

# Chapter 6

## Objective-First Nanophotonic Design

Jesse Lu and Jelena Vuckovic

**Abstract** We introduce an “objective-first” strategy for designing nanophotonic devices, and we demonstrate the design of nanophotonic coupler, cloak, and mimic devices. Simply put, our objective-first method works by prioritizing the performance of the device even above satisfying Maxwell’s equations. We show how this is accomplished starting from Maxwell’s equations, applying numerical discretization, and then solving not only for the field variables but the structure variables as well. We then demonstrate the ability to quickly produce designs for both traditional devices such as waveguide couplers, as well as more exotic devices such as optical cloaks and mimics. Finally, we point the reader to future improvements and extensions of our method.

### 6.1 Introduction

Our initial foray into design methods for nanophotonic devices began with a very simple and naive question: Could we make an inverse solver which, when given the electromagnetic fields we desire, returns the nanophotonic structure that will produce them [1]? In other words, since we already know how to solve for  $E$  and  $H$  in Maxwell’s equations, why can’t we solve for  $\epsilon$  or even  $\mu$  instead? Not surprisingly, it did not take us long to find that such a simple strategy would inevitably run into many problems.

Over the subsequent years, we were able to come up with a better solution, which we call an “objective-first” strategy for nanophotonic design, and which we present in this chapter. Although it is more advanced than our original idea, objective-first design still carries the same fundamental concept, which is to specify the electromagnetic fields, and then to solve for a structure to produce them.

One of the strengths of this approach is that it recasts the optimization problem into separably convex problems whose global optima can be reliably determined. Additionally, we show that our formulation is general enough to encompass the design of all linear nanophotonic structures and is not limited to the design of optical

---

J. Lu (✉)

Department of Electrical Engineering, Stanford University, Palo Alto, CA 94305, USA  
e-mail: [jesselu@stanford.edu](mailto:jesselu@stanford.edu)

metamaterials. While this approach typically produces designs with very high efficiencies (typically 99 %) our solutions do feature continuously varying values of the dielectric constant. Future approaches for dealing with this constraint are discussed in Sect. 6.7.3.

In this chapter, we present the simple theoretical underpinnings of objective-first design in Sects. 6.2 and 6.3—namely the numerical discretization of the electromagnetic wave equation, and the formulation of the objective-first design problem. Then, in Sects. 6.4–6.6, we show examples of the method in action in designing nanophotonic devices in three broad categories: waveguide couplers, optical cloaks, and optical mimics; the source code for which is available online [2]. In Sect. 6.7, we conclude by commenting on possible extensions of our method.

Lastly, we note that in contrast to many of the other chapters in this book, the designs proposed in this chapter are all permittivity and permeability profiles in one and two-dimensions. The additional task of choosing specific material sets to fabricate actual devices is left as an additional step.

## 6.2 The Electromagnetic Wave Equation

In this section, we outline the wave equation that is central to the application of our method, with the end-result being to show that it is separably linear (bi-linear) in the field and structure variables. We do this by first formulating this wave equation in the language of physics, and then discretizing it in order to achieve numerical solutions. We then show how one can not only obtain the solution for the fields, but also obtain the solution for the structure using simple, standard numerical tools.

### 6.2.1 Physics Formulation

First, let's derive our wave equation, starting with the differential form of Maxwell's equations,

$$\nabla \times E = -\mu_0 \frac{\partial H}{\partial t}, \quad (6.1)$$

$$\nabla \times H = J + \varepsilon \frac{\partial E}{\partial t}, \quad (6.2)$$

where  $E$ ,  $H$ , and  $J$  are the electric, magnetic and electric current vector fields, respectively,  $\varepsilon$  is the permittivity and  $\mu_0$  is the permeability, which we assume to be that of vacuum everywhere.

Assuming the time dependence  $\exp(-i\omega t)$ , where  $\omega$  is the angular frequency, these become

$$\nabla \times E = -i\mu_0\omega H, \quad (6.3)$$

$$\nabla \times H = J + i\varepsilon\omega E, \quad (6.4)$$

which we can combine to form our (time-harmonic) wave equation,

$$\nabla \times \varepsilon^{-1} \nabla \times H - \mu_0 \omega^2 H = \nabla \times \varepsilon^{-1} J. \quad (6.5)$$

In this chapter, we are only going to consider the two-dimensional form of this equation, and specifically the two-dimensional transverse electric (TE) mode [3]. In this case (6.5) is simplified because only the  $z$ -component of  $H$  is nonzero. Nevertheless, a single equation (6.5), represents all the physics which we take into account in this chapter.

## 6.2.2 Numerical Formulation

On top of the analytical formulation of the wave equation (6.5), we will now add a numerical, or discretized, formulation. This will be needed in order to solve for arbitrary structures for which there are not analytical solutions.

The salient step in order to do so is to use the Yee grid [4], which allows us to easily define the curl ( $\nabla \times$ ) operators in (6.5). Since both the individual curl operators and the equation as a whole is linear in  $H$ , we can reformulate (6.5) with a change of variables, as

$$A(p)x = b(p), \quad (6.6)$$

where  $H \rightarrow x$ ,  $\varepsilon^{-1} \rightarrow p$ ; and where

$$A(p) = \nabla \times \varepsilon^{-1} \nabla \times - \mu_0 \omega^2 \quad (6.7)$$

and

$$b(p) = \nabla \times \varepsilon^{-1} J. \quad (6.8)$$

Note that our use of  $A(p)$  and  $b(p)$  instead of  $A$  and  $b$  simply serves to clarify the dependence of both  $A$  and  $b$  on  $p$ .

Apart from using the Yee grid, which at some length scale requires that our designs conform to a rectangular grid, the only other salient implementation detail is the use of stretched-coordinate perfectly matched layers [5] where necessary, in order to prevent unwanted reflections at the boundaries of the simulation domain. The effect of such layers is to modify the curl operators, although their linear property is still maintained.

## 6.2.3 Solving for $H$

With our numerical formulation, we can now solve for the  $H$ -field (the  $E$ -field can be computed from the  $H$ -field using (6.4)) by applying general linear algebra solvers to (6.6). Recall that since we have chosen a time-harmonic formulation, solving for

$x$  in (6.6) is actually performing what is simply known as a time-harmonic or a finite-difference frequency-domain (FDFD) simulation [6]. Furthermore, since we have limited ourselves to the two-dimensional case, (6.6) is easily solved using the standard sparse solver included in Matlab on a single desktop computer.

We call the routine that solves for  $x$  in (6.6) given  $p$  a field-solver, or a simulator.

### 6.2.4 Solving for $\epsilon^{-1}$

After having built a field-solver or simulator (which finds  $x$  given  $p$ ) for our wave equation, the next step is to build a structure-solver for it. In other words, we need to be able to solve for  $p$  given  $x$ .

To do so, we return to (6.5) and remark that  $\epsilon^{-1}(\nabla \times H) = (\nabla \times H)\epsilon^{-1}$  and  $\epsilon^{-1}J = J\epsilon^{-1}$  since scalar multiplication is commutative. This allows us to rearrange (6.5) as

$$\nabla \times (\nabla \times H)\epsilon^{-1} - \nabla \times J\epsilon^{-1} = \mu_0\omega^2 H \quad (6.9)$$

which we now write as

$$B(x)p = d(x), \quad (6.10)$$

where

$$B(x) = \nabla \times (\nabla \times H) - \nabla \times J \quad (6.11)$$

and

$$d(x) = \mu_0\omega^2 H. \quad (6.12)$$

With this extremely simple trick, we have shown that we can seemingly solve for  $p$  given  $x$  with approximately the same ease as solving for  $x$  given  $p$ ! We see this because the dimensions and complexity of  $B(x)$  are basically equivalent to that of  $A(p)$ , and this implies that the same simple tools used in our field-solver should be applicable to solving (6.10). This is indeed what we find, although the later addition of constraints on  $p$  will require the use of more powerful (but just as dependable) numerical tools.

### 6.2.5 Bi-linearity of the Wave Equation

Although additional mathematical machinery must still be added to obtain a useful design tool, we have shown so far that the wave equation is separately linear in  $x$  and  $p$  (i.e., bilinear). Namely,

$$A(p)x - b(p) = B(x)p - d(x). \quad (6.13)$$

In other words, fixing  $p$  makes solving the wave equation for  $x$  a linear problem, and vice versa. Note that the joint problem, where both  $x$  and  $p$  are allowed to vary, is not linear.

The bi-linearity of the wave equation is *absolutely fundamental* in our objective-first strategy because it relies on the fact that, although simultaneously solving for  $x$  and  $p$  is very difficult, we already know how to solve linear systems ( $x$  and  $p$  separately) well. In fact, it is this very property that forms the natural division of labor which our objective-first method exploits.

### 6.3 The Objective-First Design Problem

We now describe the remaining machinery used in the objective-first method, in addition to the field-solver and the structure-solver, as previously outlined. Specifically, we introduce the idea of a design objective and a physics residual, and we reference the mathematical notion of convexity in order to motivate the need to divide the objective-first problem into two separately convex sub-problems.

#### 6.3.1 Design Objectives

Our design objective or objective function,  $f(x)$ , is simply defined as a function we wish to be minimal for the design to be produced. For instance, in the design of a device which must transmit efficiently into a particular mode, we could choose  $f(x)$  to be the negative power flow into that mode. Or, if the device was to be a low-loss resonator, we could choose  $f(x)$  to be the amount of power leaking out of the device. In general, there are multiple choices of  $f(x)$  which can be used to describe the same objective. For example,  $f(x)$  for a transmissive device may not only be the negative power transmitted into the desired output mode, but it could also be the amount of power lost to other modes, or even the error in the field values at the output port relative to the field values needed for perfect transmission. These design objectives are equivalent in the sense that, if minimized, all would produce structures with good performance. At the same time, we must consider that the computational cost and complexity of using one  $f(x)$  over another may indeed vary greatly.

#### 6.3.2 Convexity

Before formulating the design problem, we would like to add a note regarding the complexity of various optimization problems. Specifically, we want to introduce the notion of *convexity* [7] and to note the difference between problems that are convex

and those which are not. The difference is simply this: Convex problems have a single optimum point (only one local optimum, which is therefore the global optimum) which we can reliably find using existing numerical software, whereas non-convex problems typically have multiple optima and are thus much more difficult to reliably solve.

That a convex problem can be reliably solved, in this case, means that regardless of the starting guess, convex optimization software will always arrive at the globally optimal solution and will be able to numerically prove global optimality as well. Thus, the advantage in formulating a design problem in terms of convex optimization problems is to eliminate both the need to circumvent local optima and any notion of randomness. On a practical level, there exist mature convex optimization software packages among which is CVX, a convex optimization package written for Matlab [8], which we use for the examples in this chapter.

### 6.3.3 Typical Design Formulation

We now examine the typical, and most straightforward formulation of the design problem, in order to relate and contrast it to the objective-first formulation. The design problem for a physical structure is typically formulated as

$$\begin{aligned} & \underset{x,p}{\text{minimize}} \quad f(x) \\ & \text{subject to} \quad A(p)x - b(p) = 0, \end{aligned} \tag{6.14}$$

which states that we would like to vary  $x$  and  $p$  simultaneously in order to decrease  $f(x)$  while always satisfying physics (e.g., the electromagnetic wave equation).

Since solving (6.14) is quite difficult in the general sense (simultaneously varying  $x$  and  $p$  makes the problem non-convex), traditional optimization approaches, such as those described in previous chapters, have relied on either a brute-force parameter search, or a gradient-descent method utilizing first-order derivatives. In the gradient-descent case, solving (6.14) results in the well-known adjoint optimization method [9].

### 6.3.4 Objective-First Design Formulation

In contrast with the typical formulation, the objective-first formulation simply switches the roles of the wave equation and the design objective with one another:

$$\underset{x,p}{\text{minimize}} \quad \|A(p)x - b(p)\|^2 \tag{6.15}$$

$$\text{subject to} \quad f(x) = f_{\text{ideal}}. \tag{6.16}$$

Although such a switch may seem trivial, and even silly at first, we show that it fundamentally changes the nature of the design problem and actually gives us advantages in our efforts at finding a solution.

This first fundamental change, as seen from (6.15), is that we allow for a nonzero residual in the electromagnetic wave equation. This literally means that we allow for *non-physical*  $x$  and  $p$ , since  $A(p)x - b(p) \neq 0$  is permissible. And since  $A(p)x - b(p)$  can now be a nonzero entity, we choose to call it the *physics residual*. The second fundamental change is that we always force the device to exhibit ideal performance, as seen from (6.16). This, of course, ties in very closely with (6.15) since ideal performance is usually not obtainable unless one allows for some measure of error in the underlying physics (nonzero physics residual). As such, our strategy will be to iteratively vary  $x$  and then  $p$  in order to decrease the physics residual (6.15) to zero, while always maintaining ideal performance.

The primary advantage of the objective-first formulation is that, although the full problem is still non-convex, it allows us to form two convex sub-problems, as outlined below. In contrast to an adjoint method, here we can still access information regarding second-order derivatives, which decreases the time it takes to obtain a solution. An additional advantage of this approach is that our insistence that ideal performance be always attained provides a mechanism which can potentially “override” local optima in the optimization process.

To this end, we have found that such a strategy results in surprisingly non-intuitive devices which exhibit highly efficient performance, even when the starting point of the design problem is completely non-functional. Furthermore, we have found this to be true even when the physics residual fails to be completely removed.

In practice, we add an additional constraint to the original formulation, [10] which is to set hard-limits on the allowable values of  $p$ , namely  $p_0 \leq p \leq p_1$ . This is actually a relaxation of the ideal constraint, which would be to allow  $p$  to only have discrete values,  $p \in [p_0, p_1]$ , but such a constraint would essentially force us to only perform brute force trial-and-error.

Our objective-first formulation is thus:

$$\begin{aligned} & \underset{x,p}{\text{minimize}} \quad \|A(p)x - b(p)\|^2 \\ & \text{subject to} \quad f(x) = f_{\text{ideal}} \\ & \quad \quad \quad p_0 \leq p \leq p_1, \end{aligned} \tag{6.17}$$

which is still non-convex, but can be broken down into two convex sub-problems, the motivation being that each of these will be able to be easily and reliably solved.

### 6.3.5 Field Sub-problem

The first of these is the field sub-problem, which simply involves fixing  $p$  and independently optimizing  $x$ ,

$$\begin{aligned} & \underset{x}{\text{minimize}} \quad \|A(p)x - b(p)\|^2 \\ & \text{subject to} \quad f(x) = f_{\text{ideal}}. \end{aligned} \tag{6.18}$$

This problem is convex, and actually quadratic, which means that it can even be solved using standard numerical tools, in the same way as a simple least-squares problem.

The field sub-problem can be thought of as an update to  $x$  (H-field) where we try to “fit” the electromagnetic fields to the structure ( $p$ ). Of course, if it were not for the hard-constraint on the design objective, the field sub-problem would be able to perfectly fit  $x$  to  $p$ .

### 6.3.6 Structure Sub-problem

The second sub-problem is formulated by fixing  $x$  and independently optimizing  $p$ . At the same time, we use the bi-linearity property of the physics residual from (6.13) to rewrite the problem in a way that makes its convexity explicit:

$$\begin{aligned} & \underset{p}{\text{minimize}} \quad \|B(x)p - d(x)\|^2 \\ & \text{subject to} \quad p_0 \leq p \leq p_1. \end{aligned} \tag{6.19}$$

The structure sub-problem is also convex, but not quadratic because of the inequality constraints on  $p$ . However, use of the CVX package still allows us to obtain results quickly and reliably.

Note that in an analogous fashion to the field sub-problem, the structure sub-problem attempts to fit  $p$  to  $x$ , and is prevented from perfectly doing so by its own constraint. Because neither sub-problem is capable of completely reducing the physics residual to zero, they must be used in an iterative manner in order to gradually decrease the physics residual. To this end, we employ the alternating directions optimization method.

### 6.3.7 Alternating Directions

We use a simple alternating directions scheme to piece together (6.18) and (6.19), which is to say that we simply continually alternate between solving each equation



until we reach some stopping point, normally measured by how much the physics residual has decreased.

Loop:

$$\begin{aligned}
 & \underset{x}{\text{minimize}} \quad \|A(p)x - b(p)\|^2 \\
 & \text{subject to} \quad f(x) = f_{\text{ideal}}; \\
 & \underset{p}{\text{minimize}} \quad \|B(x)p - d(x)\|^2 \\
 & \text{subject to} \quad p_0 \leq p \leq p_1.
 \end{aligned} \tag{6.20}$$

The alternating directions scheme is extremely simple and does not require additional processing of  $x$  or  $p$  outside of the two sub-problems, nor does it require the use of auxiliary variables.

The advantage of the alternating directions method is that the physics residual is guaranteed to monotonically decrease with every iteration, which is useful in that no safeguards are needed to protect against “rogue” steps in the optimization procedure. Note that this robustness stems from the fact that, among other things, each sub-problem does not rely on previous values of the variable which is being optimized, but only on the variable which is held constant.

The disadvantage of such a simple scheme is that the convergence is quite slow, although we have found it to be sufficient in our cases. Related methods, such as the Alternating Directions Method of Multipliers [11], exhibit far better convergence.

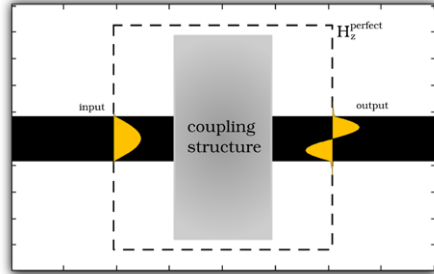
## 6.4 Waveguide Coupler Design

We first apply the objective-first formulation with the alternating directions algorithm to the design of nanophotonic waveguide couplers in two dimensions, where our goal is to couple light from a single input waveguide mode to a single output waveguide mode with as close to unity efficiency as possible. We would also like to allow the user to choose arbitrary input and output waveguides, as well as to select arbitrary modes within those waveguides (as opposed to allowing only the fundamental mode, for example).

This problem is very general and, in essence, encompasses the design of all linear nanophotonic components because the function or performance of all such components is simply to convert a defined set of input modes into a defined set of output modes. Such a broad, general problem is ideally suited for an objective-first strategy since no approximations or simplifications of the electromagnetic fields are required; we only make the simplification of working in two dimensions (transverse magnetic mode) and dealing only with a single input and output mode.

Since the electromagnetic wave equation is scale-invariant (e.g., double the length scale and half the frequency and you obtain the same equation), we state all dimensions terms of the vacuum wavelength itself. Therefore, our solutions are applicable to regions of the electromagnetic spectrum where the dielectric constants

**Fig. 6.1** Formulation of the design objective



used are achievable. In most cases, values between 1 and 12.25 are chosen because these are realizable for semiconductor devices operating at telecommunication frequencies. Finally, note that dispersive effects are ignored for all results in this chapter since we always consider device performance at a single, fixed frequency.

### 6.4.1 Choice of Design Objective

As mentioned in Sect. 6.3.1, multiple equivalent choices of design objective exist which should allow one to achieve the same device performance; however, we will choose, for generality, the following design objective:

$$f(x) = \begin{cases} x - x_{\text{perfect}} & \text{at boundary,} \\ 0 & \text{elsewhere.} \end{cases} \quad (6.21)$$

That is,  $f(x)$  simply selects the outermost values of the field in the design space and compares them to values of a perfect device.

Furthermore, we choose  $f_{\text{ideal}} = 0$  so that when placed into the objective-first problem (6.17), this will result in fixing the boundary values of the field at the edge of the design space to those of an ideal device, as shown in Fig. 6.1. In this case, we choose such an ideal device to have perfect (unity) coupling efficiency, and these ideal fields are simply obtained by using the input and output mode profiles at the corresponding ports and using values of zero at the remaining ports.

Such a design objective is general in the sense that the boundary values of the device contain all the information necessary to determine how the device will interact with its environment, when excited with the input mode in question. In other words, we only need to know the boundary field values, and not the interior field values to determine the performance of the device; and thus, it would be conceivable that such a scheme might be generally applied to linear nanophotonic devices beyond just waveguide mode couplers.

In our case, we only need to know the value of  $H_z$  and its derivative along the normal direction,  $\partial H_z / \partial n$ , along the design boundary in order to completely characterize its performance. Alternatively, one can, of course, use the outermost two layers of the  $H_z$  instead of calculating a spatial derivative.

### 6.4.2 Application of the Objective-First Strategy

Having chosen our design objective we apply the alternating directions algorithm to (6.17) which results in solving the following two sub-problems iteratively:

$$\begin{aligned} & \underset{x}{\text{minimize}} \quad \|A(p)x - b(p)\|^2 \\ & \text{subject to } x = x_{\text{perfect}}, \quad \text{at boundary;} \end{aligned} \quad (6.22)$$

$$\begin{aligned} & \underset{p}{\text{minimize}} \quad \|B(x)p - d(x)\|^2 \\ & \text{subject to } p_0 \leq p \leq p_1. \end{aligned} \quad (6.23)$$

For the results throughout this chapter, we uniformly choose  $p_0 = 1/12.25$  and  $p_1 = 1$ , corresponding to  $\varepsilon^{-1}$  of silicon and air, respectively. Additionally, since a starting value for  $p$  is initially required, we always choose to use a uniform value of  $p = 1/9$  across the entire design space. There is nothing really unique about such a choice, although we have noticed that an initial value of  $p$  near 1 often results in poor designs. Note that unlike  $p$  we do not require an initial guess for  $x$ . The only other significant value that needs to be set initially is the frequency, or wavelength of light. We use free space wavelengths in the range of 25 to 63 grid points for the results in this chapter.

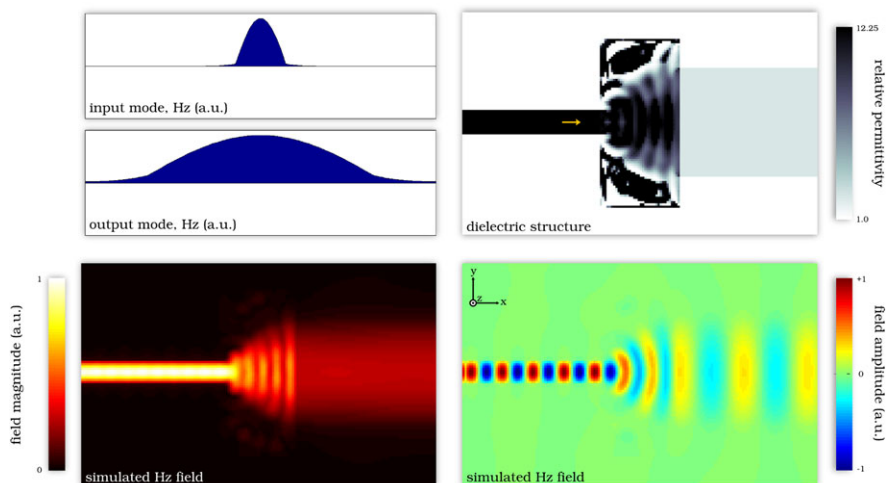
Lastly, for all the examples presented in the chapter, we run the alternating directions algorithm for 400 iterations. In terms of convergence, the physics residual never fully vanishes, and seems to asymptotically approach a nonzero value. Even so, we seem to obtain good performance from the produced designs. Note that although we do not present the convergence results here, such information can be obtained by inspecting the source code [2].

### 6.4.3 Coupling to a Wide, Low-Index Waveguide

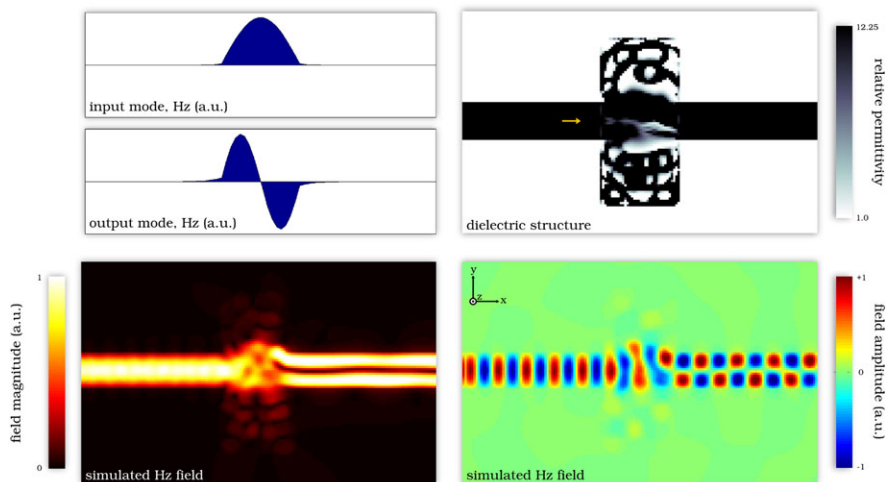
As a first example, we design a coupler from the fundamental mode of a narrow, high-index waveguide to the fundamental mode of a wide, low-index waveguide. Such a coupler would be useful for coupling from an on-chip nanophotonic waveguide to an off-chip fiber, for example.

The input and output mode profiles used as the ideal fields are shown in the upper-left corner of Fig. 6.2. The final structure is shown in the upper right plot, and the simulated  $H_z$  fields, under excitation of the input mode in this final structure, are shown in the bottom plots.

Figure 6.2 then shows that the design structure has nearly unity efficiency (99.8 %) and converts between the input and output modes within a very small footprint (roughly 1.5 square vacuum wavelengths).



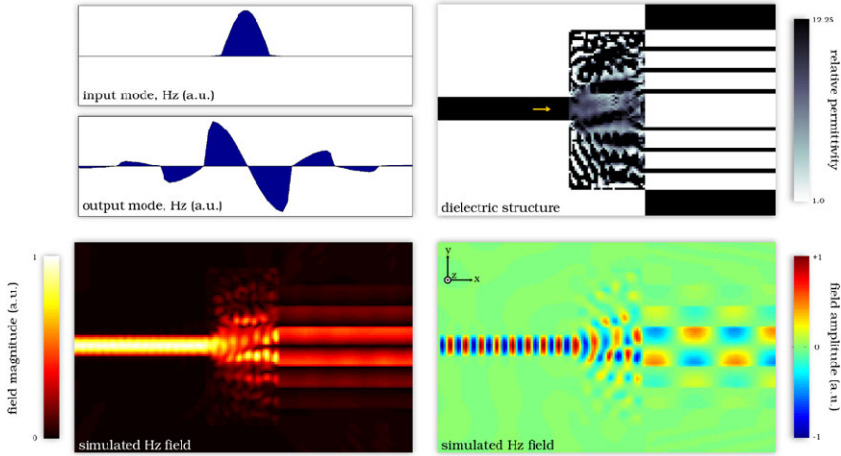
**Fig. 6.2** Coupler to a wide low-index waveguide. Efficiency: 99.8 %, device footprint:  $36 \times 76$  grid points, wavelength: 42 grid points



**Fig. 6.3** Mode converter. Efficiency: 98.0 %, device footprint:  $36 \times 76$  grid points, wavelength: 42 grid points

#### 6.4.4 Mode Converter

In addition to coupling to a low-index waveguide, we show that we can successfully apply the objective-first method to convert between modes of a waveguide. We do this by simply selecting the output mode in the design objective to be the second-order waveguide mode, as seen in Fig. 6.3. Note that the design of this coupler



**Fig. 6.4** Coupler to a wide low-index waveguide. Efficiency: 98.9 %, device footprint:  $36 \times 76$  grid points, wavelength: 25 grid points

is made challenging because of the opposite symmetries of the input and output modes. Moreover, because our initial structure is symmetric, we initially have 0 % efficiency to begin with. Fortunately, the objective-first method can still design an efficient (98.0 %) coupler in this case as well in a footprint of less than two square vacuum wavelengths.

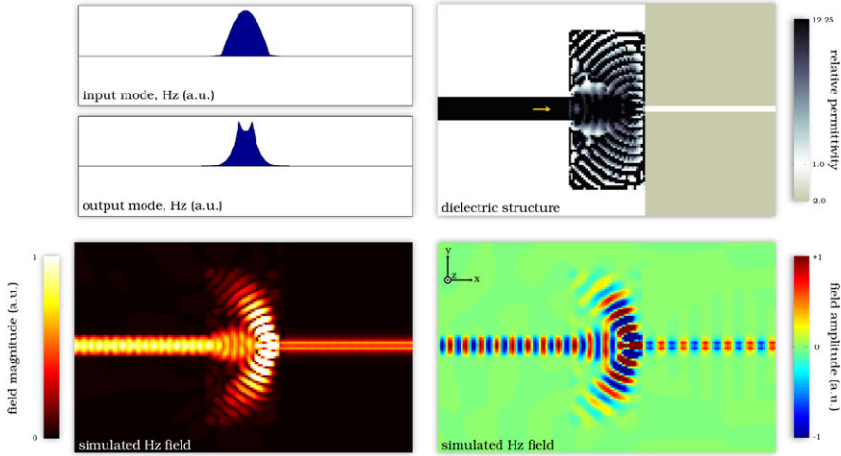
### 6.4.5 Coupling to an Air-Core Waveguide Mode

We can then continue to elucidate the generality of our method by coupling between waveguides which confine light in completely different ways. Figure 6.4 shows a high-efficiency coupling device between an index-guided input waveguide and a “air-core” output waveguide, in which the waveguiding effect is achieved using distributed Bragg reflection (instead of total internal reflection as in the input waveguide).

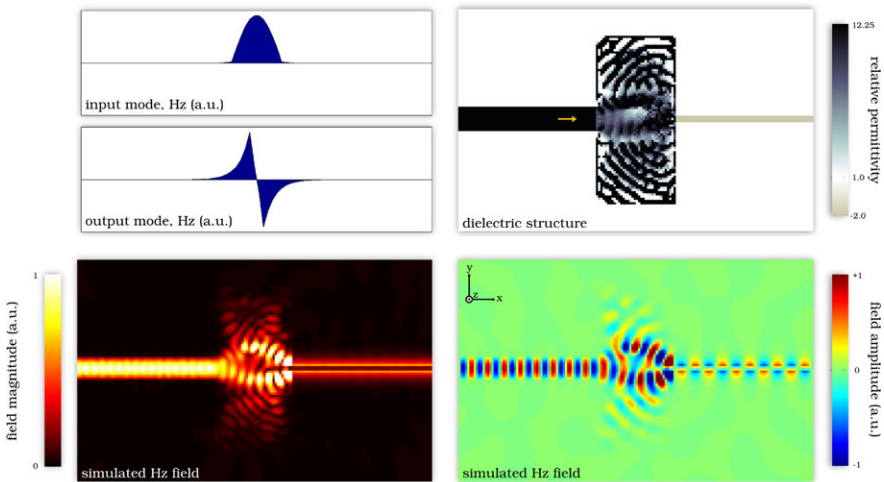
In this case, the device footprint is increased to 4.3 square vacuum wavelengths and the final efficiency is 98.9 %.

### 6.4.6 Coupling to a Metal–Insulator–Metal and Metal Wire Plasmonic Waveguides

Additionally, our design method can also generate couplers between different material systems such as between dielectric and metallic (plasmonic) waveguides, as



**Fig. 6.5** Coupler to a plasmonic metal–insulator–metal waveguide. Efficiency: 97.5 %, device footprint:  $36 \times 76$  grid points, wavelength: 25 grid points



**Fig. 6.6** Coupler to a plasmonic wire waveguide. Efficiency: 99.1 %, device footprint:  $36 \times 76$  grid points, wavelength: 25 grid points

shown in Fig. 6.5 (97.5 % efficiency). In this case, the permittivity of the metal ( $\epsilon = -2$ ) is chosen to be near the plasmonic resonance ( $\epsilon = -1$ ).

Extending this method to include plasmonic wire waveguides, Fig. 6.6 shows that efficiently coupling to this type of structure is achievable as well (99.1 % efficiency).

## 6.5 Optical Cloak Design

In the previous section, we showed that couplers between virtually any two waveguide modes could be constructed using the objective-first design method, and based on the generality of the method one can guess that it may also be able to generate designs for any linear nanophotonic device. Now, we extend the applicability of our method to the design of metamaterial devices which operate in free-space. In particular, we adapt the waveguide coupler algorithm to the to the design of optical cloaks.

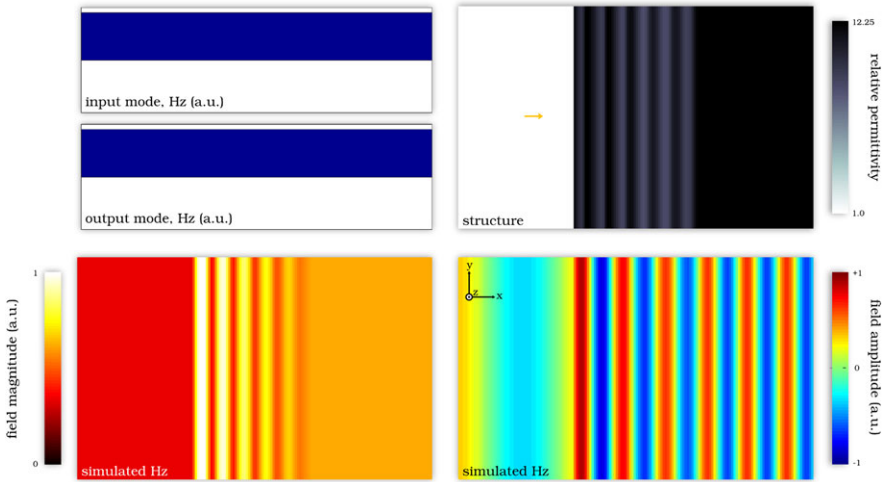
### 6.5.1 Application of the Objective-First Strategy

Adapting the method used in Sect. 6.4 to the design of optical cloaks really only requires one to change the simulation environment to allow for free-space modes. This is accomplished by modifying the upper and lower boundaries of the simulation domain from absorbing boundary conditions to periodic boundary conditions, which allows for plane-wave modes to propagate without loss until reaching the left or right boundaries, where absorbing boundary conditions are still maintained.

In terms of the design objective, we allow the device to span the entire height of the simulation domain, and thus consider only the leftmost and rightmost planes as boundary values. Specifically, for this section the input and output modes are plane waves with normal incidence, as can be expected for good cloaking devices. The achieved results all yield high efficiency, although we note that the cloaking effect is only measured for a specific input mode. That is to say, just as the waveguide couplers previously designed were single-mode devices, so the cloaks designed in this section are also “single-mode” cloaks. An additional modification, as compared to Sect. 6.4, is that we now prevent the structure from being modified in certain areas which contain the object to be cloaked. With these simple changes we continue to solve (6.17) with the alternating directions algorithm in order to design optical cloaks instead of waveguide couplers. Once again, as in Sect. 6.4, each design is run for 400 iterations with a uniform initial value of  $p = 1/9$  for the structure (where the structure is allowed to vary), and the range of  $p$  is limited to  $1/12.25 \leq p \leq 1$ , implying a dielectric cloak.

### 6.5.2 Anti-reflection Coating

As a first example, we attempt to design the simplest and most elementary “cloak-ing” device available, which, we argue, is a simple anti-reflection coating; in which case the object to be cloaked is nothing more than the interface between two dielectric materials. In this case, we use the interface between air and silicon, as shown in Fig. 6.7.



**Fig. 6.7** Anti-reflection coating. Efficiency: 99.99 %, device footprint:  $60 \times 100$  grid points, wavelength: 63 grid points

Unsurprisingly for such a simple case, we achieve a very high efficiency device. Note also that the efficiency of the device can be deduced by eye, based on the absence of reflections or standing waves in bottom two plots of Fig. 6.7.

### 6.5.3 Wrap-Around Cloak

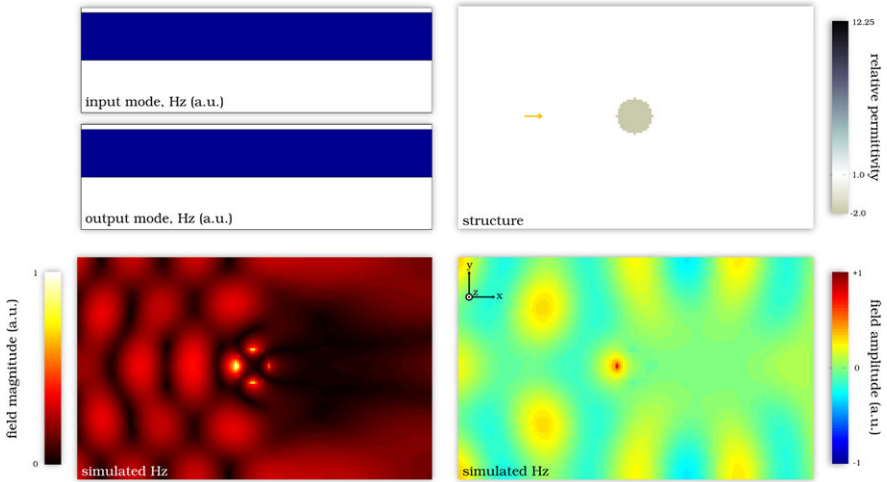
Next, we design a cloak for a plasmonic cylinder, which is quite effective at scattering light as can be seen from Fig. 6.8, where we show that the uncloaked cylinder, although sub-wavelength in size, scatters the majority of light away from the desired output (plane-wave) mode.

In designing the wrap-around cloak, we allow the structure to vary at all points within the design area except in the immediate vicinity of the plasmonic cylinder. Application of the objective-first strategy results in an efficient (greater than 99 %) device as seen in Fig. 6.9. Note that our cloak employs only isotropic, non-magnetic materials, and at the same time it is specific to a particular input and to a particular object. That is to say, it is a single-frequency, single-mode, and single-object cloaking device.

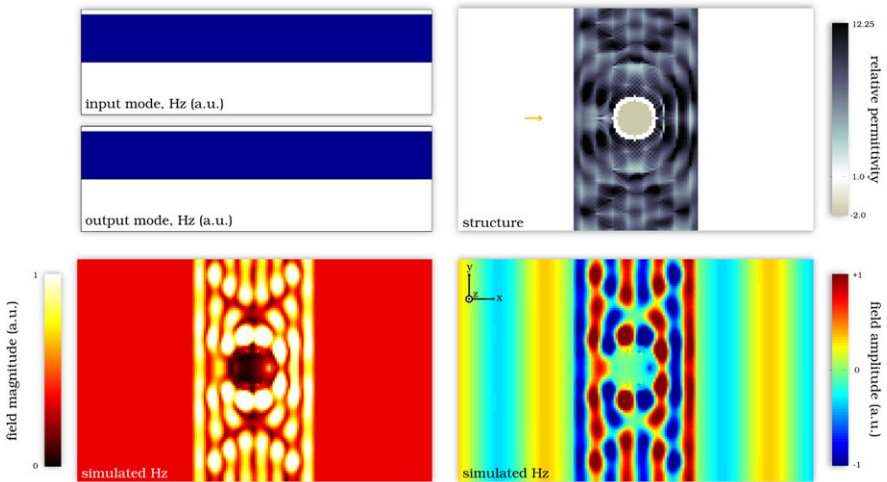
### 6.5.4 Open-Channel Cloak

With a simple modification, from the previous example, we can design a cloak which features an open channel to the exterior electromagnetic environment (Fig. 6.10).



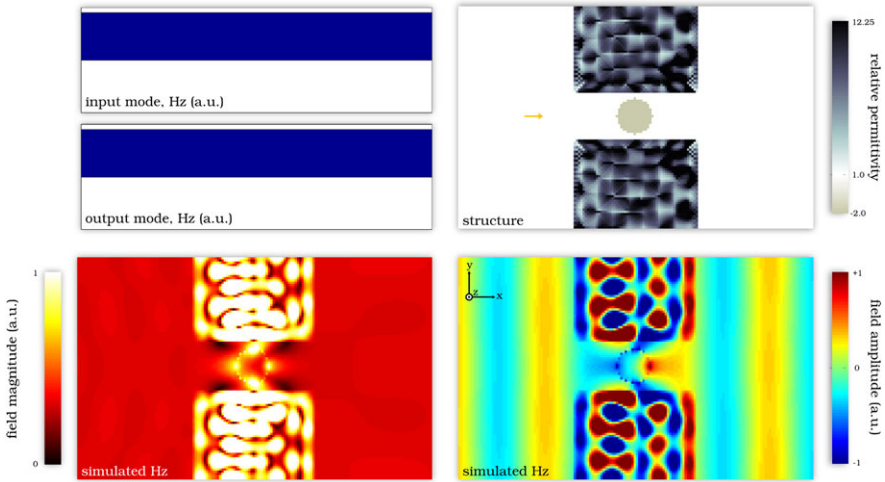


**Fig. 6.8** Plasmonic cylinder to be cloaked. 68.5 % of light is diverted away from the desired output mode



**Fig. 6.9** Wrap-around cloak. Efficiency: 99.99 %, device footprint:  $60 \times 100$  grid points, wavelength: 42 grid points

This simple modification creates an air gap that connects the cylinder to the outside world both from the front and back and may be useful in the case where one would like to remove or replace the cloaked object.



**Fig. 6.10** Open-channel cloak. Efficiency: 99.8 %, device footprint:  $60 \times 100$  grid points, wavelength: 42 grid points

Such a design is still very efficient (greater than 99 % efficiency) and demonstrates the usefulness of the objective-first strategy in cases where other methods, such as transformation optics, may require use of the entire space around the object to be cloaked.

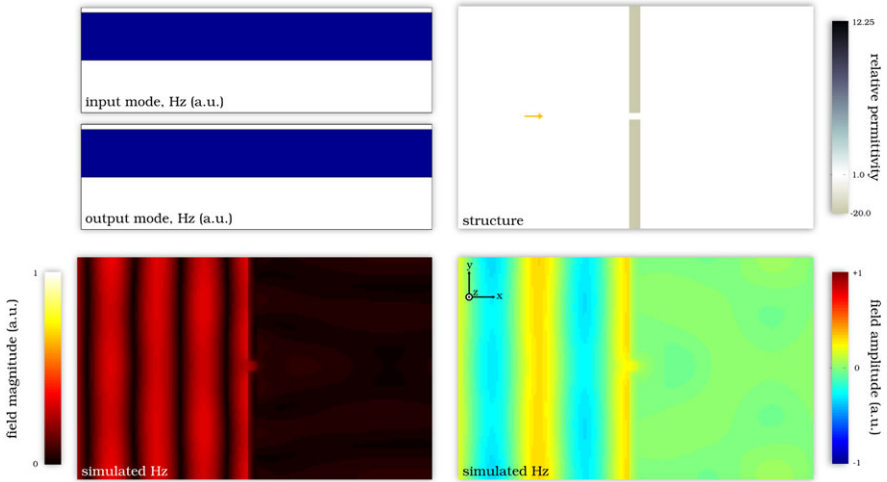
### 6.5.5 Channeling Cloak

Our last cloaking example replaces the plasmonic cylinder with a thin metallic wall in which a sub-wavelength channel is etched. Such a metallic wall is very effective at blocking incoming light (as can be seen from Fig. 6.11 where more than 99 % of the incoming light is blocked) because of its large negative permittivity ( $\epsilon = -20$ ), meaning that any cloaking device would be forced to channel all the input light into a very small aperture and then to flatten that light out into a plane wave again.

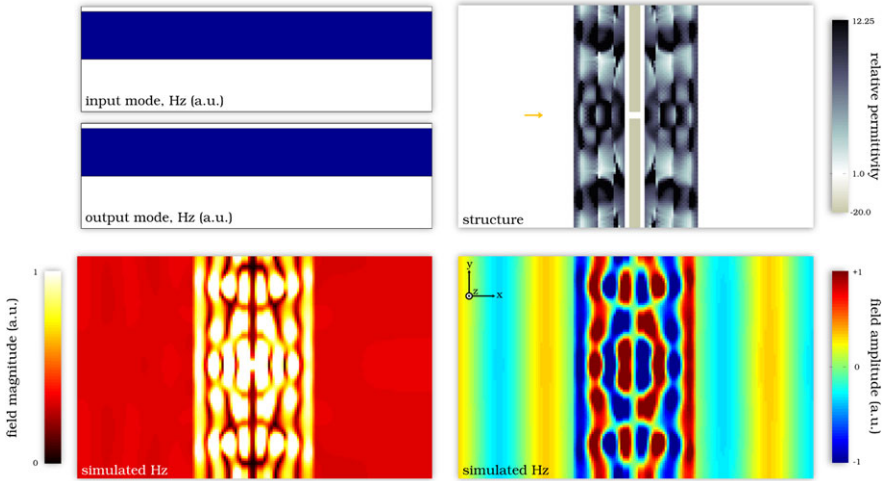
Once again, our method is able to produce a design with efficiency greater than 99 %, as shown in Fig. 6.12.

## 6.6 Optical Mimic Design

We now apply our objective-first strategy to the design of optical mimics. We define an optical mimic to be a linear nanophotonic device which mimics the output field of another device. In this sense optical mimics are anti-cloaks; where cloaks strive to make an object's electromagnetic presence vanish, mimics strive to implement an object's presence without that object actually being there.

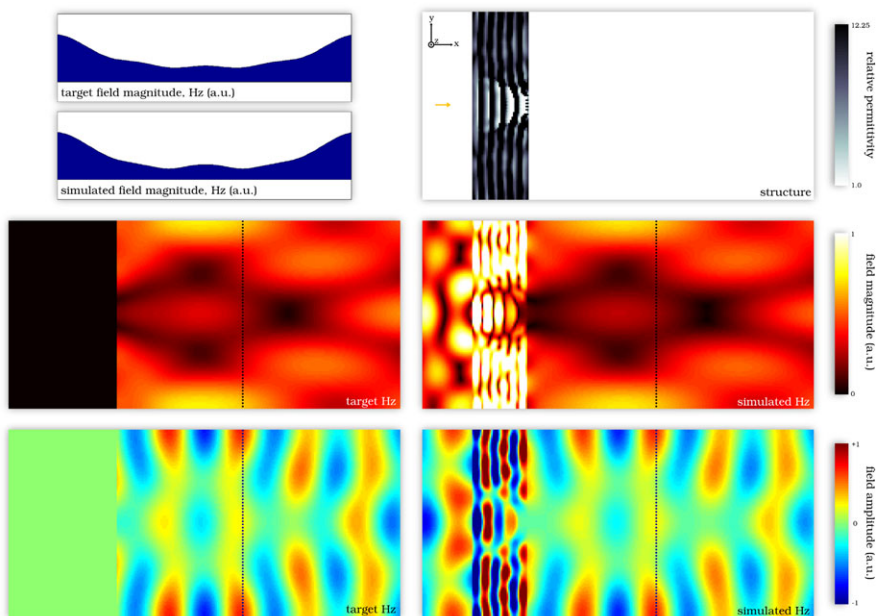


**Fig. 6.11** Metallic wall with sub-wavelength channel to be cloaked. 99.9 % of the light is blocked from the desired output plane-wave



**Fig. 6.12** Channeling cloak. Efficiency: 99.9 %, device footprint:  $60 \times 100$  grid points, wavelength: 42 grid points

As such, the design of optical mimics provides a tantalizing approach to the realization of practical metamaterial devices. That is to say, if one can reliably produce practical optical mimics, then producing metamaterials can be accomplished by simply producing an optical mimic of that material. In a more general sense, designing optical mimics is really just a recasting of the thrust of the objective-first design strategy in its purest form, namely the design of a nanophotonic device based



**Fig. 6.13** Plasmonic cylinder mimic (see Fig. 6.8 for the original object). Error: 8.1 %, device footprint:  $40 \times 120$  grid points, wavelength: 42 grid points

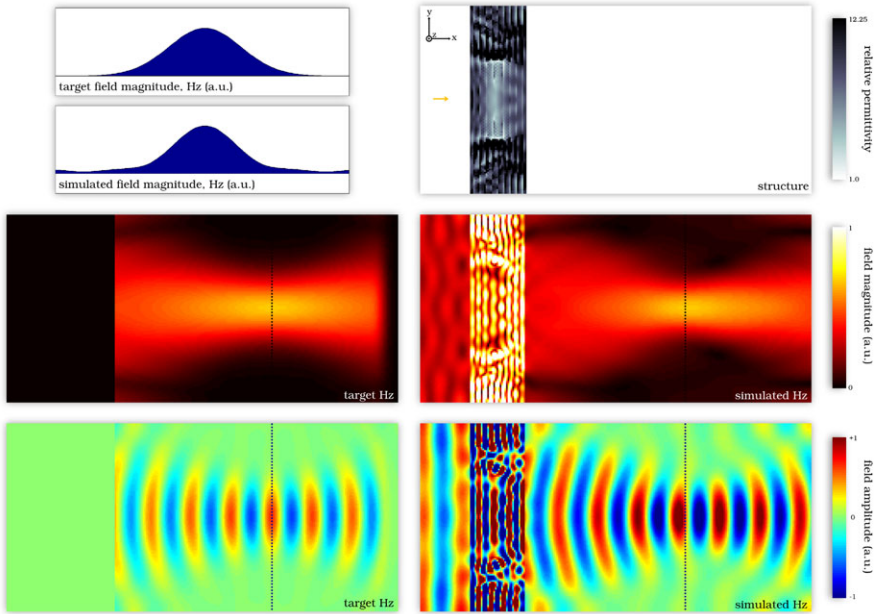
purely on the electromagnetic fields one wishes to produce. As such, devices which perform well-known optical functions (e.g., focusing, lithography) can also be designed.

### 6.6.1 Application of the Objective-First Strategy

The objective-first design of optical mimics proceeds in virtually an identical way to the design of optical cloaks, the only difference being that the output modes are specifically chosen to be those that produce the desired function. For most of the examples provided, the input illumination is still an incident plane wave. Lastly, instead of measuring efficiency, we measure the relative error of the simulated field against that of a perfect target field at a relevant plane some distance away from the device. The location of this plane is identified as a dotted line in the subsequent figures.

### 6.6.2 Plasmonic Cylinder Mimic

Our first design is simply to mimic the plasmonic cylinder which we cloaked in the previous section. Figure 6.13 shows the result of the design.



**Fig. 6.14** Full-width-half-max at focus:  $1.5 \lambda$ , focus depth: 100 grid points. Error: 12.0 %, device footprint:  $40 \times 120$  grid points ( $1.6 \lambda$  thick), wavelength: 25 grid points

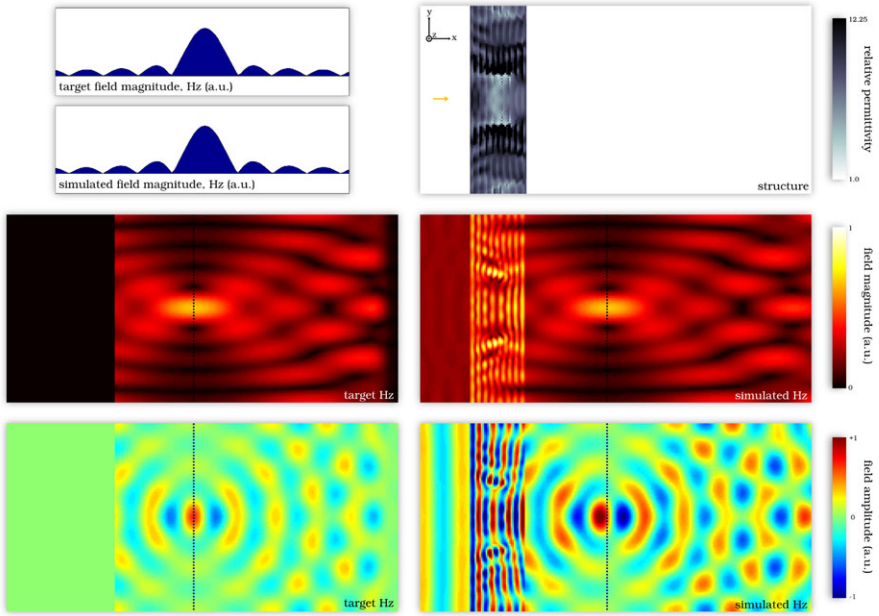
The final structure is shown in the upper right plot, while the ideal field and the simulated field are shown in the middle and bottom plots. Note that the ideal field is cut off to emphasize the fields to the right of the device (the output fields). Also, the magnitude of the fields are compared at the dotted black line at which point the relative error is also calculated. For this simple, initial mimic, the simulated field quite closely imitates that produced by a single plasmonic cylinder (8.1 % error)

### 6.6.3 Diffraction-Limited Lens Mimic

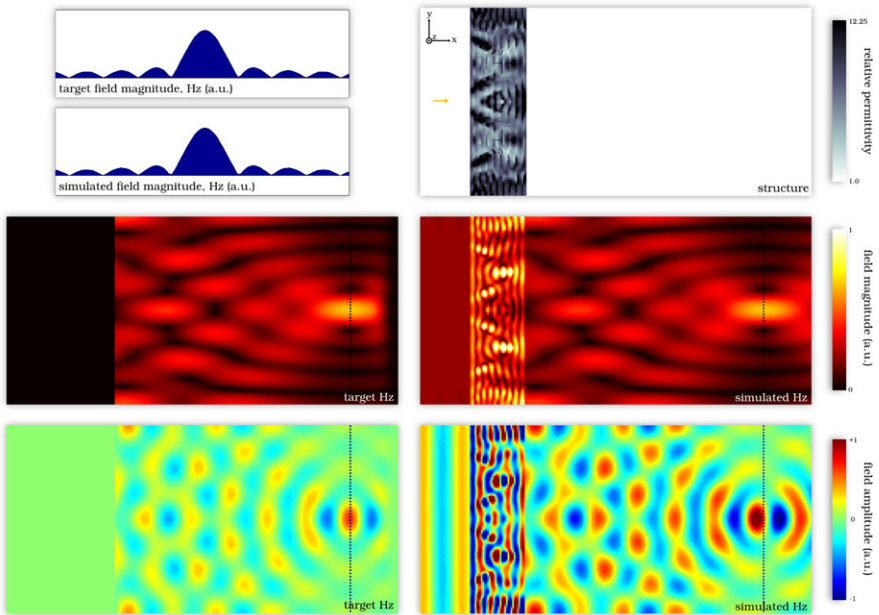
We now design a mimic for a typical diffraction-limited lens. In this case, the object which we wish to mimic does not require simulation since the fields of a lens can be readily computed. For the three figures below, the computed ideal fields are shown as the target fields.

Figure 6.14 shows the mimic of a lens with relatively moderate focusing. In such a lens, the focusing action is gradual and easily discernible by eye. The computed error in this case is 12.0 %.

In contrast, Figs. 6.15 and 6.16 are both mimics of a lens with a smaller half-wavelength spot size. Such a lens is much harder to design because of the high-frequency spatial components involved; and yet, we show that an objective-first

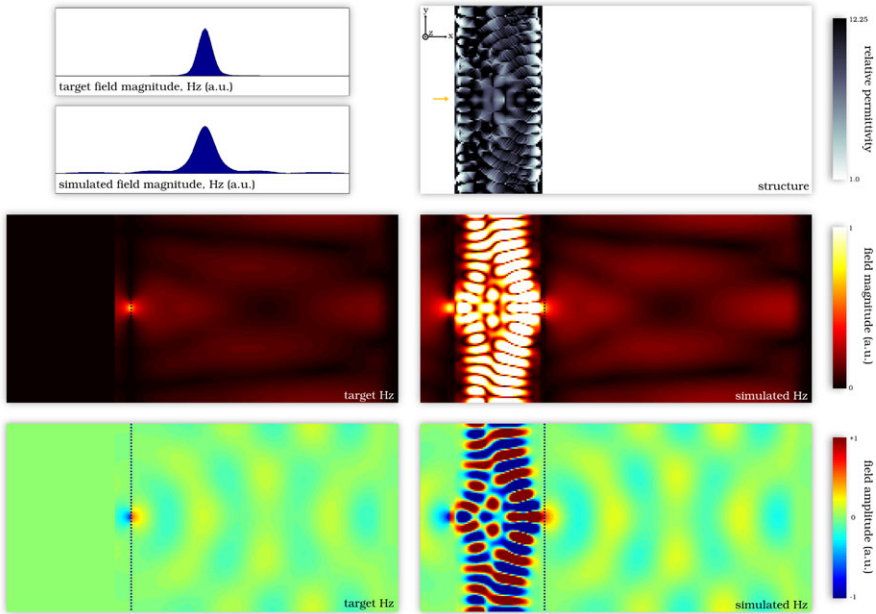


**Fig. 6.15** Full-width-half-max at focus:  $0.5 \lambda$ , focus depth: 50 grid points. Error: 5.6 %, device footprint:  $40 \times 120$  grid points ( $1.6 \lambda$  thick), wavelength: 25 grid points



**Fig. 6.16** Full-width-half-max at focus:  $0.5 \lambda$ , focus depth: 150 grid points. Error: 1.4 %, device footprint:  $40 \times 120$  grid points ( $1.6 \lambda$  thick), wavelength: 25 grid points





**Fig. 6.17** Sub-diffraction lens mimic. The target field has a full-width half-maximum of  $0.14 \lambda$ . Error: 28.6 %, device footprint:  $60 \times 120$  grid points ( $1.43 \lambda$  thick), wavelength: 42 grid points

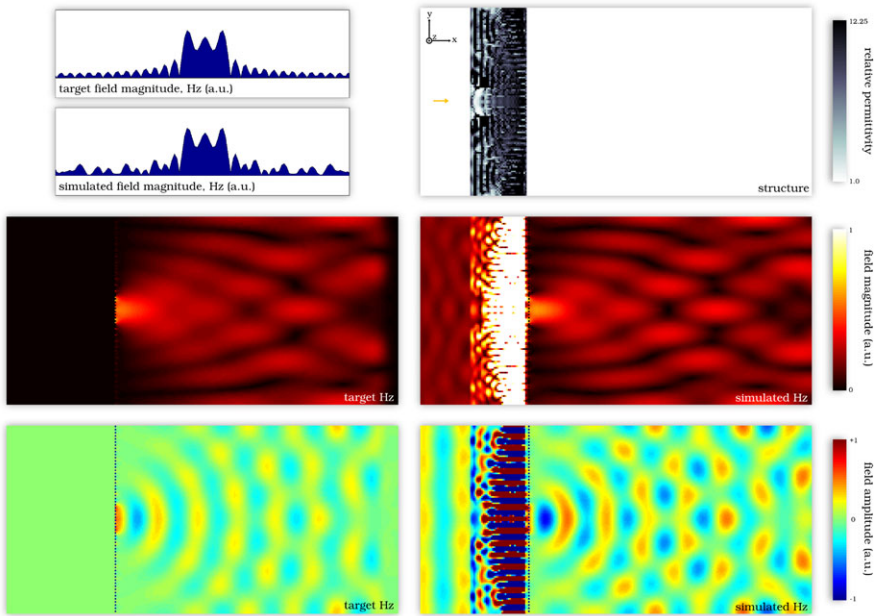
strategy can produce successful designs (5.6 % and 1.4 % error) and that this is achievable at both shorter and longer focal depths.

### 6.6.4 Sub-diffraction Lens Mimic

Our method is now employed to mimic the effect of a sub-diffraction lens. Since such a lens can be created using a negative-index material [12], this mimic can be viewed as an imitation of a negative-index material, in that the following device recreates the sub-diffraction target-field at the output plane (dotted line) when illuminated by the same target field at the input of the device. In other words, this device is an image-specific sub-diffraction imager, which is another way of saying that it is a single-mode imager.

As Fig. 6.17 shows, we are able to recreate the target field at the output, albeit with higher error (28.6 %). Although the error in this example is larger, the field produced by the device has a full-width half-maximum nearly equal to that of the target field.

Note that the target field is created simply by placing the imaging field at the output plane. Also note that, as expected, the output field decays very quickly since, for such a deeply sub-wavelength field, it is composed primarily of evanescently decaying modes.



**Fig. 6.18** Sub-diffraction optical mask. The three central peaks in the target field are each separated by  $0.28 \lambda$ . Error: 19.8 %, device footprint:  $40 \times 120$  grid points, wavelength: 25 grid points

### 6.6.5 Sub-diffraction Optical Mask

Lastly, we extend the idea of a sub-diffraction lens mimic one step further and design a sub-diffraction optical mask. Such a device takes a plane wave as its input and produces a sub-diffraction image at its output plane. Of course, akin to its lens counterpart, this output plane must lie within the near-field of the device (specifically, two computational cells away) because of its sub-wavelength nature. Figure 6.18 shows the design of a simple mask which successfully produces three peaks at its output with an error of 19.8 %.

## 6.7 Extending the Method

The objective-first method, as applied in the examples in this chapter, represents only a small foray into the area of nanophotonic design. Several key extensions to what is presented here are needed to fully address real-world nanophotonic design challenges.



### 6.7.1 Three-Dimensional Design

The first of these is the need to design fully three-dimensional structures. Doing so provides no inherent difficulties aside from the matrices in (6.6) becoming very large. This is not insurmountable as electromagnetic simulation software for three-dimensional nanophotonic structures already exists.

In fact, for certain choices of the design objective (i.e., those of low-rank) (6.18) can be efficiently solved by a small number of calls to unmodified simulation software. Of course, for general design objectives, such software will need to be modified in order to solve (6.18).

On the other hand, specialized software to solve (6.10) in any number of dimensions does not exist, although this was not a problem in two dimensions since generic linear algebra solvers are more than accurate. In three dimensions, the large size of matrix  $B(x)$  can be greatly compressed by considering only fabrication processes which modify a structure in-plane. In this way, the degrees of freedom in  $p$  can be greatly reduced and the original methods used in this chapter can still be applied. This work-around is especially appealing since in-plane structures are significantly less challenging to fabricate.

### 6.7.2 Multi-mode

A second necessary extension is to be able to consider the multiple fields that a structure produces in response to input fields of differing frequency and spatial distribution. Such an extension is straightforward in the objective-first formulation and results in the following modified problem statement,

$$\begin{aligned}
 & \underset{x_i, p}{\text{minimize}} \sum_i \|A(p)x_i - b(p)\|^2 \\
 & \text{subject to } f(x_i) = f_{i, \text{ideal}}, \quad i = 1, \dots, n, \\
 & p_0 \leq p \leq p_1,
 \end{aligned} \tag{6.24}$$

which can be separated into field and structure sub-problems as in the single-mode formulation. In the multi-mode case, this results in one structure sub-problem and  $n$  field sub-problems. Interestingly, the  $n$  field sub-problems lend themselves naturally to parallelization since they can be solved independently, leading to the possibility that a multi-mode design completing in roughly the same time as a single-mode design.

### 6.7.3 Binary Structure

Another necessary extension of our method is to force the values of  $p$  to be discrete. This is not trivial since a naive restatement of (6.17) which includes such a

constraint,

$$\begin{aligned} & \underset{x,p}{\text{minimize}} \quad \|A(p)x - b(p)\|^2 \\ & \text{subject to} \quad f(x) = f_{\text{ideal}}, \\ & \quad \quad \quad p \in \{p_0, p_1\}, \end{aligned} \tag{6.25}$$

results in a very difficult combinatorial problem.

Tractable approaches include penalizing intermediate values of  $p$  [13] or even transferring to a level-set method [9] where the distinction between materials is explicit.

### 6.7.4 Robustness

Lastly, the design of structures which are robust to both fabrication imperfections and fluctuations in environmental parameters is also a necessity for practical real-world devices.

It seems likely in this case that a heuristic approach may be most successful in this case, rather than tackling the problem head-on. For instance, to account for fluctuating material parameters induced by temperature changes one may design a device that operates over a larger bandwidth than is actually required.

## 6.8 Conclusions

We have introduced an objective-first approach to the design of nanophotonic components, and applied it to the design of waveguide couplers, optical cloaks, and optical mimics. In doing so, we hope to have exhibited both the simplicity and the breadth of our method to the design of a broad class of linear, single-mode devices. In addition to posting the source code for all the examples online [2], we have outlined the necessary extensions of our method in order to design practical, three-dimensional devices.

## References

1. J. Lu, J. Vuckovic, Objective-first design of high-efficiency, small-footprint couplers between arbitrary nanophotonic waveguide modes. *Opt. Express* **20**, 7221–7236 (2012)
2. <https://github.com/JesseLu/objective-first>
3. A. Taflov, S.C. Hagness, *Computational Electrodynamics*, 3rd edn. (Artech House, Norwood, 2005), Sect. 3.3.2
4. K.S. Yee, Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. *IEEE Trans. Antennas Propag.* **14**, 802–807 (1966)

5. S.G. Johnson, Notes on perfectly matched layers (PMLs) (2010). <http://math.mit.edu/~stevenj/18.369/pml.pdf>
6. W. Shin, S. Fan, Choice of the perfectly matched layer boundary condition for frequency-domain Maxwell's equations solvers. *J. Comput. Phys.* **231**, 3406–3431 (2012)
7. S. Boyd, L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, 2004)
8. M. Grant, S. Boyd, CVX: Matlab software for disciplined convex programming, version 1.21 (2011). <http://cvxr.com/cvx>
9. O.D. Miller, *Photonic Design: From Fundamental Solar Cell Physics to Computational Inverse Design* (U. C. Berkeley, Berkeley, 2012)
10. J. Lu, J. Vuckovic, Inverse design of nanophotonic structures using complementary convex optimization. *Opt. Express* **18**, 3793–3804 (2011)
11. S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* **3**, 1–1122 (2011)
12. J.B. Pendry, Negative refraction makes a perfect lens. *Phys. Rev. Lett.* **85**, 3966–3969 (2000)
13. M.P. Bendsoe, O. Sigmund, Material interpolation schemes in topology optimization. *Arch. Appl. Mech.* **69**, 635–654 (1999)