

Inverse-Designed Volumetric Metaoptics for Multifunctional Nanophotonics

Thesis by
Ian M Foo

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy in Applied Physics



CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2027
Defended June 11 2026

© 2027

Ian M Foo

ORCID: 0000-0001-5397-646X

All rights reserved except where otherwise noted

ACKNOWLEDGEMENTS

When I decided I wanted to be at Caltech I was about six or seven. There was no rhyme or reason to it other than the sheer randomness of seeing it mentioned somewhere in a book I was reading at the time. I had no idea what doing a degree here entailed or required, or even what a degree was— and the dream fell away very shortly after as I turned a venerable eight and moved on to other completely arbitrary life goals, now barely remembered. But it cheers me greatly that twenty-three years later I am able to say to that young child that I managed to do exactly what he wanted. I hope he is happy!

In applying to graduate school I had a very clear picture of the kind of group culture and environment I wanted. Prof. Andrei Faraon's group is everything I envisioned and more. It is a space where we all feel comfortable to suggest and pursue a myriad of intellectual paths and academic ideas. There is no sense of competition or undercutting, only a fountain of positive and sincere support from professional and exceptional people; and Andrei does an amazing job of cultivating such an environment, so rare in academia, and making it look easy. The calm and thoughtful way in which everyone in this group goes about their work was a great grounding force that removed many potential obstacles and allowed me to become the best researcher I could be. Thank you Andrei for giving me a chance to work with you; for believing in me and cheering me on; for the advice you gave and consideration you showed every step of the way; and for being a career and intellectual inspiration as only one's Ph.D supervisor can. Thank you as well to my committee: Kerry Vahala, Axel Scherer, and Harry Atwater, for being amazing mentors and role models in academia.

I have learned so much from the brilliant individuals I was lucky to work with throughout my Ph.D. No matter how much I had to learn from them— and I had a lot to learn— Conner Ballew, Neuton Li, Nia McNichols, Wenyan Chen, you were always kind and generous with your time and your lessons. I am especially grateful to Greg Roberts for teaching me everything I needed to know when I was a fresh graduate student. We all have that one senior colleague who is a mentor to us early on, and you were an endlessly patient and giving one.

I would not have been able to succeed without the help of countless people. Cecilia Gamboa, Jennifer Blankenship and Christy Jenstad keep everything running

perfectly smoothly underneath it all. I greatly appreciate the staff of The Kavli Nanoscience Institute and the Resnick High-Performance Computing Center for their efforts in keeping these facilities running, which were indispensable for my research. I am also thankful to Trevor O'Loughlin, Frederick Brady, Mohamed Sabry, and David Woolf for their insights and experience. I owe much to the entire quantum subgroup of the Faraon Group— gifted experimentalists every one— for always lending me their time, expertise, and equipment. And also to my colleagues of the metaoptics subgroup: Suki Gu and Phillippe Pearson each contain more knowledge in their heads than I ever could, and have taught me so much through the labwork and many discussions I was privileged to have with them. I am excited to see the brilliant work that Matilda Eriksson and Dennis Naughton will accomplish over the next few years. Thanks especially to Bilgehan Baspinar and Prachi Thureja, who kept me sane and present and focused through the days, weeks, and months. They are, and always will be, an endless source of inspiration to me for how I want to be as a researcher and friend.

Above all, the Ph.D has been a period of growth and heavy self-reflection. I can safely say that it was the most difficult time of my life (so far!). I had to ask myself many painful questions: What do I want in my life? What am I willing to sacrifice? How much therapy can a single human conceivably attend in a week while still showing up in the office during regular hours? But I can also safely say now that I have come out of it all more present and centered in my world than I've ever been.

Through it all, six years was a good time to hone some skills. I have been very fortunate to find teachers that fit like a long-lost puzzle piece. Molly is quite possibly the best voice teacher in the city, and most definitely the best possible teacher for me. Under her guidance I have reached more life goals than I thought imaginable. Pablo has brought me to the highest possible level of Spanish fluency that I could reach with the limited time that I had to practice. And Crystal, Sera, and Ally have taught me valuable lessons that will remain with me for a lifetime.

Spending time with friends and loved ones was the most critical thing that kept me rooted and that reminded me just what I was striving for throughout all these years. A sincere and heartfelt thank you to Allyx, Nikki, Amanda, Elodie, Christopher, Marion, Pierre, Linette, Yijie, Jieyu, Anthony, Eric, Andy, Wei Ai; and finally all of the Clowns, and the play-by-post crew, who gave me community and much-needed mirth during the pandemic and the first few years. Without you, I would have spent six years in Los Angeles without ever really discovering what living in Los Angeles

is about. I will miss the board game nights, nature trails, and our evenings spent eating good food, watching TV, playing videogames; and doing all the things that I will ultimately remember the most when all is said and done. Let the length of this document be objective proof that you mattered to me at least slightly more than that!

Thanks in particular to Zoe, for balancing out all the stress with the calmest weekends, and for reminding me to show up for myself each and every time. She introduced me to the loveliest, most genuine group of friends I've ever, or likely will ever, meet. In a way, you and I got through graduate school together, and I am heartened to know that as a therapist you will ultimately be far more useful and valuable to the world than anything I am writing down in the pages that follow.

To Inés: our journey began with me watching you type up your thesis, and now $7\frac{3161}{4872}$ years later to the day we have come full circle as I finish mine. You and I are almost unrecognizably different now, but I am very happy and grateful to have spent the intervening years changing alongside you. Thank you for all the laughter and joy. Thank you too for all the propranolol, and for never once talking to me about science, even though everybody expects that it's something we would do. I'm proud of you and, even though we've met so many excellent science communicators with upwards of a million subscribers here in LA, you remain the most talented one I've ever seen. Please hire me.

At the end of it all, I've been borne aloft by a village. If I have left out any words or any people here, please know that I have already said them to you in person. All the people that have helped me through my Ph.D, that have pushed me along so supportively, have touched my life in a not insignificant way. Every moment has been a treasure, and I'm lucky to have spent them with you.

Ian M Foo
California Institute of Technology
May 28, 2026

ABSTRACT

Optical metasurfaces have emerged as a powerful platform for compact, multifunctional light control, but the finite propagation length available within a single-layer metasurface fundamentally limits simultaneous and efficient control over multiple degrees of freedom such as phase, dispersion, polarization, and angular response. Multilayer and volumetric metaoptics overcome these limits through the use of multiple patterned surfaces or a full three-dimensional structured volume at the cost of a greatly enlarged design space and substantially increased fabrication complexity. We develop an adjoint-based framework for fabrication-aware inverse design of volumetric permittivity distributions and apply it to design and experimentally demonstrate dispersive polarizing beam-splitters. We use TiO_2 nanopillars in-scaffold to approximate a gradient refractive index spectral sorter in the mid-infrared, and separately realize ten-layer metaoptic color sorters in the visible wavelength regime using a novel Sb_2S_3 fabrication platform.

Optical metasurfaces — planar arrays of subwavelength nanostructures — have emerged as a powerful platform for compact, multifunctional light manipulation, enabling flat analogs of lenses, polarizers, holograms, and spectrometers within a single patterned layer. However, the finite propagation length available within a single-layer metasurface fundamentally limits simultaneous and efficient control over multiple degrees of freedom such as phase, dispersion, polarization, and angular response. Achieving complex optical transformations with high throughput across multiple channels can be suitably addressed by distributing the optical function across multiple layers or a full three-dimensional structured volume.

The two primary challenges in realizing multilayer, or volumetric, metaoptics are design and fabrication. We develop and present a compact, versatile software framework built around the adjoint method, which enables efficient gradient-based optimization of volumetric permittivity distributions with respect to arbitrary merit functions. The framework supports seamless switching between electromagnetic solvers, optimization schemes, and merit functions within a single workflow, and incorporates fabrication constraints directly into the optimization process, producing fabricable structures without sacrificing performance. Using this framework, we design and experimentally demonstrate dispersive polarizing beam-splitters that provide independent control of the dispersion of each polarization channel in the far field. Beam deflection angles of up to 45° are demonstrated in simulation and 30°

experimentally, with broadband operation in both achromatic and hyper-dispersive regimes. Through systematic inverse-design exploration of the device hyperparameter space, we characterize the performance of volumetric metaoptic Bayer-type color sorters under realistic fabrication constraints, and address the challenges of operating these in sensor arrays under non-collimated illumination. To expand the accessible fabrication space beyond binary geometries, we introduce a high-index-contrast pseudo-gradient-index volumetric metaoptic in TiO_2 and experimentally demonstrate spectral sorting in the mid-infrared. Finally, we present the first experimental demonstration of multilayer metaoptic color sorters operating in the visible spectrum, using a novel Sb_2S_3 -based fabrication platform to realize ten-layer devices, establishing an experimental foundation for volumetric metaoptic devices capable of multifunctional control with high efficiency under practical fabrication constraints.

PUBLISHED CONTENT AND CONTRIBUTIONS

1. Foo, I., Roberts, G., Pearson, P., Ballew, C. & Faraon, A. Volumetric Metaoptic Color Routers at Visible Wavelengths: A Performance Analysis for Practical Applications. *Manuscript submitted for publication* (2026).
I.F. planned the research effort, carried out optimizations, set up the simulations, and acquired and analyzed the data, with input from other authors. I.F. prepared the manuscript with input from all authors.
2. Foo, I., Li, N., Pearson, P. & Faraon, A. Broadband Dispersion Engineered Volumetric Metamaterial for Polarization Beam Splitting at Large Deflection Angles. *Manuscript in preparation* (2026).
I.F. provided code and updates for the inverse design framework, performed part of the optimizations, set up part of the simulations and acquired the data, fabricated and characterized the devices, and analyzed the data. I.F. prepared the manuscript with input from all authors.
3. Chen, W. *et al.* Heterogeneous Multilayered TiO_2 Metaoptics via Two-Photon Lithography. *Manuscript in preparation* (2026).
I.F. provided code for the inverse design framework of the color routers, performed optimizations of the color routers, set up the simulations and acquired the data, measured the color routers, and analyzed the data. I.F. discussed the results and contributed to the final manuscript.
4. Foo, I., Baspinar, A. B., Pearson, P., Gu, Y. & Faraon, A. Freeform Sb_2S_3 -Based Multilayered Metaoptics. *Manuscript in preparation* (2026).
I.F. provided code for the inverse design framework, performed optimizations, set up the simulations and acquired the data, measured the devices, and analyzed the data. I.F. prepared the manuscript with input from all authors.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	vi
Published Content and Contributions	ix
Table of Contents	ix
List of Illustrations	xiii
List of Tables	xv
Nomenclature	xvi
Chapter I: Introduction	1
Chapter II: Adjoint-Based Optimization and Inverse Design	14
2.1 Introduction	14
2.2 Optimization Problem Formulation	15
2.3 The Adjoint Method	16
2.3.1 The Adjoint Equations	16
2.3.2 Gradient Computation	18
2.3.3 Physical Interpretation	19
2.4 Figures of Merit	19
2.4.1 Intensity at a Target Location	20
2.4.2 Mode Coupling Efficiency	20
2.4.3 Multi-Port Transmission	20
2.5 Multi-Objective Optimization	21
2.5.1 Pareto fronts	22
2.5.2 Minimax and maxmin formulations	24
2.5.3 Weighted Chebyshev scalarization	24
2.5.4 ε -constraint method	26
2.5.5 Augmented Lagrangian method	26
2.5.6 Parallelization	29
2.5.7 Density Parameterization	31
2.5.8 Update Algorithms	31
2.6 Fabricability Constraints	33
2.6.1 Binarization	34
2.6.2 Minimum Feature Size	34
2.6.3 Binary-Constrained Update	35
2.6.4 Level-Set Refinement	35
2.7 Summary	36
2.7.1 Worked Example: Broadband Focusing Device	36
Chapter III: vipdopt: A Software Package for Volumetric Inverse Photonic Design and Optimization of Metaoptics	41
3.1 Introduction	41
3.1.1 Existing Frameworks and Their Limitations	43

	xi
3.1.2 New Technical Features	44
3.1.3 Chapter Outline	45
3.2 Background	46
3.2.1 Gradient Computation: Adjoint Method and Automatic Dif- ferentiation	46
3.3 Framework Architecture	47
3.3.1 Design Philosophy	47
3.3.2 System Overview	48
3.3.3 Design Pipeline Layer	49
3.3.4 Autodiff and Adjoint Bridge Layer	50
3.4 Optimization Schedule	51
3.4.1 Single-Stage Baseline	51
3.4.2 Multi-Stage Schedules	51
3.4.3 Comparison: FoM vs. Forward Simulations	52
3.5 Demonstration Applications	53
3.5.1 Benchmark Validation: 3D Metagrating	53
3.5.2 Spectral Reconstruction via Orthogonal Spatial Spectra	54
3.6 Discussion	55
3.6.1 Comparison to Prior Frameworks	55
3.6.2 Future Extensions	57
3.7 Conclusion	57
3.8 Minimal API Example	58
Chapter IV: Broadband Dispersion Engineered Volumetric Metamaterial for Polarization Beam Splitting at Large Deflection Angles	62
4.1 Introduction	62
4.2 Results and Discussion	65
4.2.1 Fabrication and Experiment	69
4.2.2 Performance Bounds	69
4.3 Conclusion	74
4.4 Supplementary Information	75
4.4.1 Adjoint Optimization and Projection Filters	75
4.4.2 Experimental Characterization Setup	77
4.4.3 Broader Control of Deflection and Dispersion	79
4.4.4 Calculation of Performance Bounds	82
4.4.5 Dispersions for Varying Thickness	84
4.4.6 Gradient-Index Devices	86
4.4.7 Stacking Beam Splitters	87
Chapter V: Volumetric Metaoptic Color Sorters at Visible Wavelengths: A Performance Analysis for Practical Applications	93
5.1 Introduction	93
5.2 Results	96
5.2.1 Metrics of Performance	99
5.2.2 Material Constraints for Fabricability	100
5.2.3 Three-Material Device	101
5.2.4 Crosstalk Attenuation between Color Sorter Pixels	104

5.2.5	Oblique-Incidence Optimizations	107
5.3	Conclusion	110
5.4	Supplementary Information	112
5.4.1	Comparison of Prior Color-Splitting Devices	112
5.4.2	Adjoint Optimization and Projection Filters	117
5.4.3	Changing Device Height (Layer Thickness Fixed) and Focal Length	121
5.4.4	Realistic Material Dispersion and Absorption	124
Chapter VI: Gradient-Refractive Index Multilayered TiO ₂ Metaoptics via Two-Photon Lithography		134
6.1	Introduction	135
6.2	Fabrication Procedure	137
6.3	Results and Discussion	139
6.3.1	Homogeneous Lattices	140
6.3.2	Chromatic Metalenses	140
6.3.3	Achromatic Metalenses	140
6.3.4	Scaffold Color Router	143
6.4	Conclusion	146
6.5	Experimental Section	147
6.6	Supplementary Information	150
6.6.1	Experimental Characterization Setup	150
Chapter VII: Inverse-Designed Ten-Layer Visible Color Sorter via Direct-Write Sb ₂ S ₃ Multilayer Metaoptics		155
7.1	Introduction	155
7.2	Results and Discussion	157
7.2.1	Device Design	157
7.2.2	Fabrication	160
7.3	Characterization	161
7.4	Conclusion	163
7.5	Supplementary Information	165
7.5.1	Adjoint Optimization Details	165
7.5.2	Sb ₂ S ₃ Optical Constants Across the Visible Band	167
7.5.3	Experimental Characterization Setup	168
Chapter VIII: Conclusion		173
8.1	Large-Area Metasurface Design	174
8.2	End-to-End Co-Optimization	175
8.3	Fabrication at Scale	176
8.4	Active, Tunable, and Programmable Volumetric Metaoptics	177

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
3.1 Architecture of vipdopt.	48
3.2 Example workflow of vipdopt ExperimentManager.	52
3.3 Evolution of FoM for single-stage and dual-stage optimization sched- ules.	53
3.4 Results of 3D metagrating benchmark optimization using vipdopt.	55
3.5 Spectral matching device optimization using vipdopt.	56
4.1 Comparison of different regimes of dispersion.	64
4.2 Inverse-designed zero-dispersion beam splitter. -20° for HP and $+30^\circ$ for VP.	67
4.3 Inverse-designed hyper-dispersive beam splitter. -20° for HP.	68
4.4 Simulated and measured results for zero-dispersion beam splitter in Fig. 4.2.	70
4.5 Simulated and measured results for hyper-dispersive beam splitter in Fig. 4.3.	71
4.6 Performance bounds of dispersion-engineered polarization beam split- ters.	73
4.7 Experimental characterization setup used for Fourier-plane imaging of the fabricated beam splitter.	77
4.8 [Inverse-designed hyper-dispersive beam splitter. -20° for HP and VP.	80
4.9 Inverse-designed zero-dispersion beam splitter. $+30^\circ$ for HP and -30° for VP.	81
4.10 Inverse-designed zero-dispersion beam splitter. $+45^\circ$ for HP and -45° for VP.	82
4.11 Thickness dependence of angular dispersion and transmission effi- ciency for $\pm 45^\circ$ zero-dispersion beam splitters.	85
4.12 Comparison of (separately optimized) gradient-refractive-index and binarized hyper-dispersive beam splitters.	85
4.13 Broadband polarization-splitting beam displacer formed by cascading zero-dispersion splitters.	87
5.1 Overview of the inverse-designed 10-layer RGB color sorter, includ- ing its geometry, sorting spectra, layer structure, and loss channels.	98

5.2	Hyperparameter sweeps of layer number and lateral minimum feature size, with effects on sorting spectra, peak band transmissions, and scattering losses.	100
5.3	Comparison of two- and three-material projection schemes for inverse-designed color sorters.	102
5.4	Evaluation of sidewall-based trench isolation as a strategy for reducing optical crosstalk between adjacent color sorters.	105
5.5	Reoptimization of color sorters for oblique incidence and input light of varying divergence angle with minimal loss of sorting efficiency. .	108
5.6	Selected layers of final optimized device.	118
5.7	Binarization projection used during inverse design and its effect on the FoM during optimization.	119
5.8	Examples of fabrication constraints and their effects on refractive index distribution and transmission spectra.	119
5.9	Hyperparameter sweeps of layer number (with fixed layer thickness), focal length, and device height and their effects on transmission, scattering losses, contrast, and bandwidth.	121
5.10	Comparison of idealized nondispersive simulations with single-frequency simulations using realistic material dispersion, and (n, k) refractive-index data across the visible bandwidth.	124
6.1	Conceptual overview of nano-additive manufactured multilayer TiO_2 metaoptics enabled by two-photon lithography.	136
6.2	Nano-additive manufacturing process and SEM characterization of fabricated homogeneous lattices.	138
6.3	Chromatic and achromatic TiO_2 metalenses via nano-AM, establishing the fabrication platform for multilayer metaoptics.	141
6.4	Characterization of nano-AM TiO_2 color router.	142
6.5	Metalens characterization setups.	150
7.1	Conventional and Sb_2S_3 fabrication platforms.	158
7.2	Schematic and SEM image of Sb_2S_3 color sorter with 10 layers and 9 spacers.	158
7.3	Simulated and measured performance of color sorter.	162
7.4	Optical constants n, k of Sb_2S_3 calculated from reflectance spectrometry.	167
7.5	Experimental characterization setup used to measure transmission of Sb_2S_3 color sorter	168

LIST OF TABLES

<i>Number</i>	<i>Page</i>
5.1 Comparisons of physical parameters and performance metrics of previous metaoptic-based color splitting devices.	112
5.2 Comparison of power scattering for simulations with non-dispersive and dispersive material properties.	124

NOMENCLATURE

- Dispersion.** the varying of the refractive index of materials with wavelength.
- FoM.** figure of merit, otherwise called an objective or merit function; a single scalar quantity that characterizes a device's desired performance.
- GRIN.** gradient refractive-index.
- Loss landscape.** An N -dimensional hypersurface of the merit function of an optimization, where N is the number of design parameters.
- N.A..** numerical aperture: describes the range of angles within which a lens can accept light.
- Optical axis.** a line typically passing through the shared geometric center of the elements in an optical system that defines the direction of propagation of unperturbed light in that system.

*Chapter 1***INTRODUCTION:
FROM METASURFACES TO VOLUMETRIC METAOPTICS**

The development of sight in organisms is generally held to be one of the watershed moments in evolutionary history. Across many different animal lineages, sight has independently evolved over and over again. This is simply because the ability to detect light allows an organism to detect, and interact with, entities and objects both close up and at a great distance. It is easier to find food when you can see it hanging from a tree, instead of fumbling around on your hands and knees. It is easier also to run from a predator when you can see it coming for you long before the sound of its footsteps or the smell of its fur is able to reach you.

Because of this, across both the plant and animal kingdom many organisms have developed mechanisms to control the light around them: to harness and focus light for energy; to sense light both more acutely and broadly; and to affect how their bodies interact with light, and how others see them, just to name a few of these possibilities. Some of the most fascinating examples¹ include:

- Butterfly wings contain microstructured chitin layers that act as wavelength-selective diffraction gratings, which are responsible for their intense color;
- The epidermal cells of many leaves are able to focus light by 5–20× to maximize the irradiance received by chloroplasts^{2,3};
- And of course the eye, a natural varifocal lens capable of forming sharp, multispectral images of objects and scenes at a great range of distances with an incredible field of view.

Mirroring this, we humans have gone ahead and developed an innumerable range of devices and technologies in our own efforts to control and manipulate light. All around us are optical systems containing elements such as lenses, prisms, diffraction gratings, polarizers, and spectral filters, just to name a few. They affect an incident light wave by changing or selecting certain key aspects: some of these include amplitude, propagation direction, wavelength, and polarization. However, each of these elements comes with its own shortcomings.

For example, a lens is a carefully shaped piece of material that may vary in thickness or refractive index across its plane. An incoming monochromatic plane wave, at normal incidence, has constant phase across that same plane. However, depending on the spatial position in this plane (normal to the optical axis) by passing through the lens the incoming light acquires a phase shift that depends on the thickness and refractive index of the lens at that position. For a wave, the local wavevector at every point in space is determined by the local gradient of the phase, and points toward the direction of greatest increase. We can therefore design our lens to impart a spatially varying phase profile such that the exiting light, at every point in space, has a local wavevector pointing toward a single shared focal point.

We can also achieve this effect in another way, through a technology called “metaoptics” or “optical metamaterials”. These are precise arrangements of specifically chosen nanostructures that directly affect the scattering of light. Instead of relying on the thickness and refractive index of bulk media through which the light propagates, we can choose the geometric and material properties of each nanostructure at each point in space to impart the same desired phase shift. Because these structures typically measure at below the wavelength of light, metaoptic elements have the advantage of being much more compact and can provide significant miniaturization to an optical system.

Metaoptics are also able to replace multiple optical elements at once in a system. One of the most pervasive examples⁴ of a metaoptic element is the achromatic metalens, sometimes called a broadband or zero-dispersion metalens. Conventional lenses are subject to dispersion and therefore have different focal lengths for different wavelengths. In bulk optics, this is counteracted by combining^{5,6} two (or more) lenses, called a doublet, of different refractive indices and focal lengths such that the focal length of the compound lens does not distort with wavelength. Depending on the requirements of the user, compound lens systems can sometimes contain eight or more individual lenses^{4,7,8}. In contrast, a metalens can achieve the same effect with a single, flat layer of subwavelength elements. These “meta-atoms” are chosen and arranged so as to impose the exact required phase (and amplitude) profile on light passing through, which now varies not just spatially, but also with wavelength.

Plasmonic nanoantennas, which are metallic nanostructures that support localized surface plasmon resonances, featured strongly in early optical metasurfaces^{9,10} due to their ability to impose abrupt phase discontinuities¹¹. However, such metals are significantly lossy, particularly at visible wavelengths, which motivated the rapid de-

velopment of dielectric metasurfaces. These are structures of high-index dielectric material that are able to provide strong control of optical phase while avoiding the intrinsic absorption losses of metals, paving the way for high-efficiency metaoptic devices^{12–15}. Several mechanisms exist to control phase, amplitude and/or polarization of the transmitted or reflected field^{11,16,17}. Anisotropic meta-atoms can act as local waveplates by controlling geometric, or Pancharatnam-Berry^{18–20}, phase. Rotating one of these meta-atoms changes the phase of the converted circular polarization component, allowing for full control through 2π phase coverage. The geometry of the meta-atoms can also affect the effective optical path length^{12,13} or the resonances in the scatterers^{15,16}. These mechanisms can also be combined to improve efficiency, control polarization, or engineer dispersion^{6,21,22}.

In the years since, the variety of metaoptic devices and applications has grown at a staggering rate^{10,21,22}. Metalenses have enabled compact achromatic focusing, high-NA operation, and polarization-dependent imaging^{23–25}. By controlling anisotropy and birefringence, they can transform polarization states and perform compact Stokes or Mueller polarimetry^{26,27}. Devices have also been demonstrated for holography²⁸, compact spectrometry²⁹, full-color computational imaging³⁰, and analog optical image processing³¹. More broadly, metaoptic devices have enabled a shift from simple wavefront shaping to multifunctional optical processing. Their phase profiles can be engineered with far greater local freedom than a polished surface. Because each meta-atom can be designed independently, the same device can encode multiple different optical functions into different wavelengths, polarizations, or incident angles³², such as focusing different polarizations to different locations or combining imaging with polarization analysis^{8,21–25}.

Many of these metaoptic devices are constructed as metasurfaces or “flat optics”: a single, planar array of meta-atoms in two dimensions. For such devices, as the complexity of the designed optical function increases, the primary limitation that comes to bear is the thickness³³ of the meta-atoms in that single layer. Continuing with our example of broadband achromatic focusing, the spatial phase profile required must also be different across different wavelengths. This requires control not only of phase, but also of group delay and higher-order dispersion across the entire lens aperture. If we want the lens to have high NA, we must increase the needed delay variation; if we want to further improve its field of view, we must correct more strongly for off-axis aberrations³⁴. It can become very difficult for a single-layer metasurface, with only a certain finite propagation length available, to achieve all

these things without sacrificing throughput or avoiding unwanted scattering^{6,35,36}.

In general, a single-layer metasurface cannot implement any and every desired transformation between input and output fields with perfect efficiency, arbitrary bandwidth, full polarization control, and large angular acceptance^{33,36}. For example, a single layer of nonchiral dielectric nanofins often produces a constrained Jones matrix because of symmetry³⁷. Although such structures can produce many useful polarization transformations, they cannot realize every possible transformation between arbitrary polarization states. More general polarization control requires breaking additional symmetries, introducing chirality, or otherwise increasing the number of accessible design degrees of freedom^{14,26,36}. In these situations, thicker or multilayer structures become necessary to achieve complex or multiple optical functionalities with high efficiency^{6,33,36}.

Compound or multilayer metaoptics address the limitations of single-layer devices by distributing the optical transformation across multiple patterned surfaces, or through a larger structured volume^{36,38}. Just like our previous example of a compound lens system, multiple layers of optical interactions allow the designer to correct aberrations, compensate for dispersion, improve efficiency, and access transformations that are challenging for a metaoptic device of only a single surface. Some examples of compound metaoptics include cascaded metasurfaces^{38,39}, metalens doublets^{24,40} and triplets⁴¹, double-sided metasurfaces^{40,42}, cladded multilayer stacks⁴³, folded metasurface systems²⁹, and free-standing multilayer structures⁴⁴. Metasurface doublets have been used to build compact wide-angle planar cameras corrected for monochromatic aberrations²⁴. Multilayer and hybrid metasurfaces have been explored for achromatic response, independent amplitude and phase control, and nonlocal multifunctional wavefront shaping^{36,43,45}. They also offer opportunities for active and tuneable metaoptic systems via relative motion or Moiré effects^{46–48} between individual layers³⁶. This principle has been used to design rotary varifocal devices⁴⁹ and tunable Airy-beam generators⁵⁰.

Volumetric metaoptics extends the concept of compound metaoptics to its natural limit: a fully three-dimensional design space. Instead of a single patterned surface, or a small number of stacked layers, volumetric metaoptics uses a three-dimensional refractive-index or permittivity distribution to scatter light throughout a finite volume. If there is enough distance, or sufficiently weak coupling, between each layer, the metaoptic may still be treated as a cascaded series of decoupled phase masks. Typically, however, the full multilayer device must be treated as a single sys-

tem^{27,32,33,51}, as multiple scattering and interference throughout the device volume become instrumental in implementing the target optical function. The diversity of available interactions allows for higher efficiency in the case of tasks requiring simultaneous control over multiple channels⁵¹. Spectral sorting, polarization sorting, mode conversion, multiplexing of orbital angular momentum, and compact imaging all benefit from the ability to route different components of the incident field into different output locations or modes.

As the number of design variables grows from a few geometric parameters per meta-atom to tens of thousands of voxels, forward design methods based on meta-atom libraries and intuition become insufficient for exploring the large design space to find the maximum performance possible. Instead, through shape or topology optimization^{52,53}, a merit function can be optimized with respect to all the available degrees of freedom^{54,55}. The desired function may be focusing efficiency⁵¹, channel contrast³², polarization extinction^{27,51}, spectral separation²⁷, mode overlap⁵³, or a weighted combination of multiple such objectives. Inverse design methods therefore provide the computational framework needed to exploit the full design freedom of volumetric metaoptics^{51,54,55}.

The other primary obstacle in realizing multilayer metaoptics is the complexity of fabrication. Adding layers increases the number and difficulty of fabrication steps, and demands more severe alignment tolerances. The fabrication of volumetric metaoptics requires methods capable of producing controlled three-dimensional nanoscale structures, such as multilayer lithography^{44,56–58}, two-photon polymerization^{27,59}, grayscale lithography⁶⁰, or nanoimprint-based multilayers^{58,61}. The achievable index contrast, minimum feature size, surface roughness, alignment accuracy, and material absorption all affect device performance.

Thus far, volumetric metaoptics have been used to experimentally achieve simultaneous control over multiple output channels, such as spectral routing and polarization sorting^{27,36,51}. It is currently impractical, and may remain so for some time, to fabricate three-dimensional devices of arbitrary permittivity with nanometer-scale resolution, even though such devices have been shown to be capable of almost-perfect optical efficiency and function^{62,63}. A practical volumetric metaoptic device must respect the achievable feature sizes, layer counts, aspect ratios, etch depths, sidewall profiles and layer-to-layer alignment tolerances of its fabrication platform, as well as the available materials and their dispersion and absorption^{36,54,55}.

At the time of writing, the field is transitioning strongly toward optical devices^{8,22}

that are manufacturable at scale. Translation into market technologies requires not only high optical performance, but also raises questions of compatibility with existing foundry processes and the mass production of metaoptic devices. Metaoptics must also be well-integrated with existing optoelectronic systems⁸. Some examples of market technologies that benefit directly from metaoptics include hyperspectral and/or LIDAR sensing for biological and automotive detection, wearable displays and augmented reality, computational imaging, and integrated photonics, especially where small device sizes and form factors are particularly desirable. From a more environmental perspective, metaoptics can also be useful in radiative cooling and the control of thermal emission and absorption, as well as the harvesting of solar energy.

In this thesis, solutions are presented to the problems inherent in the design and subsequent fabrication of volumetric and multilayer metaoptics. In Chapter 2, we present the theoretical framework behind the adjoint method, which is an algorithm used to efficiently compute the gradient of the merit function of a device with respect to its permittivity degrees of freedom. In Chapter 3, we discuss several different optimization schemes that make use of this gradient in order to find permittivity distributions that yield high device performance within the available design space. A compact and versatile software package is presented that is able to design metaoptic devices while seamlessly switching between electromagnetic solvers, merit functions, and optimization schemes during the course of the optimization, opening up avenues to find the highest-performing designs possible in the context of an optical function and the allowed range of device hyperparameters. Fabrication constraints are also discussed that can be directly incorporated into the optimization workflow, allowing designed devices to be successfully fabricated and still demonstrate high performance.

The primary benefit of volumetric metaoptics is multifunctionality with high efficiency. In Chapter 4, we design beam-splitters that separate light by polarization, while demonstrating the ability to control the dispersion of each polarization in the far-field. We demonstrate arbitrary control of each of the beam deflection angles, of up to 45° in simulation and 30° in experiment, and show broadband examples of both achromatic and hyper-dispersive polarizing beam-splitters.

Bayer-type color sorters, which spatially discriminate between different wavelengths of light and route them to different photodetectors, can outperform conventional Bayer filters by factors of up to $2 - 3\times$, and represent another example of multi-

functionality that is highly relevant to modern technologies. In Chapter 5, we use inverse design to explore the hyperparameter space of multilayer metaoptic-based color sorters. We detail the effects on the performance of such devices under constraints of layer counts and minimum feature sizes that are practical at the time of writing. Integrating them with modern sensors also requires that they operate in large arrays with minimal intra- and inter-pixel crosstalk. We discuss methods of minimizing this crosstalk, and of making such color sorters robust against angular incidence in the case of incident non-collimated light.

Most practical nanofabrication processes strongly favor binarized device geometries composed of only two possible materials. Gradient-refractive-index (GRIN) devices offer a richer design space by allowing the optical permittivity to vary more continuously throughout the device volume. This additional degree of freedom can improve performance by enabling finer control over phase, scattering, and dispersion. At the same time, high index contrast enables compact phase accumulation over shorter propagation distances, and expands the range of achievable scattering and dispersion responses. Despite these advantages, experimentally demonstrated GRIN-like metaoptics remain limited by fabrication complexity. In Chapter 6 we address this challenge by introducing a high-index-contrast pseudo-gradient-index TiO_2 metaoptic and experimentally demonstrate spectral sorting in the mid-infrared.

Finally, in Chapter 7 we present the first experimental demonstration of multilayer metaoptic color sorters operating in the visible regime. While single-layer metaoptics have been experimentally realized within the visible spectrum, multilayer visible metaoptics have till now remained out of reach because of the stringent fabrication requirements imposed by subwavelength feature sizes, layer-to-layer alignment tolerances, and cumulative process complexity. We use a novel Sb_2S_3 -based fabrication platform to realize ten-layer metaoptic color sorters and experimentally demonstrate spectral sorting in the visible. The work establishes both a new fabrication route for multilayer visible metaoptics, and an experimental foundation for volumetric color-routing devices beyond single-layer metasurface designs.

References

1. Vukusic, P. & Sambles, J. R. Photonic structures in biology. *Nature* **424**, 852–855. ISSN: 1476-4687. <https://www.nature.com/articles/nature01941> (2026) (Aug. 2003).

2. Vignolini, S., Moyroud, E., Glover, B. J. & Steiner, U. Analysing Photonic Structures in Plants. *Journal of the Royal Society Interface* **10**, 20130394. ISSN: 1742-5689. (2026) (Oct. 2013).
3. Vogelmann, T. C. in *Photomorphogenesis in Plants* (eds Kendrick, R. E. & Kronenberg, G. H. M.) 491–535 (Springer Netherlands, Dordrecht, 1994). ISBN: 978-94-011-1884-2. https://doi.org/10.1007/978-94-011-1884-2_18 (2026).
4. Khorasaninejad, M. & Capasso, F. Metalenses: Versatile multifunctional photonic components. *Science* **358**. ISSN: 0036-8075 (2017).
5. Born, M. *et al.* *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light* ISBN: 978-0-521-64222-4 (Cambridge University Press, Cambridge ; New York, 1999).
6. Chen, W. T., Zhu, A. Y. & Capasso, F. Flat optics with dispersion-engineered metasurfaces. *Nature Reviews Materials* **5**, 604–620. ISSN: 2058-8437. <https://www.nature.com/articles/s41578-020-0203-3> (2026) (Aug. 2020).
7. Smith, W. J. *Modern Optical Engineering, 4th Ed.* ISBN: 978-0-07-159375-5 (McGraw Hill Professional, Dec. 2007).
8. Brongersma, M. L. *et al.* The second optical metasurface revolution: moving from science to technology. *Nature Reviews Electrical Engineering* **2**, 125–143. ISSN: 2948-1201. <https://www.nature.com/articles/s44287-024-00136-4> (2026) (Feb. 2025).
9. Ni, X., Emani, N. K., Kildishev, A. V., Boltasseva, A. & Shalaev, V. M. Broad-band Light Bending with Plasmonic Nanoantennas. *Science* **335**, 427–427. <https://www.science.org/doi/10.1126/science.1214686> (2026) (Jan. 2012).
10. Chen, H.-T., Taylor, A. J. & Yu, N. A review of metasurfaces: physics and applications. *Reports on Progress in Physics* **79**, 076401. ISSN: 0034-4885. <https://doi.org/10.1088/0034-4885/79/7/076401> (2026) (June 2016).
11. Yu, N. *et al.* Light Propagation with Phase Discontinuities: Generalized Laws of Reflection and Refraction. *Science* **334**, 333–337. (2022) (Oct. 2011).
12. Lin, D., Fan, P., Hasman, E. & Brongersma, M. L. Dielectric Gradient Metasurface Optical Elements. *Science* **345**, 298–302. (2026) (July 2014).
13. Yu, Y. F. *et al.* High-transmission dielectric metasurface with 2π phase control at visible wavelengths. *Laser & Photonics Reviews* **9**, 412–418. ISSN: 1863-8899. <https://onlinelibrary.wiley.com/doi/abs/10.1002/lpor.201500041> (2026) (2015).

14. Arbabi, A., Horie, Y., Bagheri, M. & Faraon, A. Dielectric Metasurfaces for Complete Control of Phase and Polarization with Subwavelength Spatial Resolution and High Transmission. *Nature Nanotechnology* **10**, 937–943. ISSN: 1748-3395. (2026) (Nov. 2015).
15. Koshelev, K. & Kivshar, Y. Dielectric Resonant Metaphotonics. *ACS Photonics* **8**, 102–112. (2026) (Jan. 2021).
16. Bharadwaj, P., Deutsch, B. & Novotny, L. Optical Antennas. *Advances in Optics and Photonics* **1**, 438–483. ISSN: 1943-8206. <https://opg.optica.org/aop/abstract.cfm?uri=aop-1-3-438> (2026) (Nov. 2009).
17. Kildishev, A. V., Boltasseva, A. & Shalaev, V. M. Planar Photonics with Metasurfaces. *Science* **339**, 1232009. (2026) (Mar. 2013).
18. Pancharatnam, S. Generalized theory of interference, and its applications. *Proceedings of the Indian Academy of Sciences - Section A* **44**, 247–262. ISSN: 0370-0089. <https://doi.org/10.1007/BF03046050> (2026) (Nov. 1956).
19. Berry, M. The Adiabatic Phase and Pancharatnam’s Phase for Polarized Light. *Journal of Modern Optics* **34**, 1401–1407. ISSN: 0950-0340. <https://doi.org/10.1080/09500348714551321> (2026) (Nov. 1987).
20. Bomzon, Z., Biener, G., Kleiner, V. & Hasman, E. Space-variant Pancharatnam–Berry phase optical elements with computer-generated subwavelength gratings. *Optics Letters* **27**, 1141–1143. ISSN: 1539-4794. <https://opg.optica.org/ol/abstract.cfm?uri=ol-27-13-1141> (2026) (July 2002).
21. Ou, K. *et al.* Advances in Meta-Optics and Metasurfaces: Fundamentals and Applications. *Nanomaterials* **13**, 1235. ISSN: 2079-4991. <https://www.mdpi.com/2079-4991/13/7/1235> (2026) (Jan. 2023).
22. Kuznetsov, A. I. *et al.* Roadmap for Optical Metasurfaces. *ACS Photonics* **11**, 816–865. (2026) (Mar. 2024).
23. Khorasaninejad, M. *et al.* Polarization-Insensitive Metalenses at Visible Wavelengths. *Nano Letters* **16**, 7229–7234. ISSN: 1530-6984. (2026) (Nov. 2016).
24. Arbabi, A. *et al.* Miniature optical planar camera based on a wide-angle metasurface doublet corrected for monochromatic aberrations. *Nature Communications* **7**, 13682. ISSN: 2041-1723. <https://www.nature.com/articles/ncomms13682> (2026) (Nov. 2016).
25. Shrestha, S., Overvig, A. C., Lu, M., Stein, A. & Yu, N. Broadband Achromatic Dielectric Metalenses. *Light: Science & Applications* **7**, 85. ISSN: 2047-7538. (2026) (Nov. 2018).
26. Rubin, N. A., Shi, Z. & Capasso, F. Polarization in diffractive optics and metasurfaces. *Advances in Optics and Photonics* **13**, 836–970. ISSN: 1943-8206. <https://opg.optica.org/aop/abstract.cfm?uri=aop-13-4-836> (2026) (Dec. 2021).

27. Roberts, G. *et al.* 3D-patterned Inverse-Designed Mid-Infrared Metaoptics. *Nature Communications* **14**, 2768. ISSN: 2041-1723. (2025) (May 2023).
28. Ni, X., Kildishev, A. V. & Shalaev, V. M. Metasurface holograms for visible light. *Nature Communications* **4**, 2807. ISSN: 2041-1723. <https://www.nature.com/articles/ncomms3807> (2026) (Nov. 2013).
29. Faraji-Dana, M. *et al.* Compact folded metasurface spectrometer. *Nature Communications* **9**, 4196. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-018-06495-5> (2023) (Oct. 2018).
30. Colburn, S., Zhan, A. & Majumdar, A. Metasurface Optics for Full-Color Computational Imaging. *Science Advances* **4**, eaar2114. (2023) (Feb. 2018).
31. Zhou, J. *et al.* Metasurface Enabled Quantum Edge Detection. *Science Advances* **6**, eabc4385. (2026) (Dec. 2020).
32. Ballew, C., Roberts, G. & Faraon, A. Multi-dimensional wavefront sensing using volumetric meta-optics. *Optics Express* **31**, 28658–28669. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-31-18-28658> (2025) (Aug. 2023).
33. Miller, D. A. B. Why Optics Needs Thickness. *Science* **379**, 41–45. (2023) (Jan. 2023).
34. Aieta, F., Genevet, P., Kats, M. & Capasso, F. Aberrations of flat lenses and aplanatic metasurfaces. *Optics Express* **21**, 31530–31539. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-21-25-31530> (2026) (Dec. 2013).
35. Martins, A. *et al.* On Metalenses with Arbitrarily Wide Field of View. *ACS Photonics* **7**, 2073–2079. <https://doi.org/10.1021/acsphotonics.0c00479> (2026) (Aug. 2020).
36. Dorrah, A. H. Compound Meta-Optics: There Is Plenty of Room at the Top. *Nanophotonics*. ISSN: 2192-8614. (2025) (Apr. 2025).
37. Rubin, N. A., Zaidi, A., Dorrah, A. H., Shi, Z. & Capasso, F. Jones matrix holography with metasurfaces. *Science Advances* **7**, eabg7488. <https://www.science.org/doi/full/10.1126/sciadv.abg7488> (2025) (Aug. 2021).
38. Zheng, H. *et al.* Compound Meta-Optics for Complete and Loss-Less Field Control. *ACS Nano* **16**, 15100–15107. ISSN: 1936-0851. (2025) (Sept. 2022).
39. Lin, Z., Groever, B., Capasso, F., Rodriguez, A. W. & Lončar, M. Topology-Optimized Multilayered Metaoptics. *Physical Review Applied* **9**, 044030. <https://link.aps.org/doi/10.1103/PhysRevApplied.9.044030> (2026) (Apr. 2018).
40. Groever, B., Chen, W. T. & Capasso, F. Meta-Lens Doublet in the Visible Region. *Nano Letters* **17**, 4902–4907. ISSN: 1530-6984. <https://doi.org/10.1021/acs.nanolett.7b01888> (2026) (Aug. 2017).

41. Avayu, O., Almeida, E., Prior, Y. & Ellenbogen, T. Composite functional metasurfaces for multispectral achromatic optics. *Nature Communications* **8**, 14992. ISSN: 2041-1723. <https://www.nature.com/articles/ncomms14992> (2026) (Apr. 2017).
42. Arbabi, E., Arbabi, A., Kamali, S. M., Horie, Y. & Faraon, A. Controlling the sign of chromatic dispersion in diffractive optics with dielectric metasurfaces. *Optica* **4**, 625–632. ISSN: 2334-2536. <https://opg.optica.org/optica/abstract.cfm?uri=optica-4-6-625> (2024) (June 2017).
43. Malek, S. C., Overvig, A. C., Alù, A. & Yu, N. Multifunctional resonant wavefront-shaping meta-optics based on multilayer and multi-perturbation nonlocal metasurfaces. *Light: Science & Applications* **11**, 246. ISSN: 2047-7538. <https://www.nature.com/articles/s41377-022-00905-6> (2023) (Aug. 2022).
44. Dorrah, A. H., Park, J.-S., Palmieri, A. & Capasso, F. Free-standing bilayer metasurfaces in the visible. *Nature Communications* **16**, 3126. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-025-58205-7> (2025) (Apr. 2025).
45. Overvig, A. C. *et al.* Dielectric metasurfaces for complete and independent control of the optical amplitude and phase. *Light: Science & Applications* **8**, 92. ISSN: 2047-7538. <https://www.nature.com/articles/s41377-019-0201-7> (2026) (Oct. 2019).
46. Bernet, S. & Ritsch-Marte, M. Adjustable refractive power from diffractive moiré elements. *Applied Optics* **47**, 3722–3730. ISSN: 2155-3165. <https://opg.optica.org/ao/abstract.cfm?uri=ao-47-21-3722> (2026) (July 2008).
47. Colburn, S. & Majumdar, A. Simultaneous Achromatic and Varifocal Imaging with Quartic Metasurfaces in the Visible. *ACS Photonics* **7**, 120–127. <https://doi.org/10.1021/acsphotonics.9b01216> (2026) (Jan. 2020).
48. Iwami, K., Ogawa, C., Nagase, T. & Ikezawa, S. Demonstration of focal length tuning by rotational varifocal moiré metalens in an ir-A wavelength. *Optics Express* **28**, 35602–35614. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-28-24-35602> (2026) (Nov. 2020).
49. Ogawa, C., Nakamura, S., Aso, T., Ikezawa, S. & Iwami, K. Rotational varifocal moiré metalens made of single-crystal silicon meta-atoms for visible wavelengths. *Nanophotonics* **11**, 1941–1948. ISSN: 2192-8614. <https://onlinelibrary.wiley.com/doi/abs/10.1515/nanoph-2021-0690> (2026) (2022).
50. Zhang, J. C. *et al.* A 6G meta-device for 3D varifocal. *Science Advances* **9**, eadf8478. <https://www.science.org/doi/10.1126/sciadv.adf8478> (2026) (Jan. 2023).

51. Camayd-Muñoz, P., Ballew, C., Roberts, G. & Faraon, A. Multifunctional Volumetric Meta-Optics for Color and Polarization Image Sensors. *Optica* **7**, 280. ISSN: 2334-2536. (2020) (Apr. 2020).
52. Miller, O. D. *Photonic Design: From Fundamental Solar Cell Physics to Computational Inverse Design* PhD thesis (University of California, Berkeley). <http://optoelectronics.eecs.berkeley.edu/ThesisOwenMiller.pdf>.
53. Lalau-Keraly, C. M., Bhargava, S., Miller, O. D. & Yablonovitch, E. Adjoint shape optimization applied to electromagnetic design. *Optics Express* **21**, 21693. ISSN: 1094-4087. <https://www.osapublishing.org/oe/abstract.cfm?uri=oe-21-18-21693> (2020) (Sept. 2013).
54. Molesky, S. *et al.* Inverse design in nanophotonics. *Nature Photonics* **12**, 659–670. ISSN: 1749-4893. <https://www.nature.com/articles/s41566-018-0246-9> (2026) (Nov. 2018).
55. Fan, J. A. Freeform metasurface design based on topology optimization. *MRS Bulletin* **45**, 196–201. ISSN: 1938-1425. <https://doi.org/10.1557/mrs.2020.62> (2026) (Mar. 2020).
56. Zhou, Y. *et al.* Multilayer Noninteracting Dielectric Metasurfaces for Multi-wavelength Metaoptics. *Nano Letters* **18**, 7529–7537. ISSN: 1530-6984. <https://doi.org/10.1021/acs.nanolett.8b03017> (2026) (Dec. 2018).
57. Mansouree, M. *et al.* Multifunctional 25D metastructures enabled by adjoint optimization. *en. Optica* **7**, 77. ISSN: 2334-2536. <https://www.osapublishing.org/abstract.cfm?URI=optica-7-1-77> (2020) (Jan. 2020).
58. Choi, S., Zuo, J., Das, N., Yao, Y. & Wang, C. Scalable Nanoimprint Manufacturing of Functional Multilayer Metasurface Devices. *Advanced Functional Materials* **34**, 2404852. ISSN: 1616-3028. <https://onlinelibrary.wiley.com/doi/abs/10.1002/adfm.202404852> (2026) (2024).
59. Gissibl, T., Thiele, S., Herkommer, A. & Giessen, H. Two-photon direct laser writing of ultracompact multi-lens objectives. *Nature Photonics* **10**, 554–560. ISSN: 1749-4893. <https://www.nature.com/articles/nphoton.2016.121> (2026) (Aug. 2016).
60. Shanks, D. N., Wenger, T., Muller, R. E., Wallace, J. K. & Wilson, D. W. Metasurfaces with Freely Varying Height in the Visible Using Grayscale Lithography. *Nano Letters* **25**, 14090–14095. ISSN: 1530-6984. <https://doi.org/10.1021/acs.nanolett.5c03283> (2025) (Sept. 2025).
61. Einck, V. J. *et al.* Scalable Nanoimprint Lithography Process for Manufacturing Visible Metasurfaces Composed of High Aspect Ratio TiO₂ Meta-Atoms. *ACS Photonics* **8**, 2400–2409. <https://doi.org/10.1021/acsphotonics.1c00609> (2026) (Aug. 2021).

62. Catrysse, P. B., Zhao, N., Jin, W. & Fan, S. Subwavelength Bayer RGB color routers with perfect optical efficiency. *Nanophotonics* **11**. issn: 2192-8606, 2192-8614. <https://www.degruyter.com/document/doi/10.1515/nanoph-2022-0069/html> (2022) (Mar. 2022).
63. Catrysse, P. B. & Fan, S. Spectral routers for snapshot multispectral imaging. *Applied Physics Letters* **123**, 261105. issn: 0003-6951, 1077-3118. <https://pubs.aip.org/apl/article/123/26/261105/2931557/Spectral-routers-for-snapshot-multispectral> (2026) (Dec. 2023).

ADJOINT-BASED OPTIMIZATION AND INVERSE DESIGN

2.1 Introduction

The design of photonic devices has historically proceeded in a forward direction, where a geometry is proposed based on physical intuition or parametric studies, its electromagnetic response is determined, and the result is evaluated against a set of performance specifications. If the performance is inadequate, the geometry is further modified and the cycle repeats. When the number of free parameters is small (for instance, sweeping the height and width of a cylindrical pillar to identify a favorable phase and amplitude response), this approach is straightforward. However, it is nearly impossible to design the three-dimensional metaoptic devices described in the subsequent chapters of this thesis in this way. Their design spaces consist of volumetric permittivity distributions discretized across tens of thousands of independent voxels. In this regime, physical intuition is insufficient to identify permittivity distributions that satisfy multiple complex optical functionalities across several different degrees of freedom.

Inverse design addresses this in the opposite direction. The designer specifies a desired behavior through an objective function, or a figure of merit (FoM), and then searches the design space for the permittivity distribution that optimizes it. Many different optimization algorithms exist to guide the inverse design search, but they can be separated into two main groups: gradient-free and gradient-based. We will begin by talking about the latter. In a gradient-based approach, the search is guided by the gradient of the FoM with respect to every design parameter. In the case of photonic inverse design, the adjoint method allows us to efficiently compute the gradient everywhere in the design from exactly two simulations: one forward simulation driven by the physical source, and one adjoint simulation driven by a source derived from the FoM. The two simulations are independent of the number of design parameters, making gradient-based inverse design practical for the large-scale photonic devices considered here.

This chapter develops the theory and implementation of adjoint-based inverse design for linear electromagnetic systems. The optimization problem is stated precisely in Section 2.2, and the adjoint gradient is derived in Section 2.3, with a physical

interpretation of the gradient in terms of electromagnetic reciprocity provided in Section 2.3.3. Section 2.4 discusses figures of merit suited to the device classes designed in this thesis, and Section 2.5 covers the logistics of optimizing devices with multiple simultaneous objectives. Finally, Section 2.6 addresses the fabricability constraints necessary to produce physically realizable devices.

2.2 Optimization Problem Formulation

Consider a computational domain $\mathcal{P}$ containing a design region $\mathcal{D}$. At each point $\mathbf{r} \in \mathcal{D}$, the relative permittivity $\varepsilon(\mathbf{r})$ is a free parameter constrained to the interval $[\varepsilon_{\min}, \varepsilon_{\max}]$, corresponding to the two materials comprising the final device. Outside the design region the permittivity is fixed. The goal of the optimization is to find the permittivity distribution $\varepsilon(\mathbf{r})$ that maximizes a scalar figure of merit f . In practice, $\varepsilon(\mathbf{r})$ is parameterized through a continuous density variable $\rho(\mathbf{r}) \in [0, 1]$, so that the optimization takes the composed form $f(\mathbf{E}(\varepsilon(\rho)))$; the density parameterization is detailed in Section 2.5.7.

The electromagnetic response of the domain is governed by the time-harmonic Maxwell curl equations, which for a medium with a time dependence of $e^{-i\omega t}$ take the form

$$\nabla \times \mathbf{H} + i\omega\varepsilon(\mathbf{r}) \mathbf{E} = \mathbf{J}, \quad (2.1)$$

$$\nabla \times \mathbf{E} - i\omega\mu_0 \mathbf{H} = \mathbf{K}, \quad (2.2)$$

here $\mathbf{J}$ and $\mathbf{K}$ are the electric and magnetic current source densities, respectively. Throughout this chapter the medium (which in this thesis is usually dielectric) is taken to be non-magnetic ($\mu = \mu_0$) and free of magnetic sources ($\mathbf{K} = \mathbf{0}$).

We define the residuals associated with these equations as

$$\mathbf{R}_E(\mathbf{E}, \mathbf{H}, \varepsilon) \equiv \nabla \times \mathbf{H} + i\omega\varepsilon \mathbf{E} - \mathbf{J} = \mathbf{0}, \quad (2.3)$$

$$\mathbf{R}_H(\mathbf{E}, \mathbf{H}) \equiv \nabla \times \mathbf{E} - i\omega\mu_0 \mathbf{H} = \mathbf{0}. \quad (2.4)$$

The figure of merit f depends on the electromagnetic fields and, in general, on the permittivity distribution:

$$f = f(\mathbf{E}, \mathbf{H}, \mathbf{E}^*, \mathbf{H}^*, \varepsilon). \quad (2.5)$$

The fields $\mathbf{E}$ and $\mathbf{H}$ depend implicitly on ε through Maxwell's equations, so the total derivative $df/d\varepsilon$ must account for both the explicit and the implicit dependences. Computing this derivative naively—by perturbing ε at each voxel individually—requires one full electromagnetic solve per voxel, scaling as $O(N)$, N being the

number of design parameters, whereas the adjoint method reduces this computation to $O(1)$.

2.3 The Adjoint Method

To compute $df/d\varepsilon(\mathbf{r})$ at every design voxel simultaneously from two electromagnetic simulations rather than N , we introduce adjoint fields $\mathbf{v}_E$ and $\mathbf{v}_H$ determined by an auxiliary problem whose sources are derived from the figure of merit.

Consider a perturbation $\varepsilon \rightarrow \varepsilon + \delta\varepsilon$. Because Maxwell's equations must hold both before and after the perturbation, the induced field changes $\delta\mathbf{E}$ and $\delta\mathbf{H}$ must satisfy the perturbed Ampère law,

$$\nabla \times (\mathbf{H} + \delta\mathbf{H}) + i\omega(\varepsilon + \delta\varepsilon)(\mathbf{E} + \delta\mathbf{E}) = \mathbf{J}. \quad (2.6)$$

Subtracting the unperturbed equation (2.1) and discarding the second-order term $i\omega \delta\varepsilon \delta\mathbf{E}$ yields

$$\nabla \times \delta\mathbf{H} + i\omega\varepsilon \delta\mathbf{E} = -i\omega \delta\varepsilon \mathbf{E}. \quad (2.7)$$

Faraday's law (2.2) contains no ε , so its linearized form is simply

$$\nabla \times \delta\mathbf{E} - i\omega\mu_0 \delta\mathbf{H} = \mathbf{0}. \quad (2.8)$$

Because f depends on the fields through both $\mathbf{E}$ and its complex conjugate $\mathbf{E}^*$ (and likewise for $\mathbf{H}$), the first-order Taylor expansion is

$$\delta f = \int_{\mathcal{P}} \left[\frac{\partial f}{\partial \mathbf{E}} \cdot \delta\mathbf{E} + \frac{\partial f}{\partial \mathbf{E}^*} \cdot \delta\mathbf{E}^* + \frac{\partial f}{\partial \mathbf{H}} \cdot \delta\mathbf{H} + \frac{\partial f}{\partial \mathbf{H}^*} \cdot \delta\mathbf{H}^* \right] d\mathbf{r} + \frac{\partial f}{\partial \varepsilon} \delta\varepsilon. \quad (2.9)$$

For a real-valued f , the $\partial f/\partial \mathbf{E}$ term is the complex conjugate of $\partial f/\partial \mathbf{E}^*$. Dropping the explicit $\partial f/\partial \varepsilon$ term for simplicity (we will add it back in the end), we write

$$\delta f = 2 \operatorname{Re} \int_{\mathcal{P}} \left[\frac{\partial f}{\partial \mathbf{E}^*} \cdot \delta\mathbf{E}^* + \frac{\partial f}{\partial \mathbf{H}^*} \cdot \delta\mathbf{H}^* \right] d\mathbf{r}. \quad (2.10)$$

Evaluating this naively requires solving the linearized system (2.7)–(2.8) once for every voxel perturbation $\delta\varepsilon(\mathbf{r})$. In the adjoint method, we can eliminate this by expressing δf directly in terms of $\delta\varepsilon$ and the unperturbed forward fields, at the cost of solving for a single, additional field.

2.3.1 The Adjoint Equations

The $\delta\mathbf{E}$ and $\delta\mathbf{H}$ dependence in Eq. (2.10) is eliminated by forming inner products of $\mathbf{v}_E$ and $\mathbf{v}_H$ with the conjugates of Eqs. (2.7)–(2.8).

Because Eq. (2.10) involves $\delta \mathbf{E}^*$ and $\delta \mathbf{H}^*$, we first conjugate Eqs. (2.7)–(2.8). For real ε and real perturbation $\delta \varepsilon$ this gives

$$\nabla \times \delta \mathbf{H}^* - i\omega \varepsilon \delta \mathbf{E}^* = i\omega \delta \varepsilon \mathbf{E}^*, \quad (2.11)$$

$$\nabla \times \delta \mathbf{E}^* + i\omega \mu_0 \delta \mathbf{H}^* = \mathbf{0}. \quad (2.12)$$

We then form inner products of $\mathbf{v}_E$ and $\mathbf{v}_H$ with Eqs. (2.11)–(2.12) respectively, integrating over $\mathcal{P}$:

$$\int_{\mathcal{P}} \mathbf{v}_E \cdot (\nabla \times \delta \mathbf{H}^* - i\omega \varepsilon \delta \mathbf{E}^*) d\mathbf{r} = i\omega \int_{\mathcal{P}} \mathbf{v}_E \cdot \delta \varepsilon \mathbf{E}^* d\mathbf{r}, \quad (2.13)$$

$$\int_{\mathcal{P}} \mathbf{v}_H \cdot (\nabla \times \delta \mathbf{E}^* + i\omega \mu_0 \delta \mathbf{H}^*) d\mathbf{r} = 0. \quad (2.14)$$

Integrating by parts for the curl operator,

$$\int_{\mathcal{P}} \mathbf{A} \cdot (\nabla \times \mathbf{B}) d\mathbf{r} = \int_{\mathcal{P}} \mathbf{B} \cdot (\nabla \times \mathbf{A}) d\mathbf{r} - \oint_{\partial \mathcal{P}} (\mathbf{A} \times \mathbf{B}) \cdot \hat{\mathbf{n}} dS. \quad (2.15)$$

The boundary term $(\mathbf{A} \times \mathbf{B}) \cdot \hat{\mathbf{n}} dS$ generally vanishes under periodic or PML boundary conditions, which allows us to re-express (2.14) in terms of the curl of $\mathbf{v}_E$ and $\mathbf{v}_H$:

$$\int_{\mathcal{P}} \delta \mathbf{H}^* \cdot (\nabla \times \mathbf{v}_E) d\mathbf{r} - i\omega \varepsilon \int_{\mathcal{P}} \mathbf{v}_E \cdot \delta \mathbf{E}^* d\mathbf{r} = i\omega \int_{\mathcal{P}} \mathbf{v}_E \cdot \delta \varepsilon \mathbf{E}^* d\mathbf{r}, \quad (2.16)$$

$$\int_{\mathcal{P}} \delta \mathbf{E}^* \cdot (\nabla \times \mathbf{v}_H) d\mathbf{r} + i\omega \mu_0 \int_{\mathcal{P}} \mathbf{v}_H \cdot \delta \mathbf{H}^* d\mathbf{r} = 0. \quad (2.17)$$

Adding Eqs. (2.16) and (2.17) and collecting by $\delta \mathbf{E}^*$ and $\delta \mathbf{H}^*$,

$$\int_{\mathcal{P}} \delta \mathbf{E}^* \cdot (\nabla \times \mathbf{v}_H - i\omega \varepsilon \mathbf{v}_E) d\mathbf{r} + \int_{\mathcal{P}} \delta \mathbf{H}^* \cdot (\nabla \times \mathbf{v}_E + i\omega \mu_0 \mathbf{v}_H) d\mathbf{r} = i\omega \int_{\mathcal{P}} \mathbf{v}_E \cdot \delta \varepsilon \mathbf{E}^* d\mathbf{r}. \quad (2.18)$$

The right-hand side depends only on $\delta \varepsilon$. We can choose the forms of $\mathbf{v}_E$ and $\mathbf{v}_H$ such that the operator expressions multiplying $\delta \mathbf{E}^*$ and $\delta \mathbf{H}^*$ on the left are chosen to equal the corresponding partial derivatives of f in Eq. (2.10).

Setting the coefficient of $\delta \mathbf{E}^*$ in Eq. (2.18) equal to the corresponding derivative of f yields

$$\nabla \times \mathbf{v}_H - i\omega \varepsilon \mathbf{v}_E = -\frac{\partial f}{\partial \mathbf{E}}, \quad (2.19)$$

and setting the coefficient of $\delta \mathbf{H}^*$ yields

$$\nabla \times \mathbf{v}_E + i\omega \mu_0 \mathbf{v}_H = -\frac{\partial f}{\partial \mathbf{H}}. \quad (2.20)$$

Substituting accordingly,

$$-\int_{\mathcal{P}} \left[\frac{\partial f}{\partial \mathbf{E}} \cdot \delta \mathbf{E}^* + \frac{\partial f}{\partial \mathbf{H}} \cdot \delta \mathbf{H}^* \right] d\mathbf{r} = i\omega \int_{\mathcal{P}} \mathbf{v}_E \cdot \delta \varepsilon \mathbf{E}^* d\mathbf{r}. \quad (2.21)$$

Comparing with Eq. (2.10) and using the identity $\text{Re}(-iz) = \text{Im}(z)$ gives

$$\delta f = 2 \text{Re} \left[-i\omega \int_{\mathcal{P}} \mathbf{v}_E \cdot \delta \varepsilon \mathbf{E}^* d\mathbf{r} \right] = 2\omega \text{Im} \int_{\mathcal{P}} \mathbf{v}_E(\mathbf{r}) \cdot \mathbf{E}^*(\mathbf{r}) \delta \varepsilon(\mathbf{r}) d\mathbf{r}, \quad (2.22)$$

The design gradient at each voxel is therefore

$$\frac{df}{d\varepsilon(\mathbf{r})} = 2\omega \text{Im} [\mathbf{v}_E(\mathbf{r}) \cdot \mathbf{E}^*(\mathbf{r})] = 2 \text{Re} [-i\omega \mathbf{v}_E(\mathbf{r}) \cdot \mathbf{E}^*(\mathbf{r})], \quad (2.23)$$

a pointwise product of the forward electric field and adjoint electric field that requires no additional solves beyond the two already performed.

In reciprocal media (where $\varepsilon = \varepsilon^T$ and $\mu = \mu^T$ as tensors, which holds for all materials considered in this thesis), the adjoint operator equals the transpose of the forward Maxwell operator. Equations (2.19)–(2.20) therefore have the same mathematical structure as the forward Maxwell Eqs. (2.1)–(2.2), differing only in their source terms. The adjoint fields can consequently be computed by the same electromagnetic solver used for the forward simulation, with modified source distributions:

$$\mathbf{J}_{\text{adj}} = \frac{\partial f}{\partial \mathbf{E}^*}, \quad \mathbf{K}_{\text{adj}} = \frac{1}{i\omega\mu_0} \frac{\partial f}{\partial \mathbf{H}^*}. \quad (2.24)$$

These adjoint sources are determined entirely by the FoM and the forward fields recorded during the first simulation.

2.3.2 Gradient Computation

In the case where f depends on ε directly, there is an additional contribution to δf , and the total gradient is

$$\frac{df}{d\varepsilon(\mathbf{r})} = \left. \frac{\partial f}{\partial \varepsilon} \right|_{\mathbf{E}, \mathbf{H}} + 2 \text{Re} [-i\omega \mathbf{v}_E(\mathbf{r}) \cdot \mathbf{E}(\mathbf{r})]. \quad (2.25)$$

For the common case in which the FoM has no explicit permittivity dependence, the first term vanishes and the gradient reduces to the pointwise expression

$$\boxed{\frac{df}{d\varepsilon(\mathbf{r})} = 2 \text{Re} [-i\omega \mathbf{v}_E(\mathbf{r}) \cdot \mathbf{E}(\mathbf{r})]}. \quad (2.26)$$

This is the central result of the adjoint method. The gradient at every point in the design region—across all N voxels simultaneously—is obtained from the pointwise

inner product of two field distributions: the forward field $\mathbf{E}$ and the adjoint field $\mathbf{v}_E$, each requiring exactly one electromagnetic simulation.

In the case where the design parameter is not ε directly, but instead some continuous density variable ρ mapped to permittivity, for example through the linear interpolation $\varepsilon(\rho) = \varepsilon_{\min} + \rho(\varepsilon_{\max} - \varepsilon_{\min})$, the chain rule extends Eq. (2.26) to

$$\frac{df}{d\rho(\mathbf{r})} = 2 \operatorname{Re}[-i\omega (\varepsilon_{\max} - \varepsilon_{\min}) \mathbf{v}_E(\mathbf{r}) \cdot \mathbf{E}(\mathbf{r})]. \quad (2.27)$$

2.3.3 Physical Interpretation

The form of Eq. (2.26) has an intuitive electromagnetic interpretation rooted in the principle of reciprocity. Suppose the permittivity at a point $\mathbf{r}$ in the design region is increased by a small amount $\delta\varepsilon$. This perturbation introduces an additional polarization current $\delta\mathbf{J} = -i\omega \delta\varepsilon \mathbf{E}(\mathbf{r})$, which oscillates in phase with the local forward electric field. This induced current then radiates and modifies the fields wherever the FoM is sensitive to them.

To assess the effect of this induced current on the FoM, one could in principle run a separate simulation with $\delta\mathbf{J}$ as a source and observe the resulting change at the output. Done serially for every voxel, this approach recovers the finite-difference cost of $\mathcal{O}(N)$ simulations. However, electromagnetic reciprocity states that the field produced at a point $\mathbf{r}_0$ by a source at $\mathbf{r}$ is equal to the field produced at $\mathbf{r}$ by an equivalent source at $\mathbf{r}_0$. Therefore, instead of locating a source at every design voxel in turn and observing the output, one can add a source at the output location once and observe the resulting fields at every design voxel simultaneously.

The adjoint field $\mathbf{v}_E(\mathbf{r})$ is the field distribution produced throughout the design region when the FoM-derived source (the adjoint source, located at the output) is excited. The gradient in Eq. (2.26) is the overlap of the adjoint field (which carries information about how a source at $\mathbf{r}$ would affect the output) with the forward field (which carries information about the phase and amplitude of the induced dipole if $\varepsilon(\mathbf{r})$ were increased). Voxels where the forward and adjoint fields are in phase and of large amplitude are those where increasing permittivity will most improve the FoM; voxels where they are out of phase will benefit from decreasing permittivity. The result is a gradient map over the entire design region, computed in two simulations.

2.4 Figures of Merit

The generality of Eq. (2.26) means that any differentiable scalar function of the electromagnetic fields can serve as a FoM, with an associated adjoint source given by

Eq. (2.24). The following figures of merit are relevant to the device classes designed in this thesis.

2.4.1 Intensity at a Target Location

The simplest FoM is the electromagnetic intensity at a specified point $\mathbf{r}_0$:

$$f = |\mathbf{E}(\mathbf{r}_0)|^2. \quad (2.28)$$

As described in [1], due to reciprocity the corresponding adjoint source is a point dipole at $\mathbf{r}_0$ with phase and amplitude proportional to the forward electric field recorded there. The adjoint simulation amounts to exciting the structure from the target point and measuring the resulting fields throughout the design region, which by reciprocity encodes the sensitivity of the focal intensity to permittivity changes everywhere else.

2.4.2 Mode Coupling Efficiency

For devices that must couple light into a specific spatial or polarization mode—for example, a particular diffraction order or output waveguide mode—the relevant FoM is the complex amplitude of the scattered field that overlaps with a target mode $(\mathbf{E}_b, \mathbf{H}_b)$. This amplitude is evaluated as a surface overlap integral over a closed surface Ω enclosing the design region:

$$s = \frac{1}{\mathcal{N}} \iint_{\Omega} \hat{\mathbf{n}} \cdot [\mathbf{E} \times \mathbf{H}_b - \mathbf{E}_b \times \mathbf{H}] d\Omega, \quad (2.29)$$

where $\hat{\mathbf{n}}$ is the outward surface normal and $\mathcal{N} = \iint_{\Omega} \hat{\mathbf{n}} \cdot [\mathbf{E}_f \times \mathbf{H}_b - \mathbf{E}_b \times \mathbf{H}_f] d\Omega$ normalizes $|s|^2 \in [0, 1]$ to represent a coupling efficiency. The target mode $(\mathbf{E}_b, \mathbf{H}_b)$ is the desired output mode; it may differ from the input mode in both polarization and spatial profile. The adjoint source is then analogous² to the backward mode $(\mathbf{E}_b, \mathbf{H}_b)$ injected from the surface Ω .

This FoM is well-suited for optimizing focusing devices (where the target mode is a beam focused to the desired point), polarization-selective devices (where the target mode carries a specific polarization state), and spectrally demultiplexing devices (where separate target modes are defined at each wavelength).

2.4.3 Multi-Port Transmission

When the device must simultaneously route power to multiple output channels—indexed by m with associated mode amplitudes s_m —the FoM is a function of the

full set of complex amplitudes:

$$f = f(\{s_m\}, \{s_m^*\}). \quad (2.30)$$

The number of adjoint simulations required depends on whether the output ports share a common forward source. When they do—for instance, a power splitter in which a single input frequency is routed to two spatial output channels—the linearity of Eq. (2.24) allows the adjoint sources to be superposed into a single simulation:

$$(\mathbf{J}_{\text{adj}}, \mathbf{K}_{\text{adj}}) = \sum_m \frac{1}{N_m} \frac{\partial f}{\partial s_m} (\mathbf{E}_{b,m}, \mathbf{H}_{b,m}). \quad (2.31)$$

In this case only one forward and one adjoint simulation are required regardless of the number of output ports.

When output channels are driven by distinct forward sources—as is the case for a spectral demultiplexer in which each color channel is excited at a different wavelength—the forward fields differ between channels and the superposition in Eq. (2.31) is no longer valid. Each channel then requires its own forward and adjoint simulation pair, and the total number of simulations per iteration scales as $2M$, where M is the number of channels. The resulting per-channel gradients are combined using the weighted average described in Section 2.33.

2.5 Multi-Objective Optimization

Many inverse-design problems are not naturally single-objective problems. For example, a broadband color-sorting device must simultaneously route several wavelengths or wavelength bands to different target locations. Improving the performance at one wavelength may degrade the performance at another wavelength, so there is generally no single design that independently maximizes every objective. Instead, the designer must decide how the different objectives should be compared, balanced, or constrained.

Let the design variable be denoted by $\boldsymbol{\rho} \in \mathcal{D}$, where $\mathcal{D}$ is the feasible design space, and let the individual objectives be $f_j(\boldsymbol{\rho})$, $j = 1, \dots, N$. One possible design problem with multiple competing objectives is (Section 2.4.1) the intensity delivered to the target point at wavelength λ_j ,

$$f_j(\boldsymbol{\rho}) = |\mathbf{E}(\mathbf{r}_0, \lambda_j; \boldsymbol{\rho})|^2 \quad (2.32)$$

The most direct method is to define the total FoM as a weighted sum of the individual figures of merit:

$$f = \sum_{\lambda_j} w_j f_j(\boldsymbol{\rho}) \quad (2.33)$$

Each wavelength requires its own forward and adjoint simulation pair, but the resulting gradients are combined linearly according to the weights w_j as follows:

$$g(\boldsymbol{\rho}) = \sum_j w_j \frac{\partial f_j}{\partial \boldsymbol{\varepsilon}}(\boldsymbol{\rho}), \quad (2.34)$$

A natural choice is to assign greater weight to objectives that are currently underperforming. The following performance-based weighting scheme³ achieves this:

$$w_j = \max \left(\frac{2}{N_f} - \frac{f_j^2}{\sum_k f_k^2}, 0 \right), \quad (2.35)$$

FoMs that are performing well (high f_j) receive reduced weight, while poorly-performing ones dominate the update. The max operation ensures weights are never negative, which would otherwise cause a high-performing FoM to be actively degraded. Over the course of the optimization this scheme drives all objectives toward parity, preventing any single easy-to-optimize FoM from dominating at the expense of others.

Alternatively, we can choose the figure of merit as the product of the single-frequency intensities,

$$f = \prod_{\lambda_j} f_j(\boldsymbol{\rho}) \quad (2.36)$$

rather than their weighted sum. If any of the intensities is near zero, the entire product collapses. As such, this formulation implicitly penalizes severe underperformance at any single wavelength, encouraging balanced solutions without requiring explicit weight tuning. The gradient of f with respect to the design parameters follows from the product rule and remains expressible in terms of the same forward and adjoint fields computed per wavelength, though the per-wavelength contributions are now scaled by the product of all other terms rather than by fixed weights.

2.5.1 Pareto fronts

To discuss the other strategies, we first introduce the concept of the Pareto front. Instead of combining the objectives into a single quantity, here we consider the set

of designs for which no objective can be improved without sacrificing at least one other objective. This set is called the Pareto-optimal set in design space, and its image in objective space is called the Pareto front⁴.

For a maximization problem, a design $\boldsymbol{\rho}_a$ is said to dominate another design $\boldsymbol{\rho}_b$ if

$$f_j(\boldsymbol{\rho}_a) \geq f_j(\boldsymbol{\rho}_b) \quad \text{for all } j, \quad (2.37)$$

with $f_j(\boldsymbol{\rho}_a) > f_j(\boldsymbol{\rho}_b)$ for at least one f_j . A design is Pareto optimal if there is no other feasible design that dominates it. Thus, for a maximization problem, the Pareto-optimal set is

$$\mathcal{P} = \left\{ \boldsymbol{\rho}^* \in \mathcal{D} \mid \nexists \boldsymbol{\rho} \in \mathcal{D} : \begin{array}{l} f_j(\boldsymbol{\rho}) \geq f_j(\boldsymbol{\rho}^*) \quad \forall j, \\ f_k(\boldsymbol{\rho}) > f_k(\boldsymbol{\rho}^*) \quad \text{for at least one } k \end{array} \right\}. \quad (2.38)$$

By this definition, in a two-wavelength optimization, one design $\boldsymbol{\rho}_a$ may give high at one wavelength and moderate efficiency at the other, while the opposite may be true for another design $\boldsymbol{\rho}_b$. If neither design is better for both wavelength channels, both belong to the Pareto front. Here, an additional preference (e.g. equalized channel efficiency, maximum total efficiency, minimum acceptable efficiency, etc.) is usually responsible for the final decision.

In practice, the Pareto front is usually approximated rather than known exactly. A common approach is to solve the optimization problems with many different weight vectors. For example, one may maximize the simple weighted sum, Eq. 2.33, for a sweep of weight vectors $\mathbf{w}$, with normalization $\sum_j w_j = 1$. Each weight vector $\mathbf{w}$ expresses a different preference among the j objectives and gives one candidate point on the Pareto front. This method is simple and works well when the attainable objective space has a convex Pareto front. However, it cannot recover Pareto-optimal points that lie on non-convex, or unsupported, portions of the front^{5,6}. Geometrically, a weighted sum finds a point where a hyperplane with normal vector $\mathbf{w}$ touches the attainable objective set. If a portion of the Pareto front is concave, no such supporting hyperplane touches that interior concave region, and the corresponding Pareto-optimal designs are missed. Gradient-free population methods are often used when the goal is to approximate the full Pareto front rather than to optimize a single scalar objective. Two examples are NSGA-II⁷, and its successor NSGA-III⁸, which maintain a population of designs and use a diversity metric to keep the sampled solutions spread out along the Pareto front.

2.5.2 Minimax and maxmin formulations

A direct way to convert a multi-objective optimization problem into a scalar problem is to optimize the worst-performing objective. Suppose that a design variable $\boldsymbol{\rho} \in \mathcal{D}$ produces N scalar performance values

$$f_j(\boldsymbol{\rho}), \quad j = 1, \dots, N. \quad (2.39)$$

If each f_j is a loss, error, or cost to be minimized, then the pure minimax formulation is

$$\min_{\boldsymbol{\rho} \in \mathcal{D}} \max_{1 \leq j \leq N} f_j(\boldsymbol{\rho}). \quad (2.40)$$

The scalar objective is determined only by the largest component of the vector $\mathbf{f}(\boldsymbol{\rho}) = (f_1(\boldsymbol{\rho}), \dots, f_N(\boldsymbol{\rho}))$. Thus, a design is only as good as its worst channel. When the maximizer is unique and the active objective is differentiable, the gradient of the minimax objective is simply $\nabla_{\boldsymbol{\rho}} f_{j^*}(\boldsymbol{\rho})$. When multiple objectives are tied for the maximum, the objective is generally nonsmooth and nondifferentiable.⁹

If, instead, each f_j is a performance, efficiency, intensity, or figure of merit to be maximized, then the analogous worst-channel formulation is a maxmin problem,

$$\max_{\boldsymbol{\rho} \in \mathcal{D}} \min_{1 \leq j \leq N} f_j(\boldsymbol{\rho}). \quad (2.41)$$

Here the scalar objective is the smallest achieved performance across all channels. For broadband focusing, for example, if $f_j(\boldsymbol{\rho}) = |\mathbf{E}(\mathbf{r}_0, \lambda_j; \boldsymbol{\rho})|^2$, then Eq. (2.41) seeks the design whose weakest wavelength-channel intensity is as large as possible.

2.5.3 Weighted Chebyshev scalarization

Weighted Chebyshev scalarization is a structured minimax formulation applied not directly to the raw objectives, but to their weighted deviations from a reference point in objective space. Let $\mathbf{z}^* = (z_1^*, \dots, z_N^*)$ denote a reference vector, often chosen as the ideal objective vector. For a minimization problem, z_j^* is commonly the best attainable value of the j -th objective considered independently, whereas for a maximization problem it may represent the desired or independently optimized performance. In a broadband focusing problem, for instance, one may define

$$f_j(\boldsymbol{\rho}) = |\mathbf{E}(\mathbf{r}_0, \lambda_j; \boldsymbol{\rho})|^2, \quad (2.42)$$

and compare this achieved intensity to a target value z_j^* . A deviation from the target may then be written as¹.

$$d_j(\boldsymbol{\rho}) = \left| f_j(\boldsymbol{\rho}) - z_j^* \right|. \quad (2.43)$$

The weighted Chebyshev scalar objective is

$$F_{\text{TC}}(\boldsymbol{\rho}) = \max_{1 \leq j \leq N} w_j d_j(\boldsymbol{\rho}), \quad w_j > 0, \quad (2.44)$$

and the optimization problem is

$$\min_{\boldsymbol{\rho} \in \mathcal{D}} F_{\text{TC}}(\boldsymbol{\rho}) \quad \text{or} \quad \min_{\boldsymbol{\rho} \in \mathcal{D}} \max_{1 \leq j \leq N} w_j \left| f_j(\boldsymbol{\rho}) - z_j^* \right|. \quad (2.45)$$

This is a minimax problem because it minimizes the largest weighted deviation from the desired reference point $\mathbf{z}^*$ in the N -dimensional objective space. A design with excellent performance at most wavelengths but very poor performance at one wavelength is therefore penalized strongly. With suitable positive weights, weighted Chebyshev scalarization can identify Pareto-optimal points that weighted-sum scalarization may miss, especially on non-convex parts of the Pareto front⁴.

The same nonsmoothness issue from pure minimax applies here as well. If two or more channels have the same largest weighted deviation $w_j d_j(\boldsymbol{\rho})$, known as a tie point, the maximum in Eq. 2.45 is not unique and becomes nondifferentiable. To bypass this we can use a smooth approximation such as the log-sum-exp function:

$$F_\mu(\boldsymbol{\rho}) = \frac{1}{\mu} \log \sum_{j=1}^N \exp [\mu w_j d_j(\boldsymbol{\rho})], \quad \mu > 0. \quad (2.46)$$

As $\mu \rightarrow \infty$, F_μ approaches F_{TC} . For finite μ , the gradient is

$$\nabla_{\boldsymbol{\rho}} F_\mu = \sum_{j=1}^N s_j \nabla_{\boldsymbol{\rho}} [w_j d_j(\boldsymbol{\rho})], \quad s_j = \frac{\exp [\mu w_j d_j(\boldsymbol{\rho})]}{\sum_{\ell=1}^N \exp [\mu w_\ell d_\ell(\boldsymbol{\rho})]}. \quad (2.47)$$

The coefficients s_j are softmax weights over the weighted deviations. Thus, the gradient is a smooth convex combination of the channel gradients, with the largest-error channels receiving the largest weights. In the limit of large μ , the gradient has the most effect on the worst channel, recovering the behavior of the nonsmooth Chebyshev objective; for smaller μ , the optimization behaves more like a smoothed average over channels.

¹As different objectives may differ greatly in expected relative magnitude, sometimes the deviation is written $d_j(\boldsymbol{\rho}) = \left| f_j(\boldsymbol{\rho}) - z_j^* \right| / |z_j^*|$

2.5.4 ε -constraint method

The ε -constraint method is used when there are performance floors that must be met by all channels (the floor may not be the same for each channel). Instead of combining all objectives into one weighted score, it chooses one objective as the primary objective and turns the remaining objectives into constraints^{4,10}. For example, if the first wavelength is chosen as the primary objective, one may write

$$\max_{\rho \in \mathcal{D}} f_1(\rho) \quad \text{subject to} \quad f_j(\rho) \geq \varepsilon_j, \quad j = 2, \dots, N. \quad (2.48)$$

Here, ε_j is the minimum acceptable performance of the j -th constrained objective. Thus the problem becomes finding the design, out of all designs that meet their required performance floors in the other channels, which maximizes the performance in the primary channel.

The method can also be used to trace the Pareto front. One solves Eq. (2.48) repeatedly while sweeping the values of ε_j over their feasible ranges. Each choice of thresholds defines a different slice of the feasible design space, and the optimum of that constrained problem gives a candidate Pareto point. Unlike the weighted-sum method, the ε -constraint method can generate points on non-convex parts of the Pareto front because the front is sampled by constraint boundaries rather than by supporting hyperplanes^{4,11}.

For adjoint-based optimization, the constrained problem is often converted into a sequence of differentiable unconstrained or simply constrained subproblems. This can be done using exterior penalty methods, barrier methods, sequential quadratic programming, or augmented Lagrangian methods. The augmented Lagrangian approach is particularly useful because it combines constraint penalties with estimates of the Lagrange multipliers, allowing the optimization to enforce constraints without requiring the penalty parameter to become excessively large.

2.5.5 Augmented Lagrangian method

The augmented Lagrangian method is a general method for solving constrained optimization problems, including the constrained subproblems produced by the ε -constraint method^{12,13}. In general, it is used to transform a constrained problem into a sequence of easier subproblems, while still encouraging the final design to satisfy the constraints.

Consider first a minimization problem with equality constraints,

$$\min_{\rho} F(\rho) \quad \text{subject to} \quad c_i(\rho) = 0, \quad i = 1, \dots, m. \quad (2.49)$$

A simple penalty-based method would minimize

$$F(\boldsymbol{\rho}) + \frac{\beta}{2} \sum_{i=1}^m c_i(\boldsymbol{\rho})^2, \quad (2.50)$$

where $\beta > 0$ is a penalty parameter. This penalizes constraint violation. However, to enforce the constraints accurately, a pure penalty method often requires β to become very large. Large penalty parameters can make the optimization problem ill-conditioned: the objective becomes extremely steep in directions that violate the constraints and comparatively flat in other directions.

The augmented Lagrangian method improves on this by adding a Lagrange multiplier term to the quadratic penalty:

$$\mathcal{L}_A(\boldsymbol{\rho}, \boldsymbol{\eta}; \beta) = F(\boldsymbol{\rho}) + \sum_{i=1}^m \eta_i c_i(\boldsymbol{\rho}) + \frac{\beta}{2} \sum_{i=1}^m c_i(\boldsymbol{\rho})^2. \quad (2.51)$$

Here, η_i is the current estimate of the Lagrange multiplier associated with the constraint $c_i(\boldsymbol{\rho}) = 0$. The method alternates between two steps. First, for fixed multipliers and fixed penalty parameter, it approximately minimizes $\mathcal{L}_A(\boldsymbol{\rho}, \boldsymbol{\eta}; \beta)$ with respect to the design variables. Second, it updates the multipliers according to

$$\eta_i^{(k+1)} = \eta_i^{(k)} + \beta^{(k)} c_i(\boldsymbol{\rho}^{(k+1)}). \quad (2.52)$$

That is, if the current design violates a particular constraint, the multiplier is adjusted so that the next subproblem penalizes violation of that constraint more strongly.

For inequality constraints, one may write the problem in the form

$$g_i(\boldsymbol{\rho}) \leq 0, \quad i = 1, \dots, m. \quad (2.53)$$

For the ε -constraint formulation in Eq. (2.48), a constraint $f_j(\boldsymbol{\rho}) \geq \varepsilon_j$ can be written as

$$g_j(\boldsymbol{\rho}) = \varepsilon_j - f_j(\boldsymbol{\rho}) \leq 0. \quad (2.54)$$

We can introduce a nonnegative slack variable¹⁴ $s_i \geq 0$ such that the inequality constraint $g_i(\boldsymbol{\rho}) \leq 0$ becomes the equality constraint $g_i(\boldsymbol{\rho}) + s_i = 0$. Applying the equality-constrained augmented Lagrangian to this constraint gives

$$\eta_i [g_i(\boldsymbol{\rho}) + s_i] + \frac{\beta}{2} [g_i(\boldsymbol{\rho}) + s_i]^2 \quad (2.55)$$

$$= \frac{\beta}{2} \left(g_i(\boldsymbol{\rho}) + s_i + \frac{\eta_i}{\beta} \right)^2 - \frac{\eta_i^2}{2\beta}. \quad (2.56)$$

This expression is minimal when $s_i = -g_i(\boldsymbol{\rho}) - \frac{\eta_i}{\beta}$; since we defined $s_i \geq 0$, the final expression for s_i is

$$\frac{1}{2\beta} \left[\max(0, \eta_i + \beta g_i(\boldsymbol{\rho}))^2 - \eta_i^2 \right].$$

Substituting back into the expression for the augmented Lagrangian, we obtain¹⁴

$$\mathcal{L}_A(\boldsymbol{\rho}, \boldsymbol{\eta}; \beta) = F(\boldsymbol{\rho}) + \frac{1}{2\beta} \sum_{i=1}^m \left[\max(0, \eta_i + \beta g_i(\boldsymbol{\rho}))^2 - \eta_i^2 \right], \quad (2.57)$$

If any of the constraints $g_i(\boldsymbol{\rho}) \leq 0$ are violated, then $g_i(\boldsymbol{\rho}) > 0$. We aim to minimize the Lagrangian $\mathcal{L}_A$, and so in order to penalize the constraint violation, the associated KKT multiplier must satisfy $\eta_i \geq 0$. As such, the multiplier updates for the inequality-constrained augmented Lagrangian are

$$\eta_i^{(k+1)} = \max \left(0, \eta_i^{(k)} + \beta^{(k)} g_i(\boldsymbol{\rho}^{(k+1)}) \right). \quad (2.58)$$

When $\eta_i + \beta g_i(\boldsymbol{\rho}) > 0$, the maximum operator is inactive, and the inequality augmented-Lagrangian term becomes

$$\frac{1}{2\beta} \left[(\eta_i + \beta g_i(\boldsymbol{\rho}))^2 - \eta_i^2 \right] = \eta_i g_i(\boldsymbol{\rho}) + \frac{\beta}{2} g_i(\boldsymbol{\rho})^2.$$

Thus, near an active constraint or when a constraint is violated, the term has the usual multiplier-plus-quadratic-penalty form. By contrast, when $\eta_i + \beta g_i(\boldsymbol{\rho}) \leq 0$, the maximum operator returns zero and the term becomes $-\eta_i^2/(2\beta)$, which is independent of $\boldsymbol{\rho}$. In other words, the constraint has no contribution to the gradient used for the design update.

The gradient of the augmented Lagrangian has the form

$$\nabla_{\boldsymbol{\rho}} \mathcal{L}_A = \nabla_{\boldsymbol{\rho}} F(\boldsymbol{\rho}) + \sum_{i=1}^m \max(0, \eta_i + \beta g_i(\boldsymbol{\rho})) \nabla_{\boldsymbol{\rho}} g_i(\boldsymbol{\rho}). \quad (2.59)$$

Thus, the same adjoint machinery used to compute gradients of optical objectives can be reused to compute gradients of constraint functions.

We end this section by considering a multi-wavelength intensity objective using the augmented Lagrangian method and a maxmin formulation. Introduce an auxiliary scalar variable t , which represents the worst-channel focal intensity. The maxmin problem can then be written as

$$\max_{\boldsymbol{\rho} \in \mathcal{D}, t} t \quad \text{subject to} \quad f_j(\boldsymbol{\rho}) \geq t, \quad j = 1, \dots, N. \quad (2.60)$$

where each wavelength is required to perform at least as well as t . Consequently, maximizing t maximizes the minimum focal intensity across all wavelengths. At each iteration, we can then make t the prescribed floor ε_j for every wavelength. The resulting augmented Lagrangian procedure is summarized in Algorithms 1 and 2.

Algorithm 1 Maxmin multi-wavelength focusing: problem setup

Require: Wavelengths $\{\lambda_j\}_{j=1}^N$, target focal point $\mathbf{r}_0$, design space $\mathcal{D}$, initial design $\boldsymbol{\rho}^{(0)}$, initial worst-channel variable $t^{(0)}$, initial multipliers $\eta_j^{(0)} \geq 0$, initial penalty parameter $\beta^{(0)} > 0$, penalty growth factor $\gamma > 1$, tolerances τ_{opt} and τ_{con}

- 1: Define the per-wavelength focusing objectives

$$f_j(\boldsymbol{\rho}) = |\mathbf{E}(\mathbf{r}_0, \lambda_j; \boldsymbol{\rho})|^2, \quad j = 1, \dots, N.$$

- 2: Introduce an auxiliary scalar t representing the worst-channel focal intensity.
- 3: Rewrite the maxmin problem as

$$\max_{\boldsymbol{\rho} \in \mathcal{D}, t} t \quad \text{subject to} \quad f_j(\boldsymbol{\rho}) \geq t, \quad j = 1, \dots, N.$$

- 4: Convert each wavelength constraint into the standard inequality form

$$g_j(\boldsymbol{\rho}, t) = t - f_j(\boldsymbol{\rho}) \leq 0, \quad j = 1, \dots, N.$$

- 5: Set $k \leftarrow 0$
-

2.5.6 Parallelization

A complete optimization iteration consists of simulating all unique forward and adjoint sources, computing the per-FoM gradients, combining them (e.g. via Eq. 2.34), and updating the device permittivity. For an optimization with M objectives requiring independent forward and adjoint sources, up to $2M$ electromagnetic simulations must be run per iteration.

These simulations are mutually independent and can be distributed across a computing cluster: each node is assigned a single simulation, runs it to completion, and returns the result. The cluster should be configured to prevent inter-node communication during the simulation stage, since inter-node communication is substantially slower than intra-node memory access.

A useful special case arises when multiple FoMs differ only in frequency. Time-domain solvers such as FDTD respond to a short broadband pulse that simultaneously excites the full spectral range of interest; all frequency components are

Algorithm 2 Augmented Lagrangian iteration for maxmin multi-wavelength focusing

Require: Objectives $f_j(\boldsymbol{\rho})$, constraints $g_j(\boldsymbol{\rho}, t) = t - f_j(\boldsymbol{\rho}) \leq 0$, initial values $\boldsymbol{\rho}^{(0)}$, $t^{(0)}$, $\boldsymbol{\eta}^{(0)}$, $\beta^{(0)}$, γ , τ_{opt} , and τ_{con}

1: **repeat**

2: Form the augmented Lagrangian minimization objective

$$\mathcal{L}_A(\boldsymbol{\rho}, t, \boldsymbol{\eta}^{(k)}; \beta^{(k)}) = -t + \frac{1}{2\beta^{(k)}} \sum_{j=1}^N \left[\max \left(0, \eta_j^{(k)} + \beta^{(k)} [t - f_j(\boldsymbol{\rho})] \right)^2 - \left(\eta_j^{(k)} \right)^2 \right].$$

3: Approximately minimize the augmented Lagrangian:

$$(\boldsymbol{\rho}^{(k+1)}, t^{(k+1)}) \approx \arg \min_{\boldsymbol{\rho} \in \mathcal{D}, t} \mathcal{L}_A(\boldsymbol{\rho}, t, \boldsymbol{\eta}^{(k)}; \beta^{(k)}),$$

using adjoint-gradient optimization for $\boldsymbol{\rho}$.

4: Evaluate the focal intensities

$$f_j(\boldsymbol{\rho}^{(k+1)}) = \left| \mathbf{E}(\mathbf{r}_0, \lambda_j; \boldsymbol{\rho}^{(k+1)}) \right|^2, \quad j = 1, \dots, N.$$

5: Evaluate the maxmin constraint violations

$$g_j(\boldsymbol{\rho}^{(k+1)}, t^{(k+1)}) = t^{(k+1)} - f_j(\boldsymbol{\rho}^{(k+1)}), \quad j = 1, \dots, N.$$

6: Update the multipliers:

$$\eta_j^{(k+1)} = \max \left(0, \eta_j^{(k)} + \beta^{(k)} g_j(\boldsymbol{\rho}^{(k+1)}, t^{(k+1)}) \right), \quad j = 1, \dots, N.$$

7: **if** the maximum constraint violation is not decreasing sufficiently **then**

8: Increase the penalty parameter:

$$\beta^{(k+1)} \leftarrow \gamma \beta^{(k)}.$$

9: **else**

10: Keep the penalty parameter fixed:

$$\beta^{(k+1)} \leftarrow \beta^{(k)}.$$

11: **end if**

12: Set $k \leftarrow k + 1$

13: **until** the optimality residual is below τ_{opt} and

$$\max_{j=1, \dots, N} \max \left(0, t^{(k)} - f_j(\boldsymbol{\rho}^{(k)}) \right) \leq \tau_{\text{con}}$$

14: **return** $\boldsymbol{\rho}^* \leftarrow \boldsymbol{\rho}^{(k)}$

resolved in a single simulation via Fourier transformation of the time-domain response. Broadening the spectral coverage of such an objective therefore adds negligible computational cost in FDTD. When objectives differ in polarization, angle of incidence, or spatial mode profile, they cannot be batched this way and require independent simulations.

2.5.7 Density Parameterization

The design variable in the optimization is the permittivity $\varepsilon(\mathbf{r})$ at each voxel in the design region $\mathcal{D}$. For gradient-based optimization it is convenient to work with a continuous density variable $\rho(\mathbf{r}) \in [0, 1]$ that is linearly interpolated to permittivity:

$$\varepsilon(\mathbf{r}) = \varepsilon_{\min} + \rho(\mathbf{r})(\varepsilon_{\max} - \varepsilon_{\min}), \quad (2.61)$$

with $\rho = 0$ corresponding to one material (typically the surrounding medium) and $\rho = 1$ to the other (the patterned material). The design region is typically initialized with a spatially random density drawn from a Gaussian distribution centered near $\rho = 0.5$, then blurred to remove fine-scale variations that could trap the optimization in poor local minima^{15,16}.

2.5.8 Update Algorithms

Once the gradient $\partial f / \partial \rho$ is known at every voxel, the density is updated to improve the FoM. The simplest update rule is gradient ascent:

$$\rho^{(n+1)}(\mathbf{r}) = \rho^{(n)}(\mathbf{r}) + \delta \left. \frac{\partial f}{\partial \rho} \right|^{(n)}, \quad (2.62)$$

where δ is a step size chosen to be small enough that the local linearization implicit in the gradient remains valid. Gradient ascent is effective for single-objective optimizations but can exhibit slow convergence and sensitivity to the step size.

For multi-objective optimizations, the widely-used Adam algorithm¹⁷ provides substantially more robust convergence. Adam maintains exponential moving averages of both the gradient and its element-wise square, referred to as the first and second moments:

$$\mathbf{m}^{(n)} = \beta_1 \mathbf{m}^{(n-1)} + (1 - \beta_1) \nabla f^{(n)}, \quad (2.63)$$

$$\mathbf{v}^{(n)} = \beta_2 \mathbf{v}^{(n-1)} + (1 - \beta_2) \left(\nabla f^{(n)} \right)^2. \quad (2.64)$$

After bias correction $\hat{\mathbf{m}}^{(n)} = \mathbf{m}^{(n)} / (1 - \beta_1^n)$ and $\hat{\mathbf{v}}^{(n)} = \mathbf{v}^{(n)} / (1 - \beta_2^n)$, the density is updated as

$$\rho^{(n+1)} = \rho^{(n)} + \alpha \frac{\hat{\mathbf{m}}^{(n)}}{\sqrt{\hat{\mathbf{v}}^{(n)} + \eta}}, \quad (2.65)$$

where α and η are hyperparameters. The values $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\alpha = 10^{-3}$, and $\eta = 10^{-8}$ work well in practice. These adaptive step sizes allow the Adam optimizer to make progress across objectives with disparate gradient magnitudes, which is common in multi-objective photonic optimization.

Projection and continuation. An alternative binarization function to the smoothed Heaviside of Section 2.6 is the tanh projection

$$b(\rho) = \frac{\tanh(\beta\eta) + \tanh(\beta(\rho - \eta))}{\tanh(\beta\eta) + \tanh(\beta(1 - \eta))}, \quad (2.66)$$

where β controls the sharpness and η is the threshold. In order to balance landscape smoothness early in the optimization against the full desired binarization at convergence, a continuation scheme might increase β exponentially, e.g.

$$\beta(n) = 2^{7(n-1)/N}, \quad (2.67)$$

ranging from $\beta = 1$ at iteration $n = 1$ to $\beta = 128$ at $n = N$.

Covariance Matrix Adaptation Evolution Strategy. Gradient-based methods are local by definition, and eventually converge to a stationary point of f in the neighborhood of the initialization. As the number of design parameters grows, the objective function typically becomes highly multimodal. In this case, a gradient-free global search performed prior to gradient refinement can identify a favorable region of design space and reduce the risk of converging to a poor local optimum. The *Covariance Matrix Adaptation Evolution Strategy* (CMA-ES) is a derivative-free algorithm that maintains a multivariate Gaussian distribution over the density space with mean $\mathbf{m} \in \mathbb{R}^N$ and covariance matrix $\mathbf{C} \in \mathbb{R}^{N \times N}$, and iteratively updates this distribution based on the objective values of sampled candidate designs.

At each generation g , CMA-ES draws a population of λ candidate densities:

$$\rho_k^{(g)} \sim \mathcal{N}\left(\mathbf{m}^{(g)}, \left(\sigma^{(g)}\right)^2 \mathbf{C}^{(g)}\right), \quad k = 1, \dots, \lambda, \quad (2.68)$$

where $\sigma^{(g)}$ is a global step size. The top $\mu < \lambda$ candidates are selected by objective value and the mean is updated as their weighted average:

$$\mathbf{m}^{(g+1)} = \sum_{i=1}^{\mu} w_i \rho_{i:\lambda}^{(g)}, \quad \sum_i w_i = 1. \quad (2.69)$$

The covariance matrix is updated via a rank-one term that accumulates the running evolution path $\mathbf{p}_c$,

$$\mathbf{C}^{(g+1)} \leftarrow (1 - c_1)\mathbf{C}^{(g)} + c_1 \mathbf{p}_c^{(g+1)} \left(\mathbf{p}_c^{(g+1)} \right)^\top, \quad (2.70)$$

and a rank- μ term that incorporates the full selected population:

$$\mathbf{C}^{(g+1)} \leftarrow (1 - c_\mu)\mathbf{C}^{(g)} + c_\mu \sum_{i=1}^{\mu} w_i \boldsymbol{\delta}_i^{(g)} \left(\boldsymbol{\delta}_i^{(g)} \right)^\top, \quad (2.71)$$

where $\boldsymbol{\delta}_i^{(g)} = (\rho_{i:\lambda}^{(g)} - \mathbf{m}^{(g)}) / \sigma^{(g)}$. The step size σ is adapted via cumulative step-size adaptation, increasing when progress is consistent and decreasing otherwise.

CMA-ES requires one objective evaluation per candidate per generation with no gradient computations, making it compatible with any electromagnetic solver. Additionally, its default population size $\lambda \approx 4 + \lfloor 3 \ln N \rfloor$ grows slowly with the number of design parameters N . The primary limitation is the $O(N^2)$ memory cost of the full covariance matrix, which becomes impractical for $N \gg 10^3$ without low-rank approximations.

Other gradient-free strategies. Particle swarm optimization maintains a population of candidate designs that move through the design space influenced by their own best-known position and the global best, providing a different balance between exploration and exploitation than CMA-ES. Bayesian optimization builds a probabilistic surrogate model of the objective function and samples new candidates at locations predicted to be most informative, and is most effective when each function evaluation is expensive and the number of design parameters is small. In high-dimensional inverse design, gradient-free methods are most useful for an initial, early-terminated global search, with the best result then used as initialization for a subsequent gradient-based refinement stage.

2.6 Fabricability Constraints

The continuous density $\rho(\mathbf{r})$ produced by the gradient-based update is not directly fabricable: physical devices must be composed of two discrete materials (binary), and must satisfy a minimum feature size dictated by the available lithographic process. The following techniques constrain the optimization to produce designs that respect these requirements.

2.6.1 Binarization

The optimization offers no inherent preference for $\rho \in \{0, 1\}$ and is equally satisfied by any intermediate “grayscale” density that maximizes f . However, grayscale permittivity distributions cannot be practically realized in a binary material system, which is typical of most fabrication capabilities at the time of writing. To drive ρ toward binary values over the course of the optimization, a smoothed Heaviside filter is applied to the density:

$$b(\rho) = \frac{1}{1 + \exp\left[-2k\left(\rho - \frac{1}{2}\right)\right]}, \quad (2.72)$$

where k controls the sharpness of the transition at $\rho = 1/2$. For small k , $b(\rho)$ is nearly linear and the filter has little effect on the optimization landscape. As k is increased over the course of the optimization, the filter progressively enforces that intermediate densities migrate toward zero or one. A schedule in which k increases exponentially with iteration number balances the competing needs of an unconstrained initial exploration against a fully binarized final structure.

Because $b(\rho)$ is differentiable for all ρ , its gradient $b'(\rho)$ can be back-propagated through the chain of transformations from the raw density to the FoM. If the total mapping from ρ to ε is $\varepsilon(\rho) = \varepsilon_{\min} + b(\rho)(\varepsilon_{\max} - \varepsilon_{\min})$, the gradient with respect to the raw density is

$$\frac{\partial f}{\partial \rho} = \frac{\partial f}{\partial \varepsilon} \cdot (\varepsilon_{\max} - \varepsilon_{\min}) \cdot b'(\rho). \quad (2.73)$$

This chain rule structure generalizes to any sequence of differentiable operations applied to ρ prior to the electromagnetic simulation, including the Gaussian blurring described below.

2.6.2 Minimum Feature Size

Nanofabrication processes impose a minimum resolvable feature size set by the lithography system. Features below this threshold—narrow material bridges, isolated islands, or sharp corners—will either fail to print or will deviate substantially from their designed dimensions, degrading device performance. Two complementary approaches are used to prevent such features from appearing in the optimized design.

The first is Gaussian blurring of the density gradient, implemented by convolving $\partial f / \partial \rho$ with a Gaussian kernel whose standard deviation equals the desired mini-

imum feature size. This suppresses high-spatial-frequency components of the gradient, preventing the optimizer from placing features smaller than the specified scale. Applied at every iteration, or at a fixed interval of iterations, this filter smooths the density distribution and removes fine-grained structure. However, it is not guaranteed to eliminate all sub-resolution features from the final binary pattern.

A more rigorous constraint can be imposed during the level-set phase of the optimization (Section 2.6.4). In the level-set representation, fabrication penalty functions based on the local gap size and radius of curvature of the device boundary can be evaluated analytically from the first and second spatial derivatives of the level-set function. These penalties can be added as constraints, guiding the boundary away from configurations that violate the minimum feature size requirement.

2.6.3 Binary-Constrained Update

Standard gradient-based approaches—including gradient ascent and Adam—can stagnate at non-binary solutions in high index-contrast material systems, where the electromagnetic FoM gradient does not naturally point in a direction that binarizes the device. This problem becomes more acute as the refractive index contrast increases: a gradient step that improves the FoM may simultaneously push intermediate-density voxels further from the material boundaries, worsening binarization. The Heaviside filter described above mitigates this tendency but cannot fully prevent it.

One approach that addresses this more directly reformulates the permittivity update at each iteration as a constrained sub-optimization that explicitly requires the device to become more binary by some minimum amount, while still maximizing the electromagnetic FoM gradient. By leveraging the Lagrange dual of this constrained problem, the N -dimensional linear program can be reduced to a one-dimensional nonlinear optimization in a single dual variable, making the sub-optimization computationally inexpensive. This technique and its derivation are described in detail in¹⁸, and the implementation used in this thesis follows that reference.

2.6.4 Level-Set Refinement

The continuous density optimization, even with binarization constraints, produces a permittivity distribution that must eventually be thresholded to discrete material values. This thresholding step can degrade performance; the magnitude of the degradation depends on how binary the device was at the time of thresholding. Devices with binarization levels above approximately 80% prior to thresholding generally retain most of their performance, while less binary devices can suffer substantial

degradation.

The thresholded device serves as an initial seed for a level-set optimization, in which the device boundary Γ is described implicitly as the zero-level-set contour of a signed-distance function $\phi(\mathbf{r})$. The boundary evolves according to

$$\phi^{(n+1)} = \phi^{(n)} + \Delta t v_n(\mathbf{r}), \quad (2.74)$$

where v_n is the normal velocity of the boundary, computed from the electromagnetic gradient combined with any fabrication penalty terms. Because the level-set method operates on a binary structure by construction, the optimized device remains exactly binary throughout this phase. The level-set phase recovers performance lost in the thresholding step and can further enforce minimum feature size constraints analytically through curvature-based penalties.

2.7 Summary

The adjoint method has been described as a convenient tool for gradient-based inverse design of linear electromagnetic devices. Starting from frequency-domain Maxwell's equations, the adjoint method produces the gradient of a scalar FoM with respect to the permittivity at every voxel in the design region using exactly two electromagnetic simulations, independent of the number of design parameters. This gradient drives a topology optimization that iteratively improves the device. During this, we can enforce fabricability constraints, such as binarization filters to ensure the final design consists of two discrete materials, and minimum feature size enforcement to ensure the design can be reliably patterned. For devices with multiple simultaneous objectives, which is the primary concern of this thesis, gradients from independent FoMs are combined using a variety of strategies that can prevent any single objective from dominating at the expense of others.

2.7.1 Worked Example: Broadband Focusing Device

To give a concrete example of the above techniques, we trace through the full adjoint optimization of a device whose goal is to focus a normally incident broadband plane wave to a point in the focal plane. This is the simplest non-trivial test case for inverse design: the device must redistribute the phase of the incident wavefront spatially in a way that constructively concentrates power at a single target location, and it must do so over a range of wavelengths.

Setup

The design region is a rectangular volume of width w and thickness t , placed at a focal distance f above the target focal point $\mathbf{r}_0$. The permittivity at every voxel is a free parameter, initialized to a random grayscale distribution. The surrounding medium is air ($\varepsilon = 1$), and a normally incident plane wave illuminates the top surface. The optimization seeks to maximize the electric field intensity at $\mathbf{r}_0$ simultaneously for a set of discrete wavelengths $\{\lambda_j\}$ spanning the desired bandwidth. The total FoM is

$$f = \sum_j w_j |\mathbf{E}(\mathbf{r}_0, \lambda_j)|^2,$$

where the weights $w_j \propto \lambda_j^2$ compensate for the intrinsically higher intensity that shorter wavelengths produce at a diffraction-limited focus, preventing the optimization from disproportionately favoring shorter wavelengths.

Forward Simulation

For each wavelength λ_j , a forward simulation is run with the plane wave as the source. The full electric field distribution $\mathbf{E}(\mathbf{r}, \lambda_j)$ is recorded throughout the design region, and the field amplitude at the focal point $\mathbf{E}(\mathbf{r}_0, \lambda_j)$ is stored separately. In a time-domain solver, all wavelengths are excited simultaneously by a broadband pulse, and the frequency-resolved fields are obtained by Fourier transformation of the time-domain response. The forward simulation cost is therefore the same regardless of the number of wavelengths included.

Adjoint Source and Simulation

From Eq. (2.24), the adjoint source at each wavelength is

$$\mathbf{J}_{\text{adj}}(\mathbf{r}, \lambda_j) = \mathbf{E}^{\text{A}}(\mathbf{r}_0, \lambda_j) \delta(\mathbf{r} - \mathbf{r}_0).$$

This is a point dipole at the focal location $\mathbf{r}_0$, oscillating with the amplitude and phase determined by the forward field at the focal location. Therefore, for a time-domain solver a single adjoint simulation with a suitably weighted dipole source at $\mathbf{r}_0$ yields the adjoint fields $\mathbf{v}_E(\mathbf{r}, \lambda_j)$ at all wavelengths simultaneously.

Gradient and Update

The gradient of the total FoM with respect to permittivity at each voxel is assembled from Eq. (2.26):

$$\frac{df}{d\varepsilon(\mathbf{r})} = \sum_j 2w_j \operatorname{Re}[\mathbf{E}^A(\mathbf{r}, \lambda_j) \cdot \mathbf{E}(\mathbf{r}, \lambda_j)] .$$

At each voxel, this quantity is positive where increasing the permittivity improves the FoM and negative where decreasing permittivity is preferred. Assuming simple gradient descent, the density ρ is stepped in the direction of this gradient, and the process repeats.

As the optimization progresses, the device progressively redirects the incident wavefront toward the focal point. The binarization schedule gradually sharpens the continuous density until the device approaches a binary permittivity distribution. At convergence, the result is a device that focuses each wavelength to the same spatial location with high efficiency— functionality analogous to an achromatic lens, designed entirely through the optimization. The cost of two simulations per iteration is independent of the number of voxels in the design region, making this approach practical for the large three-dimensional design volumes considered in subsequent chapters.

References

1. Miller, O. D. *Photonic Design: From Fundamental Solar Cell Physics to Computational Inverse Design* PhD thesis (University of California, Berkeley). <http://optoelectronics.eecs.berkeley.edu/ThesisOwenMiller.pdf>.
2. Lalau-Keraly, C. M., Bhargava, S., Miller, O. D. & Yablonovitch, E. Adjoint shape optimization applied to electromagnetic design. *Optics Express* **21**, 21693. ISSN: 1094-4087. <https://www.osapublishing.org/oe/abstract.cfm?uri=oe-21-18-21693> (2020) (Sept. 2013).
3. Ballew, C., Roberts, G. & Faraon, A. Multi-dimensional wavefront sensing using volumetric meta-optics. *Optics Express* **31**, 28658–28669. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-31-18-28658> (2025) (Aug. 2023).
4. Miettinen, K. *Nonlinear Multiobjective Optimization* (Kluwer Academic, 1999).
5. Das, I. & Dennis, J. E. A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems. *Structural Optimization* **14**, 63–69 (1997).

6. Kim, I. Y. & de Weck, O. L. Adaptive weighted-sum method for bi-objective optimization: Pareto front generation. *Structural and Multidisciplinary Optimization* **29**, 149–158 (2005).
7. Deb, K., Pratap, A., Agarwal, S. & Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* **6**, 182–197. ISSN: 1941-0026. <https://ieeexplore.ieee.org/document/996017> (2026) (Apr. 2002).
8. Deb, K. & Jain, H. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints. *IEEE Transactions on Evolutionary Computation* **18**, 577–601. ISSN: 1941-0026. <https://ieeexplore.ieee.org/document/6600851> (2026) (Aug. 2014).
9. Clarke, F. H. in *Optimization and Nonsmooth Analysis* 24–109 (Society for Industrial and Applied Mathematics, Jan. 1990). ISBN: 978-0-89871-256-8. <https://epubs.siam.org/doi/10.1137/1.9781611971309.ch2> (2026).
10. Ehrgott, M. & Ryan, D. M. Improved ε -constraint method for multiobjective programming. *Journal of Optimization Theory and Applications* **138** (2008).
11. Mavrotas, G. Effective implementation of the ε -constraint method in multi-objective mathematical programming problems. *Applied Mathematics and Computation* **213**, 455–465 (2009).
12. Bertsekas, D. P. & Bertsekas, D. P. *Constrained Optimization and Lagrange Multiplier Methods* ISBN: 978-1-886529-04-5 (Athena Scientific, Belmont, Mass, 1996).
13. Rockafellar, R. T. The multiplier method of Hestenes and Powell applied to convex programming. *Journal of Optimization Theory and Applications* **12**, 555–562. ISSN: 1573-2878. <https://doi.org/10.1007/BF00934777> (2026) (Dec. 1973).
14. Birgin, E. G. & Martínez, J. M. in *Practical Augmented Lagrangian Methods for Constrained Optimization* 151–164 (Society for Industrial and Applied Mathematics, May 2014). ISBN: 978-1-61197-335-8. <https://epubs.siam.org/doi/10.1137/1.9781611973365.ch12> (2026).
15. Christiansen, R. E. & Sigmund, O. Inverse design in photonics by topology optimization: tutorial. *arXiv*. eprint: 2008.11816 (2020).
16. Molesky, S. *et al.* Inverse design in nanophotonics. *Nature Photonics* **12**, 659–670. ISSN: 1749-4893. <https://www.nature.com/articles/s41566-018-0246-9> (2026) (Nov. 2018).
17. Kingma, D. P. & Ba, J. *Adam: A Method for Stochastic Optimization* Jan. 2017. <https://arxiv.org/abs/1412.6980> (2025).

18. Ballew, C., Roberts, G., Zheng, T. & Faraon, A. Constraining Continuous Topology Optimizations to Discrete Solutions for Photonic Applications. *ACS Photonics* **10**, 836–844. <https://doi.org/10.1021/acsp Photonics.2c00862> (2025) (Apr. 2023).

Chapter 3

vipdopt: A SOFTWARE PACKAGE FOR VOLUMETRIC INVERSE PHOTONIC DESIGN AND OPTIMIZATION OF METAOPTICS

Abstract

We present vipdopt, a solver-agnostic software framework for nanophotonic inverse design. The framework separates the inverse design problem into five independently configurable layers: configuration and validation, design pipeline, solver backend, figure of merit, and optimization schedule. These independent layers are connected by well-defined contracts that allow any component to be replaced or extended without modifying the others. The optimization scheduler can also sequence multiple gradient-free and/or gradient-descent optimizer types within a single optimization run, and concurrent multi-device optimization via a shared thread pool allows multiple devices to be optimized simultaneously within a single run. The framework is demonstrated via inverse-design of two device classes. We design via inverse-design optimization a 3D metagrating validated against a benchmark suite described in Chen *et al.* [1], and a computational spectrometer based on spatially orthogonal spectral responses, as well as the remaining devices in this thesis.

3.1 Introduction

Nanophotonic inverse design has emerged as a prominent methodology for discovering non-intuitive device geometries that achieve electromagnetic performance beyond what conventional design approaches can reach^{2,3}. By formulating the design problem as an optimization over a continuous permittivity field in space, inverse design has produced unprecedented device performance in color routers^{4,5}, metalenses⁶, mode converters⁷, wavelength demultiplexers⁸, and other devices with large design spaces and complexities. The central role of gradient-based optimization in enabling these results is well established.

The optimization strategy chosen is critical in determining both the convergence time and performance of the final optimized result, and must be chosen carefully by the designer based on the resources available (time, memory, number of free design parameters, degree of parallelization etc.). Many nanophotonic devices are designed via gradient descent, where the gradient of the merit function with respect

to the device permittivity is computed efficiently via the adjoint method detailed in the previous Chapter 2.

The figure of merit $f(p_1, p_2, \dots, p_N)$ defines a scalar hypersurface, over the full space of N design parameters p_i , whose topology directly governs the difficulty of the optimization problem. This hypersurface is called the loss landscape as it is most often used in machine learning, where the typical figure of merit is the loss, or error, between the predicted result and the ground truth. A gradient-based optimizer traverses this surface by following the local gradient, descending (ascending) toward the nearest minimum (maximum) from wherever it is initialized. The complexity of this landscape is a function of the design problem. Simple objectives over small design domains can yield smooth, nearly convex surfaces with few local minima, while multi-objective, broadband, or highly constrained design problems over large domains tend to produce landscapes with many competing local minima separated by high barriers.

Gradient descent, despite its effectiveness, is inescapably a local method^{2,9}, searching for the nearest stationary point of the figure of merit from wherever the optimization is initialized. The final performance of the resulting optimized design therefore depends strongly on both where the initial permittivity is in the design parameter space, and on the smoothness of the loss landscape encountered along the way. For unimodal objectives, where there is a single stationary point in the entire loss landscape, the local extremum *is* the global extremum, and gradient descent works extremely well in this scenario⁹. It continues to work well⁹ for simple loss landscapes with only a few local extrema. As the loss landscape gets increasingly complex, however, which is almost invariably^{9,10} the case as the number of design parameters increase, gradient descent alone consistently converges to local minima—sometimes called “traps”—that do not have performance approaching the global optimum.

Gradient-free, or derivative-free, optimizers can be effectively used to avoid this shortcoming by exploring the design space without relying on local gradient information. They instead evaluate the figure of merit directly at candidate design points sampled from across the design space, exploring the loss landscape in ways that are not necessarily constrained to the neighborhood of the current iteration’s design. Evolutionary strategies such as CMA-ES¹¹ and particle swarm optimization maintain a population of candidate designs and update the sampling distribution based on the candidates’ relative performance. This allows the search to cross barriers

between basins of attraction, which describe regions in the loss landscape where simple gradient descent inevitably leads to the nearest local extremum with no possibility of escape. Bayesian optimization methods, which build probabilistic surrogate models of the objective function, also fall into this category. They generally require more per-iteration simulation than gradient-based methods, as each candidate requires a full forward simulation, and these methods are most effective when the number of candidates is large¹⁰. For problems with many design parameters, stochastic gradient-free methods are most useful^{10,11} in performing an initial, early-terminated global search over the design space, and then picking the best result as an initialization for a subsequent gradient-based approach to converge towards the final result.

The choice of optimizer, and the sequence in which different optimizers are applied, therefore becomes an important choice and consideration for a device optimization⁹. It is not often immediately clear^{3,10} which optimizers and optimization strategies are the best suited to a particular problem. This results in the designer trying several different optimizers in the course of solving the design problem. The choice of which optimizers, and whether or not to combine them, as well as in which order to apply them, can have significant consequences for the performance of the optimized device. While in theory any optimizer can, given the right starting conditions and hyperparameters, find the global maximum, this becomes increasingly unlikely as the number of design parameters increases. It becomes more important to find the strategy that can find *enough* sufficiently well-performing final designs.

In deciding on this, the nanophotonic designer should not have to repeatedly modify the source code of their software framework or maintain separate scripts for each strategy they want to try. The software architecture of an inverse design framework should expose the optimization strategy as an easily adjustable parameter. The strategy should be expressible in the same configuration file as all other declared parameters, so that exploring the strategy space costs no more than editing a configuration file.

3.1.1 Existing Frameworks and Their Limitations

Several software frameworks for nanophotonic inverse design have been developed in recent years, each making distinct architectural choices that reflect the problems and solvers they were designed for.

SPINS⁹ introduced a graph-based problem formulation in which the parameteriza-

tion, simulation, and objective are expressed as nodes in a computational graph, cleanly separating the physical problem from the numerical solver. By offering a well-structured framework, SPINS supports a variety of device types without requiring ad-hoc modification of individual optimization scripts for each and every device. Its primary limitation is that the optimization strategy is gradient-based only, and is limited to a single gradient descent loop with no mechanism for switching optimizer type. The framework is also tied to a specific frequency-domain finite-difference (FDFD) solver, limiting its applicability to simulation problems for which that solver is appropriate.

Hammond *et al.* [12] developed a topology optimization framework built on Meep¹³, the open-source finite-difference time-domain solver. The framework implements filter and projection transforms in the design pipeline and adjoint gradient computation via Meep’s native adjoint support, enabling fabrication-aware gradient-based optimization. Like SPINS, it runs a single gradient descent loop with no schedule mechanism, and the optimization strategy is not user-configurable without modifying the framework code.

`invrs-gym`¹⁴ takes a different approach: rather than providing an end-to-end optimization framework, it defines a suite of standardized benchmark problems and evaluation metrics against which any optimizer can be tested. However, `invrs-gym` provides no optimization loop, no schedule mechanism, or support for alternative problem scenarios, and is best used as infrastructure for benchmarking instead of a tool for designing new devices. Finally, HiLAB¹⁰ is the closest prior work to the approach taken here. It implements a fixed three-phase pipeline:

1. topology optimization produces an initial continuous design;
2. a variational autoencoder encodes this into a low-dimensional latent space;
3. Bayesian optimization searches the latent space for improved designs.

This sequence is much more capable¹⁰ of escaping local minima than a simple gradient-descent approach. However, the three phases are hard-coded, and the transition points are not configurable by the user.

3.1.2 New Technical Features

`vipdopt` addresses the limitations above through three technical features.

Composable optimization schedule. `vipdopt` introduces an intra-run optimization schedule that sequences arbitrary optimizer types within a single coherent optimization run. Within a single configuration file, the user declares the schedule and the rest of the problem specification. As such, no modifications to the code are required in order to change the optimization strategy. This generalizes the fixed multi-phase pipelines of HiLAB in a way that is applicable to any device class and any solver backend.

Multi-device concurrent optimization. `vipdopt` supports the simultaneous optimization of multiple independent devices within a single run. Each device carries its own design pipeline, simulator configuration, and figure of merit, and each runs its optimization loop in a dedicated thread via a shared thread pool. This is distinct from the prior frameworks covered in this chapter, none of which support simultaneous multi-device optimization.

Integrated hyperparameter search. `vipdopt` prioritizes hyperparameter optimization in the first stage of the optimization workflow, running multiple early-terminated optimizations via Optuna, which is built into the package’s `ExperimentManager` class. Each run has a hyperparameter configuration sampled by the TPE algorithm, and the best-performing configuration is identified before committing to a full run. All prior software frameworks require hyperparameter selection to be performed manually, outside the optimization loop. By internalizing this step, `vipdopt` reduces the iteration cost of exploring new device classes and optimization strategies.

3.1.3 Chapter Outline

The remainder of this chapter is organized as follows. Section 3.2 reviews the gradient computation background used in this work, covering the adjoint method and automatic differentiation; the inverse design problem formulation and optimization strategies are developed in Chapter 2. Section 3.3 describes the framework architecture in detail, covering each of the five layers from config validation through experiment orchestration. Section 3.4 presents the optimization schedule in detail, including the single-stage baseline, the multi-stage schedule, and a cost-normalized comparison of the two strategies. Section 3.5 demonstrates the framework on two device classes spanning different solver backends and physical operating principles. Section 3.6 discusses the framework’s position relative to prior work, its current limitations, and directions for future extension.

3.2 Background

The inverse design problem, its density-based formulation, and the topology optimization workflow are treated in full in Chapter 2; the notation adopted there is used throughout this chapter without re-derivation.

The optimization strategies used in this work, including gradient-based methods, the augmented Lagrangian optimizer, projection and continuation, and CMA-ES, are described in Section 2.5.8 of Chapter 2.

3.2.1 Gradient Computation: Adjoint Method and Automatic Differentiation

Computing the gradient with respect to the density $\partial f / \partial \theta$ is the main computational challenge of inverse design. A naive finite-difference approximation requires one forward simulation per parameter, making it prohibitively expensive for the large design spaces typical of nanophotonic problems ($N = 10^3$ – 10^6 voxels). The adjoint method and automatic differentiation are the two most common approaches to make gradient-based optimization feasible in the case of photonic inverse design.

The adjoint method. The adjoint method computes $\partial f / \partial \varepsilon$ —the gradient of the figure of merit with respect to the permittivity at every point in the design domain—using exactly two simulations, regardless of the number of design parameters. We can define the adjoint field $\mathbf{E}_{\text{adj}}$ as the solution to

$$\nabla \times \nabla \times \mathbf{E}_{\text{adj}}(\mathbf{r}) - \frac{\omega^2}{c^2} \varepsilon(\mathbf{r}) \mathbf{E}_{\text{adj}}(\mathbf{r}) = \mathbf{J}_{\text{adj}}(\mathbf{r}), \quad (3.1)$$

where the adjoint source $\mathbf{J}_{\text{adj}}$ is derived from the figure of merit: $\mathbf{J}_{\text{adj}} = -\partial f / \partial \mathbf{E}^*$, evaluated at the monitor locations. The gradient of f with respect to the permittivity at each voxel i is then

$$\frac{\partial f}{\partial \varepsilon_i} = -2\omega \operatorname{Im}[\mathbf{E}_{\text{fwd}}(\mathbf{r}_i) \cdot \mathbf{E}_{\text{adj}}(\mathbf{r}_i)], \quad (3.2)$$

where $\mathbf{E}_{\text{fwd}}$ is the forward field obtained from the original simulation with physical sources. This expression involves only the pointwise product of the forward and adjoint fields at each design voxel, which can be computed from the two field distributions without any additional simulations. The chain rule then propagates $\partial f / \partial \varepsilon$ through the design pipeline to obtain $\partial f / \partial \theta$.

The adjoint method’s cost is independent of the number of design parameters N . It requires that the solver support two types of simulation: a forward run with physical sources, and an adjoint run with sources derived from the figure of merit gradient. For natively differentiable solvers, this distinction can be handled automatically by

an autodifferentiation framework, but explicit implementation is required for otherwise external solvers.

Automatic differentiation. Automatic differentiation (AD) is an approach applicable when the forward simulation is implemented as a composition of elementary operations whose derivatives are known analytically. AD frameworks such as JAX, PyTorch, and autograd maintain a computation graph of the forward pass and apply the chain rule algorithmically in reverse (reverse-mode AD, or backpropagation) to compute the gradient of any scalar output with respect to any input. For solvers implemented in JAX (such as the RCWA backend used in this work) AD provides an alternative to the adjoint method that requires no separate adjoint implementation. Instead, the gradient is obtained by a single call to the AD framework’s gradient function, which internally traces the forward computation graph and performs differentiation. The main difference between the two approaches is the cost of implementation. The adjoint method requires deriving and implementing solver-specific adjoint equations, while AD requires only that the forward simulation be written in a differentiable programming framework.

External VJP bridge. When the solver is not natively differentiable, gradient-based computation requires an explicit bridge between the solver and the AD framework. `vipdopt` implements this via a custom vector-Jacobian product (VJP) that calls the solver twice (forward and adjoint), computes $\partial f / \partial \varepsilon$ from the field overlap formula above, propagates it through the differentiable design pipeline using the AD framework, and returns the result as if it had been computed by standard backpropagation. This approach is described in detail in Section 3.3.4.

3.3 Framework Architecture

3.3.1 Design Philosophy

`vipdopt` is organized around the central concept of independent optimization layers. In nanophotonic inverse design, a new device class may demand a different figure of merit; a new fabrication process may demand different fabrication-aware filters; and a new solver may demand a different adjoint implementation. Every such change requires rewriting the internal code of the optimizer if its framework is not decoupled and modular. `vipdopt` avoids this by interposing well-defined contracts between layers. For example, the optimizer receives and returns a flat parameter vector, which may represent either permittivity voxels, pillar heights, or waveguide widths.

Three additional design objectives follow from this philosophy. First, the framework

supports simultaneous optimization of multiple independent devices within a single run, each with its own design pipeline, simulator configuration, and figure of merit, executing concurrently via a thread pool. Second, optimization strategies are composable: the sequence of optimizer types, their hyperparameters, and the state of the design pipeline can all be specified declaratively and varied without modifying code. Third, all of the configuration is expressed in a single file so that a researcher can explore a large hyperparameter space by editing a single file and re-running. The available schemes are formalized through a schema which the configuration file is validated against prior to execution.

3.3.2 System Overview

vipdopt is organized into five layers, each with a well-defined responsibility and a minimal interface to adjacent layers. Figure 3.1 illustrates the architecture and the data flow through a single optimization iteration.




Figure 3.1: Architecture of vipdopt. The configuration file is read and used to generate both the `OptimizationManager` class and the disparate `OptimizationTasks`. The `OptimizationManager` handles parallelization and scheduling of all Tasks. Within a Task, data flow proceeds downward through the layers during each optimization iteration. The dashed upward arrow indicates gradient backpropagation through the design pipeline.

Following a single optimization iteration through this stack clarifies the role of each layer. A YAML configuration file is loaded and validated by the *Config & Validation* layer (`vipdopt.states`). Pydantic models check the shared envelope of the config, while a plugin schema registry validates deep payloads specific to each simulator, optimizer, FoM, and design pipeline type. Semantic validation then checks

cross-references and scene-expression correctness. The `OptimizeManager` instantiates runtime objects from the validated config: simulators, devices (each carrying a `DesignPipeline`), figures of merit, and `OptimizerTask` objects that bind these together. Each task is dispatched to a thread in a pool and runs its optimization loop independently. Within each task, the optimizer holds a flat parameter vector $\theta \in \mathbb{R}^n$; the design pipeline maps θ to physical parameters; the solver backend computes the objective and gradient; and the optimization schedule governs how the optimizer type and pipeline configuration change during the run.

3.3.3 Design Pipeline Layer

The design pipeline (`vipdopt.design.DesignPipeline`) is the configurable mapping from optimizer variables θ to the physical parameters consumed by the simulator. It is organized into four sequential stages.

Stage 1: Domain Transform. Maps θ to a normalized representation $\mathbf{x}$. The `bounded_01` transform is an identity for box-constrained problems; the `logit_sigmoid` transform maps an unbounded θ to $\mathbf{x} = \sigma(\theta/\tau) \in (0, 1)$.

Stage 2: Parameterization. Decodes $\mathbf{x}$ into an intermediate design representation. The `field` parameterization reshapes $\mathbf{x}$ into a density field of shape (n_x, n_y, n_z) . The `generator` parameterization passes $\mathbf{x}$ through a configurable generator transform (e.g., a multilayer perceptron) to produce a field from a low-dimensional latent vector. The `identity` parameterization passes $\mathbf{x}$ unchanged.

Stage 3: Geometry Interface. Converts the intermediate representation into a canonical continuous density field $\rho \in [0, 1]^{n_x \times n_y \times n_z}$. The `epsilon_grid` interface treats the representation as ρ directly. The `level_set` interface computes $\rho = H(-\phi)$ via a smooth Heaviside of a signed distance function ϕ . The `primitives` interface rasterizes a vector of primitive parameters into a field.

Stage 4: Transform Stack. Applies a sequence of fabrication-aware transforms to ρ to produce the final parameter field. In the most common configuration, this stack consists of three transforms: a neighbor-averaging blur that couples adjacent voxels to penalize sub-resolution features; the tanh projection of Eq. (2.66) that drives ρ toward binary values; and an affine transform $\varepsilon = \varepsilon_{\min} + (\varepsilon_{\max} - \varepsilon_{\min})\rho'$ that maps the binarized density to the physical permittivity range. All transforms are differentiable with respect to their inputs.

The pipeline also supports fabrication robustness variants. A `robustness` block specifies named variants (e.g., `nominal`, `eroded`, `dilated`), each defined by additional

transforms applied before the shared transform stack. The pipeline evaluates all variants in parallel via a `JobRunner`, and the figure of merit layer aggregates the results.

3.3.4 Autodiff and Adjoint Bridge Layer

A `Backend` object encapsulates the choice of automatic differentiation framework—one of `torch`, `autograd`, `jax`, or `numpy`—along with the corresponding array namespace. The backend is selected globally per run via the `autodiff.backend` config key, and all array operations in the design pipeline and figure of merit are dispatched through its namespace.

Two solver backends are provided by default. The RCWA backend implements a JAX-native rigorous coupled-wave analysis solver for three-dimensionally periodic structures; because it is a pure JAX function, gradients flow end-to-end via standard reverse-mode AD at a cost of approximately two to three forward passes per gradient step. The Lumerical FDTD backend connects `vipdopt` to Ansys Lumerical FDTD Solutions via subprocess dispatch, supporting full-vectorial three-dimensional simulation of arbitrary geometries. Both backends satisfy the same runtime contract: a `run(params, sim_params, backend)` method that accepts a parameter vector and returns a `sim_out` dictionary. Additional solver backends can be easily registered without modifying the optimization loop, schedule layer, or figure of merit layer.

For natively differentiable solvers (RCWA), gradients flow end-to-end via standard reverse-mode AD. For external solvers (Lumerical FDTD), the framework injects a custom VJP at the solver boundary: the forward pass runs the solver to obtain fields and objective; the backward pass runs the solver again with adjoint sources derived from the figure of merit to obtain the gradient per Eq. (??), at a cost of two forward simulations per step independent of n . This bridge is implemented in three backend-specific variants: `defvjp` for `autograd`; `jax.custom_vjp` with `jax.pure_callback` for `jax`; and `torch.autograd.Function` for `torch`. In all three cases the resulting callable is differentiable from the AD framework’s perspective, and the optimizer interacts with it identically regardless of which solver produced the gradient. The Lumerical FDTD backend has been validated against central finite-difference gradients across all device classes in this chapter, confirming adjoint correctness to within relative errors of order 10^{-4} or better.

3.4 Optimization Schedule

3.4.1 Single-Stage Baseline

The conventional approach runs a single augmented Lagrangian optimizer for 300 iterations. The projection sharpness β remains low for the first 150 iterations, during which the optimizer explores a soft, continuous design space, and then increases monotonically until full binarization is reached at iteration 300. The gradient is computed via the Lumerical adjoint at each iteration (2 forward simulations per step), for a total budget of 600 simulations. The delayed β ramp prevents binarization pressure from trapping the design in poor local minima before a good region of design space has been found. Section 2.7.1 well describes this baseline.

3.4.2 Multi-Stage Schedules

`vipdopt` addresses the local-minima problem through a multi-stage optimization schedule that sequences stochastic global search before gradient-based refinement. The full workflow proceeds in three stages.

Stage 0: Hyperparameter optimization (Optuna, ExperimentManager). The `ExperimentManager` dispatches an Optuna phase that runs many short `OptimizeManager` sub-runs with different hyperparameter configurations sampled by the TPE algorithm—learning rate, initial β , and simulator parameters—identifying the configuration under which the subsequent stages are most likely to succeed.

Stage 1: Stochastic design-space exploration (Nevergrad CMA-ES). Seeded from the best Optuna checkpoint, CMA-ES performs stochastic search over the design space without gradients. Each evaluation costs one forward simulation, allowing a large number of candidate designs to be evaluated within a modest simulation budget. The design variable at the end of this stage is the best candidate across the entire CMA-ES population.

Stage 2: Gradient-based refinement with β -continuation (AugLag). The design variable from Stage 1 is passed directly—without checkpoint serialization—to the augmented Lagrangian optimizer, which runs for $N = 300$ iterations with the exponential β schedule of Eq. (2.67). At the onset of this stage ($n = 1, \beta = 1$), minimal binarization pressure allows the optimizer to adjust the continuous density field inherited from CMA-ES. At the end ($n = N, \beta = 128$), a near-binary permittivity distribution is enforced. The pipeline override mechanism applies the β schedule by rebuilding the `DesignPipeline` transform stack at each iteration without reconstructing the optimizer or re-instantiating the manager.

The key architectural property of this schedule is that Stages 1 and 2 use the same device, the same parameter vector, and have identical (but separate) `OptimizerTasks`. A graphical representation of this workflow is given in Fig. 3.2.




Figure 3.2: Example workflow of `vipdopt` `ExperimentManager`. The `ExperimentManager` conducts several truncated `OptimizeManagers` in parallel in order to perform hyperparameter optimization, then picks the best result and performs stochastic gradient-free exploration to survey the loss landscape before finalizing the optimization through gradient-descent based methods.

3.4.3 Comparison: FoM vs. Forward Simulations

Figure 3.3 compares the figure of merit as a function of cumulative forward simulation count for the single-stage baseline and the multi-stage schedule, evaluated on a color router device optimized with the Lumerical FDTD backend. Both runs use the same device geometry, the same figure of merit, and comparable total simulation budgets. The scheduled run allocates 200 iterations to the CMA-ES stage and up to 100 to the AugLag stage; the baseline runs AugLag for 300 iterations.

The CMA-ES stage appears as a plateau in the curve: individual candidate evaluations do not monotonically improve the FoM. The AugLag stage begins with a sharp improvement, guided by the gradient-descent approach from the initialization provided by CMA-ES. The baseline shows slower initial progress, as gradient descent from a uniform initialization must first traverse the soft-projection landscape before




Figure 3.3: FoM vs. cumulative forward simulation count for the single-stage baseline and the dual-stage schedule, evaluated on the 3D metagrating in Fig. 3.4. The vertical dashed line marks the transition between the first and second stages.

reaching a competitive design region.

3.5 Demonstration Applications

`vipdopt` is demonstrated on two inverse design problems that together exercise both solver backends and the full optimization schedule:

1. A 3D metagrating that validates the RCWA backend against an established literature benchmark;
2. A computational spectrometer that demonstrates the Lumerical FDTD backend on a device with a composite figure of merit.

3.5.1 Benchmark Validation: 3D Metagrating

Physical objective. The metagrating deflects normally incident light from a silica substrate into the $(+1, 0)$ diffraction order in air, corresponding to a 50° deflection

angle. This example is adapted from Problem F of the Chen *et al.*¹ standardized benchmark suite, with reference results from multiple independent FDTD, FEM, and RCWA implementations.

Device geometry. The structure consists of a freeform silicon pattern ($\varepsilon = 3.45^2$) of thickness 325 nm on a semi-infinite silica substrate ($\varepsilon = 1.45^2$), periodic in x with period $\Lambda_x = 1371$ nm and in y with $\Lambda_y = 525$ nm. The figure of merit is the diffraction efficiency into the $(+1, 0)$ channel, defined as the power flux into that order normalized by the incident power.

Optimization setup. The RCWA backend is used for all simulations. The design pipeline applies neighbor blur followed by the tanh projection with the exponential β schedule (Eq. 2.67), mapping the continuous density to a binary silicon/air permittivity distribution. The multi-stage schedule runs Nevergrad CMA-ES for stochastic initialization followed by augmented Lagrangian gradient descent for refinement.

Results. The optimized metagrating achieves a diffraction efficiency of 88% at the design wavelength of 1050 nm, with deflection angle of 50° at this wavelength, as shown in Fig. 3.4b. This places it within the range reported by Chen *et al.* [1] for RCWA-based implementations at comparable lengthscales. The unit cell is shown in Fig. 3.4a.

3.5.2 Spectral Reconstruction via Orthogonal Spatial Spectra

Physical objective. The device functions as a compact computational spectrometer. Each spatial region of the output plane exhibits a distinct spectral response, approximately orthogonal to the responses of all other regions. An unknown input spectrum is reconstructed by measuring the transmitted intensity at each spatial region and solving a linear system $\mathbf{A}\mathbf{s} = \mathbf{y}$, where $\mathbf{A}$ is the $[N \times M]$ spectral response matrix, $\mathbf{s}$ is the unknown input spectrum, and $\mathbf{y}$ is the vector of measured intensities. The orthogonality of spatial responses minimizes the condition number of $\mathbf{A}$, which directly controls the noise sensitivity of the reconstruction.

Optimization setup. The optimization attempts to match four target spectra, establishing them in four spatially separate spectral detector regions. The figure of merit is the L2 distance between the target and obtained spectra across the device bandwidth, which is minimized throughout the optimization.

Results. The optimized device achieves the amplitudes and rough shapes of the target orthogonal spectra across the device bandwidth for all four available spectral detection regions.




Figure 3.4: (a) Permittivity distribution of the unit cell for the optimized 3D meta-grating. The 2D permittivity distribution is extruded in the z -axis. (b) Normalized intensity as a function of deflection angle and wavelength. The metagrating deflects most of the power into the $(+1, 0)$ order at an angle of 50° . Red star indicates the design wavelength 1050 nm and 50° .

3.6 Discussion

3.6.1 Comparison to Prior Frameworks

The most significant architectural distinction offered by *vipdopt* is the optimization schedule layer. SPINS employs a graph-based problem formulation that cleanly separates parameterization, simulation, and objective, but its optimization strategy is gradient-only with no mechanism for switching optimizer type or updating pipeline parameters mid-run. Hammond *et al.*'s Meep-based framework similarly runs a single gradient descent loop with filter and projection applied statically. Both are powerful tools for the problems they were designed for, but neither tool exposes optimization strategy for simple configuration by the user.

HiLAB is the closest prior work to *vipdopt* in spirit, explicitly sequencing topology optimization, a variational autoencoder, and Bayesian optimization to address the




Figure 3.5: (a) Bottom layer of device that matches orthogonal target spectra. (b) Target (dashed) and obtained (solid) spectral response curves for each spatial detector region.

local-minima problem. However, HiLaB is not as general: it has a fixed pipeline, the transition points are not user-configurable, and the approach is tightly coupled to color routers. `vipdopt` generalizes this by making the schedule a user-defined sequence of arbitrary optimizer types and pipeline configurations, applicable to any device and any solver backend.

All prior frameworks also optimize only a single device per run, whereas `vipdopt`'s thread-pool dispatch allows multiple devices to be optimized simultaneously within a single job, allowing the parallelization of independent runs without external job management. The autodifferentiation and external VJP layer is a further distinguishing feature, as no prior framework provides a unified gradient interface that transparently handles both natively differentiable solvers and closed-source external solvers via backend-specific custom VJP wrappers, across three AD frameworks (autograd, JAX, PyTorch). Finally, the ExperimentManager's Optuna integration can perform hyperparameter search as a workflow step, replacing the manual pre-process that prior frameworks require before committing to a full optimization run.

It should be noted that `vipdopt` does not yet support any eigenvalue-based objectives. The figure of merit layer operates on scalar reductions of simulator monitor outputs. Objectives requiring eigenvalue computations—such as group velocity dispersion of a photonic crystal waveguide or cavity Q -factor—are not natively supported. For example, the slow-light waveguide benchmark (Problem G) in Chen *et al.*¹ is not directly addressable with the current framework.

3.6.2 Future Extensions

The solver registry architecture is designed so that additional backends can be integrated without modifying the optimization loop, schedule layer, or figure of merit layer. Any solver that accepts a design parameter vector, runs a forward simulation, and returns monitor outputs satisfies the runtime interface contract and can be registered as a backend without modifying the framework core.

The schedule layer provides a natural extension point for adaptive strategies. Per-stage stopping criteria, i.e. terminating a stage when FoM improvement falls below a threshold rather than after a fixed iteration count, would allow the schedule to allocate simulation budget dynamically, spending more iterations where the optimizer is still making progress and transitioning earlier when it has stalled.

3.7 Conclusion

`vipdopt` is a solver-agnostic framework for nanophotonic inverse design which separates an inverse design problem into five independently configurable layers, connected by well-defined contracts that allow any component to be replaced or extended without modifying the others.

The central technical contribution is the optimization schedule layer, which sequences multiple optimizer types within a single optimization run. This generalizes the fixed multi-phase pipelines of prior frameworks, and allows the user to easily configure an optimization schedule and sequence according to their chosen device and solver backend. This is done entirely through a configuration file and, as long as no new classes are added, without the need for modification of code. The framework also offers concurrent multi-device optimization via a shared thread pool and integrated hyperparameter search via Optuna.

Two solver backends are implemented and validated. The RCWA backend provides solutions for periodic structures and calculates the gradients of the merit function via JAX automatic differentiation. The Lumerical FDTD backend provides full-wave simulation for periodic and aperiodic structures, whereby the adjoint gradient is injected into the AD computation graph via a custom vector-Jacobian product bridge. The framework is demonstrated on two device classes—a 3D metagrating and a computational spectrometer—establishing that the solver-agnostic architecture and composable optimization schedule generalize across solver backends and physical operating principles.

3.8 Minimal API Example

The following configuration file illustrates a complete vipdopt optimization run for the 3D metagrating benchmark using the RCWA backend and a two-stage schedule. The run is launched from the command line as:

```
vipdopt optimize config_metagrating.yaml --out runs/metagrating/
```

```
# config_metagrating.yaml
schema_version: 1
seed: 42
save_every: 25
autodiff:
  enabled: true
  backend: jax

simulators:
  - name: rcwa_sim
    type: rcwa
    params:
      wavelength_nm: 1050.0
      lattice:
        Lx_nm: 1371.0
        Ly_nm: 525.0
      layers:
        - {name: superstrate, material: air,
            thickness_nm: 0.0}
        - {name: design, material: silicon,
            thickness_nm: 325.0,
            is_design_layer: true}
        - {name: substrate, material: silica,
            thickness_nm: 0.0}
      diffraction_order: [1, 0]
      n_harmonics: [5, 3]

devices:
  - name: metagrating
```



```

n_params: 2700
bounds: [0.0, 1.0]
init: {kind: random, space: normalized}
design_pipeline:
  field_shape: [90, 30]
  domain: {kind: bounded_01}
  parameterization: {kind: field, cfg: {}}
  geometry_interface: {kind: epsilon_grid, cfg: {}}
  transforms:
    - type: neighbor_blur
      cfg: {iters: 1, axes: [0, 1],
            weight_self: 1.0, weight_neighbors: 1.0}
    - type: tanh_projection
      cfg: {beta: 1.0, eta: 0.5}
    - type: affine
      cfg: {scale: 10.9, bias: 2.1}

```

FoMs:

```

- name: diffraction_efficiency
  type: intensity
  cfg: {key: diffraction_order_plus1, i0: 0, i1: -1,
        weight: -1.0}

```

optimizers:

```

- name: opt_metagrating
  devices: [metagrating]
  simulator: rcwa_sim
  FoM: diffraction_efficiency
  schedule:
    - max_iters: 200
      optimizer:
        type: nevergrad
        cfg: {method: cmaes, budget: 200,
              num_workers: 1, sigma: 0.15, seed: 42}
    - max_iters: 300
      optimizer:

```

```

type: aug_lag
cfg: {lr: 0.01, mu0: 0.0, mu_factor: 1.0, steps: 1}
pipeline_overrides:
  transforms[1].cfg.beta: "2^(7*(n-1)/N)"

```

The first stage (200 evaluations) runs CMA-ES without gradient computations; each evaluation costs one RCWA forward pass. The second stage (300 iterations) runs the augmented Lagrangian optimizer with β updated each iteration according to Eq. (2.67) with $N = 300$; each iteration costs approximately two to three RCWA forward passes under JAX automatic differentiation.

References

1. Chen, M. *et al.* Validation and characterization of algorithms and software for photonics inverse design. *JOSA B* **41**, A161–A176. ISSN: 1520-8540. <https://opg.optica.org/josab/abstract.cfm?uri=josab-41-2-A161> (2025) (Feb. 2024).
2. Molesky, S. *et al.* Inverse design in nanophotonics. *Nature Photonics* **12**, 659–670. ISSN: 1749-4893. <https://www.nature.com/articles/s41566-018-0246-9> (2026) (Nov. 2018).
3. Kim, J. *et al.* Inverse design of nanophotonic devices enabled by optimization algorithms and deep learning: recent achievements and future prospects. *Nanophotonics* **14**, 121–151. ISSN: 2192-8614. <https://onlinelibrary.wiley.com/doi/abs/10.1515/nanoph-2024-0536> (2026) (2025).
4. Zou, X. *et al.* Pixel-level Bayer-type colour router based on metasurfaces. *Nature Communications* **13**, 3288. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-022-31019-7> (2026) (June 2022).
5. Kim, C. *et al.* Freeform metasurface color router for deep submicron pixel image sensors. *Science Advances* **10**, eadn9000. <https://www.science.org/doi/full/10.1126/sciadv.adn9000> (2024) (May 2024).
6. Chung, H. & Miller, O. D. High-NA achromatic metalenses by inverse design. *Optics Express* **28**, 6945–6965. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-28-5-6945> (2025) (Mar. 2020).
7. Augenstein, Y. & Rockstuhl, C. Inverse Design of Nanophotonic Devices with Structural Integrity. *ACS Photonics* **7**, 2190–2196. <https://doi.org/10.1021/acsphotonics.0c00699> (2026) (Aug. 2020).
8. Piggott, A. Y. *et al.* Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer. *Nature Photonics* **9**, 374–377. ISSN: 1749-4893. <https://www.nature.com/articles/nphoton.2015.69> (2025) (June 2015).

9. Su, L. *et al.* Nanophotonic inverse design with SPINS: Software architecture and practical considerations. *Applied Physics Reviews* **7**, 011407. ISSN: 1931-9401. <https://doi.org/10.1063/1.5131263> (2026) (Mar. 2020).
10. Marzban, R., Abiri, H., Pestourie, R. & Adibi, A. HiLAB: A Hybrid Inverse-Design Framework. en. *Small Methods* **9**, e00975. ISSN: 2366-9608. <https://onlinelibrary.wiley.com/doi/abs/10.1002/smt.202500975> (2026) (2025).
11. Hansen, N. The CMA evolution strategy: A tutorial. *CoRR* **abs/1604.00772**. arXiv: 1604.00772 tex.bibsource: dblp computer science bibliography, <https://dblp.org/tex.timestamp: Mon, 13 Aug 2018 16:47:37 +0200. http://arxiv.org/abs/1604.00772> (2016).
12. Hammond, A. M. *et al.* High-performance hybrid time/frequency-domain topology optimization for large-scale photonics inverse design. EN. *Optics Express* **30**, 4467–4491. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-30-3-4467> (2025) (Jan. 2022).
13. Oskooi, A. F. *et al.* Meep: A flexible free-software package for electromagnetic simulations by the FDTD method. en. *Computer Physics Communications* **181**, 687–702. ISSN: 0010-4655. <https://www.sciencedirect.com/science/article/pii/S001046550900383X> (2022) (Mar. 2010).
14. Schubert, M. F. *invrs-gym: a toolkit for nanophotonic inverse design research* Oct. 2024. <http://arxiv.org/abs/2410.24132> (2025).

Chapter 4

BROADBAND DISPERSION ENGINEERED VOLUMETRIC METAMATERIAL FOR POLARIZATION BEAM SPLITTING AT LARGE DEFLECTION ANGLES

This chapter is directly reproduced from [1], a paper published as a key part of the work in this thesis. As described in the abstract below, this work discusses the design, fabrication and characterization of dispersion-controlled polarization beam splitters. We experimentally demonstrate simultaneous control of polarization and dispersion in the same device, in both zero-dispersion and hyper-dispersive regimes.

Abstract

Polarization beam splitters are central components in broadband optical systems, but conventional prism- and coating-based implementations remain bulky and offer limited control over broadband chromatic response. Beam splitters based on periodic diffractive elements reduce physical footprint, but exhibit wavelength-dependent angular dispersive effects. Here we experimentally demonstrate inverse-designed transmission-mode volumetric metasurfaces that jointly control polarization splitting and chromatic dispersion over the 4.3–4.7 μm mid-infrared band in a fabrication-compatible pseudo-3D platform. We present two devices: a zero-dispersion polarization beam splitter that directs horizontally polarized light to $+30^\circ$ and vertically polarized light toward -20° with near-zero angular variation across the band; and a hyper-dispersive beam splitter that routes horizontally polarized light toward -20° while exhibiting enhanced chromatic dispersion. These results establish inverse-designed transmission-mode volumetric metastructures as a practical route to simultaneous polarization sorting and dispersion engineering in compact mid-infrared optics.

4.1 Introduction

Diffraction gratings are spatially periodic structures that enforce phase matching between an incident wave and the discrete, outgoing diffraction orders. For a fixed grating period, a broadband input beam is steered to different angles across the spectrum. As a result, any beam-splitting device that depends on periodic diffraction exhibits an inherent chromatic dispersion that has increasing magnitude with larger

bandwidths and target deflection angles. This property is advantageous in devices such as spectrometers, but is a limitation in broadband beam splitters, beam displacers, and other systems that require controlled routing of multiple wavelengths to the same angular position.

The dispersion of an optical element can be characterized relative to that of a “regular” blazed grating. For a deflection angle θ_0 at the angular frequency $\omega = 2\pi c/\lambda$, the dispersion relation is $\partial\theta/\partial\omega = -\tan(\theta_0)/\omega_0^2$. We describe regular gratings as having negative dispersion (due to the negative sign). The dispersionless device then has $\partial\theta/\partial\omega = 0$, and a positive dispersion device has a derivative opposite in sign to a “regular” grating equivalent. This is visualized in Fig. 4.1— devices of positive dispersion, such as refractive prisms, deflect light of shorter wavelengths (blue) to a greater extent than longer wavelengths (red); and devices of negative dispersion deflect light of longer wavelengths to a greater extent than shorter wavelengths. We can then describe a “hyper”-dispersive device as one that deflects light of longer wavelengths to an even greater extent than a regular grating (with the same deflection angle at the same wavelength) compared to light of short wavelengths, i.e. $|\partial\theta/\partial\omega| > |\tan(\theta)/\omega_0|$.

Precise control of dispersion is especially useful in broadband optical systems. For example, devices such as apochromatic lenses must suppress dispersion in order to maintain angular invariance across their bandwidth, whereas the resolution of spectrometers improves from enhanced dispersion. This control is challenging to maintain in the case of devices that simultaneously perform another optical function, such as polarization splitting or beam displacement. Plate beam splitters and dielectric coatings often exhibit uneven splitting ratios and nonzero wavelength dependence^{3,4} across a broad spectral band. Devices such as prism-based polarization beam splitters (PBS) can provide stronger and more robust polarization separation^{5,6}, but are typically bulky and have a substantial footprint. A single platform with a strong polarization-dependent response and broadband control of dispersive behavior would therefore enable multiple functionalities with high efficiency in integrated or miniaturized optical systems.

Metastructured optics, which consist of arrangements of subwavelength scatterers or meta-atoms, are a natural platform for this purpose. They offer precise wavefront control with far less thickness than conventional refractive or birefringent components, and have been used to realize many optical devices such as spectrometers⁷ and holograms⁸. In the context of polarization beam splitting, prior devices have




Figure 4.1: Comparison of different regimes of dispersion. (a) Positive dispersion occurs in dielectric prisms. Shorter wavelengths are deflected more than longer wavelengths. (b) Negative dispersion occurs in diffractive gratings and lenses. Longer wavelengths are deflected more than shorter wavelengths. (c) In the regime of zero dispersion / achromatic devices, light of different wavelengths is focused to the same point in the case of lenses, or deflected the same amount in the case of gratings. (d) Hyper-dispersive devices exhibit negative dispersion, but to a greater extent than a diffraction grating that achieves the same deflection angle.

shown that orthogonal polarizations can be scattered into different channels using geometric-phase^{9,10} or resonant meta-atom responses^{11,12}. Metastructured optics have also been used to achieve broadband control of chromatic dispersion, such as beam deflectors^{13–15} and achromatic lenses^{15–17}. However, many such devices rely on (locally) periodic phase gradients to achieve beam deflection^{18,19}. As such, the far-field response of these devices shifts^{15,20} or broadens^{13,21,22} under broadband illumination, or illumination at a different wavelength, in ways that are difficult to suppress while using a periodic layout of meta-atoms^{15,23,24}.

Conventional metasurfaces are typically designed as two-dimensional phase masks in a single layer, and thus have a limited number of optical modes that can be exploited to realize optical designs with high performance^{10,23,25}. This becomes a significant limitation in the case of devices with multiple functionalities, such as wavelength- or polarization-dependent behavior, where each new functionality incurs a cost in efficiency. In recent years, multifunctional devices have been implemented using three-dimensional, or volumetric, metaoptics, where multiple scattering between structures in the vertical direction allows for a greater diversity of resonances within the device, and the increased degrees of freedom due to the device's volume offer a larger design space for simultaneously addressing multiple functionalities. Simultaneous spectral and polarization sorting^{26–28}, as well as multiplexing based on orbital angular momentum²⁷, has been demonstrated using this platform. In many cases, inverse design, via either gradient-based adjoint methods^{29–37} or stochastic optimization methods^{38,39}, is well suited to navigating these large and nonintuitive design spaces. Further methods are detailed in Section 4.4.1.

Although dielectric metasurfaces have also been used to realize beam deflectors in multiple different regimes of dispersion¹³, and with at-will dispersion control¹⁵, these devices operate in reflection mode and face significant variation in efficiency across their bandwidth. In this work we investigate inverse-designed volumetric metastructures that operate in transmission and jointly tailor polarization splitting and chromatic dispersion in the mid-infrared. The devices are aperiodic, which minimizes their dependence on periodic diffractive principles and reduces the angular divergence of a broadband beam propagating from the devices; and operate in transmission mode, allowing for use in transmissive beam paths. We focus the main paper on two devices: a zero-dispersion (ZD) splitter that routes horizontally polarized (HP) light to $+30^\circ$ and vertically polarized (VP) to -20° ; and a hyper-dispersive (HD) beam deflector that directs HP light to -20° . Together they demonstrate that joint control of polarization and dispersion can be realized within the same device with high throughput.

4.2 Results and Discussion

The devices were designed in a pseudo-3D regime, where the optimized pattern is computed in the x - z plane and then extruded in the y -direction, as shown in Fig. 4.2a. This reduced representation retains longitudinal (z -axis) scattering pathways that are unavailable to a single-layer metasurface, while substantially lowering computational cost, allowing for faster and more in-depth prototyping during the design phase. The physical footprint of the device is $100 \times 100 \times 30 \mu\text{m}^3$, with a minimum feature size of $1 \mu\text{m}$ ($\sim \lambda/4$), resulting in a $100 \times 100 \times 30$ voxel design space.

A broadband Gaussian beam (FWHM $70 \mu\text{m}$) in the mid-infrared band ($\sim 4.3\text{--}4.7 \mu\text{m}$) is used to illuminate the device, composed of IP-Dip⁴⁰ ($n = 1.53$) against background air ($n = 1$), through the substrate (sapphire, Al_2O_3) from below. The objective function is calculated as the mode overlap between the transmitted far-field distribution, and pre-defined output modes in the far-field that correspond to beams propagating at the target angles for each wavelength and input polarization. Beginning with a randomized permittivity distribution as an initial seed, at each iteration, forward and adjoint simulations of the device were performed in Lumerical FDTD (ANSYS, Inc.), with perfectly-matched layer (PML) boundary conditions applied in the x -, y - and z -axes. The adjoint method was then used at each iteration of the optimization to obtain the sensitivity of the objective function with respect to the permittivity distribution, and an ADAM⁴¹-type gradient descent scheme was used to step the permittivity distribution forward. The device was then passed through

filters for binarization^{37,42} and feature-size control^{42,43} to guide the final pattern towards a fabrication-compatible binary structure.

Device optimizations were run on a high performance computing cluster (Resnick High Performance Computing Center) using 10 computing nodes, each of which was allocated 8 Intel CPU cores (mix of Skylake 2.1 GHz and Cascadelake 2.1 GHz processors), to run simulations in parallel. With these resources each simulation is completed in $\sim 2\text{--}3$ minutes, depending on the thickness of the device. Given 300 iterations, the optimizations complete in $\sim 10\text{--}15$ h of compute time.

We fabricated both devices using two-photon lithography (Nanoscribe GmbH, Germany) with the accompanying high-index-contrast polymer (IP-Dip) which has high transparency in this spectral band. Devices designed by this approach have been previously fabricated for the mid-infrared regime^{27,44} at subwavelength resolution, as well as for the microwave regime in patterned and layered polypropylene sheets.²⁶ We characterized the devices by illuminating them with a polarized beam from a tuneable light source, condensed down to $70\text{ }\mu\text{m}$ width, and imaged the back focal plane on a camera in order to visualize the far-field. Further details of the experimental setup are given in Section 4.4.2.

For the first device, an incident unpolarized beam illuminates the optimized metamaterial (depicted in Figs. 4.2a and 4.4a) from the substrate side. The horizontal polarization (HP) is deflected at an angle of $\theta_{\text{HP}} = +30^\circ$, and the vertical polarization (VP) is deflected at an angle of $\theta_{\text{VP}} = -20^\circ$. This asymmetry of deflection angles was chosen to demonstrate that the device can be optimized for arbitrary deflection angles θ_{HP} and θ_{VP} . The simulated results are displayed in Fig. 4.2. Fig. 4.2b is an overlay of the transmissions of the device for both polarizations in the far-field. For each polarization, the transmission is normalized against an input beam of the same polarization. The transmission map indicates strong directionality of the output beams, with both polarization channels remaining concentrated near their target angles throughout the design bandwidth. The angular width of the intensity peak is roughly $\pm 1.7^\circ$ (Figs. 4.2b and 4.2c), and Fig. 4.2c also visualizes the shift of the peak far-field angle across the wavelength. Notably, the change is continuously minimal across the entire bandwidth, even though in optimization only five discrete wavelength points between 4.3 and $4.7\text{ }\mu\text{m}$ (inclusive) were accounted for. In addition, the individual polarizations are separated with the same peak efficiency (60%). Relative to a regular grating with the same first-order deflection angle, the device exhibits much less absolute deviation from the target angle, and therefore less




Figure 4.2: (a) Schematic of inverse-designed zero-dispersion beam splitter, directing horizontally polarized (HP) light in one direction and vertically polarized (VP) light in the other. HP and VP polarizations can be split at arbitrary angles θ_{HP} and θ_{VP} . (b) Transmission map of the beam splitter, with HP and VP overlaid on the same plot. High intensity is seen across the entire spectrum, centered at $+30^\circ$ for HP and -20° for VP, with a spread of $\pm 0.3^\circ$ across the bandwidth. Transmission in each polarization case is normalized against the power of incident light of the same polarization on the input aperture of the device. (c) Transmission efficiency as a function of deflection angle across the device bandwidth. (d) Comparison of the dispersion of the beam splitter against a diffraction grating that directs light toward the same target angle. The deviation from the target angle at each wavelength is calculated by finding the far-field angle of maximum transmission at that wavelength. The designed beam splitter exhibits lower deviation (closer to 0°) from target angle compared to the given diffraction grating.

chromatic dispersion, across the bandwidth (Fig. 4.2d) for both polarizations.




Figure 4.3: (a) Schematic of inverse-designed hyper-dispersive beam deflector, directing horizontally polarized (HP) light with enhanced chromatic dispersion. (b) Transmission map of the beam splitter for HP polarization. High intensity is seen across the entire spectrum, centered at -20° , with a spread of $\pm 2.5^\circ$ across the bandwidth. Transmission is normalized against the power of incident light of the same polarization on the input aperture of the device. (c) Transmission efficiency as a function of deflection angle at $\lambda = 4.5 \mu\text{m}$. The position and width of the peak changes by 5.8° and 0.3° respectively across the bandwidth. (d) Comparison of the dispersion of the beam splitter against a diffraction grating that directs light toward the same target angle. The deviation from the target angle at each wavelength is calculated by finding the far-field angle of maximum transmission at that wavelength. The designed beam splitter exhibits higher deviation (farther from 0°) from target angle compared to the given diffraction grating.

The second device (Fig. 4.3a) is designed to deflect an incident horizontally polarized beam at an angle of $\theta_{\text{HP}} = -20^\circ$. Hyper-dispersive devices optimized to perform symmetrically and asymmetrically for orthogonal polarizations were designed, but their dispersive behavior in experiment was poor due to fabrication deviations; more detail on them is provided in Section 4.4.3. The simulated results for this device are displayed in Fig. 4.3. Fig. 4.3b is an overlay of the transmission of the device in the far-field, normalized against an input beam of the same polarization. The transmission map again indicates strong directionality of the output beams, with

efficiency remaining at $\sim 60\%$ over the bandwidth. The angular width of the intensity peak is roughly $\pm 3.1^\circ$ (Figs. 4.3b and 4.3c), and Fig. 4.3c also visualizes the shift of the peak far-field angle across the wavelength. There is a strong wavelength dependence of the central angle of the peak intensity. The magnitude of this dependence over the entire angular range varies depending on wavelength, but relative to a regular grating with the same first-order deflection angle, the device exhibits greater absolute deviation (Fig. 4.3d) from the target angle, and therefore higher chromatic dispersion, across the bandwidth.

4.2.1 Fabrication and Experiment

SEM pictures of the fabricated zero-dispersion (ZD) and hyper-dispersive (HD) devices are shown in Figs. 4.4a and 4.5a respectively. For visual inspection of the internal geometry, half of each device was fabricated in cross-sectional view, showing good agreement between the optimized designs and the printed internal structures. The extrusion direction of the structures (top view) is consistently straight throughout the full device.

The fabricated zero-dispersion device is shown in Fig. 4.4a, and the measured results demonstrate good agreement with simulation. In the case of both polarizations, the central angle of each beam (the median of all angular measurements taken over the entire bandwidth) is offset from the designed target angle, by -5° for HP polarization (Fig. 4.4d) and by $+3^\circ$ for VP (Fig. 4.4c). Stray transmission is noted at angles other than the central angle due to fabrication restrictions, but intensity remains high at the central angle throughout the bandwidth. Fits of the dispersion of the device (Fig. 4.4b) consistently demonstrate less deviation from the target angle, compared to a regular grating of the same deflection angle, across the wavelength span for both polarizations.

As for the fabricated hyper-dispersive device (Fig. 4.5a), the central angle of the output beam is shifted by -2° from the designed target angle. Less stray transmission is visible at angles far ($> 0^\circ$) from the central angle; this can be attributed to the device only being optimized to deflect light at a single angle, unlike the prior device. The dispersive behavior of the device is enhanced (Fig. 4.5b) relative to a regular grating of the same deflection angle.

4.2.2 Performance Bounds

For inverse design involving beam deflection or dispersion engineering, the design problem becomes increasingly difficult as one pushes toward larger angles and/or




Figure 4.4: (a) Upper left: Cross-section view of the design for zero-dispersion beam splitter after optimization. The design is binary, composed of IP-Dip (yellow regions) and air (dark blue regions). Lower left: Cross-section view of the fabricated zero-dispersion beam splitter. Only half of the device was fabricated so as to visualize the design after fabrication. Right: Top view of the fabricated zero-dispersion beam splitter. Extrusion is straight and the design does not bend along the extrusion axis. (b) Transmission map of the beam splitter in simulation. Across the spectrum, intensity is centered at $+30^\circ$ for horizontal polarization and -20° for vertical polarization, with a spread of $\pm 0.3^\circ$ across the bandwidth. (c) Experimentally characterized transmission map of the beam splitter. Transmission in the fabricated beam splitter is centered at a different angle ($+25^\circ$) than the target angle ($+30^\circ$) of the same splitter in simulation. (d) Comparison of the dispersion of the beam splitter against a diffraction grating that directs light toward the same target angle. The deviation from the target angle at each wavelength is calculated by finding the far-field angle of maximum transmission at that wavelength. The designed beam splitter exhibits lower deviation (closer to 0°) from target angle compared to the given diffraction grating.




Figure 4.5: (a) Upper left: Cross-section view of the design for hyper-dispersive beam splitter after optimization. The design is binary, composed of IP-Dip (yellow regions) and air (dark blue regions). Lower left: Cross-section view of the fabricated hyper-dispersive beam splitter. Only half of the device was fabricated so as to visualize the design after fabrication. Right: Top view of the fabricated hyper-dispersive beam splitter. Extrusion is straight and the design does not bend along the extrusion axis. (b) Transmission map of the beam splitter in simulation. Across the spectrum, intensity is centered at -20° for horizontally polarized light, with a spread of $\pm 2.5^\circ$ across the bandwidth. (c) Experimentally characterized transmission map of the beam splitter. Transmission in the fabricated beam splitter is centered at a different angle (-18°) than the target angle (-20°) of the same splitter in simulation. (d) Comparison of the dispersion of the beam splitter against a diffraction grating that directs light toward the same target angle. The deviation from the target angle at each wavelength is calculated by finding the far-field angle of maximum transmission at that wavelength. The designed beam splitter exhibits higher deviation (farther from 0°) from target angle compared to the given diffraction grating.

stronger engineered dispersion in a finite-thickness polymer structure. Since the incident beam used here is at normal incidence, imposing a different output beam angle means that the designed permittivity must, via multiple scattering through the vertical (z -) axis, change the k -vector of the light. Control of dispersion also requires that a different angular deflection must be imposed over the bandwidth of the device. In order to achieve these goals, either the design degrees of freedom must increase or a drop in overall device transmission must be accepted.

Due to constraints in fabrication discrepancy and time, the devices experimentally characterized have their deflection angle and dispersion strength reduced from what we were able to design in simulation. Examples of such devices are given in Section 4.4.3. In general, we were able to design devices with up to $\pm 45^\circ$ deflection angle for horizontal and vertical polarizations respectively, and almost zero dispersion; these devices are described in Figs. 4.9 and 4.10. We were also able to design hyper-dispersive devices that function at the same deflection angle and dispersion strength for both horizontal and vertical polarizations (Fig. 4.8). At the same time, devices with larger deflection angles and wider ranges of dispersion control could conceivably be achieved using the same geometry and design degrees of freedom. However, it is difficult to know when one has reached the global maximum in an optimization problem with a large number of design parameters such as these.

We can estimate how many design degrees of freedom are necessary by examining the performance bound of devices that allow for control of beam angle and dispersion. (The full derivation is given in Section 4.4.4.) For a conventional grating of period d , the angle of the first diffraction order is $\theta_0 = \arcsin(\lambda_0/d)$ at wavelength λ_0 . The phase required for such a first-order diffraction angle is

$$\phi(\theta_0, \omega_0, x) = -\frac{\omega_0 x}{c} \sin(\theta_0) \quad (4.1)$$

where $\omega_0 = 2\pi c/\lambda_0$. Then, a metaoptic of lateral span R that imposes dispersion $\partial\phi/\partial\omega$ obeys the relation

$$\left. \frac{\partial\phi}{\partial\omega} \right|_{\omega=\omega_0} = \frac{R}{c} \sin\theta_0 (\nu - 1) \quad (4.2)$$

Here, ν denotes a dispersion factor that indicates the dispersion of a device relative to a conventional grating of the same deflection angle θ_0 . With reference to Fig. 4.1, $\nu < 0$ applies for positive dispersion and $\nu = +1$ is the case for a conventional (regular) diffraction grating. Then, $\nu = 0$ applies for a zero-dispersion device and $\nu > +1$ for a hyper-dispersive device.

The time-bandwidth product $\Delta\omega\Delta T \leq \kappa$ sets the upper bound for the maximum phase delay achievable in a material^{22,45}. Here (derivation provided in Section 4.4.4), the time-bandwidth product obeys the relation

$$\Delta\omega \leq \frac{c\kappa}{R(\nu - 1) \sin \theta_0} \quad (4.3)$$

The fractional bandwidth $\Delta\omega/\omega$ is shown in Fig. 4.6a as a function of deflection angle θ_0 and dispersion factor ν . The closer the dispersion factor is to zero, the larger the fractional bandwidth (for example, a spread of $\Delta\lambda = 400$ nm at $\lambda = 4.5$ μm) becomes over which dispersion can be controlled, relative to a conventional grating. The same fractional bandwidth contours are overlaid over the maximum angular spread $\Delta\theta$, given in Fig. 4.6b. The maximum angular spread $\Delta\theta$ is inversely proportional to $\Delta\omega$. As such, we can increase the deflection angle range by either increasing the deflection angle θ_0 or increasing the dispersion factor ν .




Figure 4.6: (a) Contour plot of the fractional bandwidth for 4.3 μm to 4.7 μm . (b) Contour plot of the maximum achievable deflection angle range for 4.3 μm to 4.7 μm . Select contour lines of the fractional bandwidth from (a) are overlaid to illustrate the relationship between fractional bandwidth and angular range. (c) Maximum limit of fractional bandwidth $\delta\omega/\omega$ against deflection angle θ_0 . Plots for three different dispersion factors $\nu = 0, 3, 5$ are shown for devices of $100 \times 30 \mu\text{m}$, made from IP-Dip ($n = 1.53$). Smaller dispersion factors enable larger fractional bandwidths over which dispersion can be controlled. The devices displayed in this work are indicated on this plot. For each device, the larger deflection angle was used to situate them on the x-axis. Devices (1) and (3) are already near the maximum achievable bandwidth for their chosen deflection angle.

For a specific device, the practical bound of the fractional bandwidth is therefore imposed by the specific material index, device thickness, and the central intended steering angle. For the IP-Dip polymer used in the fabrication of our devices, the maximum achievable fractional bandwidth is shown as a function of the central steering angle in Fig. 4.6c.

We also plot the devices characterized in Figs. 4.4 and 4.5, as well as the devices designed in Section 4.4.3, on the same plot. Notably, the $\pm 30^\circ$ zero-dispersion device, designed for $4.0\text{--}5.0\ \mu\text{m}$ and 20° hyper-dispersive device, designed for $4.3\text{--}4.7\ \mu\text{m}$, are each located very close to the bound for their respective dispersion factor, for a device of the given geometry ($L = 30\ \mu\text{m}$, $R = 100\ \mu\text{m}$) and material index $n = 1.53$. This proximity to the theoretical bound motivates the use of volumetric inverse-designed structures in order to increase and make use of the available number of design parameters to achieve high performance when a device is required to control polarization, control, and other optical degrees of freedom.

4.3 Conclusion

The ability to select the dispersion regime and polarization response within the same device framework paves the way toward multifunctional optical elements that combine polarization, spectral, and spatial operations without sacrificing transmission efficiency. Demonstrating joint control of polarization and dispersion in transmission also expands the device landscape for compact inline optical systems, especially where beam routing must occur without introducing bulky optical elements or alignment difficulties associated with cascaded reflective optics. For example, by placing two of these devices in series, we envision a future beam displacer device, tuneable by shifting the relative distance between the two devices. Such a device could facilitate lateral translation of one polarization component of a beam while avoiding spatial separation with wavelength, and is described in Section 4.4.7. Future work would include multiplexing of circular and elliptical polarization, and engineering more exotic dispersion characteristics. We anticipate that our devices will provide the performance of conventional and bulky polarization beam splitters, with the compactness of metamaterials, and will be crucial components in future optical systems.

4.4 Supplementary Information

4.4.1 Adjoint Optimization and Projection Filters

The metastructures in this paper are all designed via inverse design using the adjoint method of photonic optimization. This approach^{26,29–31,34} defines a figure of merit correlated with device performance and efficiently computes its gradient with respect to the material permittivity distribution at every point in the design volume. Independently of the total number of design variables, this gradient requires only two electromagnetic simulations per iteration, and the design method can account for arbitrary constraints and requirements, e.g. overhead (fixed) material layers, trench enclosures, and/or platform substrate material.

The device is represented by a spatial permittivity distribution in the $x - z$ plane, where the z -axis is the direction of optical propagation. This permittivity distribution is then extruded uniformly along the y -axis. This pseudo-3D representation was adopted primarily to reduce computational cost while still allowing the design algorithm to affect the scattering of light in the direction of propagation by varying the permittivity profile in that direction.

The figure of merit (FoM) is chosen to be the mode overlap, in the far-field, between the transmitted, scattered electric and magnetic fields $\mathbf{E}$, $\mathbf{H}$ over a target area $\mathbf{S}$ and the desired output mode profiles $\mathbf{E}_m$, $\mathbf{H}_m$:

$$f(n(\vec{x})) = \frac{1}{8} \frac{\left| \int \mathbf{E} \times \overline{\mathbf{H}}_m \cdot d\mathbf{S} + \int \overline{\mathbf{E}}_m \times \mathbf{H} \cdot d\mathbf{S} \right|^2}{\int \text{Re}(\mathbf{E}_m \times \overline{\mathbf{H}}_m) \cdot d\mathbf{S}} \quad (4.4)$$

In this work, the output mode profiles are Gaussian beams in the far-field. A different target Gaussian is assigned to each input polarization state and wavelength, varying spatially in the x -direction according to the intended function of each device. For example, the zero-dispersion device is designed via five sets of simulations corresponding to five wavelength points equally distributed over the device bandwidth, and the output Gaussian for the same polarization is located at the same position. The optimizer then maximizes an aggregate objective FoM function (and gradient) that is a weighted average over the relevant design wavelengths and polarization channels. The figure of merit can be maximized over the intended target area $\mathbf{S}$ or minimized over all regions of the focal plane outside of $\mathbf{S}$, whichever is more appropriate given the device geometry and scattering environment.

For a single design wavelength and polarization, the gradient, or sensitivity $df/dn(\vec{x})$,

of the FoM with respect to the refractive index distribution can be constructed from the forward and adjoint simulation as follows:

$$\frac{df}{dn(\vec{x})} = 2n(\vec{x}) \operatorname{Re} [\mathbf{E}_{\text{fwd}} \cdot \mathbf{E}_{\text{adj}}] \quad (4.5)$$

where $\mathbf{E}_{\text{fwd}}$ is the electric field within the device region when illuminated from below the substrate by illumination of the desired mode profile and oblique incidence, and $\mathbf{E}_{\text{adj}}$ is the same electric field when the device is instead driven from the target plane by a corresponding source, sent backwards into the device. The phase and amplitude of the adjoint source are derived from the electromagnetic fields over the target area in the forward simulation^{30,31}.

The gradient (4.5) is then applied using the ADAM optimizer⁴¹:

$$\begin{aligned} m_t &= \beta_1 m_{t-1} + (1 - \beta_1) \left(\frac{df}{dn} \right)_t \\ v_t &= \beta_2 v_{t-1} + (1 - \beta_2) \left[\left(\frac{df}{dn} \right)_t \right]^2 \\ \alpha_t &= \alpha \times \frac{\sqrt{1 - \beta_2^t}}{1 - \beta_1^t} \\ n_{t+1} &= n_t + \frac{\alpha_t m_t}{\sqrt{v_t} + \epsilon} \end{aligned} \quad (4.6)$$

where the step size α is kept small, such that Eqn. 4.6 remains valid in the linear regime; and the learning rates β_1, β_2 are tuned for convergence. Upon convergence after multiple iterations, the design algorithm yields an optimized design (e.g. Figs. 4.2a, 4.3a, 4.9a, 4.8a, 4.12a) that achieves the target functionality with high efficiency.

At the end of the optimization, for fabrication purposes we require the final design to be a binary structure with only two values of refractive index n , in this case $n = 1.5$ and $n = 1$ for IP-Dip and air respectively. During the optimization, the optimizer operates on a continuous material-density field rather than directly on a binary refractive index distribution. The continuous field is mapped to the actual index distribution through a smoothed Heaviside projection filter^{26,37}:

$$\tilde{\rho} = \frac{\tanh(\beta\eta) + \tanh(\beta(\rho - \eta))}{\tanh(\beta\eta) + \tanh(\beta(1 - \eta))} \quad (4.7)$$

where ρ is the continuous density, $\eta = 0.5$ is the projection threshold, and β controls the projection sharpness. At small β the filter is approximately linear; as $\beta \rightarrow \infty$ it

approaches a step function that fully binarizes the density to either IP-Dip polymer or air. The strength β is increased gradually throughout the optimization, such that early iterations retain the flexibility of a continuous design space while late iterations are steered toward a fabricable binary structure.

Additional fabrication constraints are incorporated into the permittivity update, albeit at select iterations instead of as a differentiable continuous filter at each iteration. The nature of the extrusion of the design means that there cannot be any disconnected material islands which will fall out of the design. Every twenty iterations, disconnected material islands are identified and linked via bridges selected by Dijkstra's shortest-path algorithm, with traversal cost proportional to the distance each voxel must move in order to become fully solid. A greedy minimum spanning tree then determines the globally cheapest set of connections across all islands in the layer. Inserted bridge voxels are assigned a density of 0.75 and held fixed for the subsequent twenty iterations, giving the optimizer time to accommodate the new topology before the next patch cycle²⁷. Morphological opening (sequential erosion and dilation) is also applied every ten iterations to suppress isolated sub-resolution features, eliminating high-aspect-ratio pillars that would be mechanically fragile.

4.4.2 Experimental Characterization Setup




Figure 4.7: Experimental characterization setup used for direct- and Fourier-plane imaging of the fabricated beam splitter.

Each device was illuminated through the substrate with a collimated mid-infrared beam from a MIRcat-QT Quantum Cascade Laser (4.2–4.8 μm). A condensing lens

(L1, $f = 25$ mm) focused the beam onto the device input face through a $130\text{ }\mu\text{m}$ circular pinhole aperture, within which the $100\times 100\text{ }\mu\text{m}$ device was fabricated. Transmitted light was collected by a Thorlabs aspheric lens (L2, $f = 12.7$ mm) and a relay lens (L3, $f = 150$ mm) reimaged the back focal plane of L2 onto an FLIR A6780 MWIR thermal camera (640×512 pixels, $15\times 15\text{ }\mu\text{m}$ pixel pitch), such that the output angle of the meta-optic could be read out directly from position on the detector plane. L3 could be flipped out of the beam path to switch between Fourier and direct imaging configurations, the latter used for sample alignment. L2 was located sufficiently closely to the meta-optic that beams up to $\pm 30^\circ$ were within the numerical aperture of, and could be collected by, L2. All optics had diameter 1" and were sourced from Thorlabs.

Source output power was measured with a power meter across the bandwidth prior to BFP imaging. The total intensity summed over all pixels of the BFP image was taken to equal the power meter reading at that wavelength, and individual pixel intensities were scaled accordingly. Camera images were acquired at each wavelength with background (laser off) taken immediately afterward and subtracted, reducing errors from thermal drift of the camera housing.

In the Fourier imaging configuration, the angular coordinate was calibrated using two complementary methods. Firstly, reference gratings of spacing d , fabricated on the same substrate, were used as rulers, with known diffraction order maxima occurring at angles $\theta = \arcsin(m\lambda/d)$, providing a set of known pairs of pixel position and deflection angle. The first diffraction order was used as the primary calibration point at each wavelength, and higher orders were considered wherever their intensity was sufficient. Secondly, the origin $\theta = 0^\circ$ was established from the Airy diffraction pattern of a blank $130\text{ }\mu\text{m}$ circular aperture on the same substrate.

Due to chromatic aberration of L2, the effective Fourier plane was displaced by $\sim 40\text{ }\mu\text{m}$ across the measurement range of $4.2\text{--}4.8\text{ }\mu\text{m}$. To correct for this, the axial position of L2 was adjusted at each wavelength. The correct position at each wavelength was determined by maximizing the fringe visibility of the Airy diffraction rings produced by the blank $130\text{ }\mu\text{m}$ aperture. The corresponding micrometer position was recorded at each wavelength and applied during the device measurement throughout the bandwidth.

4.4.3 Broader Control of Deflection and Dispersion

The devices presented here document the broader simulation design space explored during the project. While they were not experimentally characterized due to difficulties and constraints with fabrication, they demonstrate that our inverse-design framework is capable of higher as well as asymmetric deflection angles, as well as control of dispersion in both polarization channels.

The first device (Fig. 4.8a) is a polarization-insensitive hyper-dispersive deflector, designed to route both HP and VP light to the same angle of $\theta = -20^\circ$ with enhanced chromatic dispersion relative to a grating with diffraction angle 20° . Unlike the single-polarization hyper-dispersive device in the main text, this device imposes the same angular dispersion on both polarization channels simultaneously, making it polarization-insensitive in both steering angle and dispersive response. The simulated results are displayed in Fig. 4.8. Figs. 4.8c and 4.8b show the transmission maps for HP and VP polarizations respectively, each normalized against an input beam of the same polarization. Both polarization channels exhibit strong directionality centered at -20° , with a spread of $\pm 10^\circ$ across the $4.0\text{--}5.0\mu\text{m}$ bandwidth. Peak efficiency is approximately $\sim 25\%$ for both polarizations (Fig. 4.8d), roughly half that of the single-polarization device in the main text. This efficiency cost arises from simultaneously addressing both polarization channels on top of dispersion response, as each additional functionality consumes design degrees of freedom that would otherwise contribute to transmission efficiency in the existing channels. Fig. 4.8e plots the dispersion of both polarization channels relative to a conventional grating of the same deflection angle, confirming that both channels exhibit greater deviation from the target angle across the full bandwidth.

The second device (Fig. 4.9) is a symmetric zero-dispersion (ZD) beam splitter, designed to deflect HP and VP light to equal and opposite angles of $\theta_{\text{HP}} = +30^\circ$ and $\theta_{\text{VP}} = -30^\circ$, over the $4.0\text{--}5.0\mu\text{m}$ band. In the main text, this was adjusted to the asymmetric $+30^\circ\text{--}20^\circ$ configuration to improve fabrication compatibility. This device demonstrates that we can achieve a balanced angular deflection between both polarization channels, and over a wider bandwidth of $4.0\text{--}5.0\mu\text{m}$ than the ZD device in the main text. Fig. 4.9b is an overlay of the transmission maps for both polarizations in the far-field, normalized against an input beam of the same polarization. Both channels remain concentrated near their target angles throughout the bandwidth, with angular peak widths of roughly $\pm 2^\circ$ (Figs. 4.9c). Peak efficiency is approximately $55\text{--}60\%$ for both polarization channels, with angular deviations




Figure 4.8: (a) Schematic of inverse-designed hyper-dispersive beam splitter, directing horizontally polarized (HP) and vertically polarized (VP) light in the same direction with high dispersion. (b, c) Transmission map of the beam splitter for VP and HP polarization respectively. High intensity is seen across the entire spectrum, centered at -20° for both polarizations, with a spread of $\pm 10^\circ$ across the bandwidth. Transmission in each polarization case is normalized against the power of incident light of the same polarization on the input aperture of the device. (d) Transmission efficiency as a function of deflection angle across the device bandwidth. Positions and widths of the peaks for each polarization change by 10° and 0.5° across the bandwidth. (e) Comparison of the dispersion of the beam splitter against a diffraction grating that directs light toward the same target angle. The deviation from the target angle at each wavelength is calculated by finding the far-field angle of maximum transmission at that wavelength. The designed beam splitter exhibits higher deviation (farther from 0°) from target angle compared to the given diffraction grating.

within roughly $\pm 1^\circ$ of the target across the band. Relative to a regular grating with the same first-order deflection angle, both polarizations exhibit less absolute deviation from the target angle across the full $1\ \mu\text{m}$ bandwidth (Fig. 4.9d).




Figure 4.9: (a) Schematic of inverse-designed zero-dispersion beam splitter, directing horizontally polarized (HP) light in one direction and vertically polarized (VP) light in the other. HP and VP polarizations can be split at arbitrary angles θ_{HP} and θ_{VP} . (b) Transmission map of the beam splitter, with HP and VP overlaid on the same plot. High intensity is seen across the entire spectrum, centered at $+30^\circ$ for HP and -30° for VP, with a spread of $\pm 2^\circ$ across the bandwidth. Transmission in each polarization case is normalized against the power of incident light of the same polarization on the input aperture of the device. (c) Transmission efficiency as a function of deflection angle across the device bandwidth. Positions and widths of the spikes for each polarization change by 1.8° and 0.4° across the bandwidth. (d) Comparison of the dispersion of the beam splitter against a diffraction grating that directs light toward the same target angle. The deviation from the target angle at each wavelength is calculated by finding the far-field angle of maximum transmission at that wavelength. The designed beam splitter exhibits lower deviation (closer to 0°) from target angle compared to the given diffraction grating.

The third device (Fig. 4.10) is a symmetric zero-dispersion beam splitter with wider angular separation, designed to deflect HP and VP light to equal and opposite angles of $\theta = \pm 45^\circ$ over the $4.3\text{--}4.7\ \mu\text{m}$ band. This device demonstrates that our inverse design framework can be pushed to significantly larger deflection angles while main-

taining near-zero chromatic dispersion. The transmission map (Fig. 4.10a), overlaying both polarization channels, shows that both remain tightly concentrated at their respective target angles across the full bandwidth, with peak efficiencies near 50% for both polarizations at the central wavelength (Fig. 4.10b). The dispersion comparison (Fig. 4.10c) shows that both channels maintain deviation below that of a reference grating of the same deflection angle across the band. As shown in Fig. 4.6c, operating at $\pm 45^\circ$ places this device close to the theoretical limit for a structure of this geometry and material index. However, optimizing for such a high deflection angle means that less device degrees of freedom are available to commit towards robustness against fabrication discrepancy, and as such this device was not characterized experimentally.




Figure 4.10: (a) Transmission map of the zero-dispersion beam splitter, with HP and VP overlaid on the same plot. High intensity is seen across the entire spectrum, centered at $+45^\circ$ for HP and -45° for VP, with a spread of $\pm 2^\circ$ across the bandwidth. Transmission in each polarization case is normalized against the power of incident light of the same polarization on the input aperture of the device. (b) Transmission efficiency as a function of deflection angle across the device bandwidth. Positions and widths of the spikes for each polarization change by 1.4° and 0.6° across the bandwidth. (c) Comparison of the dispersion of the beam splitter against a diffraction grating that directs light toward the same target angle. The deviation from the target angle at each wavelength is calculated by finding the far-field angle of maximum transmission at that wavelength. The designed beam splitter exhibits lower deviation (closer to 0°) from target angle compared to the given diffraction grating.

4.4.4 Calculation of Performance Bounds

This derivation is drawn from Arbabi et. al¹³ and Presutti & Monticone²², and is reproduced here for the reader's convenience.

For a conventional grating, the angle of the first diffraction order is given by

$$\theta_0 = \arcsin \frac{\lambda_0}{d} \quad (4.8)$$

for a grating of period d at wavelength λ . At a different wavelength λ , the deflection

is then

$$\theta' = \arcsin\left(\frac{\lambda \sin \theta_0}{\lambda_0}\right) = \arcsin\left(\frac{\omega_0 \sin \theta_0}{\omega}\right) \quad (4.9)$$

Then the difference in angular deflection is $\Delta\theta = \theta' - \theta_0$. In the linear regime for small $\Delta\theta$, $\Delta\theta \approx D(\omega - \omega_0)$ where

$$D = \left.\frac{\partial\theta}{\partial\omega}\right|_{\omega=\omega_0} = \nu D_0, \quad D_0 = \frac{-\tan \theta_0}{\omega_0} \quad (4.10)$$

To describe the strength of an arbitrary dispersion, we can express D in multiples ν of the regular (grating) above. We can then rewrite the deflected angle as

$$\theta'(\omega, \omega_0, \theta_0) = \theta_0 + \nu\Delta\theta = \theta_0 + \nu\left(\arcsin\left(\frac{\omega_0 \sin \theta_0}{\omega}\right) - \theta_0\right) \quad (4.11)$$

For a 1D blazed grating, the phase profile spatially in x required for a deflection angle θ_0 is¹³

$$\phi(\theta_0, \omega_0, x) = -\frac{\omega_0 x}{c} \sin(\theta_0) \quad (4.12)$$

and so for arbitrary θ, ω the metaoptic must impart phase

$$\phi(\theta, \omega, x) = -\frac{\omega x}{c} \sin\left(\theta_0 + \nu\left(\arcsin\left(\frac{\omega_0 \sin \theta_0}{\omega}\right) - \theta_0\right)\right) \quad (4.13)$$

Differentiating,

$$\begin{aligned} \frac{\partial\phi}{\partial\omega} &= -\frac{x}{c} \sin\left(\theta_0 + \nu\left(\arcsin\left(\frac{\omega_0 \sin \theta_0}{\omega}\right) - \theta_0\right)\right) \\ &\quad - \frac{\omega x}{c} \cos\left(\theta_0 + \nu\left(\arcsin\left(\frac{\omega_0 \sin \theta_0}{\omega}\right) - \theta_0\right)\right) \frac{\nu}{\sqrt{1 - \left(\frac{\omega_0 \sin \theta_0}{\omega}\right)^2}} \frac{-\omega \sin \theta_0}{\omega^2} \end{aligned} \quad (4.14)$$

If we evaluate the derivative at the central frequency $\omega = \omega_0$, this reduces to

$$\left.\frac{\partial\phi}{\partial\omega}\right|_{\omega=\omega_0} = \frac{-x}{c} \sin \theta_0 + \frac{x}{c} \cos \theta_0 \frac{\nu}{\sqrt{1 - \sin^2 \theta_0}} \quad (4.15)$$

$$= \frac{x}{c} \sin \theta_0 (\nu - 1) \quad (4.16)$$

The time delay-bandwidth product^{22,46} sets the upper bound for the maximum delay possible in a material:

$$\Delta\omega\Delta T \leq \kappa = \frac{\pi L}{\sqrt{3}\lambda_c} \left| \frac{\epsilon_{\max} - \epsilon_b}{\epsilon_b} \right| \quad (4.17)$$

where ΔT is the the relative group delay across the material, and L is the thickness of the device in the propagation direction.¹ If R is the lateral span across which the phase varies, we can write

$$\Delta T = \frac{\partial \phi}{\partial \omega}(x = R) - \frac{\partial \phi}{\partial \omega}(x = 0) \quad (4.18)$$

$$= \frac{R}{c} \sin \theta_0 (\nu - 1) \quad (4.19)$$

Therefore, the time-bandwidth product obeys the relation

$$\Delta \omega \leq \frac{cK}{R(\nu - 1) \sin \theta_0} \quad (4.20)$$

4.4.5 Dispersions for Varying Thickness

From Section 4.4.4, the time-bandwidth product bound (Eq. 4.17) is proportional to thickness L . A thicker device supports a larger available bandwidth $\Delta \omega$ over which dispersion can be controlled. Independently, a larger design volume (also proportional to L) makes more design degrees of freedom available, which allows the optimizer used to achieve higher transmission efficiency.

Fig. 4.11 confirms this trend for thicknesses between 10 and 50 μm . (All other devices covered in this work have thickness 30 μm .) Here, the devices used are adapted from the device displayed in Fig. 4.10. Both polarization channels show progressively narrower deviation from the target angle (Fig. 4.11a) as thickness increases, while transmission efficiency (Fig. 4.11b) remains high ($\sim 80\%$) and well-concentrated at the target angles throughout. However, at 50 μm the dispersion performance begins to worsen once again, and the transmission efficiency drops for both channels.

Partially, this may be due to the fact that the illuminating beam has a finite depth of focus; once the device thickness exceeds the Rayleigh range of the input beam, the regions near the entrance and/or exit faces of the device are poorly illuminated, depending on where the user places the focus of the input Gaussian beam. Effectively, this means that the number of voxels that the optimizer can effectively exploit is reduced. Though it is plausible that the drop in efficiency is due to absorption of the IP-Dip material as the light propagates through a thicker device, it is unlikely as the absorption is virtually zero in the bandwidth of the device⁴⁰.




Figure 4.11: (a) Comparison of the dispersions of $\pm 45^\circ$ zero-dispersion beam splitters of different thickness in the optical axis. (b) Transmission efficiency as a function of deflection angle across the device bandwidth.




Figure 4.12: (a) Cross-sectional view of the design for the hyper-dispersive beam splitter in Fig. 4.8, but without binarization or any other fabrication-compatible filters applied during the optimization. (b) Reproduction of Fig. 4.8d. Only the output corresponding to vertically polarized (VP) light is depicted here. (c) Transmission efficiency as a function of deflection angle across the device bandwidth for the design in a). Efficiency across the bandwidth is $\sim 2\times$ as high as in b). (d) Reproduction of Fig. 4.8b for VP light. (e) Transmission map of the beam splitter in a), c), only for VP light. The discretization of wavelengths across the bandwidth is much more evident.

4.4.6 Gradient-Index Devices

In photonic inverse design, optimizations that operate on a continuous (gradient-refractive-index, or GRIN) permittivity distribution consistently achieve higher figures of merit than their binarized counterparts^{35,47}. Imposing binarization means that the optimizer is restricted to a subset of the available permittivity landscape, and in highly multifunctional devices, this reduced design freedom can incur a substantial performance cost⁴². The binarization schedule used during the optimization also affects the final performance. As the Heaviside projection strength β is increased, gradients vanish in partially-binarized regions. If β is ramped too aggressively the figure of merit degrades in large discontinuous steps with little useful gradient information to guide recovery^{42,48}.

There are a few ways to bypass this issue. For example, direct level-set optimization operates on an inherently binary representation. However, it introduces a highly non-convex design space that is sensitive to initial conditions and prone to poor local optima⁴⁸. Various other techniques have been developed to mitigate this, including adaptive projection schedules⁴⁷, constrained density-based binarization⁴², dynamic binarization⁴⁹, and energy-constrained optimization⁵⁰. These approaches can approach the performance of corresponding GRIN devices, but do not generally reach it, especially for devices with a large number of simultaneous design objectives.

To quantify the efficiency cost of binarization for our devices, we designed polarizing beam splitters with the same FoMs as the $\pm 20^\circ$ hyper-dispersive device in Fig. 4.8, albeit with all binarization and fabrication-compatibility filters removed, allowing the optimizer to converge to a fully continuous GRIN structure (Fig. 4.12a). The resulting device (Fig. 4.12c) achieves roughly twice the absolute transmission efficiency of its binarized counterpart (Fig. 4.12b), with both polarization channels showing markedly stronger and more concentrated far-field peaks across the bandwidth. This suggests that the efficiencies displayed for all other devices in this work, which are binary devices, are not the highest achievable. It may be possible to use performance recovery schemes for a binary device, such as those discussed above, to approach the GRIN performance.

Relaxing these constraints also affects the dispersive behavior of the devices, as seen in Figs. 4.12d and 4.12e. The intended dispersion behavior is prescribed, as described in Section 4.4.1, by specifying the intended far-field distribution for five

¹ λ_c is the center wavelength in the background medium, and $\eta_{\max} = (\epsilon_{\max} - \epsilon_b)/\epsilon_b$ is the maximum relative permittivity contrast anywhere in the structure and band of interest.

discrete wavelengths across the bandwidth, and maximizing the mode overlap of the output field distribution with that specified. In the binarized device, the optimizer often fails to perfectly match the figure of merit at all five wavelengths simultaneously. As a result, the resulting transmission blurs smoothly across the band. In the GRIN device, however, due to the increased performance and wider permittivity space available, the optimizer is able to achieve a much closer match to the specified targets at each design wavelength, making the underlying wavelength discretization much more visible. This appears as five distinct horizontal bands across the bandwidth rather than a smooth continuum.

4.4.7 Stacking Beam Splitters




Figure 4.13: (a) Plot of E-field intensity for a beam displacer comprising a set of three copies of the device from Fig. 4.9, arranged as depicted. The input beam enters the first device on the bottom. Orthogonal polarizations are separated, and each is incident on a subsequent copy of the same device (flipped upside-down w.r.t. the optical axis). The outputs are beams of orthogonal polarization propagating in the same direction as the initial input beam. (b) Transmission map of the output of the beam displacer for horizontally polarized (HP) light. (c) Comparison of the dispersion of the beam displacer against a diffraction grating that directs light toward the same target angle (0°). The designed beam displacer exhibits lower deviation (closer to 0°) from target angle compared to the given diffraction grating.

One potential application for the polarizing beam splitters (PBS) described in this work is to combine them in series to construct compact broadband beam displacers. Figure 4.13a shows the simulated electric field intensity of such a device, comprising three copies of the $\pm 30^\circ$ zero-dispersion splitter in Fig. 4.9. The three device positions are indicated by the white dotted lines. An input beam enters the first splitter, which directs the horizontally (HP) and vertically polarized (VP) components to $+30^\circ$ and -30° respectively. Each of these beams then enters a subsequent copy of the same device, flipped upside-down with respect to the optical axis, which redirects its output back to 0° . This output beam is therefore parallel to the original input beam. The net effect is a lateral displacement of the beam, multiplexed by

polarization, without any change in propagation direction.

The transmission map (Fig. 4.13b) shows that the output (HP) beam remains tightly concentrated near 0° across the full bandwidth, with no significant broadening or distortion of the beam profile. The angular dispersion of the HP output (Fig. 4.13c) across the bandwidth remains well below that of a corresponding diffraction grating, which is consistent with the near-zero dispersion (Fig. 4.9d) of the individual splitters.

This demonstrates that the devices presented in this work can potentially serve as modular building blocks for more complex optical systems. The compactness and high efficiency of the inverse-designed splitters make them well-suited to cascaded configurations, and the preservation of dispersion performance through stacking suggests that such beam displacers could replace conventional, bulkier implementations in broadband inline optical systems.

References

1. Foo, I., Li, N., Pearson, P. & Faraon, A. Broadband Dispersion Engineered Volumetric Metamaterial for Polarization Beam Splitting at Large Deflection Angles. *Manuscript in preparation* (2026).
2. Born, M. *et al. Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light* ISBN: 978-0-521-64222-4 (Cambridge University Press, Cambridge ; New York, 1999).
3. Liu, B. *et al.* High-reliability infrared broadband thin-film polarizing beam splitter with ZnSe compensation layers. *Optics Express* **32**, 10910–10924. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-32-7-10910> (2026) (Mar. 2024).
4. Bennett, J. M. in *Handbook of Optics: Volume I - Geometrical and Physical Optics, Polarized Light, Components and Instruments* (ed Bass, M.) 3rd Edition (McGraw-Hill Education, 2010). ISBN: 978-0-07-149889-0. <https://www.accessengineeringlibrary.com/content/book/9780071498890/chapter/chapter13> (2026).
5. Chipman, R. A. in *Handbook of Optics: Volume I - Geometrical and Physical Optics, Polarized Light, Components and Instruments* (ed Bass, M.) 3rd Edition (McGraw-Hill Education, New York, 2010). ISBN: 978-0-07-149889-0. <https://www.accessengineeringlibrary.com/content/book/9780071498890/chapter/chapter15> (2026).
6. Saleh, B. E. & Teich, M. C. in *Fundamentals of Photonics* 3rd, 235–239 (Wiley, Hoboken, NJ, 2019). ISBN: 978-1-118-77007-8. (2026).

7. Faraji-Dana, M. *et al.* Compact folded metasurface spectrometer. *Nature Communications* **9**, 4196. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-018-06495-5> (2023) (Oct. 2018).
8. Wang, B. *et al.* Visible-Frequency Dielectric Metasurfaces for Multiwavelength Achromatic and Highly Dispersive Holograms. *Nano Letters* **16**, 5235–5240. ISSN: 1530-6984. <https://doi.org/10.1021/acs.nanolett.6b02326> (2026) (Aug. 2016).
9. Balthasar Mueller, J. P., Rubin, N. A., Devlin, R. C., Groever, B. & Capasso, F. Metasurface Polarization Optics: Independent Phase Control of Arbitrary Orthogonal States of Polarization. *Physical Review Letters* **118**, 113901. <https://link.aps.org/doi/10.1103/PhysRevLett.118.113901> (2026) (Mar. 2017).
10. Rubin, N. A., Zaidi, A., Dorrah, A. H., Shi, Z. & Capasso, F. Jones matrix holography with metasurfaces. *Science Advances* **7**, eabg7488. <https://www.science.org/doi/full/10.1126/sciadv.abg7488> (2025) (Aug. 2021).
11. Khorasaninejad, M., Zhu, W. & Crozier, K. B. Efficient polarization beam splitter pixels based on a dielectric metasurface. *Optica* **2**, 376–382. ISSN: 2334-2536. <https://opg.optica.org/optica/abstract.cfm?uri=optica-2-4-376> (2026) (Apr. 2015).
12. Arbabi, A., Horie, Y., Bagheri, M. & Faraon, A. Dielectric Metasurfaces for Complete Control of Phase and Polarization with Subwavelength Spatial Resolution and High Transmission. *Nature Nanotechnology* **10**, 937–943. ISSN: 1748-3395. (2026) (Nov. 2015).
13. Arbabi, E., Arbabi, A., Kamali, S. M., Horie, Y. & Faraon, A. Controlling the sign of chromatic dispersion in diffractive optics with dielectric metasurfaces. *Optica* **4**, 625–632. ISSN: 2334-2536. <https://opg.optica.org/optica/abstract.cfm?uri=optica-4-6-625> (2024) (June 2017).
14. Sell, D., Yang, J., Doshay, S., Yang, R. & Fan, J. A. Large-Angle, Multifunctional Metagratings Based on Freeform Multimode Geometries. *Nano Letters* **17**, 3752–3757. ISSN: 1530-6984. <https://doi.org/10.1021/acs.nanolett.7b01082> (2025) (June 2017).
15. McClung, A., Mansouree, M. & Arbabi, A. At-will chromatic dispersion by prescribing light trajectories with cascaded metasurfaces. *Light: Science & Applications* **9**, 93. ISSN: 2047-7538. <https://www.nature.com/articles/s41377-020-0335-7> (2024) (May 2020).
16. Chen, W. T. *et al.* A broadband achromatic metalens for focusing and imaging in the visible. *Nature Nanotechnology* **13**, 220–226. ISSN: 1748-3395. <https://www.nature.com/articles/s41565-017-0034-6> (2026) (Mar. 2018).
17. Shrestha, S., Overvig, A. C., Lu, M., Stein, A. & Yu, N. Broadband Achromatic Dielectric Metalenses. *Light: Science & Applications* **7**, 85. ISSN: 2047-7538. (2026) (Nov. 2018).

18. Lin, D., Fan, P., Hasman, E. & Brongersma, M. L. Dielectric Gradient Metasurface Optical Elements. *Science* **345**, 298–302. (2026) (July 2014).
19. Shalaev, M. I. *et al.* High-Efficiency All-Dielectric Metasurfaces for Ultracompact Beam Manipulation in Transmission Mode. *Nano Letters* **15**, 6261–6266. ISSN: 1530-6984. <https://doi.org/10.1021/acs.nanolett.5b02926> (2026) (Sept. 2015).
20. Wang, S. *et al.* A broadband achromatic metalens in the visible. *Nature Nanotechnology* **13**, 227–232. ISSN: 1748-3395. <https://www.nature.com/articles/s41565-017-0052-4> (2026) (Mar. 2018).
21. Colburn, S., Zhan, A. & Majumdar, A. Metasurface Optics for Full-Color Computational Imaging. *Science Advances* **4**, eaar2114. (2023) (Feb. 2018).
22. Presutti, F. & Monticone, F. Focusing on bandwidth: achromatic metalens limits. *Optica* **7**, 624–631. ISSN: 2334-2536. <https://opg.optica.org/optica/abstract.cfm?uri=optica-7-6-624> (2026) (June 2020).
23. Arbabi, A. & Faraon, A. Fundamental limits of ultrathin metasurfaces. *Scientific Reports* **7**, 43722. ISSN: 2045-2322. <https://www.nature.com/articles/srep43722> (2021) (Mar. 2017).
24. Sell, D. *et al.* Ultra-High-Efficiency Anomalous Refraction with Dielectric Metasurfaces. *ACS Photonics* **5**, 2402–2407. <https://doi.org/10.1021/acsphotonics.8b00183> (2026) (June 2018).
25. Miller, D. A. B. Communicating with waves between volumes: evaluating orthogonal spatial channels and limits on coupling strengths. *Applied Optics* **39**, 1681–1699. ISSN: 2155-3165. <https://opg.optica.org/ao/abstract.cfm?uri=ao-39-11-1681> (2026) (Apr. 2000).
26. Camayd-Muñoz, P., Ballew, C., Roberts, G. & Faraon, A. Multifunctional Volumetric Meta-Optics for Color and Polarization Image Sensors. *Optica* **7**, 280. ISSN: 2334-2536. (2020) (Apr. 2020).
27. Roberts, G. *et al.* 3D-patterned Inverse-Designed Mid-Infrared Metaoptics. *Nature Communications* **14**, 2768. ISSN: 2041-1723. (2025) (May 2023).
28. Ballew, C., Roberts, G. & Faraon, A. Multi-dimensional wavefront sensing using volumetric meta-optics. *Optics Express* **31**, 28658–28669. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-31-18-28658> (2025) (Aug. 2023).
29. Jensen, J. S. in *Wave Propagation in Linear and Nonlinear Periodic Media: Analysis and Applications* (eds Romeo, F. & Ruzzene, M.) 109–159 (Springer, Vienna, 2012). ISBN: 978-3-7091-1309-7. https://doi.org/10.1007/978-3-7091-1309-7_3 (2025).

30. Lalau-Keraly, C. M., Bhargava, S., Miller, O. D. & Yablonovitch, E. Adjoint shape optimization applied to electromagnetic design. *Optics Express* **21**, 21693. ISSN: 1094-4087. <https://www.osapublishing.org/oe/abstract.cfm?uri=oe-21-18-21693> (2020) (Sept. 2013).
31. Miller, O. D. *Photonic Design: From Fundamental Solar Cell Physics to Computational Inverse Design* PhD thesis (University of California, Berkeley). <http://optoelectronics.eecs.berkeley.edu/ThesisOwenMiller.pdf>.
32. Niederberger, A. C. R., Fattal, D. A., Gauger, N. R., Fan, S. & Beausoleil, R. G. Sensitivity analysis and optimization of sub-wavelength optical gratings using adjoints. *Optics Express* **22**, 12971–12981. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-22-11-12971> (2023) (June 2014).
33. Lalau-Keraly, C. *Optimizing Nanophotonics: from Photoreceivers to Waveguides* PhD thesis (University of California, Berkeley).
34. Molesky, S. *et al.* Inverse design in nanophotonics. *Nature Photonics* **12**, 659–670. ISSN: 1749-4893. <https://www.nature.com/articles/s41566-018-0246-9> (2026) (Nov. 2018).
35. Su, L. *et al.* Nanophotonic inverse design with SPINS: Software architecture and practical considerations. *Applied Physics Reviews* **7**, 011407. ISSN: 1931-9401. <https://doi.org/10.1063/1.5131263> (2026) (Mar. 2020).
36. Chung, H. & Miller, O. D. High-NA achromatic metalenses by inverse design. *Optics Express* **28**, 6945–6965. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-28-5-6945> (2025) (Mar. 2020).
37. Wang, F., Lazarov, B. S. & Sigmund, O. On projection methods, convergence and robust formulations in topology optimization. *Structural and Multidisciplinary Optimization* **43**, 767–784. ISSN: 1615-147X, 1615-1488. <http://link.springer.com/10.1007/s00158-010-0602-y> (2020) (June 2011).
38. Elsayy, M. M. R., Lanteri, S., Duvigneau, R., Fan, J. A. & Genevet, P. Numerical Optimization Methods for Metasurfaces. *Laser & Photonics Reviews* **14**, 1900445. ISSN: 1863-8880, 1863-8899. <https://onlinelibrary.wiley.com/doi/10.1002/lpor.201900445> (2021) (Oct. 2020).
39. Jiang, J. & Fan, J. A. Simulator-based training of generative models for the inverse design of metasurfaces. *Nanophotonics* **9**. arXiv:1906.07843 [physics.comp-ph], 1059–1069. ISSN: 2192-8614, 2192-8614. <http://arxiv.org/abs/1906.07843> (2026) (Feb. 2020).
40. Fullager, D. B., Boreman, G. D. & Hofmann, T. Infrared dielectric response of nanoscribe IP-dip and IP-L monomers after polymerization from 250 cm⁻¹ to 6000 cm⁻¹. *Optical Materials Express* **7**, 888–894. ISSN: 2159-3930. <https://opg.optica.org/ome/abstract.cfm?uri=ome-7-3-888> (2026) (Mar. 2017).

41. Kingma, D. P. & Ba, J. *Adam: A Method for Stochastic Optimization* Jan. 2017. <http://arxiv.org/abs/1412.6980> (2025).
42. Ballew, C., Roberts, G., Zheng, T. & Faraon, A. Constraining Continuous Topology Optimizations to Discrete Solutions for Photonic Applications. *ACS Photonics* **10**, 836–844. <https://doi.org/10.1021/acsp Photonics.2c00862> (2025) (Apr. 2023).
43. Guest, J. K. Topology optimization with multiple phase projection. *Computer Methods in Applied Mechanics and Engineering* **199**, 123–135. ISSN: 0045-7825. <https://www.sciencedirect.com/science/article/pii/S0045782509003302> (2026) (Dec. 2009).
44. Roques-Carmes, C. *et al.* Toward 3D-Printed Inverse-Designed Metaoptics. *ACS Photonics* **9**, 43–51. <https://pubs.acs.org/doi/abs/10.1021/acsp Photonics.1c01442> (2025) (Jan. 2022).
45. Tucker, R. S., Ku, P.-C. & Chang-Hasnain, C. J. Slow-Light Optical Buffers: Capabilities and Fundamental Limitations. *Journal of Lightwave Technology* **23**, 4046. <https://opg.optica.org/jlt/abstract.cfm?uri=jlt-23-12-4046> (2026) (Dec. 2005).
46. Miller, D. A. B. Fundamental Limit to Linear One-Dimensional Slow Light Structures. *Physical Review Letters* **99**, 203903. <https://link.aps.org/doi/10.1103/PhysRevLett.99.203903> (2026) (Nov. 2007).
47. Hammond, A. M. *et al.* Unifying and accelerating level-set and density-based topology optimization by subpixel-smoothed projection. *Optics Express* **33**, 33620–33642. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-33-16-33620> (2025) (Aug. 2025).
48. Christiansen, R. E. & Sigmund, O. Inverse design in photonics by topology optimization: tutorial. *JOSA B* **38**, 496–509. ISSN: 1520-8540. <https://opg.optica.org/josab/abstract.cfm?uri=josab-38-2-496> (2023) (Feb. 2021).
49. Butz, M., Abazi, A. S., Ross, R., Risse, B. & Schuck, C. Inverse design of nanophotonic devices using dynamic binarization. *Optics Express* **31**, 15747–15756. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-31-10-15747> (2025) (May 2023).
50. Zhang, G., Xu, D.-X., Grinberg, Y. & Liboiron-Ladouceur, O. Experimental demonstration of robust nanophotonic devices optimized by topological inverse design with energy constraint. *Photonics Research* **10**, 1787–1802. ISSN: 2327-9125. <https://opg.optica.org/prj/abstract.cfm?uri=prj-10-7-1787> (2025) (July 2022).

Chapter 5

VOLUMETRIC METAOPTIC COLOR SORTERS AT VISIBLE WAVELENGTHS: A PERFORMANCE ANALYSIS FOR PRACTICAL APPLICATIONS

This chapter is directly reproduced from [1], a paper published as a key part of the work in this thesis. As described in the abstract below, this work discusses the effects of hyperparameter sweeps on the performance of volumetric metaoptic color sorters. We examine also ways to improve the integration of such devices with existing image sensor arrays and enhance their performance.

Abstract

Three dimensional volumetric devices with refractive index structured at sub-wavelength scale enable custom sorting of light based on its spatial, spectral and polarization characteristics. One application is in efficient color sorting for red-green-blue color image sensors. Here we use inverse design to investigate how the performance of these color sorters depends on the complexity of the device. We take into account hyperparameters like the number of device layers, the minimum feature size, and the number of materials that can be utilized. We also analyze methods to optimize performance when light is incident from an oblique angle, and to attenuate crosstalk between sorters using trench isolation.

5.1 Introduction

Imaging systems that collect multispectral information must employ wavelength-selective elements to differentiate between light of each relevant spectral band. In many commercial image sensors that produce color images for the human eye, this is achieved by an arrangement of color-selective absorptive filters, placed above photodetector pixels, in order to quantify the amount of light in the blue, green, and red wavelength bands. Typically, this configuration consists of one red, one blue, and two green filters arranged in a 2×2 square, called a Bayer pattern². However, because light outside of each filter's passband is absorbed, the transmission efficiency of such Bayer filters can only ever reach an average of 33% under uniform broadband illumination, limiting both throughput and sensitivity.

Recent schemes involving metasurfaces have been proposed³ to improve both the ef-

efficiency and the multifunctionality of the Bayer filter arrangement, replacing the use of color filters by metaoptic-based color-splitting devices, where scattering structures accept light incident on a group of pixels and redirect each wavelength band to a different pixel. Metaoptic devices are an ideal platform for this approach as they are capable of achieving multiple different functionalities in one device⁴ with high efficiencies that are not easily reached using traditional optical elements⁵. The resulting devices can be easily adapted for monolithic integration and fabrication using well-known (e.g. CMOS) processes^{6,7}, and are well suited to optimization through inverse design⁸.

Metaoptic-based color-splitting devices can be broadly categorized into two different families. The first consists of two-dimensional, or “flat”, single-layer metasurfaces, where the single layer consists of an array of optical elements patterned in a plane, which is often orthogonal or close to orthogonal to the propagation direction of light. Prior work using these platforms^{9–15} can achieve compact color routing and performance superior to a traditional Bayer filter. However, confining the metaoptic to a single layer limits the number of optical modes available that can be exploited to realize novel optical designs with high performance^{16–19}. This is significant in the case of devices incorporating additional functions such as polarization sorting, where each new function and constraint has a corresponding cost in efficiency^{20,21}.

The second family consists of meta-optics patterned in a 3-dimensional design space, in which optical elements (or more generally, a permittivity distribution) are patterned throughout a finite volume or multiple layers instead of a single plane, expanding the design degrees of freedom and optical modes available according to the device volume^{19,22}. These devices^{19,22–27} span a great variety of geometries and configurations, and are generally able to achieve higher transmission efficiency and color contrast, but face difficulties in fabrication due to their increased complexity. A comparison of the relevant statistics of the referenced works in this study is provided in Section 5.4.1.

Finally, metaoptic-based color-splitting devices can also be distinguished between color sorters, which are devices with the detector plane located some distance from the device’s exit face, allowing for propagation of light between the exit face and detector plane; and color routers, in which light incident over the full device area is delivered directly to the intended photodetector without the need for an additional propagation distance^{3,26–28}. Color routers can, in principle, exploit nearly all incident light while simultaneously suppressing crosstalk and oblique scatter, making

them especially attractive for deep-subwavelength pixels where conventional microlenses and absorptive Bayer filters become increasingly vulnerable to diffraction and noise^{10,29}. However, this performance comes at the cost of larger design complexity due to the lack of propagation distance between the exit face and detector plane. In color sorters, that propagation distance allows for output wavefronts to separate, and as such a less extreme scattering within the metastructure is required. In many cases, color routers also require more demanding fabrication, due to potential errors in fabricated structure, layer-to-layer alignment, and vertical thickness, which are cumulative with the number of layers. At the present time, color sorters represent a compromise between optical efficiency and process compatibility that, in the near term, are able to provide improved sensitivity and reduced loss over conventional filter-based arrays.

The use of 3D platforms presents a challenging inverse design problem that is not compatible with conventional forward-design approaches due to the large amount of parameters involved and the complexity of the device response to any perturbations³⁰. Gradient-based inverse-design methods, particularly those based on the adjoint variable method, are therefore a natural framework for this problem, since they enable the design of nonintuitive permittivity distributions that use multiple scattering to realize spectral demultiplexing with high efficiency while incorporating fabrication and system-level constraints^{31–34}.

The higher number of design degrees of freedom available in 3D platforms also allow for flexibility in designing multifunctional color-splitting devices, with additional functionalities including polarization demultiplexing¹⁹, incident-angle sorting,³⁵ and Stokes polarimetry²² to enable advanced sensing applications. In this respect, inverse-design techniques are especially versatile and can be easily adapted to include these additional functionalities and constraints^{32,36–42}. Further methods are detailed in Section 5.4.2.

Previous devices designed by this approach have been fabricated^{22,43,44} at subwavelength resolution for the mid-infrared regime (3.5–5.5 μm), using two-photon lithography (Nanoscribe GmbH, Germany) with the accompanying high-index-contrast polymer (IP-Dip) which has high transparency in this spectral band; and in patterned polypropylene sheets, layered in a cube, for microwave frequencies in the K_a band¹⁹. The primary issue we have faced thus far in demonstrating multilayer metaoptics for real-world applications is the fabrication requirements for devices in the visible wavelength regime. Each layer of the multi-layer structure requires print

resolution and alignment accuracy on the order of the length of one voxel of the design, which is typically ~ 50 nm for visible wavelengths. Thus far, this is not achievable with our current technology of two-photon polymerization.

Other demonstrations of bilayer and multilayer metaoptic devices employ either the same two-photon lithography⁴⁵ or, more commonly, electron-beam lithography^{46–52} to pattern metasurfaces. The latter technique is capable of achieving small feature sizes below 100 nm⁵³, but becomes increasingly challenging to implement for more than two layers due to increasing time and alignment requirements. Where freeform 3D structures have been realized, they are often limited in the type of material available for use⁵⁴, or are confined to certain classes of structures⁵⁵. Depending on the fabrication approach, such constraints provide a practical bound on the performance of multilayer metaoptic devices that must be taken into account.

Given the large variety of geometries and device configurations present in the current literature, a systematic study of how geometry parameters, device configurations, and optimization hyperparameters affect performance is increasingly important for placing any single design in context. Kim et. al²⁵ performed one such study for 2D color sorters with RGBG subpixels in a 1D arrangement, using a genetic algorithm in order to access optimized designs close to the global maximum, and performing sweeps for parameters such as lateral device size, device thickness, number of vertical layers, and number of design pixels (cells) per layer. Most relevantly to this study, they found that for a fixed lateral device size, increasing thickness improves efficiency and then saturates; and that reducing lateral resolution (increased minimum feature size) reduces performance, which can be mitigated by increasing the number of vertical layers. Here, we validate these behaviors in the three-dimensional case, and present several different designs of 3D color filter arrays with 2×2 GBRG subpixel arrangement, while taking into account hyperparameters like the number of device layers, the minimum feature size, and the number of materials. We show that these hyperparameters have a large effect on the performance of multilayer metaoptics, and present several mitigation strategies for maximizing this performance under the constraints detailed above.

5.2 Results

To use as a baseline of comparison for the subsequent devices, we first design a 10-layer color sorting device, depicted in Fig. 5.1a, which is placed at a distance of $1.53 \mu\text{m}$ from the focal plane to focus normally incident light to different areas on

the focal plane depending on wavelength. Fig. 5.1b depicts the propagation of light of different wavelengths from the bottom (exit) face of the color sorter to regions of the focal plane that correspond to those wavelength bands. We numerically evaluate its peak transmission and color spectrum and study how this varies when different hyperparameters of the device are changed.

The device is composed of 10 vertically stacked layers, each layer with its own structure designed according to the adjoint optimization algorithm, with a randomized permittivity distribution as an initial seed. One of these layers is depicted in Fig. 5.1d. Each layer is composed of TiO_2 against background SiO_2 , two materials that at visible frequencies have an index contrast of 2.4:1.5 respectively, and are compatible with CMOS fabrication, with a minimum feature size of 51 nm ($\sim \lambda/10$), resulting in a $40 \times 40 \times 10$ voxel design space with a physical footprint of $2.04 \times 2.04 \times 2.04 \mu\text{m}^3$. The setup for evaluating the response of the color sorter is illustrated in Fig. 5.1a. The areas where light is focused into spots are determined according to the quadrants of a standard Bayer pattern (GBRG) color filter array, shown in Fig. 5.1a, with spatial boundaries matching the position and footprint of the device overhead. The device operates over the spectral band 400–700 nm, split into three bands corresponding to visible blue, green and red light with targeted central wavelengths at 430, 530 and 630 nm and spectral widths of 70, 55, and 65 nm respectively. Incident light that falls into each of these spectral bands is sorted to the corresponding quadrant in the focal plane.

At each iteration, forward and adjoint simulations of the device were performed in Lumerical FDTD (ANSYS, Inc.), with perfectly-matched layer (PML) boundary conditions applied in the x -, y - and z -axes. The ADAM optimizer⁵⁶ (learning rate 0.01) was used to apply the gradient obtained from the adjoint method to update the permittivity of the device at each iteration. Device optimizations are run on a high performance computing cluster (Resnick High Performance Computing Center) using 10 computing nodes, each of which was allocated 8 Intel CPU cores (mix of Skylake 2.1 GHz and Cascadelake 2.1 GHz processors), to run simulations in parallel. With these resources each simulation is completed in ~ 12 –15 minutes, depending on the thickness of the device. Given 300 iterations, the optimizations complete in ~ 60 –75 h of compute time.

Additional functionality (e.g. polarization sorting¹⁹, incident-angle sorting³⁵) has previously been co-optimized with color splitting devices by modifying the figure of merit and corresponding adjoint sources, but often results in an overall loss of ef-

iciency³⁵. As such, the devices examined here are solely optimized for RGB color sorting in order to maximize the sensitivity of the optimized response to design hyperparameters. The results of this study will be directly relevant to future design approaches for multifunctional color splitting devices and their applications, as the performance of the overall, multifunctional response will similarly depend on available design degrees of freedom and existing fabrication tolerances.




Figure 5.1: (a) Schematic of inverse-designed metaoptic color sorter device with 10 layers, and projected focal spots for red, green, and blue light on the focal plane. (b) Propagation of light through the color sorters is shown here from the side. When broadband light is incident on the top face (input plane) of the color sorters, light in the green (blue) band is sorted to the G (B) pixels directly under each sorter. (c) Transmission efficiency in each quadrant, defined as the power in each quadrant normalized over the power incident on the input aperture of the device. The wavelength range is divided into spectral bands corresponding to B, G, and R, and the peak transmission is recorded as the metric of performance for each band. The device is optimized for normal incidence light. (d) Selected layer of the 10-layer device after optimization. Each layer consists of TiO_2 (yellow) mesas embedded in an SiO_2 (dark blue) medium, forming a flat substrate for subsequent layers. (e) Power incident on the color sorter device can be scattered directly into power incident on the central region of the focal plane, characterized as transmission; or can be lost via the depicted scattering or loss channels. (f) Power flow analysis of 10-layer color sorter at the peak-in-band wavelengths, normalized against incident power. Numbers indicate mean values across the bandwidth.

5.2.1 Metrics of Performance

Fig. 5.1c shows the quadrant transmissions as a function of wavelength for the 10-layer device. The transmission for each quadrant is defined as the fraction of incident power reaching each quadrant, averaging over two separate simulations with x- and y-polarized incident light. This device, and all subsequent devices below, is insensitive to the polarization of the input light. We are able to achieve between 63% and 73% transmission at the peak transmission points for each quadrant. A standard color filter array can only ever achieve a maximum of 25% transmission for R and B, and 50% transmission for G, due to the absorptive filters that block out all other light for a given wavelength band.

Another important metric of the device performance is the power loss. The power lost from the device can be divided into four primary channels, defined in Fig. 5.1e: reflection, the power that escapes from the top or input aperture; side scattering, that escapes from the (vertical) sides of the color sorter and is incident upon the side faces of adjacent color sorters; and oblique scatter, which is power that leaves the exit (bottom) aperture, but is not incident on the central focal region. Instead it is directed to adjacent GBRG pixels that correspond to adjacent color sorters, modifying the intended scattering pattern there at the focal plane. A fourth loss channel is absorption. Within this wavelength range, we treat the materials that comprise the device as lossless, and assume the absorption is zero. (The effects of dispersion and absorption are discussed in Section 5.4.4.) The undesired scattered light results in optical and spectral crosstalk, contributing to image noise through such effects as desaturation and blurring^{57–63}.

All power values are normalized against the power incident on the top aperture from source. Fig. 5.1f shows the peak values of both device transmission and the above loss channels across the design wavelength spectrum. The baseline device successfully focuses 78% of the input power (averaged over wavelength) into the central region comprised of the 2×2 GBRG quadrants, while 9% is lost out the side as scattering that enters the sidewalls of adjacent color sorters. A further 11% exits the bottom aperture but does not fall into the central region, instead interfering with the focal spot pattern in adjacent pixels. Some of this effect is visible from the field norm at the focal plane depicted in Fig. 5.4d.

5.2.2 Material Constraints for Fabricability

To explore the design space and move towards more physically feasible devices for fabrication, we conducted sweeps of several hyperparameters (number of layers and feature size) and evaluated device responses (peak transmission and scattering) over these parameter sweeps. For each combination of device hyperparameters, a new, individual optimization was performed, and the response of the resulting new device was then calculated.




Figure 5.2: (a) Sorting transmission spectra in each quadrant as a function of layer number n , with total device height fixed at $2.04\mu\text{m}$, and minimum feature size 51nm . h indicates the height of each layer. (b) Peak sorting transmission values of the spectra in each quadrant as a function of layer number. (c) Mean scattering loss over the entire bandwidth of the devices in (a), as a function of layer number. Shaded regions indicate standard deviations. (d) Sorting transmission spectra in each quadrant as a function of lateral feature size, for a device with 40 device layers and total device height $2.04\mu\text{m}$. (e) Peak sorting transmission values of the spectra in each quadrant as a function of lateral feature size. Degrees of freedom available for optimization increase as the lateral feature size get smaller. (f) Scattering loss of the devices in (d) as a function of lateral feature size.

The number of layers is an important consideration in the total available permittivity voxels in the device, which scales linearly with layer number. Previous work^{19,25,26,28} has shown that total device height is an important factor in determining the efficiency and performance of multi-layer metaoptic devices. Devices with thicker overall height generally have higher transmission at their design wavelengths, while simultaneously increasing focal and side scattering due to thicker devices having

less natural angular acceptance (Section 5.4.3).

Here, we conducted sweeps of the number of layers with the total device height fixed at $2.04\text{ }\mu\text{m}$ (i.e. variable layer thickness), shown in Figs. 5.2b and 5.2c, from 5 layers to 40 layers, with the 10-layer device from the previous section as baseline, and all with a lateral minimum feature size of 51 nm. Increasing the number of layers tended to affect the peak transmission positively, since there were more degrees of freedom available in order to maximize focusing to the appropriate quadrants. The peak transmission of all bands improves sharply up to 10 layers, after which transmission becomes limited by the above loss channels, and the increase in performance flattens out. The crossover points of each spectral band also remain at the same location in terms of wavelength and transmission, which demonstrates that changing the number of device layers has the most effect on the peak transmission-in-band, and little effect out-of-band. Focal and side scattering remain roughly constant.

The number of degrees of freedom available for the optimization depends on the number of voxels in the design region, and therefore has an inverse squared relationship with the lateral minimum feature size, which we swept simultaneously in both (x,y) dimensions. All devices shown here were optimized with 40 design layers and a total device height of $2.04\text{ }\mu\text{m}$. Larger feature sizes are clearly responsible for a monotonic decrease in peak transmission (Fig. 5.2e), while increasing losses (Fig. 5.2f) through unwanted scattering channels. As the lateral feature sizes get larger, the peak transmission drops even farther, but the spectral shape and the overlap points stay roughly constant. For instance, when going from 51 nm features to 85 nm features, all three spectral bands experience a performance drop as shown in Fig. 5.2e: between 3–15% depending on the spectral band. If there is a requirement to use feature sizes larger than 51 nm, and still maintain the same amount of performance and transmission, using more device layers can alleviate this requirement. In the other direction, if there is a need to reduce the number of device layers (for example to simplify the fabrication process), decreasing the minimum feature size and thereby increasing the number of device voxels is a viable strategy to maintain performance.

5.2.3 Three-Material Device

Our color sorters are fully binary devices composed of two materials (up to this point a TiO_2 / SiO_2 matrix). The devices obtained through the standard ADAM optimization are of continuous grayscale density, and we enforce a binary permittivity




Figure 5.3: (a) Sigmoidal projection filter that projects device density to refractive index distribution for two materials, shown in Eq. (5.1). Red: at weak strength, the filter merely scales the density to the refractive index bounds. Blue: with increasing strength, the filter is an increasingly sharp Heaviside function. (b) Three-tiered sigmoidal projection filter (5.2) projecting device density to refractive index distribution for three materials. (c) Fully-projected refractive index distribution for layer 2 of the resulting 3-material device. (d) Sorting transmission spectrum into B, G, R quadrants for a 5-layer device with 51 nm lateral feature size. Grey lines represent the overall device transmission, which is the sum of transmissions into the B, G, R quadrants. (e) Sorting transmission spectrum into R,G,B quadrants for a 5-layer device with 85 nm lateral feature size. (f) Scattering and loss channels for power outflow from the device for 5 layers, 85 nm lateral feature size and 2(3) materials. Numbers indicate mean values across the bandwidth.

by passing this density through a smoothened Heaviside projection filter²⁶:

$$\tilde{\rho} = \frac{\tanh(\beta\eta) + \tanh(\beta(\rho - \eta))}{\tanh(\beta\eta) + \tanh(\beta(1 - \eta))} \quad (5.1)$$

where the density ρ is projected toward 0 for $\rho < \eta$ and toward 1 for $\rho > \eta$ with increasing projection strength β . For a two-material filter, we set $\eta = 0.5$ i.e. exactly halfway between the density bounds of 0 and 1. At low values $\beta \rightarrow 0$, the projection filter is linear (Fig. 5.3a, red line) and simply scales the density directly to the intended refractive index range, but remains continuous. At high values $\beta \rightarrow \infty$, the projection filter approximates a step function (Fig. 5.3a, blue line) and binarizes all intermediate density values to 0/1.

Significant performance drops are observed whenever the strength β of the binarization filter is ramped up (Fig. 5.7c), as projecting the device's grayscale density

through the Heaviside filter perturbs the loss landscape^{19,22}. A particular device configuration that was previously in a local minimum of the prior loss landscape is unlikely to be in a similar minimum of the new, perturbed loss landscape, which typically reduces the device's performance as a result. This is only a temporary problem as the optimizer, in most cases, is able to quickly recover its previous performance within a few iterations. Since volumetric inverse-designed devices function based on multiple scattering between the separate and individual component layers, allowing for a greater number and diversity of materials in the design improves the diversity of resonances and scattering that can be achieved within the device. From an optimization standpoint, because the optimizer has more options in the three-material case, it has greater versatility in choosing a material distribution that can scatter light of different wavelengths in the precise way that is optimal for the required figure of merit. This increased performance is especially significant when considering color sorters of low layer number and coarse feature size, which is necessary depending on the limitations of the lithography method being used to fabricate such devices.

To change this implementation from two materials to three materials, we change the shape of the projection filter (Fig. 5.3b), given below:

$$\tilde{\rho} = \frac{\tanh(\beta\eta) + \frac{1}{2} [\tanh(\beta(2\rho - \eta)) + \tanh(\beta(2\rho - 3\eta))]}{\tanh(\beta\eta) + \tanh(\beta(1 - \eta))} \quad (5.2)$$

Here the permittivity of the third material is taken to be halfway between the initial two, but can be straightforwardly extended to any number and type of materials.

We first directly compare two color sorters, designed with five layers and 51 nm lateral feature sizes, with two and three materials respectively. In both cases the devices are composed of TiO_2 ($n = 2.4$), SiO_2 ($n = 1.5$), and in the latter case a third intermediate material which we have set at a refractive index halfway in between (i.e. 1.95), but that can be adjusted as necessary to fit fabrication compatibilities. Several materials, including SiN, HfO_2 , and AlN, are CMOS-compatible and have refractive indices $\sim 1.9 \leq n \leq \sim 2.1$ over the device bandwidth, and could be used for this purpose. The corresponding refractive index distribution is shown in Fig. 5.3c. The transmission spectrum (Fig. 5.3d) of the three-material device features a 3-12% absolute increase in the peak transmission depending on the spectral band of interest. There is also improved out-of-band suppression, which directly improves the color contrast of the device. Scattering and loss channels are largely unaffected by the inclusion of additional materials to the projection filter, and the average device transmission increases only by 2%.

This trend is observed also for color sorters designed with five layers and coarser 85 nm lateral feature sizes. Here the absolute increase is 5-8% depending on the spectral band, as shown in Fig. 5.3e. Fig. 5.3f shows the changes in the transmission and scattering channels, averaged across the design wavelength range. The former increases by about 5% in this case, and the latter channels are largely unaffected.

In projecting to 3 materials, in-band transmission is increased, and off-band transmission is suppressed at the peak wavelengths of other bands. The in-band amplification is most clear for devices with 51 nm lateral feature size, whereas for devices with 85 nm feature size the primary improvement is in off-band suppression. The improvements in performance are also clearest for the blue wavelength band. By including a third material, we introduce new interfaces within the device of different refractive indices, at which shorter wavelengths (blue) experience larger angles of refraction from the normal than longer wavelengths (red). As such, the scattering effects that can be leveraged by the inverse-design optimization have a wider range, leading to a higher figure of merit (in-band transmission).

5.2.4 Crosstalk Attenuation between Color Sorter Pixels

Optical and spectral crosstalk between adjacent devices is a persistent problem in pixel design^{57–59} that leads to spectral overlap and signal noise in the color information as a result of a reduced ability to distinguish between spectral bands. This has subsequent impacts on image resolution and can contribute to the formation of artifacts after processing. The types of crosstalk relevant to our color sorter design are detailed above in Fig. 5.1e and the accompanying discussion.

In many conventional CMOS image sensors, deep trench isolation is a technique used to attenuate optical crosstalk between adjacent pixels^{64–67}. Trenches are etched in between pixel stacks, or unit cells, and these trenches are filled with a material that either absorb or reflect electromagnetic fields within each pixel stack and prevents leakage to adjacent stacks. Here, we can implement a similar strategy by including a sidewall material in a rectangular configuration surrounding each cuboid color sorter device. An example of such a configuration is shown in Fig. 5.4a. In order to minimize scattering between the sides of adjacent color sorters, the sidewalls completely cover the sides of each color sorter. They also extend down towards the focal plane to block oblique scattering from the exit face of one color sorter to the detector regions of adjacent color sorters. A side view of the color sorter and sidewall configuration is given in Fig. 5.4b.




Figure 5.4: (a) Proposed implementation of air sidewalls, placed in a rectangular configuration around the color sorter. Here only three of the four sidewalls are shown for visual clarity. Depending on existing fabrication capabilities, the trenches may not necessarily extend all the way down to the focal plane. (b) Propagation of light through the color sorters, with air sidewalls present, is shown here from the side. Wavy arrows indicate possible paths by which the included air sidewalls confine lossy light into the correct pixel stack. (c) Test configuration for modelling the effect of adjacent devices on the color sorter scattering pattern at the focal plane. Nine identical color sorters are arranged in a 3×3 pattern, and incident light is focused tightly into the top face of the central sorter. (d) Field norm at the focal plane for a 3×3 array of color sorters, with incident light focused only into the central sorter. (e) Field norm at the focal plane for the same but optimized with air sidewalls in between each color sorter, in the configurations depicted in (a) and (b). Note the lower intensity in the pixels to the north, south, east, and west. (f) Power flow analysis of 10-layer color sorter with air sidewalls of thickness 85 nm and depth 1.148 μm , and feature size 51 nm. Numbers indicate average values over the wavelength range. (g) Evolution of mean focal region power and scattering losses as sidewall depth is changed from 62.5%-100% of the focal length (1.53 μm). Dashed lines are the same for an identical device, except optimized with no sidewalls. (h) Mean focal region power and scattering losses for different material fills of sidewalls, of thickness 85 nm and depth 1.53 μm .

In order to model the effect of side and oblique scattering on the performance of a color sorter, a supercell of nine identical devices was assembled as shown in Fig. 5.4c, in a 3×3 array, above a corresponding 3×3 array of pixels at the focal plane. Each color sorter in this array is optimized for normal incidence. The central color sorter was excited from above with a Gaussian wave source at normal incidence, with waist diameter restricted to the side length ($2.04 \mu\text{m}$) of the color sorter such that most of the incoming light is focused primarily on the central color sorter. An aperture directly above the central color sorter blocks all light that might be incident on adjacent devices in the supercell, so as to ensure they receive no additional excitation. The effect of side and oblique scattering on the performance of adjacent devices is demonstrated in Fig. 5.4d. For all bands, asymmetry is observed in the focal spots; for the green and red bands, a significant amount of light is scattered outside of the central pixel region. Spectral crosstalk is especially evident for the green wavelength band, where light is present in the subpixels corresponding to other wavelength bands.

To maximize the optical confinement within each pixel stack, the trench material is typically of a lower refractive index than the background material^{68,69}. Air ($n = 1$) sidewalls have a high index contrast against the background material (SiO_2 , $n = 1.5$) of our setup. As such, we chose air as the medium for sidewalls between color sorters and designed new devices, identical to the device first presented in Fig. 5.1a, but incorporating the above sidewall configuration. The inclusion of air sidewalls was seen to reduce side and oblique scattering by 2.8% and 6.7% respectively, allowing for 81% focal transmission from the optimized device, as shown in Fig. 5.4f. Fig. 5.4e shows that much of the power incident on the focal plane is now confined to the central pixel and the shape of the focal spots in each subpixel is more isotropic.

The depth of the sidewalls is a constraint that depends heavily on the fabrication capabilities available. For example, although an air sidewall that extends all the way down to the focal plane would give the maximum possible confinement of light to each pixel stack, the presence of electronic interconnects may prevent the etching of a sidewall to that depth. We explored the effect of sidewall depth on focal transmission and scattering losses, which is shown in Fig. 5.4g. Extending the trench depth from 75% of the focal length (the distance between the exit face of the color sorter and the focal plane) to 100% of the focal length gives only a 4.2% increase in focal transmission, and decreases side scattering and oblique scattering by 2.1% and 1% respectively. As such, increasing trench depth beyond 75% of the focal length (here

1.48 μm) offers little benefit.

Finally, the material used to fill the trenches is critically important in determining the consequent effects on device performance. Often, the material used is chosen for its optical properties as well as appropriateness in fabrication. Although air has previously been used for trench isolation with success in attenuating crosstalk^{70–72}, in general using solid dielectric trench fills (such as poly-Si⁷³) and/or conductive liners is more common. Fig. 5.4h shows the scattering distribution for devices designed in the same way as in Fig. 5.4a, but with air replaced by a different material. Tungsten (W) has been previously explored²⁴ as a sidewall material for a color router device, while copper (Cu) provides another example of the effects of a metal sidewall relative to a bare device: side scatter and reflection are minimized, and power is mainly directed out of the device exit face. Of all these, air provides the highest amount of power transmitted directly to the focal region, whereas other materials have non-zero extinction coefficients in the device bandwidth and impose power losses due to absorption. Depending on the material used, the use of sidewalls to achieve deep trench isolation for metaoptic color sorters is a viable strategy to confine power within pixel stacks and prevent unwanted noise between pixels at the focal plane.

5.2.5 Oblique-Incidence Optimizations

The performance and calculations performed above are in consideration of a single color sorter, operating in isolation. However, each device only sorts light for a single 2×2 Bayer pixel. If these color sorters are to be implemented in an actual device, they must be implemented as an array, placed directly above a corresponding array of sensor pixels. Under focused light from a preceding imaging lens, this arrangement receives a range of angled components that have to be accounted for depending on the spatial position in the device array and numerical aperture of the lens. In other words, different color sorters in the array receive differently angled components that have to be accounted for depending on their spatial position in the array and the numerical aperture of the lens. Under normal circumstances there is a significant performance drop-off (Fig. 5.5a) the more oblique the incident angle, due to the increasing difficulty of mode-matching higher angular components into light that is aimed downwards and focused into the correct quadrants. The angular range (here defined as the angle at which the transmission drops to 50% of that at normal incidence) is about $\theta = \pm 7^\circ$, which informs the number of re-optimizations that are necessary to ensure consistent performance across an array of such devices.




Figure 5.5: (a) Transmission into each quadrant, normalized against input power, for a device optimized at normal incidence. Each data point is the peak transmission for a quadrant over its spectral band, evaluated for a source with incident angle θ deviated from the normal, with no re-optimization performed. The polar angle θ is defined as shown in the inset of (c). (b) Transmission spectra for two devices optimized for incident light at $\theta = 0^\circ$ and $\theta = 30^\circ$. The differences in peak performance for the blue, green, and red transmission spectra, as well as overall transmission, are indicated above the figure. (c) Comparison of peak transmission in the green band as the incidence angle θ of incoming light is deviated from the original angle of optimization. The polar angle θ and azimuthal angle ϕ are defined as shown in the inset. Here we compare nine devices optimized at the same polar angle $\theta = 5^\circ$, and each with different azimuthal angles ϕ spaced at 40° increments between 0° – 360° . The green band is bounded by the maximum and minimum transmission values of the nine devices as the incident polar angle is varied. Across all devices, there is only a small deviation in device performance and angular width. A device optimized at $\theta = 0^\circ$ is shown in black for comparison. (d–f) Peak transmission efficiencies for the blue (400–490 nm), green (490–565 nm) and red (565–700 nm) wavelength bands as functions of the divergence angle of the incident source. Each line corresponds to a device optimized to function at the indicated source divergence angle.

Here, we extend the concept to design devices that operate specifically at non-normal angles of incidence and have robust performance with respect to angular deviation. We have successfully used the adjoint method of optimization to design splitters that function with high performance over a broad range of angles, centered on an oblique angle of incidence. The sorting efficiency of these devices is shown in simulation in Fig. 5.5b for central incident angles of 0° and 30° . The latter device optimized at $\theta = 30^\circ$ has comparable performance to the former device optimized at normal incidence, with a small drop in peak transmission of -2% and -1% for the blue and green band, and improved peak transmission of +4% in the red band. Overall, the average transmission across the entire spectrum goes up by 2.5%. The devices are also robust against variation of the azimuthal angle ϕ (Fig. 5.5c) at $\theta = 5^\circ$. A set of nine devices were optimized at 40° increments of ϕ around a cone with polar angle $\theta = 5^\circ$ relative to normal incidence. Here, the polar angle is the angle between the wavevector of the incident source and the z-axis, and the azimuthal angle the angle between the x-axis and the projection of the incident wavevector onto the xy-plane, as shown in the inset of Fig. 5.5c. At the intended incidence angle $\theta = 5^\circ$, the lowest-performing device has only 4.6% lower transmission (averaged across the operating bandwidth) than the highest-performing device. This performance gap widens to a maximum 13.7% across the entire angular range explored in Fig. 5.5c. Across all devices, there is only a small deviation in device performance and angular width, although we expect this deviation to increase as the polar angle θ increases, which may necessitate more azimuthal re-optimizations to maintain the same invariance.

We are also able to adapt our device to a wide variety of arbitrary imaging systems without difficulty. In simulation, we set the input source to a finite Gaussian beam, changing the divergence angle to simulate focused spots from lenses of different numerical apertures contingent with typical camera smartphone lenses, which feature divergence angles of roughly $13 \pm 5^\circ$ in the background material SiO_2 of our color sorter device. Figs. 5.5d to 5.5f show the angular robustness of devices optimized to each function at four different values of divergence angles, 8° , 13° , 18° , and 23° . The divergence angle of the incident Gaussian source is then swept for that device to simulate its performance robustness under imaging systems of different, suboptimal numerical apertures. For all three wavelength bands, the maximum transmission recoverable for such re-optimizations ranges between 93%-107% of the peak transmission at divergence angle 13° , and the (divergence) angular range for all devices is $\sim \pm 14^\circ$. This ability to tune the central angle of the device without noticeably sacrificing performance, for a wide variety of preceding imaging systems, directly

enables the implementation of color sorting devices in practical scenarios.

5.3 Conclusion

A compact three-dimensional color sorter for RGB spectral splitting in the visible wavelength regime (400-700 nm) is thoroughly investigated. The splitter has lateral footprint $2.04 \times 2.04 \mu\text{m}$ and height $2.04 \mu\text{m}$. Although such color sorters are capable of high efficiency and multifunctionality, in order to both meet the limitations of current lithography methods and at the same time surpass the current transmission limit (25% for blue and red, and 50% for green) per spectral band inherent to a color-filter-based array, there must be severe constraints on the dimensions and lithographic resolution of 3D metaoptic color sorters. Primary among these are the number of design layers and the minimum feature size of the voxels within each layer. Our inverse design method is able to achieve 50%, 52%, and 59% transmission for blue, green, and red spectral bands respectively, for a device of 5 layers and 51 nm feature size, which are near the limits of currently existing lithography methods^{52,74}. By extending the number of allowed materials in the color sorter to three or more, this performance goes up to 63%, 58%, and 62%, and could be achieved by any fabrication method that could accommodate materials of three or more refractive indices in a device, or make use of gradient refractive indices.

Separately, in order for these color sorters to be used in a realistic imaging system, cross-talk between adjacent sorters must be minimized to preserve imaging quality. An array of such sorters above a Bayer-type sensor array must be able to robustly accept a range of incident angles from any preceding optics, and this range may vary depending on the position of each sorter in the array. We are able to straightforwardly extend our inverse design method to accommodate the presence of dielectric trenches in between pixels to minimize such cross-talk, with unchanged performance in the optimized device's ability to split colors. Incoming light will also impact an array of pixels at different incident angles and divergence angles depending on the preceding optics, and our inverse design method demonstrates similar levels of performance when optimizing for incoming light of a wide variety (up to 30°) of incident angles and a wide range (8° - 23°) of divergence angles.

In practice, the most plausible routes toward wide-scale manufacturing of these devices are layer-by-layer fabrication processes³. At each step, a single layer of a design is formed by depositing and patterning high-index material laterally, then filling in the background with a lower-index material and planarizing. This flat layer then

forms the substrate for subsequent layers of the design. In schemes like these, while additional layers increase the number of available degrees of freedom and improve device performance, they also increase the difficulty of the multilayer fabrication process through higher required alignment precision, especially between multiple layers, and cumulative thickness discrepancies in the vertical direction. Figs. 5.2b and 5.9a also show that performance gains begin to saturate at higher layer counts.

Future work will involve experimental realization of these designs via novel freeform fabrication processes for multilayer lithography and characterization of their fabrication tolerance and performance, as well as improved control of the placement of both features of different sizes and materials. Image sensors and other devices that are capable of sorting light based on a wide variety of properties lend themselves well to novel applications such as optical microscopy, remote sensing and computational imaging. The results presented here support the realization and practical use of these devices in modern technology.

5.4 Supplementary Information

5.4.1 Comparison of Prior Color-Splitting Devices

Table 5.1: The below tables are comparisons of physical parameters and performance metrics of previous metaoptic-based color splitting devices referenced in this work. In the columns describing simulated and measured performance, all values are peak values within each band unless stated otherwise:

^a: average over band; ^s: sum or integral over band; ^p: value at single wavelength point.

Ref	Layers	Thickness (μm)	Pixel Pitch (μm)	Lateral Size (μm)	Focal Plane Distance (μm)	Min Feature Size Lateral (nm)	Material / Background
[15]	1	0.6	25.0	50×50	110	50 (pillar width)	GaN / air on sapphire
[13]	1	1.25	1.6	448×448	4	400 (meta- atom period)	SiN / air on fused silica
[9]	1	0.6	1.0	2.0×2.0	—	125	Si ₃ N ₄ / air on quartz
[12]	1	0.3	0.8	1.6×1.6	2	80	TiO ₂ / air on SiO ₂
[22]	6	18	15.075	30.15×30.15	25	750	IP-Dip / air on sapphire
[10]	1	0.6	0.6	1.2×1.2	0.6	50	SiN / air on glass
[11]	1	0.998	1.12	2.24×2.24	4	150	Si ₃ N ₄ / air on SiO ₂
[75]	1	0.6	1.1	2.2×2.2	4	100	TiO ₂ / air on SiO ₂
[24]	50	1	0.7	1.4×1.4	0	20	TiO ₂ / SiO ₂
[14]	1	0.6	1.1	3.3×1.1	4	100	TiO ₂ / air on SiO ₂

Ref	Channels	Added Functionality	Opt. Algorithm	Power Norm.	Simulation (%)	Measured (%)
[15]	2×2 Bayer (RGBG)	—	Library Design	Incident power on device	—	R 15.9 ^p / G 1 37.9 ^p G 2 27.6 ^p / B 38.3 ^p
[13]	2×2 Bayer (RGBG)	—	Library Design	Incident power on all pixels	R 65 ^s / G 50 ^s B 25 ^s / T 77 ^a	R 48 ^s / G 48 ^s B 12 ^s / T 65 ^a
[9]	2×2 Bayer (RGBG)	—	Genetic algorithm	Incident power on device	R 61 / G 57 B 50 / T 93 ^a	R 58 / G 59 B 49 / T 84 ^a
[12]	2×2 Bayer (RGBG)	—	Path-information inverse design	Incident power on device	R 58.3 ^a / G 52.6 ^a B 69.6 ^a / T 86 ^a	R 49.2 ^a / G 43.7 ^a B 52.5 ^a / T 81.4 ^a
[22]	2×2 Bayer (3 mid-IR bands)	Polarization sorting in middle band	Adjoint gradient-based inverse design	Incident power at focal plane	HP: 78, 63, 73 VP: 78, 63, 74	HP: 47, 38, 46 VP: 45, 39, 44
[10]	2×2 Bayer (RGBG)	—	Latent-matrix topology opt. + autodiff (ADAM)	Incident power on all pixels	R 60 / G 57 B 65 / T 85 ^a	R 42 ^p / G 52 ^p B 33 ^p
[11]	2×2 Bayer (RGBG)	—	Particle Swarm Optimization (PSO)	Incident power on all pixels	R 64.3 / G 70.4 B 54.6	R 49.5 ^a / G 55.3 ^a B 37.7 ^a
[75]	2×2 Bayer (RGB+NIR)	—	Genetic algorithm + resets based on PSO	Incident power on device	R 49.7 / G 55.4 / B 43.2 NIR 61.9 / T 89 ^a	R 49.7 / G 55.4 / B 43.2 NIR 61.9 / T 89 ^a
[24]	2×2 Bayer (RGBG)	—	Adjoint optimization with MMA	Incident power on all pixels	R 75.4 / G 81.5 / B 63.1	—
[14]	3×1 RGB	—	NSGA-II	Incident power on all pixels	R 59.1 / G 52.4 / B 49.4	—

Ref	Layers	Thickness (μm)	Pixel Pitch (μm)	Lateral Size (μm)	Focal Plane Distance (μm)	Min Feature Size Lateral (nm)	Material / Background
[19]	5	2	1.0	2.0×2.0	1.5	60	TiO ₂ / SiO ₂
[26]	50	2	0.3	0.64×0.64	0	10	TiO ₂ / SiO ₂
[26]	50	2	0.3	0.64×0.64	0	10	TiO ₂ / SiO ₂
[27]	50	2	0.3	0.99×0.66	0	10	TiO ₂ / SiO ₂ on Si
[27]	50	6	1.0	2.97×2.97	0	30	TiO ₂ / SiO ₂ on InGaAs
[76]	1	0.6	1.0	2.0×2.0	2	50	Si ₃ N ₄ / air on SiO ₂
[25]	4	1.5	0.25	1	0.5	31	$n_2 = 2.0 / n_1 = 1.5$ (dispersionless)
[28]	200	2	0.4	1.6 width	0	10	Si ₃ N ₄ + SiO ₂ on Si

Ref	Channels	Added Functionality	Opt. Algorithm	Power Norm.	Simulation (%)
[19]	2×2 Bayer (RGBG)	Polarization sorting in G	Adjoint topology optimization	Incident power on device	R 68.8 ^a / G 51.2 ^a / B 65.4 ^a
[26]	2×2 Bayer (RGBG)	—	Adjoint optimization with MMA	Incident power on device	R 99 ^p / G 99 ^p / B 99 ^p
[27]	3×2 visible channels	—	Adjoint optimization with MMA	Incident power on device	All channels >99
[27]	3×3 SWIR channels	—	Adjoint optimization with MMA	Incident power on device	All channels >98
[76]	2×2 Bayer (RGBG)	—	Adjoint level-set optimization	Incident power on device	R 53 / G 70 / B 47
[25]	4×1 RGBG	—	Genetic algorithm	Incident power on device	R 70.1 / G 57.2 / B 77.4
[28]	^{4×1} RGB+NIR	—	Adjoint L-BFGS	Incident power on device	R 99 ^p / G 99 ^p / B 99 ^p / NIR 99 ^p

5.4.2 Adjoint Optimization and Projection Filters

The adjoint method inverse design approach^{32,36–39,41} involves defining a figure of merit function that is correlated with device performance, and efficiently computing the gradient of this function with respect to the discrete permittivity profile within the design region to iteratively evolve the index distribution. Its versatility allows the technique to be easily and quickly adapted to additional functionalities and constraints, as well as to arbitrary environmental conditions including but not limited to platform substrate, overhead material layers and sidewall enclosures.

We elaborate on this method using the example of an image sensor that splits light of different wavelength bands and different polarizations into different spatial areas on a plane (Fig. 5.6a). The sensor comprises a 3D scattering structure that can be represented by a refractive index distribution over space $n(\vec{x})$, which at the end of the optimization we require to be a binary structure with only two values of n being the refractive indices of TiO_2 and SiO_2 . In order to achieve the target functionality, the figure of merit is selected as the mode overlap between the electric and magnetic fields $\mathbf{E}$, $\mathbf{H}$ over a target area $\mathbf{S}$, and the desired mode profiles $\mathbf{E}_m$, $\mathbf{H}_m$:

$$f(n(\vec{x})) = \frac{1}{8} \frac{\left| \int \mathbf{E} \times \overline{\mathbf{H}}_m \cdot d\mathbf{S} + \int \overline{\mathbf{E}}_m \times \mathbf{H} \cdot d\mathbf{S} \right|^2}{\int \text{Re}(\mathbf{E}_m \times \overline{\mathbf{H}}_m) \cdot d\mathbf{S}} \quad (5.3)$$

which serves as a proxy³² for focusing efficiency of light of the desired wavelength bands and polarizations into their corresponding quadrants in the focal plane. The figure of merit can be maximized over the intended area $\mathbf{S}$ or minimized over all areas in the focal plane outside of $\mathbf{S}$, either of which may be more appropriate depending on the device parameters and optical scattering environment. Both the figure of merit and all mentioned fields are also functions of wavelength and can be modulated to fit the desired spectral profile in each area.

The gradient, or sensitivity, of the figure of merit with respect to the refractive index distribution $df/dn(\vec{x})$ can be constructed from the electromagnetic fields within the device region, of one forward and one adjoint simulation:

$$\frac{df}{dn(\vec{x})} = 2n(\vec{x}) \text{Re} [\mathbf{E}_{\text{fwd}} \cdot \mathbf{E}_{\text{adj}}] \quad (5.4)$$

where $\mathbf{E}_{\text{fwd}}$ is the electric field within the device region when illuminated from above by illumination of the desired mode profile and oblique incidence, and $\mathbf{E}_{\text{adj}}$ is the same electric field, but when the device region is illuminated from below by the same desired mode profile, sent backwards into the device. The phase and amplitude of

$\mathbf{E}_{\text{adj}}$ are derived from the electromagnetic fields over the target area in the forward simulation.

This gradient can then be applied using an appropriate optimization method such as the ADAM optimizer:

$$\begin{aligned}
 m_t &= \beta_1 m_{t-1} + (1 - \beta_1) \left(\frac{df}{dn} \right)_t \\
 v_t &= \beta_2 v_{t-1} + (1 - \beta_2) \left[\left(\frac{df}{dn} \right)_t \right]^2 \\
 \alpha_t &= \alpha \times \frac{\sqrt{1 - \beta_2^t}}{1 - \beta_1^t} \\
 n_{t+1} &= n_t + \frac{\alpha_t m_t}{\sqrt{v_t} + \epsilon}
 \end{aligned} \tag{5.5}$$

where the step size α is kept small so as to keep perturbations to the refractive index small enough such that Eq. 5.5 remains valid in the linear regime; and the learning rates β_1, β_2 are adjusted to ensure successful convergence. After multiple iterations, the algorithm approaches convergence and yields an optimized design (Fig. 5.6b) that achieves the target functionality with high efficiency.




Figure 5.6: (a) Schematic of inverse-designed color sorter device, and projected focal spots on the focal plane. (b) Selected layers of a 40-layer device after optimization, with binarization imposed. Each layer is comprised of regions of TiO_2 embedded in a background SiO_2 medium, which forms a flat substrate for subsequent layers.

As detailed in Eq. (5.1), we impose a binary permittivity on the continuous grayscale devices obtained from the above through the use of a smoothened Heaviside projection filter²⁶. Figure 5.7b shows the corresponding weights that are applied to the

inverse-design gradient as a result of these filters, color-coded by increasing strength from red to green to blue.




Figure 5.7: (a) Sigmoidal projection filter that projects device density to refractive index distribution for two materials, shown in Eq. 5.1. Red: at weak strength, the filter merely scales the density to the refractive index bounds. Blue: with increasing strength, the filter is an increasingly sharp Heaviside function. (b) Corresponding gradient weights that are applied to the adjoint sensitivity function during back-propagation. (c) Figure of merit for a two-material device as it proceeds through the optimization from 0–300 iterations. Dips in the figure of merit for ≥ 150 iterations are due to increasing strength of binarization filter presented in Eq. 5.1.




Figure 5.8: (a) Fully-binarized refractive index distribution for layer 0 (bottom) of a 5-layer device with feature size 85 nm and lateral size $4\mu\text{m}$. (b) In a separate optimization with all other parameters kept the same, a two-pixel border (174 nm) is constrained to have background index and receive no sensitivity updates throughout the optimization so as to ensure a continuous write area limited to $4\times 4\mu\text{m}^2$. (c) Morphological opening (consecutive erosion and dilation filters) is applied to the pattern every 10 iterations to eliminate islands of high permittivity below a certain minimum size so as to avoid fabricating high-aspect-ratio pillars with large risk of crumbling. (d) Comparison between transmission spectra for a., b., c. shows little change in overall performance.

We can also enforce other necessary fabrication restrictions by the application of carefully tuned filters, chosen such that they are differentiable and continuous so as

to avoid a sharply varying gradient across the device region. For example, in a device written through an electron-beam lithography process, due to the finite size of the writing field it is typically necessary to limit the continuous write area to about $4\text{ }\mu\text{m} \times 4\text{ }\mu\text{m}$. This constraint is easily achieved by setting to zero the sensitivity passed to the device in a thin (174 nm) border region in the lateral dimension, depicted in Fig. 5.8b. Finally, to avoid the appearance of high-aspect-ratio and unsupported pillars, we use consecutive erosion and dilation filters to remove islands of a minimum size in each layer, shown in Fig. 5.8c.

When we compare the performances of a device designed with no constraints (Fig. 5.8b.), and devices with the above constraints applied during optimization (Figs. 5.8b, 5.8c), there is very little noticeable performance difference, as shown in Fig. 5.8d. Other fabrication restrictions, as well as tolerance to fabrication errors, can be arbitrarily built into the inverse design process as needed.

The method used above took $\sim 12\text{--}20$ minutes on a high performance computing cluster, as mentioned in the paper, depending on the device thickness and general size of the simulation region. This results in prohibitively large simulation times when scaling to larger device arrays. For individual iterations, GPU acceleration could help reduce simulation runtime. Partitioning the simulation region (domain decomposition) into subregions that are then simulated in parallel may also offer a possible route toward reducing runtime. Reducing the iteration number is also an option, though the schedules of filters must be carefully adjusted to allow the figure of merit to recover (Fig. 5.7c) whenever filters are strengthened. This may be addressed through the implementation of improved optimization strategies, such as coarse-to-fine refinement^{77,78}, or the staged use of different optimizers during different phases of the optimization^{79,80}.

5.4.3 Changing Device Height (Layer Thickness Fixed) and Focal Length




Figure 5.9: (a) Peak transmission values of the spectra in each quadrant as a function of layer number, with layer thickness fixed at 204 nm. Here, the devices have feature size 51 nm and a focal length of 1.5 μm . (b) Scattering losses of the same device as a function of layer number, with layer thickness fixed. Loss channels are as defined in Fig. 5.1f. (c) Contrast of the transmission spectra in each quadrant as a function of focal length, defined as the distance between the bottom aperture of the device and the image plane. Here, the devices, separate from a., have 10 layers and feature size 51 nm. (d) Peak transmission values of the spectra in each quadrant as a function of focal length. (e) Scattering losses of the same device as a function of focal length. (f) Intensity cross-section of the R quadrant across the x-dimension at $\lambda = 630 \text{ nm}$ (R band). Gaussian fits are taken and the FWHM increases while peak intensity decreases with greater focal length.

Depending on the method of fabrication, another form of variation that can be important in affecting device performance is layer number, but with the layer thickness fixed, i.e. changing total device height. This differs from the analysis presented in Section 5.2, where the total device height is fixed and so the parameter that changes is the layer thickness. This consideration becomes relevant when using a fabrication method where consecutive layers are stacked on top of each other, and the thickness of each layer is not controllable. In this case, the limits to the device stack height achievable becomes fabrication time and alignment precision. In the case of fixed layer thickness, increasing the number of layers tended to affect the peak transmission positively (Fig. 5.9a), since there were more degrees of freedom available in order to maximize focusing to the appropriate quadrants. Increasing the device stack height also provides incoming light with a longer propagation distance within

the device, allowing for better scattering control and therefore focusing of different colors of light to the focal plane. As device height increases, the area of the device sides also increases, allowing for greater amounts of side scattering. (Fig. 5.9b). Oblique scattering and reflection remained roughly the same.

The overall thickness of the color sorter becomes relevant also when considering the integration of the sorter with microlenses to improve performance, just as traditional Bayer color filters are combined with microlenses to enhance performance, which have been found to provide a $1.5 - 2\times$ increase in efficiency depending on wavelength⁸¹. In general, a thick router or sorter competes with the space budget normally occupied by the combined microlens and color-filter stack. Depending on the architecture, the thickness of a sorter also affects the effective propagation distance to the photodiode. Conventional microlenses do not scale well into the deep sub-micron regime²⁹, and as such for small pixels, it becomes preferable to design color splitter devices as replacements for the full microlens-and-color-filter stack^{10,13,26,28}. However, for thin single-layer or few-layer devices, co-integration with reduced or modified microlenses may still be plausible.

We also found that the device performance depends strongly on the distance of the focal plane relative to the device height. Efficient focusing requires that the color sorter is optimized to direct light of each wavelength band into target focal spots of adequately tight radius. This focusing is achieved entirely by multiple scattering within the volume of the device region. Similar to a conventional lens, the closer the propagation distance to the focal plane, the more tightly focused the light must be when leaving the lens in order to converge at the appropriate spot on the focal plane. This principle is demonstrated in Fig. 5.9c. Although the peak transmission for each band is generally higher at shorter focal lengths (Fig. 5.9d), since the image plane is closer, the light is not as tightly focused. The contrast between the in-band and out-of-band intensity is lower, which compromises the quality of color information in any image formed. Without adding more degrees of freedom to the design region, higher-NA devices therefore lose performance at shorter focal lengths. Conversely, for longer focal lengths, the effective numerical aperture for a fixed device footprint starts to lead to larger focal spot sizes (Fig. 5.9f) and diffraction begins to limit the focal spot radius that can be achieved until it becomes comparable to the quadrant width. Oblique scattering also increases with focal length (Fig. 5.9e), as tighter focusing is needed for outgoing light to land in the focal region.

This trend also informs as to what happens when pixel size, or lateral device size,

is reduced. As the effective numerical aperture decreases, diffraction during propagation and the associated lateral spreading of the focal spots increasingly compete with the desired spectral separation. The focal spot radius can become comparable to the subpixel width, and efficiency gains are progressively offset by reduced color contrast and higher inter-pixel crosstalk. This is especially true for color sorter devices that rely partially on the propagation distance between exit face and detector plane; even if the focal distance is reduced in order to keep the same NA, the spectrally sorted light has less room to separate before reaching the detector plane. For subpixel sizes approaching the extreme sub-micron regime, or the wavelength scale, color router approaches would start to become more promising, at least in terms of pure efficiency, as they fundamentally do not suffer from this effect. In practice, the choice of using a color sorter or color router architecture ultimately depends on the tradeoff between device performance and available fabrication methods and tolerances.

5.4.4 Realistic Material Dispersion and Absorption




Figure 5.10: (a) Sorting transmission spectrum into B, G, R quadrants for a 10-layer device with 51 nm lateral feature size. Faded lines are calculated from a broadband simulation where dispersion is ignored. Solid lines are calculated from multiple single-frequency simulations where dispersion is taken into account. There is a 6–13% performance drop when taking dispersion into account, depending on the spectral band. (b) Refractive index data for TiO₂ and SiO₂ across the device bandwidth. Horizontal dotted lines indicate the assumed n for each material when ignoring dispersion.

Scattering Channel	422 nm		538 nm		664 nm	
	ND	D	ND	D	ND	D
Sorted Color	72.51%	59.01%	63.48%	56.21%	69.83%	63.74%
Focal	85.00%	73.20%	80.29%	67.60%	84.30%	80.08%
Oblique	7.55%	10.05%	7.59%	7.54%	7.52%	5.66%
Reflection	1.84%	7.56%	2.26%	7.67%	1.81%	2.72%
Side	5.61%	8.16%	9.86%	11.46%	6.37%	8.56%

Table 5.2: Power scattered into each channel at 422 nm, 538 nm, and 664 nm. “ND” indicates the case where dispersion is ignored, calculated with a single broadband simulation; “D” indicates the case where realistic dispersion is used, calculated with multiple single-frequency simulations.

The results in this study are performed assuming that the materials used in the device (TiO₂ and background SiO₂) are non-dispersive and have the same refractive index across the device bandwidth of 400–700 nm. The effects of absorption are also ignored. Fig. 5.10b shows the actual refractive index n and extinction coefficient k of both materials across the device bandwidth⁸². While the n, k values for SiO₂ do not significantly vary or deviate from the assumed value of $n = 1.5, k = 0$ across this wavelength range, the refractive index of TiO₂ does deviate from the assumed

$n = 2.4$, most starkly at lower wavelengths. The deviation decreases with increasing wavelength.

Fig. 5.10a superimposes the same device, which is the baseline device with binary $\text{TiO}_2/\text{SiO}_2$, 10 layers, and 51 nm lateral feature size, displayed in Fig. 5.1a. One set of results (the faded line) is performed with a broadband simulation, assuming that the refractive indices of the device materials do not vary with wavelength, and are the same results as in Fig. 5.1a. This device is optimized, and simulated post-optimization, with these assumptions. The solid line is calculated using the same device design, optimized without dispersion, but simulated post-optimization taking dispersion into account. This is done using ten single-frequency simulations, with the refractive index of the device materials updated accordingly for each frequency. The shifts in transmission and scattering channel values are given in Table 5.2 for select wavelengths. In general, the shift in peak wavelength for each color channel is small, with the discrepancy increasing with wavelength; meanwhile, the drop in performance is 6-13% depending on the spectral band. The performance reduction is greater for shorter wavelengths, and this can be attributed to the high deviation from $n = 2.4$ of the refractive index of TiO_2 at shorter wavelengths.

The discrepancy introduced by the present approximation is small and as such the results and conclusions for these two materials in the wavelength range studied remain broadly valid. A more rigorous treatment of material dispersion and absorption would require a significantly more computationally expensive optimization procedure. Specifically, rather than performing a single broadband simulation at each iteration, multiple single-frequency simulations would be required to accurately represent the wavelength dependence of the material response. This approach can be partially mitigated through parallelization, but would still increase the computational cost beyond a practical level for the present study. The same improvements discussed in Section 5.4.2 for improving scalability would also help enable this workflow for future studies.

References

1. Foo, I., Roberts, G., Pearson, P., Ballew, C. & Faraon, A. Volumetric Metaoptic Color Routers at Visible Wavelengths: A Performance Analysis for Practical Applications. *Manuscript submitted for publication* (2026).
2. Bayer, B. E. US3971065A. <https://patents.google.com/patent/US3971065A/en> (2025)(1976).

3. Park, C. *et al.* Towards subwavelength pixels: nanophotonic color routers for ultra-compact high-efficiency CMOS image sensors. *Journal of Optics* **26**, 093002. ISSN: 2040-8986. <https://doi.org/10.1088/2040-8986/ad5dcc> (2026) (July 2024).
4. Zhou, Y. *et al.* Multifunctional metaoptics based on bilayer metasurfaces. *Light: Science & Applications* **8**, 80. ISSN: 2047-7538. <https://www.nature.com/articles/s41377-019-0193-3> (2026) (Sept. 2019).
5. Dainese, P. *et al.* Shape optimization for high efficiency metasurfaces: theory and implementation. *Light: Science & Applications* **13**, 300. ISSN: 2047-7538. <https://www.nature.com/articles/s41377-024-01629-5> (2026) (Oct. 2024).
6. Li, N. *et al.* Large-area metasurface on CMOS-compatible fabrication platform: driving flat optics from lab to fab. *Nanophotonics* **9**, 3071–3087. ISSN: 2192-8614. <https://onlinelibrary.wiley.com/doi/abs/10.1515/nanoph-2020-0063> (2026) (2018).
7. Zuo, J. *et al.* Chip-integrated metasurface full-Stokes polarimetric imaging sensor. *Light: Science & Applications* **12**, 218. ISSN: 2047-7538. <https://www.nature.com/articles/s41377-023-01260-w> (2023) (Sept. 2023).
8. Colburn, S. & Majumdar, A. Inverse design and flexible parameterization of meta-optics using algorithmic differentiation. *Communications Physics* **4**, 1–11. ISSN: 2399-3650. <https://www.nature.com/articles/s42005-021-00568-6> (2024) (Mar. 2021).
9. Zou, X. *et al.* Pixel-level Bayer-type colour router based on metasurfaces. *Nature Communications* **13**, 3288. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-022-31019-7> (2026) (June 2022).
10. Kim, C. *et al.* Freeform metasurface color router for deep submicron pixel image sensors. *Science Advances* **10**, eadn9000. <https://www.science.org/doi/full/10.1126/sciadv.adn9000> (2024) (May 2024).
11. Shao, Y. *et al.* Pixelated Bayer spectral router based on a sparse meta-atom array. *Chinese Optics Letters* **23**, 073601. <https://opg.optica.org/col/abstract.cfm?uri=col-23-7-073601> (2026) (July 2025).
12. Li, J. *et al.* Single-Layer Bayer Metasurface via Inverse Design. *ACS Photonics* **9**, 2607–2613. <https://doi.org/10.1021/acsp Photonics.2c00990> (2026) (Aug. 2022).
13. Miyata, M., Nemoto, N., Shikama, K., Kobayashi, F. & Hashimoto, T. Full-color-sorting metalenses for high-sensitivity image sensors. *Optica* **8**, 1596. ISSN: 2334-2536. <https://www.osapublishing.org/abstract.cfm?URI=optica-8-12-1596> (2021) (Dec. 2021).

14. Chen, M. *et al.* Full-color nanorouter for high-resolution imaging. *Nanoscale* **13**, 13024–13029. ISSN: 2040-3372. <https://pubs.rsc.org/en/content/articlelanding/2021/nr/d1nr02166d> (2026) (Aug. 2021).
15. Chen, B. H. *et al.* GaN Metalens for Pixel-Level Full-Color Routing at Visible Light. *Nano Letters* **17**, 6345–6352. ISSN: 1530-6984, 1530-6992. <https://pubs.acs.org/doi/10.1021/acs.nanolett.7b03135> (2022) (Oct. 2017).
16. Lin, Z., Groever, B., Capasso, F., Rodriguez, A. W. & Lončar, M. Topology-Optimized Multilayered Metaoptics. *Physical Review Applied* **9**, 044030. <https://link.aps.org/doi/10.1103/PhysRevApplied.9.044030> (2026) (Apr. 2018).
17. Rubin, N. A., Zaidi, A., Dorrah, A. H., Shi, Z. & Capasso, F. Jones matrix holography with metasurfaces. *Science Advances* **7**, eabg7488. <https://www.science.org/doi/full/10.1126/sciadv.abg7488> (2025) (Aug. 2021).
18. Arbabi, A. & Faraon, A. Fundamental limits of ultrathin metasurfaces. *Scientific Reports* **7**, 43722. ISSN: 2045-2322. <https://www.nature.com/articles/srep43722> (2021) (Mar. 2017).
19. Camayd-Muñoz, P., Ballew, C., Roberts, G. & Faraon, A. Multifunctional Volumetric Meta-Optics for Color and Polarization Image Sensors. *Optica* **7**, 280. ISSN: 2334-2536. (2020) (Apr. 2020).
20. Yuan, Q. *et al.* Recent advanced applications of metasurfaces in multi-dimensions. *Nanophotonics* **12**, 2295–2315. ISSN: 2192-8614. <https://onlinelibrary.wiley.com/doi/abs/10.1515/nanoph-2022-0803> (2026) (2023).
21. Shim, H., Kuang, Z., Lin, Z. & Miller, O. D. Fundamental limits to multifunctional and tunable nanophotonic response. *Nanophotonics* **13**, 2107–2116. ISSN: 2192-8614. <https://onlinelibrary.wiley.com/doi/abs/10.1515/nanoph-2023-0630> (2026) (2024).
22. Roberts, G. *et al.* 3D-patterned Inverse-Designed Mid-Infrared Metaoptics. *Nature Communications* **14**, 2768. ISSN: 2041-1723. (2025) (May 2023).
23. Johlin, E. Nanophotonic color splitters for high-efficiency imaging. *iScience* **24**, 102268. ISSN: 2589-0042 (Apr. 2021).
24. Lee, S. *et al.* Inverse design of color routers in CMOS image sensors: toward minimizing interpixel crosstalk. *Nanophotonics* **13**, 3895–3914. ISSN: 2192-8614. <https://www.degruyterbrill.com/document/doi/10.1515/nanoph-2024-0269/html> (2025) (Aug. 2024).
25. Kim, S., Park, C., Kim, S., Chung, H. & Jang, M. S. Design parameters of free-form color splitters for subwavelength pixelated image sensors. *iScience* **26**, 107788. ISSN: 2589-0042. <https://www.sciencedirect.com/science/article/pii/S2589004223018655> (2026) (Oct. 2023).

26. Catrysse, P. B., Zhao, N., Jin, W. & Fan, S. Subwavelength Bayer RGB color routers with perfect optical efficiency. *Nanophotonics* **11**, 2192–8606, 2192–8614. <https://www.degruyter.com/document/doi/10.1515/nanoph-2022-0069/html> (2022) (Mar. 2022).
27. Catrysse, P. B. & Fan, S. Spectral routers for snapshot multispectral imaging. *Applied Physics Letters* **123**, 261105. ISSN: 0003-6951, 1077-3118. <https://pubs.aip.org/apl/article/123/26/261105/2931557/Spectral-routers-for-snapshot-multispectral> (2026) (Dec. 2023).
28. Zhao, N., Catrysse, P. B. & Fan, S. Perfect RGB-IR Color Routers for Sub-Wavelength Size CMOS Image Sensor Pixels. *Advanced Photonics Research* **2**, 2000048. ISSN: 2699-9293, 2699-9293. <https://onlinelibrary.wiley.com/doi/10.1002/adpr.202000048> (2022) (Mar. 2021).
29. Huo, Y., Fesenmaier, C. C. & Catrysse, P. B. Microlens performance limits in sub-2 μ m pixel CMOS image sensors. *Optics Express* **18**, 5861–5872. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-18-6-5861> (2026) (Mar. 2010).
30. Chang, Z. C., Huang, R., Chen, P. & Li, G. Shaping light with multilayer metasurfaces: Design, fabrication, and applications. *Journal of Physics D: Applied Physics*. ISSN: 0022-3727. <http://iopscience.iop.org/article/10.1088/1361-6463/ae0e80> (2025) (2025).
31. Shen, Z. & Huang, D. A Review on Metasurface Beam Splitters. *Nanomanufacturing* **2**, 194–228. ISSN: 2673-687X. <https://www.mdpi.com/2673-687X/2/4/14> (2025) (Dec. 2022).
32. Lalau-Keraly, C. M., Bhargava, S., Miller, O. D. & Yablonovitch, E. Adjoint shape optimization applied to electromagnetic design. *Optics Express* **21**, 21693. ISSN: 1094-4087. <https://www.osapublishing.org/oe/abstract.cfm?uri=oe-21-18-21693> (2020) (Sept. 2013).
33. Piggott, A. Y. *et al.* Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer. *Nature Photonics* **9**, 374–377. ISSN: 1749-4893. <https://www.nature.com/articles/nphoton.2015.69> (2025) (June 2015).
34. Sell, D., Yang, J., Doshay, S., Yang, R. & Fan, J. A. Large-Angle, Multifunctional Metagratings Based on Freeform Multimode Geometries. *Nano Letters* **17**, 3752–3757. ISSN: 1530-6984. <https://doi.org/10.1021/acs.nanolett.7b01082> (2025) (June 2017).
35. Ballew, C., Roberts, G. & Faraon, A. Multi-dimensional wavefront sensing using volumetric meta-optics. *Optics Express* **31**, 28658–28669. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-31-18-28658> (2025) (Aug. 2023).

36. Jensen, J. S. in *Wave Propagation in Linear and Nonlinear Periodic Media: Analysis and Applications* (eds Romeo, F. & Ruzzene, M.) 109–159 (Springer, Vienna, 2012). ISBN: 978-3-7091-1309-7. https://doi.org/10.1007/978-3-7091-1309-7_3 (2025).
37. Miller, O. D. *Photonic Design: From Fundamental Solar Cell Physics to Computational Inverse Design* PhD thesis (University of California, Berkeley). <http://optoelectronics.eecs.berkeley.edu/ThesisOwenMiller.pdf>.
38. Niederberger, A. C. R., Fattal, D. A., Gauger, N. R., Fan, S. & Beausoleil, R. G. Sensitivity analysis and optimization of sub-wavelength optical gratings using adjoints. *Optics Express* **22**, 12971–12981. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-22-11-12971> (2023) (June 2014).
39. Lalau-Keraly, C. *Optimizing Nanophotonics: from Photoreceivers to Waveguides* PhD thesis (University of California, Berkeley).
40. Molesky, S. *et al.* Inverse design in nanophotonics. *Nature Photonics* **12**, 659–670. ISSN: 1749-4893. <https://www.nature.com/articles/s41566-018-0246-9> (2026) (Nov. 2018).
41. Su, L. *et al.* Nanophotonic inverse design with SPINS: Software architecture and practical considerations. *Applied Physics Reviews* **7**, 011407. ISSN: 1931-9401. <https://doi.org/10.1063/1.5131263> (2026) (Mar. 2020).
42. Chung, H. & Miller, O. D. High-NA achromatic metalenses by inverse design. *Optics Express* **28**, 6945–6965. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-28-5-6945> (2025) (Mar. 2020).
43. Roques-Carmes, C. *et al.* Toward 3D-Printed Inverse-Designed Metaoptics. *ACS Photonics* **9**, 43–51. <https://pubs.acs.org/doi/abs/10.1021/acsphotonics.1c01442> (2025) (Jan. 2022).
44. Fan, Z. *et al.* Holographic multiplexing metasurface with twisted diffractive neural network. *Nature Communications* **15**, 9416. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-024-53749-6> (2025) (Oct. 2024).
45. Pan, C.-F. *et al.* 3D-printed multilayer structures for high-numerical aperture achromatic metalenses. *Science Advances* **9**, eadj9262. <https://www.science.org/doi/10.1126/sciadv.adj9262> (2025) (Dec. 2023).
46. Melzer, J. E. & McLeod, E. Assembly of multicomponent structures from hundreds of micron-scale building blocks using optical tweezers. *Microsystems & Nanoengineering* **7**, 45. ISSN: 2055-7434. <https://www.nature.com/articles/s41378-021-00272-z> (2025) (June 2021).

47. Jin, X. & Zentgraf, T. Independent Wavefront Multiplexing with Metasurfaces via Non-Injective Transformation. *Advanced Materials* **n/a**, e11823. ISSN: 1521-4095. <https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.202511823> (2025).
48. Chang, T. *et al.* Universal Metasurfaces for Complete Linear Control of Coherent Light Transmission. *Advanced Materials* **34**, 2204085. ISSN: 1521-4095. <https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.202204085> (2025) (2022).
49. Zheng, H. *et al.* Compound Meta-Optics for Complete and Loss-Less Field Control. *ACS Nano* **16**, 15100–15107. ISSN: 1936-0851. (2025) (Sept. 2022).
50. Mirzapourbeinekalaye, B., McClung, A. & Arbabi, A. General Lossless Polarization and Phase Transformation Using Bilayer Metasurfaces. *Advanced Optical Materials* **10**, 2102591. ISSN: 2195-1071. <https://onlinelibrary.wiley.com/doi/abs/10.1002/adom.202102591> (2025) (2022).
51. Kim, M. *et al.* Multilayer all-polymer metasurface stacked on optical fiber via sequential micro-punching process. *Nanophotonics* **12**, 2359–2369. ISSN: 2192-8606. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11501433/> (2025).
52. Dorrah, A. H., Park, J.-S., Palmieri, A. & Capasso, F. Free-standing bilayer metasurfaces in the visible. *Nature Communications* **16**, 3126. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-025-58205-7> (2025) (Apr. 2025).
53. Liu, L., Holzer, A., Raney, N. & Babicheva, V. E. Nanoantennas and metasurfaces tailored by electron beam lithography and substrate conductivity. *Scientific Reports* **15**, 29339. ISSN: 2045-2322. <https://www.nature.com/articles/s41598-025-12996-3> (2025) (Aug. 2025).
54. Qin, N., Qian, Z.-G., Zhou, C., Xia, X.-X. & Tao, T. H. 3D electron-beam writing at sub-15 nm resolution using spider silk as a resist. *Nature Communications* **12**, 5133. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-021-25470-1> (2025) (Aug. 2021).
55. Zeng, S., Tian, T., Oh, J., Lin, Z.-H. & Shih, C.-J. Direct nanopatterning of complex 3D surfaces and self-aligned superlattices via molecular-beam holographic lithography. *Nature Communications* **16**, 3436. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-025-58651-3> (2025) (Apr. 2025).
56. Kingma, D. P. & Ba, J. *Adam: A Method for Stochastic Optimization* Jan. 2017. <http://arxiv.org/abs/1412.6980> (2025).
57. Theuwissen, A. J. P. CMOS image sensors: State-of-the-art. *Solid-State Electronics. Papers Selected from the 37th European Solid-State Device Research Conference - ESSDERC'07* **52**, 1401–1406. ISSN: 0038-1101. <https://www>

- w.sciencedirect.com/science/article/pii/S0038110108001317 (2025) (Sept. 2008).
58. Hirakawa, K. *Cross-talk explained in 2008 15th IEEE International Conference on Image Processing* (Oct. 2008), 677–680. <https://ieeexplore.ieee.org/document/4711845> (2025).
 59. Blockstein, L. & Yadid-Pecht, O. Crosstalk quantification, analysis, and trends in CMOS image sensors. *Applied Optics* **49**, 4483. ISSN: 0003-6935, 1539-4522. <https://opg.optica.org/abstract.cfm?URI=ao-49-24-4483> (2025) (Aug. 2010).
 60. Wu, X. & Zhang, X. Joint Color Decrosstalk and Demosaicking for CFA Cameras. *IEEE Transactions on Image Processing* **19**, 3181–3189. ISSN: 1941-0042. <https://ieeexplore.ieee.org/document/5482075> (2025) (Dec. 2010).
 61. Vaillant, J., Mornet, C., Decroux, T., Hérault, D. & Schanen, I. *Characterization of pixel crosstalk and impact of Bayer patterning by quantum efficiency measurement in Digital Photography VII* **7876** (SPIE, Jan. 2011), 362–371. <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/7876/787613/Characterization-of-pixel-crosstalk-and-impact-of-Bayer-patterning-by/10.1117/12.879603.full> (2026).
 62. Kim, J., Jeong, K. & Kang, M. G. Crosstalk Correction for Color Filter Array Image Sensors Based on Lp-Regularized Multi-Channel Deconvolution. *Sensors (Basel, Switzerland)* **22**, 4285. ISSN: 1424-8220. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9185327/> (2025) (June 2022).
 63. Goossens, T. Crosstalk elimination by rearranging thin-film filters. *Optics Letters* **47**, 3920. ISSN: 0146-9592, 1539-4794. <https://opg.optica.org/abstract.cfm?URI=ol-47-15-3920> (2025) (Aug. 2022).
 64. Gilardi, G. *et al.* Deep Trenches for Thermal Crosstalk Reduction in InP-Based Photonic Integrated Circuits. *Journal of Lightwave Technology* **32**, 4262–4268. <https://opg.optica.org/jlt/abstract.cfm?uri=jlt-32-24-4262> (2025) (Dec. 2014).
 65. Tseng, Y.-T., Chang, Y.-H., Liu, M. C., Liu, S.-C. & Tsai, C.-S. US20150295005A1. <https://patents.google.com/patent/US20150295005A1/en> (2025)(2015).
 66. Yoo, D. *et al.* Automotive 2.1 μm Full-Depth Deep Trench Isolation CMOS Image Sensor with a 120 dB Single-Exposure Dynamic Range. *Sensors* **23**, 9150. ISSN: 1424-8220. <https://www.mdpi.com/1424-8220/23/22/9150> (2025) (Jan. 2023).
 67. Choi, S., Lee, S., Koo, C., Nah, S., Kim, J. H., *et al.* World smallest 200Mp CMOS Image Sensor with 0.56 μm pixel equipped with novel Deep Trench Isolation structure for better sensitivity and higher CG. *International Image*

- Sensor Society*. <https://imagesensors.org/papers/10.60928/sel5-hn0e> (2025) (2023).
68. Okawa, T. *et al.* A 1/2inch 48M All PDAF CMOS Image Sensor Using 0.8 μ m Quad Bayer Coding 2 $\times$ 2OCL with 1.0lux Minimum AF Illuminance Level in 2019 IEEE International Electron Devices Meeting (IEDM) (Dec. 2019), 16.3.1–16.3.4. <https://ieeexplore.ieee.org/document/8993499> (2025).
 69. Chang, Y.-C. *et al.* 0.8 μ m Color Pixels with Wave-Guiding Structures for Low Optical Crosstalk Image Sensors. *Electronic Imaging* **33**, 1–5. ISSN: 2470-1173. <https://library.imaging.org/ei/articles/33/7/art00009> (2025) (Jan. 2021).
 70. Kunikiyo, T. US8884319B2. <https://ppubs.uspto.gov/api/patents/html/8884319?source=USPAT&requestToken=eyJzdWIiOiJmYjc2ZDQyOC02ZjBhLTQwNzctOTFiOC0wYjMwNmRhYTE2MWMiLCJ2ZXIiOiIION2I1NzI4Ny05MWMwLTQ5ODctOWY3OS1mMWZhMGZiNTEyYzIiLCJleHAiOiJbB9> (2026)(2014).
 71. Bak, H. *et al.* Advanced Color Filter Isolation Technology for Sub-Micron Pixel of CMOS Image Sensor in 2022 International Electron Devices Meeting (IEDM) (Dec. 2022), 37.6.1–37.6.4. <https://ieeexplore.ieee.org/document/10019484> (2026).
 72. Park, S. *et al.* A 64Mpixel CMOS Image Sensor with 0.56 μ m Unit Pixels Separated by Front Deep-Trench Isolation in 2022 IEEE International Solid-State Circuits Conference (ISSCC) **65** (Feb. 2022), 1–3. <https://ieeexplore.ieee.org/document/9731750> (2025).
 73. Kim, D. *et al.* 6.10 A 1/1.56-inch 50Mpixel CMOS Image Sensor with 0.5 μ m pitch Quad Photodiode Separated by Front Deep Trench Isolation in 2024 IEEE International Solid-State Circuits Conference (ISSCC) **67** (Feb. 2024), 118–120. <https://ieeexplore.ieee.org/document/10454448> (2026).
 74. Wang, W. *et al.* Grayscale Electron Beam Lithography Direct Patterned Antimony Sulfide Jan. 2024. <http://arxiv.org/abs/2401.13427> (2025).
 75. Yan, Z. *et al.* Pixel-Level Spectral Routers for RGB-IR Sensing. *ACS Photonics* **12**, 6374–6380. <https://doi.org/10.1021/acsp Photonics.5c01866> (2025) (Nov. 2025).
 76. Wang, L., Liu, X., Peng, Y. & Xiao, J.-J. Nano-router metasurface inverse design by multi-body shape optimization. *Optics Letters* **50**, 4786–4789. ISSN: 1539-4794. <https://opg.optica.org/ol/abstract.cfm?uri=ol-50-15-4786> (2026) (Aug. 2025).
 77. Wang, F., Lazarov, B. S. & Sigmund, O. On projection methods, convergence and robust formulations in topology optimization. *Structural and Multidisciplinary Optimization* **43**, 767–784. ISSN: 1615-147X, 1615-1488. <http://link.springer.com/10.1007/s00158-010-0602-y> (2020) (June 2011).

78. Christiansen, R. E. & Sigmund, O. Inverse design in photonics by topology optimization: tutorial. *JOSA B* **38**, 496–509. ISSN: 1520-8540. <https://opg.optica.org/josab/abstract.cfm?uri=josab-38-2-496> (2023) (Feb. 2021).
79. Ballew, C., Roberts, G., Zheng, T. & Faraon, A. Constraining Continuous Topology Optimizations to Discrete Solutions for Photonic Applications. *ACS Photonics* **10**, 836–844. <https://doi.org/10.1021/acsp Photonics.2c00862> (2025) (Apr. 2023).
80. MacLellan, B. *et al.* Inverse Design of Photonic Systems. *Laser & Photonics Reviews* **18**, 2300500. ISSN: 1863-8899. <https://onlinelibrary.wiley.com/doi/abs/10.1002/lpor.202300500> (2026) (2024).
81. Crocherie, A., Vaillant, J. & Hirigoyen, F. Three-dimensional broadband FDTD optical simulations of CMOS image sensor. *arXiv.org*. <https://arxiv.org/abs/2310.10302v1> (2026) (Oct. 2023).
82. Polyanskiy, M. N. Refractiveindex.info database of optical constants. *Scientific Data* **11**, 94. ISSN: 2052-4463. <https://www.nature.com/articles/s41597-023-02898-2> (2026) (Jan. 2024).

GRADIENT-REFRACTIVE INDEX MULTILAYERED TiO₂ METAOPTICS VIA TWO-PHOTON LITHOGRAPHY

This chapter is directly reproduced from [1], a paper published as a key part of the work in this thesis. As described in the abstract below, in this work we design, fabricate and characterize several metaoptic devices, including a gradient-index color router, using a novel two-photon-lithography platform to realize TiO₂/air volumetric metaoptics.

Figures 6.1 to 6.3 reprinted with permission from the copyright holder, Wenyuan Chen.

Abstract

Flat optics based on subwavelength metasurfaces enable compact wavefront control for imaging, sensing, and communications, but the functionality of single-layer architectures is fundamentally constrained by the limited degrees of freedom available within a single optical interaction. Multilayer and volumetric metaoptics offer a pathway to overcome these limitations; however, their realization using conventional top-down nanofabrication is hindered by fabrication complexity, alignment requirements, and limited scalability. Here, we establish a nanoscale additive manufacturing platform for volumetric metaoptics based on high-index crystalline titanium dioxide (TiO₂) by integrating scaffold-supported voxel architectures with a shrinkage-mitigated two-photon lithography workflow and ceramic conversion. The approach enables large-area ($\geq 90\ \mu\text{m}$ lateral dimension), mechanically continuous three-dimensional TiO₂ architectures with thicknesses of $10\ \mu\text{m}$ – $20\ \mu\text{m}$ and post-conversion feature sizes approaching $100\ \text{nm}$, fabricated in a single lithographic process without layer-to-layer alignment. Leveraging this capability, we demonstrate mid-infrared metaoptics operating across the $4\ \mu\text{m}$ – $5\ \mu\text{m}$ band, including monolayer chromatic metalenses, heterogeneous multilayer achromatic metalenses with numerical apertures up to 0.49, and an inverse-designed volumetric color router that performs color routing within a compact ceramic volume. Together, these results establish nanoscale additive manufacturing of high-index ceramics as a viable route to heterogeneous, multilayer metaoptics and, more broadly, position it as a promising pathway toward volumetric architectures that support multifunctionality within

compact three-dimensional optical systems for imaging, sensing, and spectral routing.

6.1 Introduction

Over the past decade, planar optical metasurfaces have matured into a widely adopted paradigm for wavefront engineering^{2–6}. By arranging subwavelength scatterers (meta-atoms) across a two-dimensional plane, metasurfaces enable precise control over the phase, amplitude, and polarization of light, leading to compact optical components such as flat metalenses^{6–8}, polarization routers^{9–11}, and holographic plates¹². Their compatibility with wafer-scale fabrication has further enabled integration into consumer electronics and on-chip photonic systems¹³. A central appeal of metasurfaces lies in their potential for multifunctionality¹⁴. However, the planar, single-interaction geometry imposes intrinsic constraints: limited degrees of freedom within a thin layer and reliance on subtractive nanofabrication processes such as lithography and etching. As performance demands increase—encompassing high numerical aperture, broadband operation, and multi-channel functionality—these constraints increasingly motivate architectures that distribute optical interactions through depth rather than confining them to a single plane.

Extending metasurfaces into the third dimension through multilayer stacks or fully volumetric architectures opens the regime of metaoptics, where multiple optical interactions can be engineered within a compact footprint. Demonstrations to date include cascaded metasurfaces^{15,16}, double-sided devices¹⁷, bilayer architectures^{18,19}, and inverse-designed volumetric structures²⁰. Despite these advances, scalability remains a central challenge. Top-down multilayer implementations typically require repeated planar processing steps with stringent overlay alignment and planarization. As the number of layers increases, accumulated misalignment and process complexity limit yield, device area, and practical heterogeneity across the aperture. In an industrial context, realizing even a modest number of aligned layers often demands infrastructure comparable to semiconductor foundries, making such approaches viable only when multilayer optics deliver transformative functionality beyond conventional components.

Additive manufacturing offers a compelling alternative by enabling voxel-by-voxel construction of arbitrary three-dimensional architectures. Two-photon lithography (TPL) allows volumetric patterning with submicron resolution by confining polymerization to a localized voxel via nonlinear two-photon absorption^{21–24}. In polymer-




Figure 6.1: Conceptual overview of nano-additive manufactured multilayer TiO_2 metaoptics enabled by two-photon lithography. (a) Multilayer TiO_2 nanoarchitectures fabricated via two-photon lithography, comprising heterogeneous unit-cell designs arranged on square lattices. The resulting structures achieve overall thicknesses of up to $20 \mu\text{m}$ and lateral dimensions ranging from $40 \mu\text{m}$ to beyond $90 \mu\text{m}$. (b) Rendering of a homogeneous monolayer chromatic metalens designed for operation at $4.5 \mu\text{m}$. (c) Rendering of a heterogeneous bilayer achromatic metalens engineered for broadband operation across the $4 \mu\text{m}$ – $5 \mu\text{m}$ spectral band. (d) Rendering of a heterogeneous multilayer color router obtained through inverse design, enabling spatial separation of multiple wavelength bands.

based systems such as IP-Dip, direct-write compatibility and geometric flexibility have enabled stacked and heterogeneous metaoptics, including high-aspect-ratio meta-atoms for orbital angular momentum holography¹², achromatic and multilayer metalenses^{16,25,26}, and inverse-designed wavefront sorters and volumetric optical elements^{20,27}. However, most available TPL resists remain polymeric or hybrid, with relatively low refractive indices ($n \lesssim 1.6$). While high-index inorganic dielectrics can overcome these limitations, their realization via TPL is considerably more challenging. Hybrid organic–inorganic photoresists typically require post-print thermal conversion to form crystalline oxides such as TiO_2 ^{28,29}, a process that introduces shrinkage-induced stress, adhesion failure, and geometry distortion—effects particularly detrimental to large-area and spatially heterogeneous architectures. Areal scalability further compounds these challenges as stitching across multiple TPL write fields can introduce mechanical discontinuities that act as stress concentrators during calcination if not carefully engineered.

Here, we establish a nano-additive manufacturing framework for high-index TiO_2 metaoptics that addresses both ceramic-conversion robustness and large-area scalability. The methodology integrates (1) a minimal-contact anchoring geometry that enables clean substrate release without sacrificial layers while promoting isotropic shrinkage, and (2) a controlled-overlap stitching strategy that preserves mechanical continuity across multiple TPL write fields. As illustrated in Fig. 6.1a, we demonstrate high-resolution TiO_2 nanoarchitectures with device footprints exceeding $\sim 90\ \mu\text{m}$ and thicknesses up to $20\ \mu\text{m}$ using a custom metal–organic hybrid photoresist and optimized post-processing. Building on this fabrication platform, we establish a complete design-to-fabrication pipeline for multifunctional, multilayer metaoptics operating in the $4\ \mu\text{m}$ – $5\ \mu\text{m}$ band using scaffold-assisted three-dimensional meta-atom designs. We demonstrate chromatic TiO_2 metalenses spanning numerical apertures of 0.13–0.49 (Fig. 6.1b), heterogeneous multilayer achromatic metalenses at $\text{NA} = 0.25$ (9 layers) and $\text{NA} = 0.49$ (15 layers) across $4\ \mu\text{m}$ – $5\ \mu\text{m}$ (Fig. 6.1c), and an inverse-designed 20-layer scaffolded color router that spatially separates multiple wavelength bands within the same spectral window (Fig. 6.1d). Collectively, these demonstrations chart a progression from homogeneous monolayer devices to heterogeneous multilayer architectures, where design complexity is leveraged to achieve enhanced functionality within a compact volumetric form factor achievable through additive manufacturing.

6.2 Fabrication Procedure

The nano-additive manufacturing (nano-AM) process developed in this work enables fabrication of large-area TiO_2 nanostructures with minimal structural deformation through a coordinated design of contact-layer geometry, print strategy, and post-processing (Fig. 6.2). We employ an acrylate-functionalized TiO_2 precursor resin (Fig. 6.2a) synthesized via ligand exchange between titanium(IV) ethoxide and acrylic acid, following a formulation previously reported by Vyatskikh *et al.*²⁸ (see Experimental Section). The resin is confined between a glass slide and a silanized substrate spaced by $\sim 100\ \mu\text{m}$, enabling two-photon lithography of three-dimensional lattice architectures. After printing and solvent development, precursor lattices are mechanically released from the growth substrate using a micromanipulation-assisted process and transferred to a receiving fused silica substrate prior to calcination (Fig. 6.2b). Ceramic conversion is carried out using slow thermal ramps in an open-atmosphere furnace, yielding crystalline TiO_2 architectures with phase compositions controlled by the peak temperature.




Figure 6.2: Nano-additive manufacturing process and SEM characterization of fabricated homogeneous lattices. (a) Illustration of the two-photon lithography setup, where the TiO₂ precursor resin is sandwiched between a glass chip and the sample substrate. The main chemical constituents of the resin, based on a titanium(IV) acrylate formulation, are shown. Schematics of the contact-layer design and writing strategy, featuring minimal-contact anchoring geometry and a staggered printing approach, are also illustrated. (b) Post-print processing sequence, during which the printed structures are transferred to a receiving substrate using a micro-manipulation needle, followed by calcination to obtain crystalline TiO₂ nanostructures. (c) 52° tilted-view SEM image of a calcined TiO₂ lattice. The lattice was intentionally flipped to expose the contact-layer geometry for imaging. (Scale bar: 20 μm) (d) Magnified SEM image of the corner region highlighted in (c). Blue and red insets denote the contact-layer and bulk-layer unit-cell geometries, respectively. (Scale bar: 2 μm) (e–f, g–h, i–j) Top-view SEM images of representative homogeneous lattice geometries: octahedron, octet, and isotruss, respectively. (Scale bar: e, g, i: 2 μm; f, h, j: 500 nm)

For device footprints exceeding a single two-photon lithography write field, stitching is required to assemble continuous large-area, multilayer lattices. While adjacent fields can be directly aligned in polymeric systems, stitched interfaces in ceramic precursors act as stress concentrators during thermal conversion, often leading to seam-localized fracture and yield loss. To mitigate this effect, we implement a controlled-overlap stitching strategy based on an interleaved ABAB tiling scheme incorporated directly into both the CAD geometry and the DeScribe writing code. Alternating layers are laterally offset by half a tile width in both in-plane directions, producing a checkerboard-style arrangement in which stitch boundaries do not align vertically across layers. This interleaving distributes mechanical stress throughout the lattice volume, improves structural continuity across tile interfaces, and enables reliable fabrication of large-area, multilayer TiO_2 lattices after calcination. Unlike the native stitching routine implemented in the DeScribe software—which requires global preprocessing of the full STL geometry prior to writing, we implement stitching through explicit geometric partitioning at the CAD and layer level. This approach reduces the computational burden associated with generating and processing large, complex three-dimensional nanoarchitected STL files by operating on smaller subsections, while simultaneously enabling layer-by-layer control of stitch placement and offsets required to enforce interleaved stitching across multilayer lattices.

During ceramic conversion, volumetric shrinkage arises from organic removal and densification of the titanium-containing network. If precursor structures remain rigidly anchored to the substrate, suppressed in-plane contraction induces stress gradients that can result in lateral distortion, base warping, delamination, or fracture³⁰. While prior approaches—including sacrificial pedestals, spring suspensions, and release layers—have been demonstrated to mitigate these effects^{31–34}, they are difficult to scale to large-area, multilayer architectures or are incompatible with complex device integration. Here, we implement a minimal-contact anchoring strategy by modifying the contact layer such that each unit cell interfaces with the substrate at a single discrete point (Fig. 6.2a). This point-contact geometry provides sufficient mechanical stability during printing while substantially reducing interfacial constraint during calcination, enabling near-isotropic shrinkage and preserving lattice fidelity.

6.3 Results and Discussion

The following results in Sections 6.3.1 to 6.3.3 were contributed by co-authors and establish the fabrication platform and baseline device performance that underpin

the color router work in Section 6.3.4.

6.3.1 Homogeneous Lattices

To validate the nano-AM fabrication process, homogeneous TiO_2 lattices with three representative unit cell geometries—octet, octahedron, and isotruss—were fabricated over $450\ \mu\text{m} \times 450\ \mu\text{m} \times 25\ \mu\text{m}$ footprints using the interleaved stitching strategy described in Section 2. Following calcination at $900\ ^\circ\text{C}$, structures shrank with near-perfect isotropy (volumetric shrinkage factor $\sim 5\times$), yielding rutile TiO_2 lattices with final lateral dimensions of $\sim 89\ \mu\text{m}$. Mechanical and geometric continuity was preserved across all stitched write fields with no observable fracture, confirming that the nano-AM process supports large-area, mechanically robust oxide lattice fabrication.

6.3.2 Chromatic Metalenses

Chromatic singlet metalenses were designed by mapping a hyperboloidal phase profile onto the discretized meta-atom lattice at a design wavelength of $4.5\ \mu\text{m}$. A series of lenses spanning NAs of 0.13–0.49 were fabricated and characterized (Fig. 6.3a–e). Measured focal spots were near diffraction-limited across all characterized designs, with focusing efficiencies of 47–57%, in close agreement with simulated values of 55–60%. These results confirm that the nano-AM TiO_2 platform supports high-fidelity chromatic wavefront shaping in the mid-infrared.

6.3.3 Achromatic Metalenses

Achromatic operation requires simultaneous control of phase and group delay across the operating bandwidth—a constraint that homogeneous monolayer designs cannot satisfy due to their limited dispersion tunability. The nano-AM platform addresses this through vertical heterogeneity: stacking layers of distinct pillar diameters substantially expands the accessible phase–dispersion space beyond homogeneous configurations (Fig. 6.3f–g). Heterogeneous bilayer achromatic metalenses at $\text{NA} = 0.25$ (9 layers) and $\text{NA} = 0.49$ (15 layers) were fabricated and demonstrated maximum focal shifts of 9% and 15% respectively over an 800 nm bandwidth (Fig. 6.3h–i). This principle of vertical heterogeneity as a design degree of freedom extends directly to the multilayer color router in Section 6.3.4, where full voxel-by-voxel control of the three-dimensional refractive index distribution enables wavelength-dependent wavefront sorting within a compact ceramic volume.




Figure 6.3: Chromatic and achromatic TiO₂ metalenses via nano-AM, establishing the fabrication platform for multilayer metaoptics. (a) Schematic of a meta-atom unit cell, consisting of a central TiO₂ pillar of diameter D supported by scaffold beams of diameter d_s , arranged on a square lattice of period a . The stacking number N controls the overall height of the unit cell, resulting in a vertically uniform homogeneous monolayer. (b) SEM image of the calcined TiO₂ chromatic metalens. (NA= 0.74, Scale bar: 2 μm) (c) Measured and simulated focusing efficiencies as a function of NA. (d) Measured axial (xz) and focal-plane (xy) intensity distributions for a chromatic metalens with NA= 0.49 at the design wavelength of 4.5 μm . (e) Calculated phase at the center wavelength ($\lambda = 4.5 \mu\text{m}$) and group delay over the target bandwidth $\Delta\omega$ (corresponding to $\lambda = 4.5 \mu\text{m}$) for homogeneous monolayer meta-atoms as a function of the number of stacked layers N . For each N , data points corresponds to pillar diameters varying from 100 to 900 nm. (f) Calculated phase and group delay for heterogeneous bilayer configurations with $N = 9$. Each color denotes a distinct bilayer configuration $k^2_{(x+y)}$, with an example configuration shown in the inset. The heterogeneous bilayer set encloses the phase–dispersion space of homogeneous monolayers and homogeneous bilayers. This bilayer design principle may be generalized to the multilayer architecture employed in the color router (Section 6.3.4). (g) Measured far-field intensity distributions of the NA = 0.25 achromatic metalens. Normalized axial (xz) intensity cross sections are shown for selected wavelengths spanning 4–5 μm . (h, i) Measured focal length and lateral full width at half maximum (FWHM) as a function of wavelength for the NA= 0.25 and NA= 0.49 achromatic metalenses. Maximum relative focal shifts are 9% and 15% respectively over the 800 nm bandwidth.




Figure 6.4: Characterization of nano-AM TiO₂ color router. (a) Schematic of inverse-designed metaoptic color router device with 20 layers and width 40×40 μm, and projected focal spots for red, green, and blue light on the focal plane. We distinguish between the two opposite quadrants accepting green light as G1 and G2. (b) Selected layers of the 20-layer device after optimization. Each voxel in each layer is assigned a refractive index value $n \in [1, 2.4]$ where $n = 2.4$ is the assumed refractive index of TiO₂. Layer 0 is the top layer. (c) Phase shift of a voxel with refractive index n is matched to the corresponding width of a pillar (1 μm high) in-scaffold that produces the same phase shift. (d) Tilted view of optical microscopy (yellow, pre-ceramic) and scanning electron microscopy (post-ceramic) images of the fabricated device. Inset shows the top layer of the designed device. (Scale bar: Top: 100 μm; Bottom left: 50 μm; Bottom right: 20 μm) (e) Sorting efficiency in each quadrant, defined as the power in each quadrant normalized over the total power incident on all four quadrants. The wavelength range is divided into spectral bands corresponding to B, G and R. The CI color router device is optimized for normal incidence light with no scaffold present. The displayed values are obtained from simulations performed post-optimization. (f) Experimentally measured sorting efficiencies of the scaffold color router over each quadrant area across the wavelength range. (g) Simulated intensity at the focal plane using monochromatic incident light at the given wavelengths for the scaffold color router. (h) Experimentally obtained image at the focal plane using monochromatic incident light for the scaffold color router. The given wavelengths are the values at which sorting efficiency for the relevant quadrant is maximal.

6.3.4 Scaffold Color Router

Beyond conventional wavefront shaping, the key advantage of AM and TPL lies in enabling volumetric architectures that support multifunctional optical responses which would be challenging to fabricate using conventional top-down lithography techniques. Color routers, or spectral routers, are an example of such multifunctional meta-optic devices that provide both spectral and spatial multiplexing in a compact footprint. As such, they have been suggested as spectral routers for use in image-sensor pixels^{35–43}. In conventional Bayer-type color-filter arrays, the filters used are absorptive and only transmit at most 25% (red or blue) or 50% (green) of incident light^{44,45}. Metaoptic-based color routers provide avenues to overcome these limitations by redirecting incident light instead of absorbing it, thus providing higher transmission and signal-to-noise ratio with the same pixel pitch. Building on this, additional degrees of freedom can be added to the existing spatial and spectral functionality, such as control of dispersion⁴⁶, phase⁴⁷, and polarization⁴⁸. Multi-layer or volumetric meta-optics have been shown to be an effective route towards this increased multifunctionality^{36,48–50}, offering more scattering interactions between multiple layers that provide control over optical degrees of freedom^{47,51} not possible with single-layer devices. Fabrication of such devices typically requires precise multi-layer alignment and is improved by the use of materials with high index contrast (with respect to the background) in the relevant wavelength range such as TiO₂⁵². Nano-AM is therefore well-positioned as a fabrication platform to realize multilayer color-routers made of TiO₂.

Here, we combine computational inverse design⁵³ with the nano-AM TiO₂ platform to realize a fully three-dimensional metaoptic capable of wavelength-dependent wavefront sorting. Specifically, we demonstrate an inverse-designed, scaffolded color router operating across the 4 μm –5 μm band (Fig. 6.4a). The device occupies a $40 \times 40 \times 20 \mu\text{m}$ volume discretized on a cubic $1 \mu\text{m}^3$ voxel grid, resulting in a 20-layer architecture, each of which varying laterally while remaining uniform in the vertical direction. Each voxel is assigned a refractive index value $n \in [1, 2.4]$, where $n = 2.4$ corresponds to dense TiO₂, and the distribution of refractive index is continuous across the device. The optimization targets routing of normally incident monochromatic light at different wavelengths into distinct spatial quadrants at a prescribed output plane. Two opposite quadrants (G1 and G2 in Fig. 6.4a) are designated to simultaneously accept wavelengths in the green spectral band. Selected layers of the optimized 20-layer device are shown in Fig. 6.4b, and details of the inverse-design methodology are given below.

To implement the optimized refractive index distribution using the nano-AM platform, each voxel value is mapped to an equivalent scaffolded pillar geometry. Specifically, the phase delay imparted by a voxel of refractive index n is matched to the phase shift produced by a 1 μm -tall TiO_2 pillar of corresponding width embedded within the scaffold (Fig. 6.4c). This phase-matching procedure enables translation of a continuous index profile into a realizable voxelized pillar architecture while preserving the intended optical response. To avoid confusion, from this point onwards we denote the inverse-designed, continuous-index device as the “continuous-index” (CI) color router, and the pillar-in-scaffold implementation as the scaffold color router.

Optical microscopy of the pre-ceramic structure and scanning electron microscopy (SEM) of the post-ceramic device confirm successful fabrication of the full 20-layer voxelized geometry with high structural fidelity (Fig. 6.4d). The inset of Fig. 6.4d allows comparison between the pattern of the top layer of the CI color router, and the visible top-layer structures of the scaffold color router (pre- and post-ceramic), showing qualitative agreement in the spatial layout produced by the pillar phase-matching.

We examine the performance of the optimized routers by the sorting efficiency, which is defined for a target quadrant as the ratio of power incident on that quadrant to the total power incident on the full 2×2 -quadrant pixel array in the focal plane. Fig. 6.4e compares the simulated sorting efficiency spectra of the CI and scaffold color routers. For the CI color router, we achieve peak sorting efficiencies (80% for blue, 90% for green, 78% for red) across the three spectral bands with good spectral separation. The target transmission/sorting spectrum used to define band edges and desired routing behavior, as well as a discussion of the crosstalk, is given in the Supplement. In contrast, the scaffold color router exhibits markedly different behavior: the B and R bands show reduced sorting efficiency (57% and 55% respectively), and the intended symmetric behavior in the green band between the G1 and G2 quadrants is broken. Instead, G2 approximately follows the expected spectral behavior while G1 becomes relatively flat across the bandwidth, which compromises spectral selectivity and increases crosstalk. These effects are directly visible in Fig. 6.4g, which displays the simulated focal-plane intensity distributions over each quadrant at the center of their respective spectral band. In each spectral band, although the majority of intensity resides in the desired quadrant, a significant amount of power is directed to undesired quadrants, particularly near the diagonal boundaries.

We attribute this discrepancy primarily to the limitations of the phase-matching procedure applied above, where we assume that the response of each individual voxel in the CI router can be represented by an independent phase delay, whereas the actual pillar-in-scaffold device supports both multiple scattering and near-field coupling, within each layer and also between layers. Such interactions are not captured when using a voxel-to-pillar library approach, and the resulting spectral behavior of the scaffold device is perturbed greatly from the CI device as a result. Correcting this would require optimizing the actual pillar-in-scaffold geometry using full-wave simulation of the entire scaffold color router structure at each iteration, and performing the optimization directly with respect to the scaffold degrees of freedom, i.e. the spatially varying pillar width distribution $p(x, y, z)$ and stepping the device using the computed gradients $\partial F / \partial p$, where F is the figure of merit. This approach was not used due to limits of computational complexity and time, but is expected to recover performance comparable to or potentially exceeding that of the CI color router. Additionally, the bare scaffold (with no pillars) behaves as a strongly diffractive structure with intensity patterns at the focal plane that exhibit highly diagonal ($\times$ -like) symmetry, as shown in Fig. S5. When considering the traditional Bayer arrangement, the desired spatial routing pattern means that areas of high intensity are often located close to the diagonals as well. Because our device is optimized using voxels of continuous index, with no scaffold present, any undesired scaffold-induced diffraction cannot be accounted for until after the optimization is complete. This increases apparent crosstalk from the intended meta-optic routing response, which is consistent with the above observation in Fig. 6.4g that undesired areas of high intensity tend to cluster near the diagonal boundaries. To mitigate this, we rotated the quadrant arrangement of our color router by 45° , as shown in Fig. 6.4a, relative to the conventional pixel geometry which tends to center high areas of intensity in a “+” arrangement rather than an “ $\times$ ”. This rotation improves contrast between the scaffold’s intrinsic diffraction features and the meta-optic device’s intended routing response, and has little to no effect on the device’s performance as spatial and spectral routing are largely unaffected by scrambling of the routing regions⁵⁴.

The fabricated scaffold color router was characterized using the mid-IR optical setup shown in Fig. 6.5. Fig. 6.4f shows the measured sorting-efficiency spectrum. Notably, the same asymmetry in the behavior of the G1 and G2 quadrants associated with the simulated scaffold color router is seen experimentally, supporting the interpretation that these effects originate from scaffold-mediated scattering. Other reasons for the discrepancy between simulation and experiment include residual

porosity or incomplete densification after calcination, as well as the finite numerical aperture of the measurement system. In general, however, the device displays clear spatial separation of the B, G, and R spectral bands, achieving peak measured sorting efficiencies of 48%, 45%, and 51% respectively, which surpass the limits imposed by traditional absorptive color-filter Bayer pixel arrays. This result demonstrates that inverse-designed volumetric meta-optic devices can be robustly implemented using a fully printed, high-index ceramic architecture. This result highlights the potential of the nano-AM TiO_2 platform for realizing compact, multifunctional and three-dimensional meta-optic elements.

6.4 Conclusion

In this work, we established nano-additive manufacturing of high-index TiO_2 as a general and scalable platform for realizing volumetric metaoptics. By integrating scaffold-supported voxel architectures with a shrinkage-mitigated two-photon lithography workflow and transfer-enabled ceramic conversion, this approach enabled the fabrication of large-area, mechanically continuous three-dimensional oxide nanostructures in a single lithographic process without layer-to-layer alignment. The resulting platform supported near-100 nm feature sizes while preserving geometric fidelity and optical quality through development, transfer, and high-temperature calcination, addressing key limitations that have historically constrained the broader application of additive manufacturing of high-index ceramic photonic structures.

Building on this fabrication capability, we demonstrated a sequence of mid-infrared metaoptical devices with increasing functional complexity operating across the 4 μm –5 μm band. Homogeneous lattices and monolayer chromatic metalenses established a baseline for phase-only wavefront shaping in additively manufactured TiO_2 , achieving near-diffraction-limited focusing over numerical apertures up to 0.49. We then extended this framework to heterogeneous multilayer metalenses, exploiting vertical stacking as an additional design degree of freedom to independently tailor phase and group delay. These achromatic metalenses suppressed chromatic focal shift over a 1 μm bandwidth while maintaining stable focal positions and diffraction-limited spot sizes at both moderate ($\text{NA} = 0.25$) and high ($\text{NA} = 0.49$) numerical apertures. Finally, we demonstrated a fully volumetric, inverse-designed scaffold-based color router that performed wavelength-dependent wavefront sorting within a compact ceramic volume, illustrating the ability of the nano-AM platform to implement nonlocal and multifunctional optical responses beyond the reach of planar metasurfaces.

More broadly, this work positions nano-additive manufacturing of high-index ceramics as a promising route toward multilayer, volumetric metaoptics that support multifunctionality within compact three-dimensional architectures. Such capabilities open pathways toward miniaturized optical systems for imaging, sensing, and spectral routing, and provide a foundation for future extensions incorporating polarization control, nonlocal responses, and fully inverse-designed volumetric photonic devices.

6.5 Experimental Section

Materials (contributed by co-authors): Titanium(IV) ethoxide (>97%), acrylic acid (anhydrous, 99%), 2-methoxyethanol (anhydrous, 99.8%), pentaerythritol triacrylate (technical grade), dichloromethane (DCM; anhydrous, $\geq 99.8\%$), and propylene glycol monomethyl ether acetate (PGMEA; >99.5%) were purchased from Sigma-Aldrich. 7-diethylamino-3-thenoylcoumarin was purchased from Exciton. All chemicals were used as received.

Resin Preparation (contributed by co-authors): The hybrid photoresist preparation approach was adapted from Vyatskikh *et al.*²⁸. Acrylic acid (576.5 mg) was slowly added to titanium(IV) ethoxide (456.2 mg) inside a glovebox and manually agitated. The solution was then placed in the vacuum antechamber of the glovebox for 30 min to remove ethanol produced as a result of the ligand-exchange reaction. The resulting solution was subsequently removed from the glovebox and mixed with 2-methoxyethanol (300 mg) and pentaerythritol triacrylate (1.2 g), agitated for 30 s using a vortex mixer, and transferred to a yellow-light environment. Next, 7-diethylamino-3-thenoylcoumarin (30 mg) was dissolved in dichloromethane (DCM; 150 mg) and added to the mixture, followed by an additional 30 s of vortex mixing. The resulting photoresist was allowed to rest for 5 min to release air bubbles trapped during agitation.

Additive Manufacturing (contributed by co-authors): Before printing, silicon substrates were sonicated in acetone and isopropyl alcohol (IPA) for 5 min each, and silanized to promote structure adherence. Two-photon lithography was performed on a Photonic Professional GT (Nanoscribe GmbH) using a Zeiss Plan-Apochromat 63 $\times$ /1.4 NA oil immersion objective. Typical writing conditions used 20–25 mW laser power at the sample and a writing speed of 30 mm/s. Samples were printed in a “sandwich” configuration with the resin confined between a glass slide and the substrate using Kapton tape spacers, producing a total gap of 100 μ m and enabling

prints up to this height. Print sections of the pillars and scaffold lattices were generated in a commercially available 3D CAD software (Fusion 360). The models were batched processed with the commercial software Describe (Nanoscribe GmbH) to generate machine code for TPL. After printing, the samples were developed in 2-methoxyethanol overnight, followed by 20min in PGMEA and 3min in filtered IPA. The samples were then transferred to a critical point dryer (Autosamdri-931).

Release, Transfer and Calcination (contributed by co-authors): Following CPD, dried precursor lattices were mechanically released from the silicon substrate using a micro-manipulation needle attached to a commercially available wire bonder (West-Bond 7476D-79) and transferred to a fused silica receiving substrate for calcination. Thermal conversion to crystalline TiO_2 was performed in an open-atmosphere tube furnace (MTI OTF-1500X). Samples were ramped from room temperature to the target calcination temperature (typically 550–900°C) at 0.5°C/min, held at peak temperature for 5 min, and cooled to room temperature at 1°C/min.

Inverse Design of Color Router: The color routers presented here are designed using adjoint-based inverse design^{55,56}, which enables gradient-based optimization over the large design space presented by the use of a multilayer voxelized grid. The (wavelength-dependent) figure of merit used was the mode overlap integral between the obtained electromagnetic field over each target pixel region, and the desired spatial field distribution over that region, with intensity varying spectrally as needed for the color router functionality. The gradient of the figure of merit with respect to the permittivity was obtained via a corresponding adjoint simulation, which was then used to step the refractive index profile of the color router via an ADAM-based⁵⁷ optimizer.

Numerical Simulations of Color Router: Numerical simulations of the continuous-index color routers were performed using Lumerical FDTD. Here PMLs were used along x -, y -, and z -directions, and the color routers were excited by normally incident Gaussian illumination polarized across the x and y axes and focused on the routers' input faces, with beam waist 3/4 of the device lateral size.

Optical Characterizations for Color Router Metaoptics: The experimental setups used to measure the sorting efficiencies and image the color router's focal plane are shown in the Supplement. The color router was printed in a square aperture. There is a change in focal length of the device due to dispersion of the measuring optics. As such, we translated the imaging objective (L3 in Fig. 6.5) axially in order to measure planes of differing distance (normal to the optical axis). Each device

was printed in a square aperture of diameter $50\text{ }\mu\text{m}$, and for each value of the laser wavelength, we imaged an identical square aperture, with no devices printed in it. The device's axial distance was tuned until the empty square aperture was in focus, and the micrometer value was noted to characterize the corresponding focal length shift for each laser wavelength. To measure the power incident on the input face of the device, we measured the power through an empty square aperture with a power meter. Then, to measure the power incident on each pixel of the 2×2 pixel array at the focal plane, we placed an aperture of the same shape and size as the pixel at the focal plane and measured the power passing through that aperture.

6.6 Supplementary Information

6.6.1 Experimental Characterization Setup




Figure 6.5: **Metalens characterization setups.** **a)** Experimental setups used to measure the focusing efficiency of the metalenses and color routers. L4 is inserted when characterizing color routers to weakly focus light onto the device. **b)** Setup used to capture focal spot images of the metalenses.

References

1. Chen, W. *et al.* Heterogeneous Multilayered TiO_2 Metaoptics via Two-Photon Lithography. *Manuscript in preparation* (2026).
2. Yu, N. & Capasso, F. Flat optics with designer metasurfaces. *Nature Materials* **13**, 139–150. ISSN: 1476-1122 (2014).
3. Kildishev, A. V., Boltasseva, A. & Shalaev, V. M. Planar Photonics with Meta-surfaces. *Science* **339**, 1232009. ISSN: 0036-8075 (2013).
4. Lin, D., Fan, P., Hasman, E. & Brongersma, M. L. Dielectric gradient meta-surface optical elements. *Science* **345**, 298–302. ISSN: 0036-8075 (2014).
5. Yu, N. *et al.* Light propagation with phase discontinuities: generalized laws of reflection and refraction. *Science* **334**, 333–7. ISSN: 0036-8075 (2011).
6. Khorasaninejad, M. & Capasso, F. Metalenses: Versatile multifunctional photonic components. *Science* **358**. ISSN: 0036-8075 (2017).
7. Wang, S. *et al.* Broadband achromatic optical metasurface devices. *Nature Communications* **8**, 187 (2017).
8. Chen, W. T. *et al.* A broadband achromatic metalens for focusing and imaging in the visible. *Nature Nanotechnology* **13**, 220–226. ISSN: 1748-3387 (2018).

9. Arbabi, A., Horie, Y., Bagheri, M. & Faraon, A. Dielectric metasurfaces for complete control of phase and polarization with subwavelength spatial resolution and high transmission. *Nature Nanotechnology* **10**, 937–943. ISSN: 1748-3387 (2015).
10. Arbabi, E., Kamali, S. M., Arbabi, A. & Faraon, A. Full-Stokes Imaging Polarimetry Using Dielectric Metasurfaces. *ACS Photonics* **5**, 3132–3140. ISSN: 2330-4022 (2018).
11. Rubin, N. A. *et al.* Matrix Fourier optics enables a compact full-Stokes polarization camera. *Science* **365**. ISSN: 0036-8075 (2019).
12. Ren, H. *et al.* Complex-amplitude metasurface-based orbital angular momentum holography in momentum space. *Nature Nanotechnology* **15**, 948–955. ISSN: 1748-3387 (2020).
13. Pitruzzello, G. Metaoptics for the consumer market. *Nature Photonics* **17**, 6–7. ISSN: 1749-4885 (2023).
14. Kamali, S. M., Arbabi, E., Arbabi, A. & Faraon, A. A review of dielectric optical metasurfaces for wavefront control. *Nanophotonics* **7**, 1041–1068 (2018).
15. Zhou, Y. *et al.* Multifunctional metaoptics based on bilayer metasurfaces. *Light: Science & Applications* **8**, 80. ISSN: 2095-5545 (2019).
16. Pan, C.-F. *et al.* 3D-printed multilayer structures for high-numerical aperture achromatic metalenses. *Science Advances* **9**, eadj9262 (2023).
17. Arbabi, A. *et al.* Miniature optical planar camera based on a wide-angle metasurface doublet corrected for monochromatic aberrations. *Nature Communications* **7**, 13682. eprint: 1604.06160 (2016).
18. Li, Z. *et al.* Heterogeneous Freeform Metasurfaces: A Platform for Advanced Broadband Dispersion Engineering. *arXiv*. eprint: 2412.12028 (2024).
19. Dorrah, A. H., Park, J.-S., Palmieri, A. & Capasso, F. Free-standing bilayer metasurfaces in the visible. *Nature Communications* **16**, 3126 (2025).
20. Roberts, G. *et al.* 3D-patterned inverse-designed mid-infrared metaoptics. *Nature Communications* **14**, 2768 (2023).
21. Jaiswal, A. *et al.* Two decades of two-photon lithography: Materials science perspective for additive manufacturing of 2D/3D nano-microstructures. *iScience* **26**, 106374. ISSN: 2589-0042 (2023).
22. Harinarayana, V. & Shin, Y. Two-photon lithography for three-dimensional fabrication in micro/nanoscale regime: A comprehensive review. *Optics & Laser Technology* **142**, 107180. ISSN: 0030-3992 (2021).
23. Farsari, M. & Chichkov, B. N. Two-photon fabrication. *Nature Photonics* **3**, 450–452. ISSN: 1749-4885 (2009).

24. Skliutas, E. *et al.* Multiphoton 3D lithography. *Nature Reviews Methods Primers* **5**, 15 (2025).
25. Ren, H. *et al.* An achromatic metafiber for focusing and imaging across the entire telecommunication range. *Nature Communications* **13**, 4183. eprint: 2201.07158 (2022).
26. Balli, F., Sultan, M., Lami, S. K. & Hastings, J. T. A hybrid achromatic metalens. *Nature Communications* **11**, 3892. eprint: 1909.07941 (2020).
27. Roques-Carmes, C. *et al.* Toward 3D-Printed Inverse-Designed Metaoptics. *ACS Photonics* **9**, 43–51. ISSN: 2330-4022 (2022).
28. Vyatskikh, A., Ng, R. C., Edwards, B., Briggs, R. M. & Greer, J. R. Additive Manufacturing of High-Refractive-Index, Nanoarchitected Titanium Dioxide for 3D Dielectric Photonic Crystals. *Nano Letters* **20**, 3513–3520. ISSN: 1530-6984 (2020).
29. Zhang, W. *et al.* Printing of 3D photonic crystals in titania with complete bandgap across the visible spectrum. *Nature Nanotechnology* **19**, 1813–1820. ISSN: 1748-3387 (2024).
30. Mori, T. *et al.* Pick and place process for uniform shrinking of 3D printed micro- and nano-architected materials. *Nature Communications* **14**, 5876 (2023).
31. Brigo, L. *et al.* 3D Nanofabrication of SiOC Ceramic Structures. *Advanced Science* **5**, 1800937. ISSN: 2198-3844 (2018).
32. Vyatskikh, A. *et al.* Additive manufacturing of 3D nano-architected metals. *Nature Communications* **9**, 593 (2018).
33. Liu, Y. *et al.* Structural color three-dimensional printing by shrinking photonic crystals. *Nature Communications* **10**, 4340. eprint: 1905.05913 (2019).
34. Portela, C. M. *et al.* Supersonic impact resilience of nanoarchitected carbon. *Nature Materials* **20**, 1491–1497. ISSN: 1476-1122 (2021).
35. Wang, P. & Menon, R. Ultra-high-sensitivity color imaging via a transparent diffractive-filter array and computational optics. *Optica* **2**, 933–939. ISSN: 2334-2536. <https://opg.optica.org/optica/abstract.cfm?uri=optica-2-11-933> (2026) (Nov. 2015).
36. Camayd-Muñoz, P., Ballew, C., Roberts, G. & Faraon, A. Multifunctional Volumetric Meta-Optics for Color and Polarization Image Sensors. *Optica* **7**, 280. ISSN: 2334-2536. (2020) (Apr. 2020).
37. Miyata, M., Nemoto, N., Shikama, K., Kobayashi, F. & Hashimoto, T. Full-color-sorting metalenses for high-sensitivity image sensors. *Optica* **8**, 1596. ISSN: 2334-2536. <https://www.osapublishing.org/abstract.cfm?URI=optica-8-12-1596> (2021) (Dec. 2021).

38. Ren, J. *et al.* Wavelength-dependent multifunctional metalens devices via genetic optimization. *Optical Materials Express* **11**, 3908–3916. ISSN: 2159-3930. <https://opg.optica.org/ome/abstract.cfm?uri=ome-11-11-3908> (2026) (Nov. 2021).
39. Catrysse, P. B., Zhao, N., Jin, W. & Fan, S. Subwavelength Bayer RGB color routers with perfect optical efficiency. *Nanophotonics* **11**. ISSN: 2192-8606, 2192-8614. <https://www.degruyter.com/document/doi/10.1515/nanoph-2022-0069/html> (2022) (Mar. 2022).
40. Li, J. *et al.* Single-Layer Bayer Metasurface via Inverse Design. *ACS Photonics* **9**, 2607–2613. <https://doi.org/10.1021/acsp Photonics.2c00990> (2026) (Aug. 2022).
41. Zou, X. *et al.* Pixel-level Bayer-type colour router based on metasurfaces. *Nature Communications* **13**, 3288. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-022-31019-7> (2026) (June 2022).
42. Kim, C. *et al.* Freeform metasurface color router for deep submicron pixel image sensors. *Science Advances* **10**, eadn9000. <https://www.science.org/doi/full/10.1126/sciadv.adn9000> (2024) (May 2024).
43. Hsu, W.-L. *et al.* Simplest but Efficient Design of a Color Router Optimized by Genetic Algorithms. *ACS Photonics* **12**, 1402–1408. <https://doi.org/10.1021/acsp Photonics.4c01905> (2026) (Mar. 2025).
44. Bayer, B. E. US3971065A. <https://patents.google.com/patent/US3971065A/en> (2025)(1976).
45. Palum, R. *Image sampling with the bayer color filter array in PICS 2001: Image processing, image quality, image capture, systems conference, montréal, quebec, canada, april 22-25, 2001* tex.bibsource: dblp computer science bibliography, <https://dblp.org> tex.timestamp: Tue, 15 Jul 2014 19:48:31 +0200 (IS&T - The Society for Imaging Science and Technology, 2001), 239–245.
46. Chen, W. T. *et al.* Dispersion-engineered metasurfaces reaching broadband 90% relative diffraction efficiency. en. *Nature Communications* **14**, 2544. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-023-38185-2> (2026) (May 2023).
47. Mansouree, M. *et al.* Multifunctional 25D metastructures enabled by adjoint optimization. en. *Optica* **7**, 77. ISSN: 2334-2536. <https://www.osapublishing.org/abstract.cfm?URI=optica-7-1-77> (2020) (Jan. 2020).
48. Roberts, G. *et al.* 3D-patterned Inverse-Designed Mid-Infrared Metaoptics. *Nature Communications* **14**, 2768. ISSN: 2041-1723. (2025) (May 2023).
49. Ballew, C., Roberts, G., Camayd-Muñoz, S., Debbas, M. F. & Faraon, A. Mechanically reconfigurable multi-functional meta-optics studied at microwave frequencies. *Scientific Reports* **11**, 11145. ISSN: 2045-2322. <http://www.nature.com/articles/s41598-021-88785-5> (2021) (Dec. 2021).

50. Ballew, C., Roberts, G. & Faraon, A. Multi-dimensional wavefront sensing using volumetric meta-optics. *Optics Express* **31**, 28658–28669. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-31-18-28658> (2025) (Aug. 2023).
51. Li, Z., Pestourie, R., Lin, Z., Johnson, S. G. & Capasso, F. Empowering Meta-surfaces with Inverse Design: Principles and Applications. *ACS Photonics* **9**, 2178–2192. <https://doi.org/10.1021/acsp Photonics.1c01850> (2023) (July 2022).
52. Kischkat, J. *et al.* Mid-infrared optical properties of thin films of aluminum oxide, titanium dioxide, silicon dioxide, aluminum nitride, and silicon nitride. *Applied Optics* **51**, 6789–6798. ISSN: 2155-3165. <https://opg.optica.org/ao/abstract.cfm?uri=ao-51-28-6789> (2026) (Oct. 2012).
53. Li, Z. *et al.* Inverse design enables large-scale high-performance meta-optics reshaping virtual reality. *Nature Communications* **13**, 2409. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-022-29973-3> (2022) (May 2022).
54. Ballew, C., Roberts, G. & Faraon, A. Multi-dimensional wavefront sensing using volumetric meta-optics. *Optics Express* **31**, 28658 (2023).
55. Hughes, T. W., Minkov, M., Williamson, I. A. D. & Fan, S. Adjoint Method and Inverse Design for Nonlinear Nanophotonic Devices. *ACS Photonics* **5**, 4781–4787. <https://doi.org/10.1021/acsp Photonics.8b01522> (2022) (Dec. 2018).
56. Zhang, G., Xu, D.-X., Grinberg, Y. & Liboiron-Ladouceur, O. Experimental demonstration of robust nanophotonic devices optimized by topological inverse design with energy constraint. *Photonics Research* **10**, 1787–1802. ISSN: 2327-9125. <https://opg.optica.org/prj/abstract.cfm?uri=prj-10-7-1787> (2025) (July 2022).
57. Kingma, D. P. & Ba, J. *Adam: A Method for Stochastic Optimization* Jan. 2017. <http://arxiv.org/abs/1412.6980> (2025).

INVERSE-DESIGNED TEN-LAYER VISIBLE COLOR SORTER VIA DIRECT-WRITE Sb_2S_3 MULTILAYER METAOPTICS

This chapter is directly reproduced from [1], a paper published as a key part of the work in this thesis. As described in the abstract below, this work introduces the first experimental demonstration of multi-layer metaoptic color routers in the visible, enabled by a novel Sb_2S_3 fabrication platform.

Figures 7.1 and 7.2c reprinted with permission from the copyright holder, A. Bilgehan Baspinar.

Abstract

Color imaging systems require spectral discrimination at the pixel level, conventionally achieved through absorptive Bayer color filter arrays. Each pixel accepts only its designated color band and discards the remainder, imposing an efficiency ceiling on each color channel. Metaoptic color sorters address this by spatially sorting different spectral components of light into their designated pixel regions instead of absorbing them. Single-layer implementations have been experimentally demonstrated in the visible, but multilayer architectures—which offer substantially greater design freedom and multifunctionality—have not previously been fabricated at visible wavelengths. Here we report the first experimental realization of a multilayer color sorter operating in the visible regime: a ten-layer Sb_2S_3 /HSQ inverse-designed device occupying a $4 \times 4 \times 6.325 \mu\text{m}^3$ volume, compatible with a $2 \mu\text{m}$ pixel pitch, the fabrication of which is achieved through an Sb_2S_3 direct-write platform requiring only two process steps per patterned layer, bypassing the repeated deposition and etch cycles that have made multilayer visible metaoptics impractical to date. Measured sorting efficiencies of 32.1%, 58.0%, and 30.5% surpass the absorptive Bayer limit across all three channels, and establish the Sb_2S_3 direct-write platform as a practical route to high-layer-count visible metaoptics.

7.1 Introduction

Color discrimination at the pixel level is central to digital imaging. In the dominant sensor architecture, a Bayer color filter array places an absorptive filter over each photodetector, restricting the detected light to the red, green, or blue band assigned to

that pixel while absorbing the remaining spectral components.² The result is a fixed per-pixel efficiency ceiling: red and blue photodetectors receive no more than one quarter of the photons incident on their aperture, and green detectors no more than one half.³ As pixel pitches in modern sensors approach the single-micron or sub-micron level, the shrinking collection area limits the available signal per exposure, making this additional photon loss increasingly significant for low-light imaging performance.

Metaoptic color sorters have been proposed as a means of recovering these discarded photons. Rather than blocking out-of-band light, a color sorter redistributes each wavelength component into its designated pixel, so that the full incident photon flux is utilized across the sensor. Single-layer implementations based on subwavelength-patterned dielectric metasurfaces have been fabricated and experimentally characterized, employing design strategies that include phase-gradient deflection⁴, inverse-designed freeform patterns^{5–7}, and evolutionary optimization^{6,8,9}. These fabricated devices have shown that routing efficiencies above the absorptive Bayer limit are achievable, and have established the viability of metaoptic color sorting in the visible.

Single-layer color sorters, however, face intrinsic design constraints. A single patterned plane must simultaneously provide spectral discrimination and selectivity, high transmission, and suppression of inter-channel crosstalk. Because of the finite thickness and fundamental limitations of a single-layer metasurface^{10,11}, the design degrees of freedom available are limited and it becomes difficult to achieve all these at once.

Volumetric or multilayer architectures, which distribute the optical interaction across multiple patterned layers (typically in the direction of optical propagation), offer a substantially expanded design space and have been predicted in simulation to improve routing efficiency, spectral selectivity, and crosstalk suppression relative to single-layer devices.^{12–15}

Multilayer metaoptic color routing has been experimentally demonstrated at mid-infrared wavelengths, where the longer operating wavelength regime places the required minimum feature sizes and layer-to-layer alignment tolerance at scales accessible by existing fabrication tools such as two-photon lithography^{16,17}. To date, however, no multilayer color sorter operating in the visible has been experimentally fabricated. Metaoptics operating in the visible typically require sub-150 nm feature resolution at each layer^{4,6}, as well as fabrication alignment accuracy between layers

at comparable scales^{18,19}, and a layer count high enough to provide sufficient additional degrees of freedom for high performance^{12,15}. Conventional multilayer dielectric metasurface processes, such as stacked amorphous-silicon or silicon-nitride platforms, generally require repeated cycles of material deposition, lithographic patterning, reactive-ion etching, mask removal, spacer deposition, and planarization steps for each patterned layer^{19,20}. As such, the complexity of fabrication and the effects of alignment errors scale cumulatively with layer count, making multilayer visible metaoptics impractical to fabricate via standard means.

This need can be suitably addressed by a direct-write electron-beam lithography approach based on an antimony trisulfide (Sb_2S_3) precursor²¹. In this fabrication platform, the precursor is spin-coated and patterned directly by electron-beam exposure, which triggers in-situ conversion to high-index Sb_2S_3 . A planarization layer of hydrogen silsesquioxane (HSQ) is then applied before the next layer is written. Each patterned layer requires only these two steps, bypassing the material deposition and plasma-etch cycles that account for the majority of process steps in standard multilayer workflows, while retaining interlayer alignment accuracy below 10 nm. This platform has previously been used to demonstrate three-layer resonant filter devices in the visible and near-infrared²¹, but its applicability to volumetric metaoptic devices with higher layer counts has not yet been established.

Here we extend the Sb_2S_3 direct-write platform to ten patterned layers and demonstrate the first experimentally fabricated multilayer color sorter operating in the visible. The device is designed by adjoint-based inverse design to route light across three separate spectral bands into a 2×2 Bayer-type quadrant arrangement. We present the design methodology, fabrication process, and transmission-mode optical characterization, and report measured sorting efficiencies that surpass the absorptive Bayer limit across all three color channels. The ten-layer architecture demonstrated here exceeds the three-layer device count of the prior demonstration on this platform²¹ by more than threefold, and to our knowledge represents the highest layer count realized for a direct-write metaoptic color sorter in the visible regime.

7.2 Results and Discussion

7.2.1 Device Design

The device occupies a $4 \times 4 \times 6.325 \mu\text{m}^3$ volume, comprising ten Sb_2S_3 patterned layers each 340 nm thick, separated by nine HSQ spacer layers each 325 nm thick, on a silicon carbide substrate. The $4 \times 4 \mu\text{m}$ lateral footprint accommodates a 2×2




Figure 7.1: (a) Conventional nanofabrication platforms would require the displayed steps, repeated for each layer, in order to fabricate a multilayer metaoptic device. (b) Novel Sb_2S_3 platform that reduces the number of steps per layer. The SB-BDCA precursor is spun onto the previous layer and directly patterned via EBL. It is then developed via immersion in isopropanol and annealing. HSQ is then spun on and planarized.




Figure 7.2: (a) Schematic of inverse-designed metaoptic color sorter device with 10 layers and 9 spacers, and projected focal spots for red, green, and blue light on the focal plane. (b) Topmost layer of the 10-layer device after optimization. Each Sb_2S_3 patterned layer (magenta) is embedded in a HSQ (white) medium, forming a flat substrate for subsequent layers. (c) SEM image of top layer of device, the same layer as displayed in b).

pixel array at a $2\text{ }\mu\text{m}$ pixel pitch. The design material is patterned Sb_2S_3 ($n = 2.0\text{--}2.2$ across the visible band, discussed further in Section 7.5.2) against a background of HSQ, which serves as both the interlayer spacer and the surrounding medium ($n = 1.4$). The color sorter is designed using inverse design via the adjoint method^{lalau2013, molesky2018}, which computes the gradient of a figure of merit with respect to the permittivity at every point in the design volume from only two electromagnetic simulations per iteration, independent of the total number of design variables. The figure of merit is the mode overlap between the simulated transmitted field over each target quadrant at the focal plane and the desired spatial intensity distribution, evaluated at the three target wavelengths. The device is optimized to route normally incident light at 550 nm (green), 650 nm (red), and 750 nm (near-infrared or NIR) into the four quadrants of a 2×2 Bayer-type pixel array at a focal distance of $7.04\text{ }\mu\text{m}$, with green routed to two diagonally opposite quadrants and red and NIR each occupying one quadrant. Target wavelengths were selected at 550 nm, 650 nm, and 750 nm instead of the standard RGB values of 475 nm, 550 nm, and 650 nm, as the silicon carbide substrate exhibits particularly high optical absorption at wavelengths below $\sim 500\text{ nm}$, introducing a systematic loss in transmission²² that degrades device performance in the conventional blue channel. The figure of merit is maximized as a weighted average over the three wavelength channels. The device was illuminated in simulation by a normally incident Gaussian beam with a waist radius of $85\text{ }\mu\text{m}$, sufficiently large relative to the $4\text{ }\mu\text{m}$ device footprint to approximate uniform plane-wave excitation across the device. Transmission is normalized against the simulated transmission through an empty aperture of equivalent dimensions as the device, albeit with no device present. In simulation (Fig. 7.3a) for the three bands green, red, and NIR, we obtain efficiency peaks of 55.9%, 66.4%, and 42.8% at 593 nm, 659 nm and 732 nm respectively. These peaks are redshifted from the design targets of 550 nm, 650 nm, and 750 nm by approximately +43 nm, +9 nm, and -18 nm; however, they well surpass the limits of a conventional absorptive-based Bayer color filter. Fig. 7.3c displays the intensities at the focal plane at the representative peak wavelengths. At each band, good contrast is observed between the designated quadrant(s) for that band, and all other quadrants.

The initial seed used was a randomized permittivity distribution. Forward and adjoint simulations were performed in Lumerical FDTD (ANSYS, Inc.) with perfectly-matched layer (PML) boundary conditions applied in all three axes. Using the adjoint method we obtained the sensitivity of the figure of merit with respect to the permittivity distribution, stepping it forward with the ADAM²³ gradient descent scheme

(learning rate 0.01). We enforce a binary $\text{Sb}_2\text{S}_3/\text{HSQ}$ distribution throughout the optimization via a Heaviside projection scheme^{24,25} applied within the optimization loop, with the projection sharpness increased progressively with iteration, such that the final design is a binary structure suited for fabrication. A minimum feature size of 115 nm is enforced^{25,26} for the designed pattern throughout the optimization.

Device optimizations were run on a high performance computing cluster (Resnick High Performance Computing Center) using 10 computing nodes, each of which was allocated 8 Intel CPU cores (mix of Skylake 2.1 GHz and Cascadelake 2.1 GHz processors), to run simulations in parallel. With these resources each simulation was completed in ~ 12 – 15 minutes depending on the thickness of the device. Given 300 iterations, the optimizations completed in ~ 60 – 75 h of compute time.

7.2.2 Fabrication

The ten-layer $\text{Sb}_2\text{S}_3/\text{HSQ}$ color sorter is realized using the direct-write EBL platform introduced by Baspinar et al.²¹, based on an antimony–butyldithiocarbamic acid (Sb–BDCA) precursor that undergoes in situ decomposition into Sb_2S_3 upon electron-beam exposure. Conventional multilayer EBL processes for high-index dielectric nanostructures typically involve approximately nine steps per layer, including material deposition, resist coating, exposure, development, hard-mask deposition and lift-off, plasma etching, mask removal, and planarization. The present approach requires four steps with no deposition or etch stages, and interlayer alignment accuracy limited only by the overlay accuracy of electron-beam lithography, which is approximately 5 – 10 nm²⁷. As such, the process overhead per-layer is substantially reduced.

The Sb–BDCA precursor is synthesized by combining antimony(III) oxide dissolved in ethanol with butyldithiocarbamic acid, itself formed by reacting 1-butylamine with carbon disulfide. The resulting complex is stable over several weeks and is applied by spin-coating directly onto the substrate immediately prior to each patterning step. Devices are fabricated on a silicon carbide (SiC) substrate on which gold alignment markers are defined by a standard lift-off process prior to any device patterning. This alignment marker set is referenced for all ten layers.

For each patterned layer, the precursor solution is diluted in ethanol and spin-coated in two successive passes to achieve the target 340 nm film thickness. The coated substrate is then exposed using a 100 kV EBL system at a beam current of 20 nA. A multipass strategy distributes the total accumulated dose of 500 mC/cm² across ten

sequential passes of 50 mC/cm^2 each, promoting spatially uniform decomposition of the precursor. Following exposure, the substrate is immersed in isopropanol for ten minutes to remove both the unexposed precursor and the organic decomposition byproducts produced in the exposed regions, then annealed at 180°C for two minutes on a hotplate under nitrogen atmosphere to drive complete conversion of the exposed material to Sb_2S_3 . Subsequently, planarization of HSQ is performed by spin-coating HSQ to a thickness of 325 nm over the patterned Sb_2S_3 structures, followed by a partial-cure anneal at 180°C for two minutes. This sequence is repeated for each of the ten layers. Scanning electron microscopy imaging of the completed device (Fig. 7.2c) shows successful realization of the design pattern at each layer, with Sb_2S_3 features resolved to 115 nm.

7.3 Characterization

The fabricated device is characterized in transmission using a visible-wavelength optical setup (Fig. 7.5). Broadband illumination is normally incident on the device from below through the SiC substrate, provided by a SuperK COMPACT supercontinuum laser (NKT Photonics). A condensing lens focuses the beam onto the device input face through a $150 \mu\text{m}$ square aperture fabricated on the substrate, within which the device is located.

The focal plane of the device passes through a relay system comprising two lenses of focal length 150 mm, with $75\times$ magnification such that the $4\times 4 \mu\text{m}^2$ device aperture corresponds to a $300\times 300 \mu\text{m}^2$ field in the image plane. A $150 \mu\text{m}$ square aperture is then stepped across the four quadrant positions of the Bayer array in the focal plane, which allows the optical power directed into each quadrant to be measured independently. The sorting efficiency for each quadrant is defined as the fraction of the total power incident on the 2×2 quadrant array which is directed into that quadrant. The incident power reference is established at each wavelength by measuring transmission through an identical $150 \mu\text{m}$ square aperture with no device present. Camera images are acquired at each wavelength immediately followed by a background frame with the laser blocked. The background is subtracted prior to analysis in order to reduce contributions from ambient illumination and detector thermal drift.

Measured sorting efficiencies reach 32.1% for green, 58.0% for red, and 30.5% for NIR at the respective peak wavelengths (Fig. 7.3b), with efficiency peaks observed at 570 nm, 630 nm, and 732 nm for the green, red, and NIR channels respectively.




Figure 7.3: (a) Sorting efficiency in each quadrant, defined as the power in each quadrant normalized over the total power incident on all four quadrants. The wavelength range is divided into three distinct spectral bands denoted B, G and R, centered at 550 nm, 650 nm and 750 nm respectively. The green spectral band is duplicated across quadrants G1 and G2 for increased sensitivity. The given wavelengths are the values at which sorting efficiency contrast for the relevant quadrant is maximal. (b) Experimentally measured sorting efficiencies of the color sorter over each quadrant area across the wavelength range. (c) Simulated intensities at the focal plane for the color router.

These values surpass the conventional limits of an absorptive Bayer filter across all three channels. Focal-plane intensity images acquired at each peak wavelength confirm spatial separation of the three color channels into their designated quadrants, with the majority of transmitted power concentrated in the target region at each wavelength (Fig. 7.3d). Residual intensity near quadrant boundaries constitutes the primary source of inter-channel crosstalk.

The measured efficiencies fall below the simulated values, which is consistent with discrepancies observed in prior volumetric metaoptic fabrications and attributable to several sources. Due to fabrication discrepancies, there are deviations between the actual dimensions of Sb_2S_3 features and their intended dimensions, which perturbs the sorting efficiency and introduces asymmetry between channels. Cumulative electron-beam exposure of the HSQ planarization layers during successive patterning steps also increases the refractive index of deposited spacer material, which is not accounted for during the design and optimization. Cumulative electron-beam exposure of the HSQ planarization layers during successive patterning steps can further crosslink or increase the density of HSQ, modifying its optical properties relative to the unexposed spacer assumed in simulation^{28,29}. Surface roughness at each of the Sb_2S_3 /HSQ interfaces also causes additional scattering losses that are not accounted for in the lossless simulation model.

7.4 Conclusion

We have demonstrated the first experimentally fabricated multilayer color sorter operating in the visible regime. The device consists of ten inverse-designed Sb_2S_3 /HSQ layers occupying a $4 \times 4 \times 6.325 \mu\text{m}^3$ volume consistent with a $2 \mu\text{m}$ pixel pitch, routing 550 nm, 650 nm, and 750 nm light into a standard 2×2 Bayer quadrant arrangement at a focal distance of $7.04 \mu\text{m}$. Measured sorting efficiencies of 32.1%, 58.0%, and 30.5% surpass the throughput ceiling of absorptive Bayer color filters across all three channels.

Natural extensions of this work include devices with higher layer counts for improved spectral selectivity and reduced inter-channel crosstalk,^{12,15} as well as integration of polarization sorting as an additional multiplexing dimension within the same volumetric framework.^{16,30}

Natural extensions of this work include devices with higher layer counts^{12,31} for improved spectral selectivity and reduced inter-channel crosstalk, integration of polarization sorting as an additional multiplexing dimension^{16,30} within the same volu-

metric framework. These results position the Sb_2S_3 direct-write platform as a practical route to high-layer-count visible metaoptics and establish inverse-designed volumetric color sorters as experimentally viable components for next-generation image sensors and compact hyperspectral imaging systems.

7.5 Supplementary Information

7.5.1 Adjoint Optimization Details

Inverse design of the color sorter is performed using the adjoint method^{32,33}, which provides the gradient of a scalar figure of merit with respect to the permittivity at every point in the design volume from exactly two electromagnetic simulations per iteration, regardless of the number of design variables.

The figure of merit f quantifies the degree to which the device routes each target wavelength into its designated quadrant at the focal plane. For N_λ wavelengths and N_c color channels, it is defined as a weighted average of mode overlap integrals:

$$f(\mathbf{n}) = \frac{1}{N_\lambda N_c} \sum_{\lambda, c} \frac{\left| \int_{S_c} (\mathbf{E}_\lambda \times \mathbf{H}_{m,\lambda,c}^* + \mathbf{E}_{m,\lambda,c}^* \times \mathbf{H}_\lambda) \cdot d\mathbf{S} \right|^2}{2 \int_{S_c} \text{Re}(\mathbf{E}_{m,\lambda,c} \times \mathbf{H}_{m,\lambda,c}^*) \cdot d\mathbf{S}}, \quad (7.1)$$

where $\mathbf{E}_\lambda$, $\mathbf{H}_\lambda$ are the simulated transmitted electric and magnetic fields at wavelength λ ; $\mathbf{E}_{m,\lambda,c}$, $\mathbf{H}_{m,\lambda,c}$ are the target mode profiles for color channel c at wavelength λ ; and S_c is the surface corresponding to the target quadrant of channel c at the focal plane. The target mode profiles are Gaussian distributions within their respective $2\text{ }\mu\text{m} \times 2\text{ }\mu\text{m}$ quadrant regions, with $1\text{ }\mu\text{m}$ FWHM so as to confine them within the boundary of each region. These profiles exist at the device focal plane, defined $7.04\text{ }\mu\text{m}$ beyond the device output face. Three wavelength bands are used, centered at 550 nm , 650 nm , and 750 nm , with the 650 nm channel optimized over two diagonally opposite quadrants. The aggregate figure of merit f may be maximized over the target quadrant regions or minimized over non-target regions to suppress crosstalk, as appropriate for each channel.

For each wavelength, a forward simulation computes the fields $\mathbf{E}_{\text{fwd}}$, $\mathbf{H}_{\text{fwd}}$ throughout the device volume under normally incident Gaussian excitation with waist radius $85\text{ }\mu\text{m}$, entering from below through the SiC substrate. The adjoint simulation then drives the device from the appropriate region in the focal plane, depending on the wavelength band, using a source whose amplitude and phase are derived from the fields over the target quadrant in the forward simulation³². The adjoint fields $\mathbf{E}_{\text{adj}}$, $\mathbf{H}_{\text{adj}}$ within the design volume are then recorded.

The sensitivity of f with respect to the refractive index at each point $\mathbf{x}$ in the design volume is assembled from the forward and adjoint fields as³²:

$$\frac{df}{dn(\mathbf{x})} = 2n(\mathbf{x}) \operatorname{Re}(\mathbf{E}_{\text{fwd}}(\mathbf{x}) \cdot \mathbf{E}_{\text{adj}}(\mathbf{x})) . \quad (7.2)$$

The gradient is accumulated as a weighted sum across all wavelengths and channels. The refractive index distribution is updated using the ADAM algorithm²³ with learning rate $\alpha = 0.01$ and standard moment decay parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$:

$$\begin{aligned} m_t &= \beta_1 m_{t-1} + (1 - \beta_1) \nabla f_t, \\ v_t &= \beta_2 v_{t-1} + (1 - \beta_2) (\nabla f_t)^2, \\ n_{t+1} &= n_t + \alpha \frac{\hat{m}_t}{\sqrt{\hat{v}_t + \varepsilon}}, \end{aligned} \quad (7.3)$$

where $\hat{m}_t = m_t / (1 - \beta_1^t)$ and $\hat{v}_t = v_t / (1 - \beta_2^t)$ are bias-corrected moment estimates and $\varepsilon = 10^{-8}$.

The optimizer operates on a continuous material density field $\rho \in [0, 1]$, which is mapped to a permittivity distribution through a smoothed Heaviside projection filter:

$$\tilde{\rho} = \frac{\tanh(\beta\eta) + \tanh(\beta(\rho - \eta))}{\tanh(\beta\eta) + \tanh(\beta(1 - \eta))}, \quad (7.4)$$

with projection threshold $\eta = 0.5$. The projected density $\tilde{\rho}$ is linearly interpolated between the permittivities of HSQ ($\varepsilon_{\text{HSQ}} = 1.4^2$) and Sb_2S_3 ($\varepsilon_{\text{Sb}_2\text{S}_3} = \bar{n}^2$, with $\bar{n}$ taken at the wavelength of interest):

$$\varepsilon(\tilde{\rho}) = \varepsilon_{\text{HSQ}} + \tilde{\rho}(\varepsilon_{\text{Sb}_2\text{S}_3} - \varepsilon_{\text{HSQ}}) . \quad (7.5)$$

The sharpness parameter β is initialized at $\beta = 10^{-3}$ and increased progressively throughout the optimization. At low β , the projection is nearly linear, preserving a continuous design space in early iterations; as β increases, the filter approaches a step function, driving the density field toward a binary Sb_2S_3 /HSQ distribution. The gradient of f with respect to ρ is obtained by composing Eq. (7.2) with the derivative of Eqs. (7.4)–(7.5) via the chain rule.

To ensure compatibility with the 115 nm minimum feature size of the fabrication process, a morphological opening operation (sequential erosion followed by dilation) using a 115 nm structuring element is applied to the density field at regular

intervals during the optimization. This suppresses isolated sub-resolution features that would be mechanically or lithographically unstable, without requiring a differentiable constraint in the objective. The filter is applied every ten iterations; the density values of affected voxels are adjusted and held fixed for the subsequent two iterations before being released back to the optimizer.

7.5.2 Sb_2S_3 Optical Constants Across the Visible Band




Figure 7.4: (a) Refractive index n and extinction coefficient k of Sb_2S_3 . Values are calculated from samples annealed at 180°C and 300°C. For further details of the annealing step in fabrication, see Section 7.2.2. (b) Reflectance spectroscopy performed on sample of Sb_2S_3 annealed at 180°C. Fit performed using a transfer-matrix-method thin film model. (c) Same as (b), but annealed at 300°C.

The simulations presented in the main text treat Sb_2S_3 as a lossless dielectric with no dispersion across the bandwidth of the device. This section characterizes the actual optical constants of the fabricated Sb_2S_3 across the 550 nm–750 nm operating band, and estimates the resulting efficiency penalty relative to the lossless model.

The complex refractive index $\tilde{n} = n + ik$ of the Sb_2S_3 films produced by the direct-write EBL process is extracted by fitting reflectance measurements performed on a continuous Sb_2S_3 film deposited under identical processing conditions on a silicon substrate. Reflectance spectra are acquired using a broadband supercontinuum source (the same one described in Section 7.5.3) over the wavelength range 450 nm–950 nm and fit using a transfer-matrix-method (TMM) thin-film model, in which the dispersive dielectric function is parameterized using a Lorentz oscillator:

$$\hat{\varepsilon}(\omega) = \varepsilon_{\infty} + \sum_j \frac{s_j}{\omega_{0,j}^2 - \omega^2 - i\Gamma_j\omega}, \quad (7.6)$$

where ε_{∞} captures contributions from higher-energy transitions, and s_j , $\omega_{0,j}$, Γ_j are the oscillator strength, resonance frequency, and linewidth of the j -th oscillator, respectively. Model parameters are determined by minimizing the residual between the simulated and measured reflectance via a particle-swarm optimization algorithm²¹. The resulting $n(\lambda)$, $k(\lambda)$ across the visible band are shown in Fig. 7.4.

7.5.3 Experimental Characterization Setup




Figure 7.5: .
Experimental characterization setup.

The characterization setup operates in two configurations sharing a common illumination and collection arm: a power-measurement configuration for quantifying sorting efficiency, and an imaging configuration for acquiring focal-plane intensity maps. The setup is shown in Fig. 7.5.

Broadband illumination is provided by a SuperK COMPACT supercontinuum laser (NKT Photonics), which generates a spatially coherent white-light beam spanning the visible band. The output beam passes through a 90:10 (R:T) beamsplitter, with 10% of the incident power directed to a reference detector for continuous power monitoring and 90% routed into the main optical path. A 50:50 beamsplitter further divides the beam: one arm feeds a direct-imaging camera used for sample alignment and beam profiling, and the other arm constitutes the sample illumination path. A

linear polarizer (LP1) and half-waveplate (HWP) sets the polarization state incident on the device. The beam is directed onto the input face of the device through the silicon carbide substrate at normal incidence. The device is fabricated within a 150 μm square aperture that suppresses stray light from adjacent substrate structures.

An objective lens (Olympus 100 $\times$, NA 0.8) placed below the substrate collects the transmitted light. Bandpass filters for the visible regime, placed after LP2, may be used to isolate a narrow spectral band from the broadband output light for imaging. The beam is relayed to two detection paths: a camera in the visible for imaging of the focal plane, and a spectrometer (Princeton Instruments) for spectrum collection.

The focal plane is reimaged with a total magnification of 75 $\times$, mapping the 4 $\times$ 4 μm device aperture to a 300 $\times$ 300 μm field. Each 2 $\times$ 2 μm quadrant subtends 150 $\times$ 150 μm in the reimaged plane, matched to the dimensions of a 150 μm square physical aperture placed at this plane. The aperture is stepped across the four quadrant positions and the per-quadrant optical power spectrum is recorded. The incident power reference at each wavelength is established by translating the substrate to position an identical 150 μm aperture containing no device into the beam path and recording the transmitted power.

Using the bandpass filters to isolate the desired spectral band, camera images are then acquired with a background frame taken immediately afterward with the beam blocked; the background is subtracted prior to analysis.

References

1. Foo, I., Baspinar, A. B., Pearson, P., Gu, Y. & Faraon, A. Freeform Sb_2S_3 -Based Multilayered Metaoptics. *Manuscript in preparation* (2026).
2. Bayer, B. E. US3971065A. <https://patents.google.com/patent/US3971065A/en> (2025)(1976).
3. Palum, R. *Image sampling with the bayer color filter array* in *PICS 2001: Image processing, image quality, image capture, systems conference, montréal, quebec, canada, april 22-25, 2001* tex.bibsource: dblp computer science bibliography, <https://dblp.org> tex.timestamp: Tue, 15 Jul 2014 19:48:31 +0200 (IS&T - The Society for Imaging Science and Technology, 2001), 239–245.
4. Miyata, M., Nemoto, N., Shikama, K., Kobayashi, F. & Hashimoto, T. Full-color-sorting metalenses for high-sensitivity image sensors. *Optica* **8**, 1596. ISSN: 2334-2536. <https://www.osapublishing.org/abstract.cfm?URI=optica-8-12-1596> (2021) (Dec. 2021).

5. Li, Z. *et al.* Inverse design enables large-scale high-performance meta-optics reshaping virtual reality. *Nature Communications* **13**, 2409. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-022-29973-3> (2022) (May 2022).
6. Zou, X. *et al.* Pixel-level Bayer-type colour router based on metasurfaces. *Nature Communications* **13**, 3288. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-022-31019-7> (2026) (June 2022).
7. Kim, C. *et al.* Freeform metasurface color router for deep submicron pixel image sensors. *Science Advances* **10**, eadn9000. <https://www.science.org/doi/full/10.1126/sciadv.adn9000> (2024) (May 2024).
8. Hsu, W.-L. *et al.* Simplest but Efficient Design of a Color Router Optimized by Genetic Algorithms. *ACS Photonics* **12**, 1402–1408. <https://doi.org/10.1021/acsp Photonics.4c01905> (2026) (Mar. 2025).
9. Yan, Z. *et al.* Pixel-Level Spectral Routers for RGB-IR Sensing. *ACS Photonics* **12**, 6374–6380. <https://doi.org/10.1021/acsp Photonics.5c01866> (2025) (Nov. 2025).
10. Miller, D. A. B. Communicating with waves between volumes: evaluating orthogonal spatial channels and limits on coupling strengths. *Applied Optics* **39**, 1681–1699. ISSN: 2155-3165. <https://opg.optica.org/ao/abstract.cfm?uri=ao-39-11-1681> (2026) (Apr. 2000).
11. Arbabi, A. & Faraon, A. Fundamental limits of ultrathin metasurfaces. *Scientific Reports* **7**, 43722. ISSN: 2045-2322. <https://www.nature.com/articles/srep43722> (2021) (Mar. 2017).
12. Catrysse, P. B., Zhao, N., Jin, W. & Fan, S. Subwavelength Bayer RGB color routers with perfect optical efficiency. *Nanophotonics* **11**. ISSN: 2192-8606, 2192-8614. <https://www.degruyter.com/document/doi/10.1515/nanoph-2022-0069/html> (2022) (Mar. 2022).
13. Catrysse, P. B. & Fan, S. Spectral routers for snapshot multispectral imaging. *Applied Physics Letters* **123**, 261105. ISSN: 0003-6951, 1077-3118. <https://pubs.aip.org/apl/article/123/26/261105/2931557/Spectral-routers-for-snapshot-multispectral> (2026) (Dec. 2023).
14. Lee, S. *et al.* Inverse design of color routers in CMOS image sensors: toward minimizing interpixel crosstalk. *Nanophotonics* **13**, 3895–3914. ISSN: 2192-8614. <https://www.degruyterbrill.com/document/doi/10.1515/nanoph-2024-0269/html> (2025) (Aug. 2024).
15. Chi, C., Fan, Z., Song, Q., Cai, W. & Duan, H. Multilayer topological metasurface-based color routers. *EN. Chinese Optics Letters* **24**, 010501. ISSN: 1671-7694. <https://www.researching.cn/articles/OJ2083715bc7f2a5f5> (2026) (Nov. 2025).

16. Roberts, G. *et al.* 3D-patterned Inverse-Designed Mid-Infrared Metaoptics. *Nature Communications* **14**, 2768. ISSN: 2041-1723. (2025) (May 2023).
17. Chen, W. *Additive Nanofabrication of 3D TiO₂ Metaoptics via Two-Photon Lithography* en. phd (California Institute of Technology, 2026). <https://resolver.caltech.edu/CaltechTHESIS:02072026-200423437> (2026).
18. Yoon, G. *et al.* Fabrication of three-dimensional suspended, interlayered and hierarchical nanostructures by accuracy-improved electron beam lithography overlay. en. *Scientific Reports* **7**, 6668. ISSN: 2045-2322. <https://www.nature.com/articles/s41598-017-06833-5> (2026) (July 2017).
19. Zheng, Y. *et al.* Bilayer optical metasurfaces with multiple broken symmetries for nonlinear wavelength generation. en. *npj Metamaterials* **2**, 6. ISSN: 3059-3727. <https://www.nature.com/articles/s44455-025-00016-3> (2026) (Feb. 2026).
20. Yang, W., Zhou, J., Tsai, D. P. & Xiao, S. Advanced manufacturing of dielectric meta-devices. en. *Photonics Insights* **3**, R04. ISSN: 2791-1748. <https://www.spiedigitallibrary.org/journals/photonics-insights/volume-3/issue-2/R04/Advanced-manufacturing-of-dielectric-meta-devices/10.3788/PI.2024.R04> (2026) (June 2024).
21. Baspinar, A. B., Pearson, P. & Faraon, A. Multilayer Q-BIC-Like Optical Filters with High Throughput Direct-Write Multilayer Lithography. en. *Advanced Optical Materials* **n/a**. eprint: <https://advanced.onlinelibrary.wiley.com/doi/pdf/10.1002/adom.7128e71282>. ISSN: 2195-1071. <https://onlinelibrary.wiley.com/doi/abs/10.1002/adom.71282> (2026).
22. Ci, S. *et al.* Photoluminescence properties of N-B Co-doped fluorescent SiC. *Semiconductor Science Technology* **35**. ADS Bibcode: 2020SeScT..35e5009C, 055009. ISSN: 0268-1242. <https://ui.adsabs.harvard.edu/abs/2020SeScT..35e5009C> (2026) (May 2020).
23. Kingma, D. P. & Ba, J. *Adam: A Method for Stochastic Optimization* Jan. 2017. <http://arxiv.org/abs/1412.6980> (2025).
24. Wang, F., Lazarov, B. S. & Sigmund, O. On projection methods, convergence and robust formulations in topology optimization. *Structural and Multidisciplinary Optimization* **43**, 767–784. ISSN: 1615-147X, 1615-1488. <http://link.springer.com/10.1007/s00158-010-0602-y> (2020) (June 2011).
25. Ballew, C., Roberts, G., Zheng, T. & Faraon, A. Constraining Continuous Topology Optimizations to Discrete Solutions for Photonic Applications. *ACS Photonics* **10**, 836–844. <https://doi.org/10.1021/acsp Photonics.2c00862> (2025) (Apr. 2023).
26. Guest, J. K. Topology optimization with multiple phase projection. *Computer Methods in Applied Mechanics and Engineering* **199**, 123–135. ISSN: 0045-7825. <https://www.sciencedirect.com/science/article/pii/S0045782509003302> (2026) (Dec. 2009).

27. Dawant, R., Seils, R., Ecoffey, S., Schmid, R. & Drouin, D. Hybrid cross correlation and line-scan alignment strategy for CMOS chips electron-beam lithography processing. *Journal of Vacuum Science & Technology B* **40**, 012601. ISSN: 2166-2746. <https://doi.org/10.1116/6.0001278> (2026) (Dec. 2021).
28. Choi, S., Word, M. J., Kumar, V. & Adesida, I. Comparative study of thermally cured and electron-beam-exposed hydrogen silsesquioxane resists. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena* **26**, 1654–1659. ISSN: 1071-1023. <https://doi.org/10.1116/1.2960565> (2026) (Sept. 2008).
29. Zhang, Y. *et al.* Exposure of Hydrogen Silsesquioxane in Electron Beam Lithography. *ACS Applied Materials & Interfaces* **17**, 38720–38730. ISSN: 1944-8244. <https://doi.org/10.1021/acsami.5c07013> (2026) (July 2025).
30. Camayd-Muñoz, P., Ballew, C., Roberts, G. & Faraon, A. Multifunctional Volumetric Meta-Optics for Color and Polarization Image Sensors. *Optica* **7**, 280. ISSN: 2334-2536. (2020) (Apr. 2020).
31. Wang, L., Liu, X., Peng, Y. & Xiao, J.-J. Nano-router metasurface inverse design by multi-body shape optimization. *Optics Letters* **50**, 4786–4789. ISSN: 1539-4794. <https://opg.optica.org/ol/abstract.cfm?uri=ol-50-15-4786> (2026) (Aug. 2025).
32. Lalau-Keraly, C. M., Bhargava, S., Miller, O. D. & Yablonovitch, E. Adjoint shape optimization applied to electromagnetic design. *Optics Express* **21**, 21693. ISSN: 1094-4087. <https://www.osapublishing.org/oe/abstract.cfm?uri=oe-21-18-21693> (2020) (Sept. 2013).
33. Molesky, S. *et al.* Inverse design in nanophotonics. *Nature Photonics* **12**, 659–670. ISSN: 1749-4893. <https://www.nature.com/articles/s41566-018-0246-9> (2026) (Nov. 2018).

CONCLUSION AND OUTLOOK: FUTURE DIRECTIONS FOR VOLUMETRIC METAOPTICS

The development of metaoptics has been driven by a simple but powerful idea: using subwavelength structures as the primary mechanism for controlling optical response. In planar metasurfaces, this idea has enabled thin optical elements that impose spatially varying phase, amplitude, and polarization responses with a form factor far smaller than conventional bulk optics. While single-layer metasurfaces can be extremely powerful, they are constrained by the fact that light interacts with the structure over a limited propagation distance. This imposes tradeoffs between diffraction efficiency, bandwidth, field of view, numerical aperture, dispersion engineering, and polarization control. Multilayer and volumetric architectures allow light to scatter, interfere, couple, and accumulate phase across a structured volume, relaxing these constraints and enabling transformations that are difficult or impossible for a single metasurface.

This thesis addresses the two key challenges, design and fabrication, in the experimental realization of volumetric metaoptics. It presents results that establish that these devices are physically realizable, and can be experimentally demonstrated across multiple wavelength regimes. In Chapter 3, we presented a modular inverse design framework built to optimize volumetric permittivity distributions. It can implement dynamic and versatile optimization schemes that are easily tuned by the user, and incorporates fabrication constraints to ensure the final results are compliant with fabrication capabilities. Crucially, the framework developed here is general enough that it does not need to be rebuilt for each new application, as is often currently the case in nanophotonic inverse design, greatly simplifying the workload of the designer.

Using this framework, Chapter 4 emphasized the multifunctional capabilities of volumetric metastructures by featuring dispersive polarizing beam-splitters, with independent control over the chromatic dispersion of each separated polarization channel. The beam-splitters were capable of large splitting angles up to 45° and broadband operation in both achromatic and hyper-dispersive regimes. Chapter 5 addressed the restrictions in experimental realization of volumetric metasurfaces

that are imposed by limited fabrication capabilities. In it we conducted a systematic hyperparameter exploration of volumetric Bayer-type color sorters under realistic fabrication constraints, characterizing performance as a function of layer count, pixel pitch, and material contrast. We also discussed several methods for adapting and integrating such metaoptic devices with on-sensor arrays. In general, many metaoptic devices must be adapted or re-designed in order to function seamlessly with external optophotonic systems.

Chapter 5 also showed that many realistic fabrication constraints cause a substantial reduction in performance. This issue can be addressed with improved computational schemes, but finding a way to relax some of those constraints is another viable method. In Chapter 6 we removed a common fabrication requirement of binary-material devices and used a novel fabrication platform to realize a pseudo-GRIN volumetric metaoptic device. Finally, in Chapter 7 we used the methods of Chapter 3 and the lessons of Chapter 5 to design a fabricable multilayer metaoptic device in the visible region. Such devices were previously impractical to fabricate, but a new Sb_2S_3 -based fabrication platform enabled the first realization of a ten-layer metaoptic color sorter, capable of spectral sorting in the visible regime. The results in this thesis show that the challenges of design and fabrication present in volumetric metaoptics can be overcome, giving us a path towards making them usable in application.

Large-Area Metasurface Design

Many of the most compelling applications of metaoptics require devices which span areas much larger than the micron scale. Imaging, display optics, augmented and virtual reality, automotive sensing, free-space communication, spectroscopy, photovoltaics, and wafer-level sensor integration all require devices with effective apertures on the scale of hundreds of microns, millimeters, centimeters, or larger^{1,2}. A centimeter-scale metasurface in the visible can contain hundreds of millions of subwavelength features. It is computationally prohibitive to use conventional electromagnetic solvers to simulate domains that span that many wavelengths. This restriction is usually bypassed by the use of local periodic approximations. However, for high-NA focusing, rapid phase variation, strong nonlocality or aperiodicity, or in the case of volumetric coupling, such approximations can fail. This design problem remains one of the central problems in the development of practical and applicable large-area metasurfaces¹.

Progress on this front must require a hierarchy of simulation and optimization methods. At the smallest scale, full-wave solvers will remain necessary for simulating local environments and the generation of training data. At intermediate scales, domain decomposition³, locally periodic models^{4,5}, and aperiodic Fourier methods⁶ can reduce the cost of evaluating large simulation domains. At the largest scales, surrogate models and reduced-order solvers may be necessary^{1,2}.

One promising direction is the use of such neural surrogate solvers^{7,8}. Given a simulation problem, they allow the prediction of its resulting electromagnetic field distributions while avoiding repeated full-wave simulations, which is especially valuable during optimization. Recent work on large-scale metasurface inverse design⁹ has shown that learned forward predictors can account for local coupling between neighboring meta-atoms. This successfully addresses one of the main problems inherent in simple library-based design, which assumes that neighboring structures are independent. In reality, coupling between structures and nonlocal responses often alters the actual electromagnetic response. Graph neural networks¹⁰, operator-learning methods⁸, and hybrid physics-informed models⁷ are all natural candidates. However, the minimum useful size of the required training data sets is often extremely large^{11,12}; and if the designer is not careful, the training model must be continually updated to reflect fabrication discrepancies¹³ in order to avoid error. This places constraints on the reusability of such models¹¹ to suit different design environments and requirements.

End-to-End Co-Optimization

Computational imaging represents another important potential direction. In conventional imaging and analysis systems, the optics and the computation are often treated as separate. The imaging system captures data, which is then transferred to and computationally processed by software, with no communication in between either system. Metaoptics can be used to instead create systems where the optical elements are optimized end-to-end together with the necessary reconstruction algorithms. In these co-optimizations, a metaoptic may deliberately encode depth, spectrum, polarization, or phase information into a measurement that is non-intuitive or not directly interpretable by humans, but can be reconstructed computationally and often with better fidelity¹⁴. This is especially valuable in sensing systems where the components are limited by size, weight, and power^{2,15,16}, where metaoptics can implement compact and multi-faceted optical transformations in the place of bulk optics. Volumetric metaoptics show particular promise in this regard, being able

to simultaneously discriminate intensity, wavelength or polarization^{17,18}, angle¹⁹, depth, or phase^{14,20} for such multimode sensing.

Metaoptics can also be used for analog optical computing^{21,22} and image processing²³, sometimes replacing the computational software entirely. Spatial differentiation²⁴, edge enhancement²⁵, and other linear operations can also be implemented optically by designing the angular or spatial transfer function of the metasurface. By performing operations at the speed of light, these devices reduce the amount of downstream electronic computation. However, they also emphasize the need for broadband and polarization-robust behavior if they are to be used outside laboratory conditions^{2,26,27}.

Considering the integration of metaoptics with image sensors as an example, the metaoptic can be placed directly on or near photodetectors, LEDs, and other devices. At the environment level, the metaoptic should be optimized to fit restrictions of the surrounding environment or electronics, such as pixel pitch, substrate (if any), passivation layers, the presence of microlenses, trench isolation or metal interconnects, or photodetector spectral responsivity. At the optical level, the device should be optimized for the intended illumination (e.g. bandwidth or divergence angle), field of view, and working distance. If any of these parameters are flexible, they may be co-optimized alongside the nanophotonic structure to produce higher overall performance^{16,17}.

Fabrication at Scale

Constraints imposed by fabrication restrictions have been a central theme in this thesis, but unintended discrepancies in fabrication are also especially damaging to the performance of volumetric metaoptics. The use of multiple layers requires accurate registration between each patterned layer, and the optical response can be adversely affected by lateral shifts and rotations, wafer bow, nonuniform shrinkage, tilted sidewall angles and accumulated surface roughness. Local errors due to stitching fields or the variation of exposure dose can produce unintended scattering across the area of the device. One potential way to resolve this is by including fabrication variability in the optimization workflow. Metaoptic devices can be made robust against such variability by optimizing geometrically eroded and dilated variants, that emulate systematic lithographic and etch bias errors, along with the original geometry. This has been demonstrated for topology-optimized metagratings²⁸, metasurface color routers²⁹, and other metaoptic devices^{30–32}.

To produce metaoptics at scale, these problems must be solved, but using foundry processes. Techniques such as electron-beam lithography and focused ion beam milling are poorly suited for large-volume or large-area manufacturing^{1,33}. Future metaoptics will need to rely increasingly on scalable processes such as deep ultra-violet lithography³⁴, nanoimprint lithography³³, roll-to-roll nanoimprinting³⁵, step-and-repeat imprinting³⁶, and transfer printing³⁷.

Active, Tunable, and Programmable Volumetric Metaoptics

Most demonstrated metaoptics are passive and static, having a fixed optical response once fabricated. This is appropriate for many applications, but it limits the ability of a device to adapt to changing wavelength, field of view, focus distance, illumination condition, or computational task, to name just a few. A major long-term direction is therefore the development of active, tunable, and programmable volumetric metaoptics.

Tunable devices can be implemented^{38,39} through mechanical motion⁴⁰, phase-change materials⁴¹, liquid crystals⁴², electro-optic materials⁴³, MEMS actuation⁴⁴, thermo-optic tuning⁴⁵, carrier injection⁴⁶, or strain⁴⁰. In a volumetric architecture, tunability could be implemented not only by perturbing the response of a single layer, but by changing the relative alignment, spacing, orientation, or material state of multiple scattering planes, bearing in mind that throughout this, the volumetric metaoptic must still be treated as a single scattering structure in simulation and design. This creates opportunities for varifocal lenses, switchable color routers, reconfigurable polarization optics, adaptive wavefront correctors, and programmable optical encoders.

Programmable volumetric metaoptics represent an even broader and more exciting possibility. Diffractive optical computing has already been used⁴⁷ to demonstrate neural networks and learning tasks like image recognition over multiple wavelength regimes⁴⁸. If each layer becomes reconfigurable, similarly to spatial light modulators in multi-plane light conversion⁴⁹, such devices could perform adaptive imaging, dynamic beam forming, compressive sensing, or other optical neural-network tasks⁵⁰. So far, this has been demonstrated to a limited extent through physical rearrangement by hand^{51,52}, and through phase-change materials in simulation⁵³, for the near-infrared, microwave, and telecom wavelength regimes.

Conclusion

Not all of the above challenges and considerations are specific to volumetric metaoptics, and they appear more broadly across the field of nanophotonic inverse design. However, volumetric metasurfaces stand out in their ability to perform complex and multifunctional optical functions in the same device with high efficiency. They are a promising platform for compact imaging, spectroscopy, sensing, display, communication, and optical information processing, and will enable new optical components and capabilities that transform many of the systems and technologies we all use today.

References

1. Choi, M. *et al.* Realization of high-performance optical metasurfaces over a large area: a review from a design perspective. *npj Nanophotonics* **1**, 31. ISSN: 2948-216X. <https://www.nature.com/articles/s44310-024-00029-2> (2026) (Aug. 2024).
2. Kuznetsov, A. I. *et al.* Roadmap for Optical Metasurfaces. *ACS Photonics* **11**, 816–865. (2026) (Mar. 2024).
3. Skarda, J. *et al.* Low-overhead distribution strategy for simulation and optimization of large-area metasurfaces. *npj Computational Materials* **8**, 1–6. ISSN: 2057-3960. <https://www.nature.com/articles/s41524-022-00774-y> (2023) (Apr. 2022).
4. Pestourie, R. *et al.* Inverse design of large-area metasurfaces. EN. *Optics Express* **26**, 33732–33747. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-26-26-33732> (2024) (Dec. 2018).
5. Wang, S., Zhang, S., Song, N. & Xue, D. Large-Scale Metasurface Simulation Using Local-Segmented Approach. *Materials* **18**, 649. ISSN: 1996-1944. <https://www.mdpi.com/1996-1944/18/3/649> (2026) (Jan. 2025).
6. Phan, T. *et al.* High-efficiency, large-area, topology-optimized metasurfaces. *Light: Science & Applications* **8**, 48. ISSN: 2047-7538. <https://www.nature.com/articles/s41377-019-0159-5> (2024) (May 2019).
7. Chen, M. *et al.* High Speed Simulation and Freeform Optimization of Nanophotonic Devices with Physics-Augmented Deep Learning. *ACS Photonics* **9**, 3110–3123. <https://doi.org/10.1021/acsphotonics.2c00876> (2026) (Sept. 2022).
8. Augenstein, Y., Repän, T. & Rockstuhl, C. Neural Operator-Based Surrogate Solver for Free-Form Electromagnetic Inverse Design. *ACS Photonics* **10**, 1547–1557. <https://doi.org/10.1021/acsphotonics.3c00156> (2026) (May 2023).

9. Xu, B. *et al.* Deep-Learning-Enabled Inverse Design of Large-Scale Metasurfaces With Full-Wave Accuracy (Laser Photonics Rev. 20(5)/2026). *Laser & Photonics Reviews* **20**, e71006. ISSN: 1863-8899. <https://onlinelibrary.wiley.com/doi/abs/10.1002/lpor.71006> (2026) (2026).
10. Khoram, E., Wu, Z., Qu, Y., Zhou, M. & Yu, Z. Graph Neural Networks for Metasurface Modeling. *ACS Photonics* **10**, 892–899. <https://doi.org/10.1021/acsp Photonics.2c01019> (2026) (Apr. 2023).
11. Li, Z., Pestourie, R., Lin, Z., Johnson, S. G. & Capasso, F. Empowering Metasurfaces with Inverse Design: Principles and Applications. *ACS Photonics* **9**, 2178–2192. <https://doi.org/10.1021/acsp Photonics.1c01850> (2023) (July 2022).
12. Fan, Z. *et al.* Transfer-Learning-Assisted Inverse Metasurface Design for 30% Data Savings. *Physical Review Applied* **18**, 024022. <https://link.aps.org/doi/10.1103/PhysRevApplied.18.024022> (2026) (Aug. 2022).
13. Peng, R., Ren, S., Malof, J. & Padilla, W. J. Transfer learning for metamaterial design and simulation. *Nanophotonics* **13**, 2323–2334. ISSN: 2192-8614. <https://onlinelibrary.wiley.com/doi/abs/10.1515/nanoph-2023-0691> (2026) (2024).
14. Lin, Z. *et al.* End-to-end nanophotonic inverse design for imaging and polarimetry. *Nanophotonics* **10**, 1177–1187. ISSN: 2192-8606 (2020).
15. Colburn, S., Zhan, A. & Majumdar, A. Metasurface Optics for Full-Color Computational Imaging. *Science Advances* **4**, eaar2114. (2023) (Feb. 2018).
16. Brongersma, M. L. *et al.* The second optical metasurface revolution: moving from science to technology. *Nature Reviews Electrical Engineering* **2**, 125–143. ISSN: 2948-1201. <https://www.nature.com/articles/s44287-024-00136-4> (2026) (Feb. 2025).
17. Camayd-Muñoz, P., Ballew, C., Roberts, G. & Faraon, A. Multifunctional Volumetric Meta-Optics for Color and Polarization Image Sensors. *Optica* **7**, 280. ISSN: 2334-2536. (2020) (Apr. 2020).
18. Roberts, G. *et al.* 3D-patterned Inverse-Designed Mid-Infrared Metaoptics. *Nature Communications* **14**, 2768. ISSN: 2041-1723. (2025) (May 2023).
19. Ballew, C., Roberts, G. & Faraon, A. Multi-dimensional wavefront sensing using volumetric meta-optics. *Optics Express* **31**, 28658–28669. ISSN: 1094-4087. <https://opg.optica.org/oe/abstract.cfm?uri=oe-31-18-28658> (2025) (Aug. 2023).
20. Liu, J., Wang, H., Li, Y., Tian, L. & Paiella, R. Asymmetric metasurface photodetectors for single-shot quantitative phase imaging. *Nanophotonics* **12**, 3519–3528. ISSN: 2192-8614. <https://onlinelibrary.wiley.com/doi/abs/10.1515/nanoph-2023-0354> (2026) (2023).

21. Silva, A. *et al.* Performing Mathematical Operations with Metamaterials. *Science* **343**, 160–163. <https://www.science.org/doi/10.1126/science.1242818> (2026) (Jan. 2014).
22. Zangeneh-Nejad, F., Sounas, D. L., Alù, A. & Fleury, R. Analogue computing with metamaterials. *Nature Reviews Materials* **6**, 207–225. ISSN: 2058-8437. <https://www.nature.com/articles/s41578-020-00243-2> (2022) (Mar. 2021).
23. Hu, X. *et al.* Metasurface-based computational imaging: a review. *Advanced Photonics* **6**, 014002. ISSN: 2577-5421. <https://www.spiedigitallibrary.org/journals/advanced-photonics/volume-6/issue-01/014002/Metasurface-based-computational-imaging-a-review/10.1117/1.AP.6.1.014002> (2026) (Feb. 2024).
24. Wan, L. *et al.* Optical analog computing of spatial differentiation and edge detection with dielectric metasurfaces. *Optics Letters* **45**, 2070–2073. ISSN: 1539-4794. <https://opg.optica.org/ol/abstract.cfm?uri=ol-45-7-2070> (2026) (Apr. 2020).
25. Pearson, P., Roberts, G. & Faraon, A. Inverse-designed metasurfaces for multifunctional spatial frequency filtering. *Optica* **12**, 1090–1099. ISSN: 2334-2536. <https://opg.optica.org/optica/abstract.cfm?uri=optica-12-7-1090> (2026) (July 2025).
26. Zhou, J. *et al.* Metasurface Enabled Quantum Edge Detection. *Science Advances* **6**, eabc4385. (2026) (Dec. 2020).
27. Ou, K. *et al.* Advances in Meta-Optics and Metasurfaces: Fundamentals and Applications. *Nanomaterials* **13**, 1235. ISSN: 2079-4991. <https://www.mdpi.com/2079-4991/13/7/1235> (2026) (Jan. 2023).
28. Wang, E. W., Sell, D., Phan, T. & Fan, J. A. Robust design of topology-optimized metasurfaces. *Optical Materials Express* **9**, 469–482. ISSN: 2159-3930. <https://opg.optica.org/ome/abstract.cfm?uri=ome-9-2-469> (2026) (Feb. 2019).
29. Kim, C. *et al.* Freeform metasurface color router for deep submicron pixel image sensors. *Science Advances* **10**, eadn9000. <https://www.science.org/doi/full/10.1126/sciadv.adn9000> (2024) (May 2024).
30. Chen, M., Jiang, J. & Fan, J. A. Design Space Reparameterization Enforces Hard Geometric Constraints in Inverse-Designed Nanophotonic Devices. *ACS Photonics* **7**, 3141–3151. <https://doi.org/10.1021/acsp Photonics.0c01202> (2025) (Nov. 2020).
31. Wen, F., Jiang, J. & Fan, J. A. Robust Freeform Metasurface Design Based on Progressively Growing Generative Networks. *ACS Photonics* **7**, 2098–2104. <https://doi.org/10.1021/acsp Photonics.0c00539> (2026) (Aug. 2020).

32. Jenkins, R. P., Campbell, S. D. & Werner, D. H. Establishing exhaustive meta-surface robustness against fabrication uncertainties through deep learning. *Nanophotonics* **10**, 4497–4509. ISSN: 2192-8606. <https://www.scopus.com/pages/publications/85118988955> (2026) (Dec. 2021).
33. Oh, D. K. *et al.* Nanoimprint lithography for high-throughput fabrication of metasurfaces. *Frontiers of Optoelectronics* **14**, 229–251. ISSN: 2095-2767. <https://doi.org/10.1007/s12200-021-1121-8> (2026) (June 2021).
34. Park, J.-S. *et al.* All-Glass 100 mm Diameter Visible Metalens for Imaging the Cosmos. *ACS Nano* **18**, 3187–3198. ISSN: 1936-0851. <https://doi.org/10.1021/acsnano.3c09462> (2026) (Jan. 2024).
35. Hoang, T. *et al.* 300-unit-per-second roll-to-roll manufacturing of visible metalenses. *Nature* **652**, 1188–1194. ISSN: 1476-4687. <https://www.nature.com/articles/s41586-026-10369-y> (2026) (Apr. 2026).
36. Choi, S., Zuo, J., Das, N., Yao, Y. & Wang, C. Scalable Nanoimprint Manufacturing of Functional Multilayer Metasurface Devices. *Advanced Functional Materials* **34**, 2404852. ISSN: 1616-3028. <https://onlinelibrary.wiley.com/doi/abs/10.1002/adfm.202404852> (2026) (2024).
37. Yeh, T.-W. *et al.* Nanotransfer Printed Dual-Layer Metasurfaces for Infrared Cut-off Applications. *ACS Applied Nano Materials* **7**, 25593–25602. <https://doi.org/10.1021/acsanm.4c04734> (2026) (Nov. 2024).
38. Shaltout, A. M., Shalaev, V. M. & Brongersma, M. L. Spatiotemporal light control with active metasurfaces. *Science* **364**, eaat3100. <https://www.science.org/doi/10.1126/science.aat3100> (2026) (May 2019).
39. Santonocito, A., Patrizi, B. & Toci, G. Recent Advances in Tunable Metasurfaces and Their Application in Optics. *Nanomaterials* **13**, 1633. ISSN: 2079-4991. <https://www.mdpi.com/2079-4991/13/10/1633> (2026) (Jan. 2023).
40. Pan, X. *et al.* Recent progress in mechanically reconfigurable metasurfaces. *APL Photonics* **10**, 080901. ISSN: 2378-0967. <https://doi.org/10.1063/5.0274326> (2026) (Aug. 2025).
41. Peimyoo, N. *et al.* Electrical tuning of optically active interlayer excitons in bilayer MoS₂. *Nature Nanotechnology* **16**, 888–893. ISSN: 1748-3395. <https://www.nature.com/articles/s41565-021-00916-1> (2026) (Aug. 2021).
42. Kim, J. *et al.* Tunable metasurfaces towards versatile metalenses and metaholograms: a review. *Advanced Photonics* **4**, 024001. ISSN: 2577-5421. <https://www.spiedigitallibrary.org/journals/advanced-photonics/volume-4/issue-02/024001/Tunable-metasurfaces-towards-versatile-metalenses-and-metaholograms-a-review/10.1117/1.AP.4.2.024001> (2026) (Mar. 2022).

43. Shirmanesh, G. K., Sokhoyan, R., Wu, P. C. & Atwater, H. A. Electro-optically Tunable Multifunctional Metasurfaces. *ACS Nano* **14**, 6912–6920. ISSN: 1936-0851, 1936-086X. <https://pubs.acs.org/doi/10.1021/acsnano.0c01269> (2021) (June 2020).
44. Arbabi, E. *et al.* MEMS-tunable dielectric metasurface lens. *Nature Communications* **9**, 812. ISSN: 2041-1723. <https://www.nature.com/articles/s41467-018-03155-6> (2026) (Feb. 2018).
45. Zangeneh Kamali, K. *et al.* Electrically programmable solid-state metasurfaces via flash localised heating. *Light: Science & Applications* **12**, 40. ISSN: 2047-7538. <https://www.nature.com/articles/s41377-023-01078-6> (2026) (Feb. 2023).
46. Iyer, P. P., Pendharkar, M., Palmstrøm, C. J. & Schuller, J. A. III–V Heterojunction Platform for Electrically Reconfigurable Dielectric Metasurfaces. *ACS Photonics* **6**, 1345–1350. <https://doi.org/10.1021/acsp Photonics.9b00178> (2026) (June 2019).
47. Lin, X. *et al.* All-optical machine learning using diffractive deep neural networks. *Science* **361**, 1004–1008. <https://www.science.org/doi/10.1126/science.aat8084> (2022) (Sept. 2018).
48. Mengu, D., Luo, Y., Rivenson, Y. & Ozcan, A. Analysis of Diffractive Optical Neural Networks and Their Integration With Electronic Neural Networks. *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1–14. ISSN: 1558-4542 (Jan. 2020).
49. Kupianskyi, H., Horsley, S. A. R. & Phillips, D. B. High-dimensional spatial mode sorting and optical circuit design using multi-plane light conversion. *APL Photonics* **8**, 026101. ISSN: 2378-0967. <https://doi.org/10.1063/5.0128431> (2026) (Feb. 2023).
50. Abou-Hamdan, L. *et al.* Programmable metasurfaces for future photonic artificial intelligence. *Nature Reviews Physics* **7**, 331–347. ISSN: 2522-5820. <https://www.nature.com/articles/s42254-025-00831-7> (2026) (June 2025).
51. He, C. *et al.* Pluggable multitask diffractive neural networks based on cascaded metasurfaces. *Opto-Electronic Advances*, 2300005–9. ISSN: 2096-4579. <http://www.chinaoceanengin.cn/en/article/doi/10.29026/oea.2024.2300005> (2023) (July 2023).
52. Yu, Z. *et al.* Multifunctional movable-type coding metasurface enabling reconfigurable diffractive neural networks. *Light: Science & Applications* **15**, 127. ISSN: 2047-7538. <https://www.nature.com/articles/s41377-026-02216-6> (2026) (Feb. 2026).

53. Zarei, S. Programmable diffractive deep neural networks enabled by integrated rewritable metasurfaces. *Scientific Reports* **15**, 35624. ISSN: 2045-2322. <https://www.nature.com/articles/s41598-025-19638-8> (2026) (Oct. 2025).

INDEX

F

figure of merit, 14

M

merit function, 14

O

objective function, 14

P

Pareto front, 22

