGER, Institute of Science and Technology Austria, Austria

WOLFGANG HEIDRICH, King Abdullah University of Science and Technology, Kingdom of Saudi Arabia

BERND BICKEL, Institute of Science and Technology Austria, Austria

ACM Reference Format:

Thomas Auzinger, Wolfgang Heidrich, and Bernd Bickel. 2018. Computational Design of Nanostructural Color for Additive Manufacturing – Supplemental Material. *ACM Trans. Graph.* 37, 4, Article 159 (August 2018), 9 pages. <https://doi.org/10.1145/3197517.3201376>

1	Diffraction Grating Theory	1
1.1	Diffraction Efficiency	1
1.2	Reciprocity	2
2	Colorization Objective	3
2.1	Color Difference	3
2.2	Difference Aggregation	3
2.3	Wavelength Discretization	3
2.4	Viewing Direction Discretization	3
3	Source Definition	3
4	Phase Alignment	4
5	Intensity Interpolation	4
6	Monitor Plane Intensity	5
7	Result Analysis	6
8	Result Statistics	6
9	Measurement Setups	6
	References	7

1 DIFFRACTION GRATING THEORY

We design nanostructures as tilings of fundamental unit cells. Assuming sufficiently many repetitions, the optical behavior of such structures is well approximated by an infinite tiling of the unit cell. A fundamental result from electromagnetic theory – commonly referred to as the *Bloch Theorem* [Bloch 1929] – states that the fields around such a periodic structure exhibit the *same* periodicity [Eckart 1933]. As an immediate consequence, the possible propagation directions of certain light waves are restricted to a discrete set. This effect explains the behavior of our design in the terms of *diffraction gratings*, which split incoming light into various beams traveling in different directions. These directions can be fully characterized by the lateral size L of the unit cell and do *not* depend on the actual structures; as illustrated in [Figure 1](#), they are given by the *grating equation*

$$\sin(\theta_m) - \sin(\theta_i) = \frac{m\lambda}{L} \quad (1)$$

where for a given diffraction order $m \in \{\dots, -2, -1, 0, 1, 2, \dots\}$ and incident light at angle θ_i , the m -th order beam direction is given by the angle θ_m as shown in [Figure 1 \(left\)](#). Note that this is a general result and we do *not* employ scalar diffraction theory in this work. Diffraction gratings act as a *dispersive* optical element since the diffracted beam angles vary with the wavelength λ of the incident light. It is this property, which gives rise to the rainbow colors that can be observed when a beam of white light is diffracted; see [Figure 1 \(center\)](#) for an illustration.

Depending on θ_i , λ , and – most importantly – L , only a subset of orders m can satisfy [Equation \(1\)](#), i.e., we have

$$\left\lceil -\frac{L}{\lambda}(1 + \sin(\theta_i)) \right\rceil \leq m \leq \left\lfloor \frac{L}{\lambda}(1 - \sin(\theta_i)) \right\rfloor \quad (2)$$

with the ceiling and floor functions $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$. Thus, the number of possible diffracted directions increases with size L of the unit cell and the frequency $f = c/\lambda$ of the incident light.

1.1 Diffraction Efficiency

While the simple relationship in [Equation \(1\)](#) determines the angles of the diffracted orders, the actual amount of light that is transmitted into the various orders depends on the actual unit cell structures and requires a much more elaborate treatment. It is usually measured as the ratio of the incident spectral power $P_i(\theta_i, \lambda)$ and the spectral power $P(\theta_m, \lambda)$ of light diffracted into order m and it is commonly denoted as (*absolute*) *diffraction efficiency* η , i.e.,

$$\eta(\theta_i, \theta_m, \lambda) = \frac{P(\theta_m, \lambda)}{P_i(\theta_i, \lambda)} \quad (3)$$

Authors' addresses: Thomas Auzinger, Institute of Science and Technology Austria, Am Campus 1, 3400, Klosterneuburg, Austria, thomas.auzinger@ist.ac.at; Wolfgang Heidrich, King Abdullah University of Science and Technology, Al Khwarizmi Building 1, 23955-6900, Thuwal, Kingdom of Saudi Arabia, wolfgang.heidrich@kaust.edu.sa; Bernd Bickel, Institute of Science and Technology Austria, Am Campus 1, 3400, Klosterneuburg, Austria, bernd.bickel@ist.ac.at.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3197517.3201376>.




Fig. 1. Diffraction grating theory. (Left) Depending on the incident angle θ_i and the ratio of grating period L and wavelength λ , incident light is split into a varying number of diffraction orders with different directions (e.g., θ_1 for order 1). (Center) The amount and direction of diffraction orders vary with the incident wavelength and create the typical rainbow patterns; here, arrow lengths indicate the orders and carry no physical meaning. (Right) Due to the reciprocity property of diffraction gratings, the grating efficiencies (resp. transmittance ratios) for the standard setting $\eta(\theta_i, \theta_3, \lambda)$ and the reciprocal setting $\eta(\theta_3, \theta_i, \lambda)$ are identical.

with $0 \leq \eta \leq 1$. A diffraction grating is fully specified by the knowledge of η . Any incident spectral light distribution $P_i(\theta_i, \lambda)$ can be transformed to an outgoing spectral light distribution $P(\theta, \lambda)$ via

$$\begin{aligned} P(\theta, \lambda, P_i) &= \int \eta(\theta_i, \theta, \lambda) P_i(\theta_i, \lambda) d\theta_i \\ &= \sum_{m \in M(\theta, \lambda)} \eta(\theta_i(m, \theta, \lambda), \theta, \lambda) P_i(\theta_i(m, \theta, \lambda), \lambda) \end{aligned} \quad (4)$$

where the summation is over those orders $M(\theta, \lambda)$ such that the m -th order exists and coincides with direction θ . $\theta_i(m, \theta, \lambda)$ denotes the corresponding incident direction according to Equation (1). Note that $P(\theta, \lambda, P_i)$ describes the full spectral composition of diffracted light into direction θ . The corresponding observable color $c(\theta, P_i)$ can be readily computed with the help of standard color matching functions $\mathbf{m}(\lambda)$, i.e.,

$$c(\theta, P_i) = \int_{\lambda} \mathbf{m}(\lambda) P(\theta, \lambda, P_i) d\lambda \quad (5)$$

In this and the following sections we assume unpolarized light, i.e., all quantities are averaged over all polarizations; details on the treatment of the polarized components can be found in Section 5 of the main document.

1.2 Reciprocity

For a given structure, we aim to compute the observable colors for a set of outgoing directions. For this, we approximate the diffraction efficiency function η by sampling both the directions and wavelengths. To achieve this, one chooses a set \mathcal{O} of outgoing directions and computes $\eta(\theta_i, \theta, \lambda)$ for all $\theta \in \mathcal{O}$ and all compatible incident angles θ_i according to Equation (1). For each of these computations,

the interaction between the incident light and the structure has to be simulated. As described in Section 5 of the main document, the 100s of nanometer-scale features of our structures require wave and polarization effects to be taken into account and we employ full electromagnetic simulation. In such a setting, different incident light directions have to be simulated separately. However, the generally large number of compatible incident angles θ_i requires an impractically large amount of simulations.

By exploiting a fundamental property of diffraction gratings, this issue can be mitigated by exchanging the role of incident and outgoing light. For the materials and light conditions of our setting, diffraction gratings exhibit a fundamental property – commonly referred to as *reciprocity* –, which states that the diffraction efficiency η is symmetric in incident and outgoing angle. As shown in Figure 1 (right), this can be visualized by comparing two settings: *i*) the standard setting with incident light direction θ_i and outgoing light direction θ_m of the m -th diffraction order; and *ii*) the reciprocal setting with incident light direction θ_m instead of θ_i . In the latter case, there exists a $-m$ -th diffraction order with direction θ_i and reciprocity states that the efficiencies of both settings are *identical*.

Reciprocity allows us to directly compute efficiencies $\eta(\theta_i, \theta, \lambda)$ for outgoing direction θ and all compatible incident angles θ_i by simulating light incident from direction θ and with power $P_i(\theta, \lambda)$ and by measuring the powers $P(\theta_i, \lambda)$ for all θ_i . This leads not only to a significant reduction of required simulation, but also avoids problematic incident light directions for our simulation method (see Section 5.2.2 of the main document).

2 COLORIZATION OBJECTIVE

Since we aim to realize a desired surface colorization, our design objective is stated in terms of color differences. In a continuous setting, a desired colorization can be stated as a target color function $\check{c}(\theta)$ given for all possible viewing directions θ . Formulating this design task as an optimization problem, we aim to find a nanostructure, whose simulated colors $\mathbf{c}(\theta)$ reproduce the target colors $\check{c}(\theta)$ as close as possible. In other words, we want to minimize the total color difference $d(\mathbf{c}(\theta), \check{c}(\theta))$ over all viewing directions. This setting poses four immediate design choices: *i)* which distance measure to choose to quantify the difference between colors; *ii)* how to aggregate these individual differences over all viewing directions; *iii)* how to discretize the wavelengths; and *iv)* how to discretize the viewing directions to enable a numerical treatment. In the following, we describe our design rationale for each of these choices.

2.1 Color Difference

Since we perform electromagnetic simulations, we obtain spectral power distributions as results. The color of such a spectral power distribution is determined by its *tristimulus values* X , Y , and Z , which mimic the human eye's response to light of such spectral distribution. For the three types of cone cell in the human eye, their corresponding spectral sensitivity functions $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$ quantify the cells' responsivity for each wavelength of the observed light. Thus, the tristimulus values X , Y , Z are obtained as

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \frac{K}{\int_{\lambda} I(\lambda) \bar{y}(\lambda) d\lambda} \int_{\lambda} \tau(\lambda) I(\lambda) \begin{pmatrix} \bar{x}(\lambda) \\ \bar{y}(\lambda) \\ \bar{z}(\lambda) \end{pmatrix} d\lambda, \quad (6)$$

where the $\tau(\lambda)$ denotes the spectral transmittance of our nanostructure, $I(\lambda)$ the spectral power distribution of the illuminant, and K a scaling factor. Considering the tristimulus values as components of a three-dimensional vector gives rise to the CIEXYZ color space (as long as one of the *CIE standard observers* is used as spectral sensitivity functions – in which case they are also referred to as *color matching functions*).

Through transformation functions, it is possible to transform the CIEXYZ values into coordinates of one of the many existing color spaces. Generally, these color spaces have specialized purposes, such as the sRGB color space for electronic displays, the CMYK color space for printing, or the perceptually uniform color spaces CIELAB and CIELUV. In the latter, Euclidean distances between different color values correlate well with the perceived color differences of a human observer. More advanced perceptually correct color differences were developed over the decades and the latest standard is given by the CIELAB Delta E* variant CIEDE2000.

In the context of our method, all these color spaces and color difference functions can be used. The only requirement is the existence of a derivative of all associated transforms and difference computation. This is fulfilled by all aforementioned color spaces and differences function as they are given as analytic expressions. In our reported results, we employed simple Euclidean distance in the sRGB color space out of convenience. With increased fabrication precision, the benefit of using more advanced variants will become perceivable.

2.2 Difference Aggregation

Given the color differences $d(\theta)$ for each viewing direction θ , an aggregated value φ is required to judge the overall quality of the associated nanostructure (in term of the colorization objective). We use the 1-norm for this purpose (i.e., $\varphi = \int_{\theta} d(\theta) d\theta$) but any other suitable L^p norm can be used.

2.3 Wavelength Discretization

In a numeric setting, the integrals in Equation (6) have to be approximated by finite sums. Thus, we need to find a suitable sets of wavelengths for which the spectral transmissions of the nanostructures are simulated. Due to the low-index dielectric materials used for our structures and due to the absence of photonic crystal behavior, we do not observe strong resonances across the visible spectrum. Accordingly, the spectral power distributions are sufficiently smooth to be well approximated by a few equidistant samples across the spectrum. In all our results, we used a spectral sample spacing of 20 nm.

2.4 Viewing Direction Discretization

In order to treat the aforementioned color differences in a computational setting, the range of viewing angles has to be discretized. For the range of viewing directions, for which target colors were defined, we used equidistantly placed samples. Thus, the difference aggregation is approximated by a finite sum, i.e.,

$$\varphi = \sum_{\theta_k} \|\check{c}(\theta_k) - \mathbf{c}(\theta_k)\| \approx \int_{\theta} d(\theta) d\theta.$$

A detailed formulation of the objective function is given as Equation (4) of the main document. The number of angular samples is given by N_{target} in Table 4.

For visualization purposes – such as the results in Figure 7 in the main document –, we might also be interested in obtaining colors for viewing directions that do not coincide with on the target directions θ_k . For this purpose, we do not interpolate the colors but the spectral power distributions instead. This is necessary due to specific properties of our simulation setting, where different wavelengths travel in different directions. See Sections 3 and 5 for details.

3 SOURCE DEFINITION

Light is injected into our simulation along planar surfaces above and below the unit cell (see Ω_S in Figure 4 of the main document). Conventionally, short light pulses are utilized to simulate the structure's response to a wide range of frequencies. In this setting, the recorded signals are transformed into the frequency domain via some variant of a discrete Fourier transform (DFT). As explained in Section 5.2.4 of the main document, this approach is not suitable for our setting due to memory limitations and/or spectral leakage issues. Since we need only a small number of wavelengths for the final color transform – as described in Section 2.3 –, we use a set of superimposed *continuous wave* sources, which excite each of these frequencies separately. During simulation, these superposed fields are propagated simultaneously.




Fig. 2. Numerical grid dispersion. The Bloch boundary conditions allow the choice of a propagation direction $\theta_{i,\text{ref}}$ for *only a single* wavelength (here 550 nm). The finite-difference time-domain (FDTD) method then forces light waves with a different wavelength to propagate along different directions $\theta_i(\lambda)$.

For each frequency, we inject light as a linearly polarized plane wave source at an incident angle θ_i . This is achieved by defining a time-varying source current \mathbf{J}_E or \mathbf{J}_H – depending on the polarization – that mimics the effect of an oblique plane wave passing through the planar source area at $y = y_S$, i.e., $\mathbf{J}_E(x, y_S, z, t) = \mathbf{J}_H(x, y_S, z, t) = (0, 0, e^{2\pi i k_x x} e^{-i\omega t})$. The corresponding source intensities are $\eta/8$ and $1/(8\eta)$ for TE and TM polarization, due to $|\mathbf{E}|^2 = (\eta/2)^2 |\mathbf{J}_E|^2$ and $|\mathbf{H}|^2 = 1/(2\eta)^2 |\mathbf{J}_H|^2$. These intensities supply the denominator for the transmittance definition in Equation 3 of the main document.

Since we assume an infinite tiling of unit cells across the substrate, we utilize lateral periodic boundary conditions during the simulation. For plane waves with oblique incidence, there exists a phase shift across the unit cell when looking at the wave’s phase on the left and the right boundary of the simulated unit cell; for a later unit cell length L the phase shift is given by $e^{2\pi i k_x L}$ where k_x is the x -component of the wave’s wave vector \mathbf{k} .

A consequence of using such a (periodic) Bloch boundary condition in an finite-difference time-domain (FDTD) simulation is that all k_x are effectively set to the same value. This can be considered a limitation of the FDTD method and it leads to an undesired deflection of waves. Although k_x is held fixed for waves of all frequencies, the length of their wave vectors \mathbf{k} is still determined by their individual frequencies f , i.e., $\|\mathbf{k}\| = k = 2\pi f/c$, where k denotes the wavenumber. Thus, the y -component k_y of the wavevectors have to adapt according to $k = \|(k_x, k_y)\|$.

The result is an *angular discrepancy* between plane waves of different frequency that are simulated together. We are free to choose a *reference* plane wave source with frequency f_{ref} that propagates in its intended direction $\theta_{i,\text{ref}}$, i.e.,

$$k_x = \frac{2\pi f_{\text{ref}}}{c} \sin(\theta_{i,\text{ref}}).$$

All other plane waves sources with frequencies f are forced to emit their light in the direction

$$\theta_i(f) = \arcsin\left(\frac{f_{\text{ref}}}{f} \sin(\theta_{i,\text{ref}})\right) \quad (7)$$

as shown in Figure 2. This so-called *numeric grid dispersion* is a fundamental limitation of the FDTD method and far-field intensities have to be angle-corrected by a suitable interpolation.

The broadband fixed-angle source technique (BFAST) would allow multiple frequencies to propagate in the same direction by reformulating the FDTD algorithm itself [Liang et al. 2014]; however we forgo its use because it becomes unstable in certain settings and it would further complicate the gradient computation.

Another downside of FDTD stems from the incompatibility of periodic boundary conditions and the absorbing perfectly matched layer (PML). At grazing incidences, the PMLs reflect light back into the unit cell which corrupts the intensity results. Our strategy of using a reciprocal simulation setting becomes vital in this aspect: as explained later in Section 7 of the main document, we limit our colorization objective – and thus our reciprocal incidences – to viewing directions in an unproblematic range of $\approx \pm 45^\circ$ around the surface normal. This would not be possible in a non-reciprocal setting with wide-angle illumination.

Further care has to be taken that PML reflections are not unintentionally caused by the aforementioned angle discrepancy. If, for example, blue light at 450 nm is injected at 40° , red light at 700 nm would propagate in the same simulation at a grazing angle of 89° ! This poses an essential problem for our signal reconstruction approach as described in Appendix A: the contribution due to the PML reflection of the red light corrupt recorded fields, which causes the reconstruction to fail. Thus, every light component that would cause such reflections must not be injected into the simulation in the first place. Therefore, selective *frequency culling* has to be performed for each simulation to ensure that only frequencies of interest are injected. This concept is illustrated in Figure 3 where no samples are placed in the red exclusion regions.

4 PHASE ALIGNMENT

Regarding the computation of $\mathbf{A}^\top \mathbf{E}$, we have to take care of a subtle issue regarding the relative phases of both fields. While the global phase is irrelevant for computing φ from the near-field E_z or H_z , this is not the case for the gradient. If we record the fields after a total simulation time T_{total} , each frequency component $\mathbf{A}(\mathbf{x}, f)$ of the adjoint field will be recorded for a current source $e^{-i\omega T_{\text{total}}} \mathbf{J}_A(\mathbf{x}, f)$ instead of $\mathbf{J}_A(\mathbf{x}, f)$. This phase difference corrupts the gradient and as a remedy, we simply correct the individual frequency components of current sources with a phase correction term $e^{i\omega T_{\text{total}}}$. Out of convenience, we apply this correction to all sources – even in the forward setting where it does not matter.

5 INTENSITY INTERPOLATION

In order to compute the objective function φ in the reciprocal setting, we require the knowledge of the outgoing spectral light distributions $L(\vec{\theta}_k, \lambda)$ for a set of outgoing directions $\vec{\theta}_k$ along which we compute the colors (see Equation 4 of the main document). However, due to the spectral angle discrepancy as given in Equation (7), it is




Fig. 3. Angular interpolation. We interpolate the transmittances at the target direction $\check{\theta}_k$ (black crosses) from the spectral samples (colored dots). The propagation directions $\theta_i(f)$ for individual simulations are depicted as gray lines and we maximize their reuse for all spectral components. Note that our interpolation avoids PML reflections (red regions). See Section 5 for details.

not possible to propagate all light frequencies in the same direction when simulating them simultaneously in one FDTD run. This leaves us with two options:

- (1) For each outgoing direction $\check{\theta}_k$, we simulate all frequencies separately. Since each simulation would only propagate monochromatic light, we can choose its propagation direction exactly. Note that this approach does not require the aforementioned frequency culling. The amount of required simulations is the product of the number of outgoing directions and the number of simulated frequencies.
- (2) We select a set of alternative outgoing directions and simulate for each direction all those frequencies that do not need to be culled. Afterward, we interpolate the results of these simulations to obtain the desired outgoing light distributions.

Option 1 is impractical due to the vast number of required simulations – up to several hundred in our setting. Thus, we choose the alternative option and employ a simple scheme to reduce the amount of required simulations as shown in Figure 3:

- For the largest wavelength λ_{\max} , we choose outgoing directions $\vartheta_{\kappa}(\lambda_{\max})$ that lie halfway between the target directions $\check{\theta}_k$. After simulation, the resulting outgoing spectral light distributions $L(\vartheta_{\kappa}(\lambda_{\max}), \lambda_{\max})$ can be used to obtain $L(\check{\theta}_k, \lambda_{\max})$ through linear interpolation from its nearest neighbors. This is visualized at the $\lambda = \lambda_{\max} = 700$ nm line of Figure 3, where the values at the target directions $\check{\theta}_k$ (i.e., the black crosses) can be interpolated from the simulated values at the simulated directions $\vartheta_{\kappa}(\lambda_{\max})$ (i.e., the red dots).
- For the next smaller wavelength λ , we reuse all previous ϑ_{κ} . This does not require frequency culling since the actual propagation angles decrease together with the wavelength. Additional outgoing directions are added if the original target directions $\check{\theta}_k$ are not covered anymore (e.g., the left- and right-most additional directions on the $\lambda = 500$ nm line of Figure 3).

When simulating one of these additional directions, the previous wavelengths are culled, since they are not needed for the interpolation and they might be located in the exclusion region where they cause PML reflections.

- This process is iterated until all spectral samples are covered and all required outgoing spectral light distributions $L(\check{\theta}_k, \lambda)$ are interpolated from the corresponding $L(\vartheta_{\kappa}, \lambda)$. Note that all ϑ_{κ} and the corresponding interpolation weights are independent of the structure and can be computed in a preprocess.

In the context of Figure 3, each simulation corresponds to one outgoing direction (gray line) and propagates all unculted spectral components (colored dots) along it. Note that due to the varying number of frequencies for different simulation, an optimal reconstruction has to be precomputed for each frequency count (see Appendix A of the main document).

6 MONITOR PLANE INTENSITY

The total intensity $I_{\text{inf, total}}(f)$ of the near field as mentioned in Section 5.1.2 of the main document is obtained by

$$I_{\text{inf, total}}(f) = \int \langle S_y \rangle(x, f) dx = \begin{cases} \frac{1}{2\omega\mu} \Re \left(\int E_z(x, f) \mathcal{F} \left(k_y \mathcal{F}^{-1} (E_z(\cdot, f)) \left(\frac{k_x}{2\pi} \right) \right) (x) dx \right) \\ \frac{1}{2\omega\epsilon} \Re \left(\int H_z(x, f) \mathcal{F} \left(k_y \mathcal{F}^{-1} (H_z(\cdot, f)) \left(\frac{k_x}{2\pi} \right) \right) (x) dx \right) \end{cases} \quad (8)$$

where the integration is taken over the near-field plane in a unit cell; the two cases are for TE and TM polarization.

As explained in the following, this is a consequence of the Maxwell equation when stated in the spatial frequency domain. Using the Fourier transform on the spatial component of the phasor $\mathbf{E}(\mathbf{x}, f)$, we obtain the spatial-frequency phasor $\hat{\mathbf{E}}(\mathbf{k}/2\pi, f)$ as

$$\hat{\mathbf{E}} \left(\frac{\mathbf{k}}{2\pi}, f \right) = \int_{\mathbb{R}^2} \mathbf{E}(\mathbf{x}, f) e^{i\mathbf{k}\cdot\mathbf{x}} dx = \mathcal{F}^{-1}(\mathbf{E}(\cdot, f)) \left(\frac{\mathbf{k}}{2\pi} \right)$$

for which the Maxwell equations are given as

$$\begin{aligned} \mathbf{k} \times \hat{\mathbf{E}} \left(\frac{\mathbf{k}}{2\pi}, f \right) &= \omega\mu_0 \hat{\mathbf{H}} \left(\frac{\mathbf{k}}{2\pi}, f \right) - i\hat{\mathbf{J}}_H \left(\frac{\mathbf{k}}{2\pi}, f \right) \\ \mathbf{k} \times \hat{\mathbf{H}} \left(\frac{\mathbf{k}}{2\pi}, f \right) &= -\omega\hat{\epsilon} \left(\frac{\mathbf{k}}{2\pi} \right) * \hat{\mathbf{E}} \left(\frac{\mathbf{k}}{2\pi}, f \right) + i\hat{\mathbf{J}}_E \left(\frac{\mathbf{k}}{2\pi}, f \right) \end{aligned} \quad (9)$$

where $*$ denotes convolution.

In source-free regions with constant permittivity – such as the far-field above and below our structures –, it follows from Equation (9) that

$$\begin{aligned} \hat{\mathbf{H}} \left(\frac{\mathbf{k}}{2\pi}, f \right) &= \frac{1}{\omega\mu_0} \mathbf{k} \times \hat{\mathbf{E}} \left(\frac{\mathbf{k}}{2\pi}, f \right) \\ \hat{\mathbf{E}} \left(\frac{\mathbf{k}}{2\pi}, f \right) &= \frac{1}{\omega\epsilon} \mathbf{k} \times \hat{\mathbf{H}} \left(\frac{\mathbf{k}}{2\pi}, f \right) \end{aligned}$$

with the x -component given as

$$\begin{aligned} \hat{H}_x \left(\frac{\mathbf{k}}{2\pi}, f \right) &= \frac{1}{\omega\mu_0} k_y \hat{E}_z \left(\frac{\mathbf{k}}{2\pi}, f \right) \\ \hat{E}_x \left(\frac{\mathbf{k}}{2\pi}, f \right) &= \frac{1}{\omega\epsilon} k_y \hat{H}_z \left(\frac{\mathbf{k}}{2\pi}, f \right). \end{aligned} \quad (10)$$

With the integrated y -component of the Poynting vector

$$\int \langle S_y \rangle(x, f) dx = \frac{1}{2} \Re \left(\int E_z(\mathbf{x}) \overline{H_x(\mathbf{x})} - E_x(\mathbf{x}) \overline{H_z(\mathbf{x})} dx \right),$$

Equation (8) is a direct consequence of the identities of Equation (10).

7 RESULT ANALYSIS

In Tables 1 and 2, we provide quantitative results for all optimized designs in the main document. When optimizing for a colorization along the normal viewing direction, we give in Table 2 the color values for the desired target colorization and the resulting colorization using our method. Both colorizations are given in sRGB and CIEXYZ color space and their difference is given in terms of the CIEDE2000 standard [Sharma et al. 2004].

For designs, which were optimized for wide-angle viewing directions, we sampled the color differences on a dense set of viewing directions. For each viewing direction, where a target colorization was defined, we evaluated its CIEDE2000 color difference with the colorization of our results. As shown in Figure 7 of the main document, the resulting colorizations vary with different viewing directions. In Table 1, we characterize the distributions of these differences by providing the mean color difference and standard deviation for each distribution together with the minimal and maximal color difference that were encountered.

In Table 3, the colorizations of the fabricated samples are analyzed. For this, we utilized the sample images obtained with macrophotography (see Figures 9 and 10 of the main document). Since each sample consists of a multitude of small $50 \mu\text{m} \times 50 \mu\text{m}$ patches, three quantities were computed:

Per-patch color differences. To quantify the stability of the colorization inside each patch, we the average color for each patch. This allows the computation of per-patch CIEDE2000 color difference between all the pixels of the patch and its

average color. These differences were aggregated into the average color differences per patch. Table 3 gives the mean color differences across all patches and their standard deviations.

Color difference between patches. While the previous quantity computes the color differences with respect to a per-patch average color, we also computed the overall color differences with respect to the average color of the *whole sample*. This allows us to judge the colorization stability across all patches.

Comparison with the computed result. As a final step, we repeat the previous computation but instead of the average sample color, we compare the individual pixels with the predicted color of our design method.

From Table 3, we see that the per-patch colorization is very stable, which is characterized by the small color differences inside each patch. The color differences across patches are significantly higher, however. This is caused by the interface finder problem described in Section 8 of the main document. While the yellow colorization of Figure 9 could be reproduced well, we encountered a systematic error with the blue colorization objective of Figure 10. This is characterized by a large average color difference when comparing the blue color of the fabricated sample with the cyan color as predicted by our design method. At the same time, the standard deviation of the color differences is small across the sample, which confirms the visual impression of having a uniform colorization.

8 RESULT STATISTICS

In Table 4, we give a quantitative structure description for each result in the main document as well as various characteristics of the optimization process. Note that reinforcement leads to an increased focus count, which is not shown here. The number of L-BFGS-B iteration is given by N_{iter} . N_{sim} denotes the number of simulation that is required for the evaluation of the objective function. The additional computation of its gradient requires the same amount of simulation. The number of target directions (and corresponding target colors) is given by N_{target} . Note that **Runtime** and N_{iter} are not perfectly correlated since individual L-BFGS-B iterations evaluate the objective function and its gradient a varying amount of times.

9 MEASUREMENT SETUPS

The feature sizes of our samples span several orders of magnitude as they cover an area of square millimeters while the smallest features have sub-micrometer size. Thus, we employ a range of measurement setups.

Electron Microscopy. We used a *Merlin VP compact* by Zeiss at 2 kV for detailed structure characterizations.

Brightfield Transmission Microscopy. To investigate the overall structure of $50 \mu\text{m} \times 50 \mu\text{m}$ patches we used a *Nikon Eclipse* inverted widefield microscope with 10x, 20x, and 60x objectives. We ensured correct Köhler illumination for all measurements.

Macrophotography. To test the real-world behavior of our samples, we used a uniform light-emitting diode (LED) backlight manufactured by *Metaphase* and a *Nikon 5100D* digital single-lens reflex camera (DSLR) with a *Tamaron 90mm F/2.8 1:1* macro objective. As

Table 1. Result analysis. For each objective, the color differences between the optimized result and the desired target colorization are characterized by their mean and standard deviation. Furthermore, the minimal and maximal color difference for each objective is shown in the last two columns.

Figure	Result	CIEDE2000		
		Mean \pm deviation	Min	Max
Figure 7	red	0.299 \pm 0.028	0.258	0.357
	green	0.374 \pm 0.061	0.271	0.471
	blue	0.337 \pm 0.047	0.263	0.433
	yellow	0.288 \pm 0.037	0.228	0.403
	magenta	0.270 \pm 0.025	0.213	0.319
	cyan	0.203 \pm 0.045	0.103	0.279
Figure 8	split	0.241 \pm 0.028	0.205	0.317
	split	0.349 \pm 0.040	0.305	0.431
	split	0.302 \pm 0.041	0.235	0.351
Figure 1	foreground	0.374 \pm 0.066	0.269	0.555
	background	0.378 \pm 0.039	0.329	0.428

Table 2. Result analysis. For each objective, the target color, as well as the predicted color of the reinforced structure are given in sRGB and CIEXYZ color space. Their differences according to the CIEDE2000 standard is shown in the last column.

Figure	Result	sRGB		CIEXYZ		CIEDE2000
		Target color	Result color	Target color	Result color	
Figure 6 (top)	red	(1.000, 0.000, 0.000)	(0.687, 0.310, 0.437)	(0.436, 0.223, 0.014)	(0.240, 0.161, 0.128)	0.274
	green	(0.000, 1.000, 0.000)	(0.332, 0.697, 0.362)	(0.385, 0.717, 0.097)	(0.226, 0.345, 0.121)	0.205
	blue	(0.000, 0.000, 1.000)	(0.204, 0.433, 0.739)	(0.143, 0.061, 0.714)	(0.148, 0.151, 0.377)	0.228
	yellow	(1.000, 1.000, 0.000)	(0.881, 0.791, 0.318)	(0.821, 0.939, 0.111)	(0.566, 0.594, 0.127)	0.139
	magenta	(1.000, 0.000, 1.000)	(0.729, 0.331, 0.796)	(0.579, 0.283, 0.728)	(0.334, 0.209, 0.442)	0.112
	cyan	(0.000, 1.000, 1.000)	(0.008, 0.841, 0.886)	(0.528, 0.777, 0.811)	(0.369, 0.531, 0.608)	0.095
Figure 6 (bottom)	red	(1.000, 0.000, 0.000)	(0.914, 0.813, 0.818)	(0.436, 0.223, 0.014)	(0.687, 0.669, 0.526)	0.369
	green	(0.000, 1.000, 0.000)	(0.734, 0.801, 0.749)	(0.385, 0.717, 0.097)	(0.525, 0.576, 0.438)	0.284
	blue	(0.000, 0.000, 1.000)	(0.784, 0.823, 0.849)	(0.143, 0.061, 0.714)	(0.598, 0.632, 0.563)	0.536
	yellow	(1.000, 1.000, 0.000)	(0.913, 0.891, 0.820)	(0.821, 0.939, 0.111)	(0.742, 0.771, 0.542)	0.259
	magenta	(1.000, 0.000, 1.000)	(0.779, 0.524, 0.770)	(0.579, 0.283, 0.728)	(0.419, 0.330, 0.427)	0.166
	cyan	(0.000, 1.000, 1.000)	(0.564, 0.772, 0.830)	(0.528, 0.777, 0.811)	(0.430, 0.501, 0.526)	0.179
Figure 9	yellow	(1.000, 1.000, 0.000)	(0.855, 0.777, 0.629)	(0.821, 0.939, 0.111)	(0.574, 0.583, 0.317)	0.244
Figure 10	blue	(0.000, 1.000, 1.000)	(0.511, 0.875, 0.930)	(0.528, 0.777, 0.811)	(0.504, 0.631, 0.681)	0.110

Table 3. Result analysis. For the fabricated samples, the average color differences are quantified for three aspects: *i*) the color stability inside the patches; *ii*) the color stability across different patches; and *iii*) the color difference when compared with the predicted color for the reinforced structure.

Figure	Result	Average CIEDE2000 differences		
		Per patch	Between patches	Compared with result
Figure 9	yellow	0.011 ± 0.005	0.047 ± 0.019	0.076 ± 0.030
Figure 10	blue	0.007 ± 0.003	0.020 ± 0.015	0.218 ± 0.013

shown in Figure 4, the sample and the backlight were mounted on a rotating table that enabled us to simulate varying viewing directions. A zero aperture iris, as depicted in Figure 5, was used to simulated varying incident light cones.

REFERENCES

- Felix Bloch. 1929. Über die Quantenmechanik der Elektronen in Kristallgittern. *Zeitschrift für Physik* 52, 7-8 (jul 1929), 555–600. <https://doi.org/10.1007/bf01339455>
- Carl Eckart. 1933. A General Derivation of the Formula for the Diffraction by a Perfect Grating. *Physical Review* 44, 1 (jul 1933), 12–14. <https://doi.org/10.1103/physrev.44.12>
- Bin Liang, Ming Bai, Hui Ma, Naiming Ou, and Jungang Miao. 2014. Wideband Analysis of Periodic Structures at Oblique Incidence by Material Independent FDTD Algorithm. *IEEE Transactions on Antennas and Propagation* 62, 1 (jan 2014), 354–360. <https://doi.org/10.1109/tap.2013.2287896>
- Gaurav Sharma, Wencheng Wu, and Edul N. Dalal. 2004. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application* 30, 1 (2004), 21–30. <https://doi.org/10.1002/col.20070>

Table 4. Result statistics.

Figure	Result	Focus count	Objective function		Runtime	N_{iter}	N_{sim}	N_{target}
			Initial	Final				
Figure 6 (top)	red	36	0.159	0.023	03:27:50	67	1	1
	green	22	0.144	0.023	03:05:13	62	1	1
	blue	32	0.159	0.020	03:02:46	48	1	1
	yellow	39	0.090	0.009	01:31:29	39	1	1
	magenta	30	0.098	0.028	00:43:50	21	1	1
	cyan	17	0.067	0.006	01:03:34	10	1	1
Figure 6 (bottom)	red	36	0.232	0.210	00:30:24	7	1	1
	green	29	0.241	0.117	03:20:38	54	1	1
	blue	29	0.249	0.187	01:06:11	15	1	1
	yellow	36	0.114	0.090	02:21:39	18	1	1
	magenta	37	0.126	0.029	04:23:59	55	1	1
	cyan	34	0.121	0.023	02:43:08	43	1	1
Figure 7	red	37	0.175	0.118	12:22:47	35	23	9
	green	35	0.175	0.139	05:59:52	20	23	9
	blue	39	0.168	0.075	03:47:42	40	23	9
	yellow	31	0.103	0.067	06:51:09	34	23	9
	magenta	28	0.103	0.082	01:51:37	17	23	9
	cyan	28	0.096	0.056	05:02:50	24	23	9
Figure 8	split	26	0.142	0.095	08:52:12	22	23	9
	split	26	0.154	0.135	02:07:08	8	23	9
	split	37	0.152	0.107	03:28:53	16	23	9
Figure 9	yellow	16	0.083	0.062	01:00:44	15	1	1
Figure 10	blue	11	0.074	0.047	00:13:50	10	1	1
Figure 1	foreground	35	0.167	0.134	08:37:16	22	33	15
	background	30	0.161	0.148	03:02:17	9	33	15

Table 5. Macrophotography setup – legend of Figures 4 and 5.

Annotation	Description
1	Digital single-lens reflex camera (DSLR) with a macro objective.
2	Sample holder with mounted sample. The fabricated structure is visible as a bright spot in the center of the glass substrate in Figure 5.
3	Zero aperture iris.
4	Uniform light-emitting diode (LED) backlight
5	Rotating table. Several translation and rotation stages are used to place the sample exactly on the rotation axis of the table.




Fig. 4. Macrophotography setup – top view. See Table 5 for a description of the annotations.




Fig. 5. Macrophotography setup – top view. See Table 5 for a description of the annotations.