

Heuristic methods and performance bounds for photonic design

GUILLERMO ANGERIS,*  JELENA VUČKOVIĆ, AND STEPHEN BOYD

Department of Electrical Engineering, Stanford University, Stanford, California 94305, USA

**angeris@stanford.edu*

Abstract: In the photonic design problem, a scientist or engineer chooses the physical parameters of a device to best match some desired device behavior. Many instances of the photonic design problem can be naturally stated as a mathematical optimization problem that is computationally difficult to solve globally. Because of this, several heuristic methods have been developed to approximately solve such problems. These methods often produce very good designs, and, in many practical applications, easily outperform ‘traditional’ designs that rely on human intuition. Yet, because these heuristic methods do not guarantee that the approximate solution found is globally optimal, the question remains of just how much better a designer might hope to do. This question is addressed by performance bounds or impossibility results, which determine a performance level that no design can achieve. We focus on algorithmic performance bounds, which involve substantial computation to determine. We illustrate a variety of both heuristic methods and performance bounds on two examples. In these examples (and many others not reported here) the performance bounds show that the heuristic designs are nearly optimal, and can be considered globally optimal in practice. This review serves to clearly set up the photonic design problem and unify existing approaches for calculating performance bounds, while also providing some natural generalizations and properties.

© 2021 Optical Society of America under the terms of the [OSA Open Access Publishing Agreement](#)

1. Introduction

An important part of photonics, and many other scientific and engineering fields, is the design and construction of physical devices. A ‘physical device’, as used in this paper, includes anything as simple as a spherical lens, where a scientist can easily find an optimal lens for a given application by using basic algebra and ray optics, all the way to potentially very complicated devices and applications such as range detection and mapping using LiDAR [1], where building an ‘optimal’ device, in nearly any practically useful sense, is an open research problem. Though the two examples we give are in the field of photonics (and this is, indeed, the focus of our review), the term refers to any device that can, at least theoretically, be built, and whose desired behavior can be mathematically specified.

Traditionally, the design of physical devices was done by an engineer or scientist, whom we will generally call a designer, for a specific application. The designer would have a library, either physically or through experience, of well-understood components or materials, each performing a specific function. These components would then be carefully pieced together, often with a good amount of ingenuity, in order to perform the desired task. In many cases, the resulting designs could then be modified in part or in whole, or combined with other designs, in order to perform even more complicated functions. This procedure, while effective in practice, is time consuming and sometimes even tangential to the final application of the design itself.

A second approach to constructing physical devices was initially explored in the early 1960s within the field of electrical engineering, originally for the purpose of recognizing printed letters [2] and has since been extended to many other fields [3–24], with sometimes very surprising results [25]. In this approach, the designer specifies a mathematical objective function which, given a design, outputs a number representing how well the input design matches the desired

specifications; the lower this number is, the better the design. This function is then fed into an optimization algorithm, which attempts to minimize this objective function by finding a design that is good in the sense specified by the designer. In almost all practical cases, the algorithm will fail to find the best possible design, and, except in very specific scenarios, may never do so even when left to run for a very long time. But, in many applications, the resulting designs found often have much better performance than any design found by humans. This approach can be seen as a declarative approach to design: the user specifies *what* they want, while ceding control of *how* it should be done to the optimization algorithm. This idea has many names in different fields, with field-specific connotations and denotations; these include ‘automated design’, ‘computational design’, ‘inverse design’ (in photonics and aerospace engineering), ‘shape design’ and ‘generative design’ (in mechanical engineering), ‘topology design’ (in several fields), or ‘synthesis’ (in hardware design), among many others. We will simply call this optimization problem the ‘physical design problem’, with the understanding that many, if not all, of the previously mentioned problems are instances of the physical design problem.

2. Physical design problem

The usual physical design problem can be formally stated in many ways. In this review, we focus on a simple but general formulation, which, as we show in this section, includes many important problems in photonic design.

2.1. Physics

The design problem starts with a physical theory that describes the behavior of the *field* (which we will write as z) under some *excitation* (which we will write as b). The field z and excitation b are vectors in some (typically infinite dimensional) vector space. We focus here on the case when the physical theory is linear, in which case we can write the *physics equation* as

$$Az = b, \quad (1)$$

where A is a linear operator. Problems in physical design are governed by physics equations such as Maxwell’s equations, Helmholtz’s equation, the heat equation, the Schrödinger equation, among many others, which are linear in many important applications, and therefore of the form of (1).

Electromagnetic wave equation. For example, a common way of writing the electromagnetic (EM) wave equation in terms of the electric field E and the currents J , for a monochromatic wave with angular frequency ω is [26, §2],

$$-\nabla \times \nabla \times E + \omega^2 \mu_0 \varepsilon E = -i\omega \mu_0 J,$$

where ε denotes the permittivities at each point in space and μ_0 is the magnetic permeability, which we assume to be constant throughout space, in this example. We can then make the following correspondences:

$$\underbrace{(-\nabla \times \nabla \times \cdot + \omega^2 \mu_0 \varepsilon)}_A \underbrace{E}_z = \underbrace{-i\omega \mu_0 J}_b, \quad (2)$$

which naturally leads to an equation of the form of (1).

Discretization. We will work with an appropriate discretization of the field, excitation, and physics Eq. (1). We will overload notation to use the same symbols for their discretized versions. In the sequel, the field z will be a vector in \mathbf{R}^n , the excitation b will be a vector in \mathbf{R}^m , and the linear operator A will be a matrix in $\mathbf{R}^{m \times n}$. The physics Eq. (1) is then a set of m linear equations

in n scalar variables. Complex fields and excitations can be reduced to the real case by separating them into their real and imaginary parts.

Solutions and simulations. For a fixed A and excitation b , we will call any z which satisfies (1) a *solution* of the physics equation. In general there can be a unique solution, many solutions, or no solution. We will focus on the case when there is a unique solution, *i.e.*, $m = n$ and A is invertible, so $z = A^{-1}b$.

We refer to computing the field $z = A^{-1}b$, given A and b , as a *simulation*. There are many simulation methods, including generic methods for solving linear equations such as sparse-direct methods [27,28] or iterative methods [29,30], and custom methods crafted specifically for the particular physics equations [31–33]. We note that, in practical photonic design, the resulting linear systems can be very large, with the number of variables, n , often in the millions or tens of millions. Solving systems in the upper end of this range often requires the use of large-scale linear solvers [34].

Approximate solutions and physics residual. Some of the methods we will see work with approximate solutions of the physics equation. For any field vector z and excitation b , we define the *physics residual* as $r_{\text{phys}} = Az - b$. A reasonable numerical measure of the size of the residual is $\|Az - b\|/\|b\|$, where $\|\cdot\|$ is a norm, typically the Euclidean norm $\|\cdot\|_2$. Simulations, especially those that use iterative methods, produce fields with small physics residuals.

Modes. A simple trick can be used to represent the modes of a system as a solution to (1). Suppose the original (discretized) physics equation is $H z = \lambda z$, where λ is the eigenvalue and z is an associated mode. Directly expressing this as $Az = b$ with $A = H - \lambda I$ and $b = 0$ yield a physics equation that is not invertible and has multiple solutions (including, of course, $z = 0$). To fix a unique solution, we use a linear normalization and insist that $c^T z = 1$, where $c \in \mathbf{R}^n$ is some nonzero vector. We then represent the mode equation and normalization as $Az = b$ with

$$A = \begin{bmatrix} H - \lambda I \\ c^T \end{bmatrix}, \quad b = \begin{bmatrix} 0_n \\ 1 \end{bmatrix}. \quad (3)$$

This has a unique solution, provided c is not an eigenvector of H with eigenvalue λ and that λ is a simple eigenvalue. Note that the linear normalization in (3) differs from the usual choice of normalization, $\|z\|_2 = 1$, where $\|\cdot\|_2$ is the Euclidean norm.

2.2. Design parameters

In physical design, the designer is able to change the system physics Eq. (1), by choosing some parameters that affect the physics, *i.e.*, A and b . Thus A and b depend on some *design parameters* $\theta \in \mathbf{R}^d$. For example, in photonic design, θ is generally a variable that controls the permittivities inside of the device. We can then write the physics Eq. (1) with explicit dependence on the design parameters as

$$A(\theta)z = b(\theta). \quad (4)$$

When $A(\theta)$ is invertible we have $z(\theta) = A(\theta)^{-1}b(\theta)$; *i.e.*, the field also depends (implicitly) on the design parameters θ .

Affine physics design. In many practical cases (and in all of the examples we will show), A and b are *affine functions* of the design parameters; *i.e.*,

$$A(\theta) = A_0 + \sum_{i=1}^d \theta_i A_i, \quad b(\theta) = b_0 + \sum_{i=1}^d \theta_i b_i,$$

where $A_i \in \mathbf{R}^{m \times n}$ and $b_i \in \mathbf{R}^m$. These A_i are usually sparse matrices and vectors; that is, each design parameter θ_i affects just a few entries of A and b . Because Eq. (4) is affine in θ when

holding z fixed, and affine in z when holding θ fixed, this type of equation is sometimes called ‘bi-affine’ or ‘multilinear’ in θ and z .

Diagonal physics design. A common and useful special case of (4) is when $b(\theta)$ is a constant, $b(\theta) = b_0$, and $A(\theta)$ can be written as

$$A(\theta) = A_0 + \mathbf{diag}(\theta), \quad (5)$$

where $\mathbf{diag}(\theta)$ is a matrix whose diagonal entries contain the elements of the vector θ and is zero elsewhere. In other words, $A_i = E_{ii}$, where E_{ii} is the matrix with only one nonzero entry (which is one), in the i, i entry. We will call this special case of (4) the *diagonal physics equation*.

EM wave equation in diagonal form. A specific example of a diagonal physics equation are Maxwell’s Eqs. (2), whenever a designer is allowed to vary the permittivities. In this case, we can define θ to be proportional to the permittivities (with proportionality constant $\mu_0\omega^2$) such that the following correspondences can be made:

$$\underbrace{(-\nabla \times \nabla \times \cdot + \mu_0\omega^2 \varepsilon)}_{A_0} \underbrace{E}_{z} = \underbrace{-i\omega\mu_0 J}_b.$$

This correspondence results in an equation of the form of (5). (We can also similarly write the more general case where the designer is allowed to vary both the permittivities and permeabilities, in this form. See, *e.g.*, the appendix in [35].)

Low-rank updates. Whenever $A(\theta)$ is sparse (which is very common in practice), it is often possible to compute an explicit factorization of $A(\theta)$ (*e.g.*, the sparse Cholesky factorization, when $A(\theta)$ is positive semidefinite) which makes evaluating $A(\theta)^{-1}b(\theta)$ inexpensive, after the factorization. Additionally, whenever the matrices A_i are also low rank (as in the diagonal case (5), for example), updates to the factorization of $A(\theta)$ can be efficiently computed [28]. This implies that we can also efficiently evaluate $A(\theta')^{-1}b(\theta')$, when only a small number of entries of θ' differ from those of θ , given the factorization for $A(\theta)$.

Parameter constraints. In general, a designer has constraints on the parameters that can be chosen. Because of this, we will define the *feasible parameter set*, $\Theta \subseteq \mathbf{R}^d$, such that only design parameters satisfying $\theta \in \Theta$ are feasible or valid. In many applications, Θ is a hyperrectangle (or box) indicating that each component of θ must lie in some interval given by $\theta^{\min}, \theta^{\max} \in \mathbf{R}^d$, *i.e.*,

$$\Theta = \{\theta \in \mathbf{R}^d \mid \theta^{\min} \leq \theta \leq \theta^{\max}\},$$

where the inequalities are elementwise.

In photonic design, it is often not possible to vary the permittivities along an interval but are instead allowed to be one of two possible values. This leads to another common set of parameter constraints, where each component of θ is constrained to be exactly one of two elements (we will call this class of constraints *Boolean constraints*):

$$\Theta = \{\theta \in \mathbf{R}^d \mid \theta_i \in \{\theta_i^{\min}, \theta_i^{\max}\} \text{ for } i = 1, \dots, d\}.$$

In this class of constraints, the total number of parameters that are feasible is $|\Theta| = 2^d$. The feasible parameter set Θ can also include fabrication constraints such as minimum possible feature sizes, among other possibilities [36], but we will focus on the common cases of box or Boolean constraints.

Normalization. We can re-parametrize the design parameters to lie between -1 and 1 (or any other limits). So, without loss of generality, we can always consider the upper and lower bounds to be $\theta^{\min} = -\mathbf{1}$ and $\theta^{\max} = \mathbf{1}$, where $\mathbf{1}$ is the vector with all entries equal to one. To do this, we introduce a new parameter $\delta \in \mathbf{R}^d$ and define

$$\theta = \bar{\theta} + \rho \circ \delta,$$

where $\bar{\theta} = (\theta^{\max} + \theta^{\min})/2$ is the parameter midpoint, while $\rho = (\theta^{\max} - \theta^{\min})/2$ is the parameter radius, and \circ denotes the elementwise (Hadamard) product. The constraint $\theta \in \Theta$ becomes

$-1 \leq \delta \leq 1$, where the inequalities are elementwise, in the box-constrained case, or $\delta \in \{-1, 1\}^d$ in the Boolean case. We then have $A(\theta) = \tilde{A}(\delta)$, $b(\theta) = \tilde{b}(\delta)$, with \tilde{A} and \tilde{b} affine, and

$$\tilde{A}_0 = A(\bar{\theta}), \quad \tilde{b}_0 = b(\bar{\theta}),$$

while

$$\tilde{A}_i = \rho_i A_i, \quad \tilde{b}_i = \rho_i b_i, \quad i = 1, \dots, d.$$

2.3. Optimization problem

The design objective is often written as a function of the fields, specifying how well the resulting field matches the desired objective. This objective function could specify the power in a given direction, the field overlap (*i.e.*, the inner product between the current field and a desired one), or the total energy, all of which can be written as functions depending only on the field, that the designer may wish to optimize. For example, the designer may wish to maximize the power transmitted through a specific port of a device at a given frequency [37], or to maximize the focusing efficiency of a lens within a specific region [38,39]. Finding a device that best matches this objective can be directly phrased as a mathematical optimization problem.

Objective function. In other words, we seek to find a design whose field optimizes an *objective function* $f : \mathbf{R}^n \rightarrow \mathbf{R} \cup \{+\infty\}$. The function's input is a field z (generated by some design $\theta \in \Theta$) and its output is a number that specifies how good or bad this field is, or how well the field matches the designer's specification. Without loss of generality, we will assume that a higher number is worse (*i.e.*, a designer wishes to minimize f), but we can just as well maximize f by, equivalently, minimizing its negative, $-f$. We allow the objective function f to take on infinite values to denote hard constraints on the desired field: if $f(z) = +\infty$ for some field z , then z is not a feasible field.

Problem statement. We can then compactly write the problem that a designer wishes to solve (or approximately solve), which we will call the *physical design problem*:

$$\begin{aligned} & \text{minimize} && f(z) \\ & \text{subject to} && A(\theta)z = b(\theta) \\ & && \theta \in \Theta, \end{aligned} \tag{6}$$

where the problem variables are the fields $z \in \mathbf{R}^n$ and design parameters $\theta \in \mathbf{R}^d$, while the problem data include the matrix $A(\theta) \in \mathbf{R}^{m \times n}$, the excitation $b(\theta) \in \mathbf{R}^m$, and the parameter constraint set $\Theta \subseteq \mathbf{R}^d$. We will call the special case of (6) where the physics equation is the diagonal physics Eq. (5) the *diagonal physical design problem*. (As a reminder, in photonic design, θ is usually proportional to the permittivities, while $A(\theta)$ is the operator corresponding to the electromagnetic wave Eq. (2).)

Problem attributes. Problem (6) has several important properties. In many practical cases, the function f is a convex function and the set Θ is a convex set, which implies that problem (6) is a convex problem in the variable z , when holding θ fixed, and is a convex problem in θ , when holding z fixed. This property leads to some useful heuristics for approximately solving problem (6); *cf.*, [40,41]. Additionally, problem (6) often has a smooth (differentiable) objective function f , while Θ can almost always be represented as a number of smooth equality and inequality constraints, which happens in many practical applications and all examples presented in this review. In this case, we can apply general nonlinear optimization solvers such as IPOPT [42] directly to problem (6).

Computational hardness. On the other hand, it is not difficult to show that even finding a feasible design and field for problem (6) is, in general, a computationally difficult problem (*i.e.*, it is NP-hard) even when the design parameters θ are unconstrained; that is, even if $\Theta = \mathbf{R}^n$. To do

this, we will reduce the *subset sum problem* [43], a problem known to be NP-hard, to an instance of (4). This would imply that, if we could efficiently solve problem (4), then we could efficiently solve the subset sum problem, which is widely believed to be computationally hard to solve. (See, e.g., [44] for a good overview of P vs. NP and its implications.)

The subset sum problem asks: given $c \in \mathbf{R}^n$, is there a nonzero binary vector $x \in \{0, 1\}^n$ such that $c^T x = 0$? We will show that, given c , we can answer this question by finding a field z and a design θ that satisfy the constraints of (6), which will imply that, in general, problem (6) is computationally difficult.

First, note that we can write the following conditions on z and θ ,

$$c^T z = 0, \quad \theta_{n+1} \mathbf{1}^T z = 1, \quad \theta_i (z_i - 1) = 0, \quad z_i = \theta_i, \quad i = 1, \dots, n,$$

as an instance of (4), by appropriately stacking the conditions into a matrix form. Now, the last two conditions are true if, and only if, $z_i(z_i - 1) = 0$, which also happens only when $z_i \in \{0, 1\}$ for $i = 1, \dots, n$. The second condition implies that $\mathbf{1}^T z \neq 0$, and, when combined with the first condition, this statement is true if, and only if, there exists a nonzero solution to the subset sum problem, with the provided vector c . This, in turn, shows that finding a feasible design θ and field z such that the constraints of (6) are satisfied must, in general, be a problem that is at least as hard as the subset sum problem.

Because problem (6) is likely to be computationally difficult to solve exactly when the number of parameters is large, we will focus on heuristics which approximately solve the problem for the remainder of this paper.

Multi-scenario design. A common design task is to find a single device that has good performance across many different scenarios. For example, the device might need to be robust against temperature variations, or the device might be required to filter out a number of specific wavelengths, while allowing others through. In this case, we will assume that the device satisfies S instances of the physics Eq. (1), where the design parameters θ are held fixed across the S instances, but the physics equation or the excitation, is allowed to vary:

$$A^s(\theta) z^s = b^s(\theta), \quad s = 1, \dots, S.$$

Here $A^s(\theta) \in \mathbf{R}^{m_s \times n_s}$, $b^s(\theta) \in \mathbf{R}^{m_s}$ and $z^s \in \mathbf{R}^{n_s}$ for $s = 1, \dots, S$. This leads to the multi-scenario physical design problem:

$$\begin{aligned} & \text{minimize} && f(z^1, \dots, z^S) \\ & \text{subject to} && A^s(\theta) z^s = b^s(\theta), \quad s = 1, \dots, S \\ & && \theta \in \Theta, \end{aligned} \quad (7)$$

where the variables are the fields $z^s \in \mathbf{R}^{n_s}$ in each of the $s = 1, \dots, S$ scenarios and the design parameters $\theta \in \mathbf{R}^d$, while the objective function $f: \mathbf{R}^{n_1} \times \dots \times \mathbf{R}^{n_S} \rightarrow \mathbf{R}$ can depend on the fields at any of the S scenarios.

In fact, it turns out that we can write any instance of the multi-scenario physical design problem (7) as an instance of (6). To do this, we collect all of the individual physical equations into a single constraint by placing all of the $A^s(\theta)$ along the diagonal of a larger matrix $A(\theta)$, and stacking the excitations $b^s(\theta)$ and fields z^s . More specifically, define

$$A(\theta) = \begin{bmatrix} A^1(\theta) & 0 & \dots & 0 \\ 0 & A^2(\theta) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & A^S(\theta) \end{bmatrix}, \quad b(\theta) = \begin{bmatrix} b^1(\theta) \\ b^2(\theta) \\ \vdots \\ b^S(\theta) \end{bmatrix}, \quad z = \begin{bmatrix} z^1 \\ z^2 \\ \vdots \\ z^S \end{bmatrix}.$$

So we can write the S distinct physical equations, $A^s(\theta)z^s = b^s(\theta)$ for scenarios $s = 1, \dots, S$, as the single equation, $A(\theta)z = b(\theta)$, with dimensions $A(\theta) \in \mathbf{R}^{m \times n}$, $z \in \mathbf{R}^n$, and $b(\theta) \in \mathbf{R}^m$, where $m = m_1 + \dots + m_S$ and $n = n_1 + \dots + n_S$. Overloading notation slightly, such that $f(z) = f(z_1, \dots, z_S)$, then reduces problem (7) to one of the form of (6). In other words, it suffices to only consider a problem of the form of (6).

Eliminating the field variables. A simple (and relatively common) equivalent formulation of problem (6) is to note that, because we have assumed $A(\theta)$ is invertible, we can write it as a problem that depends only on the design variables; *i.e.*, since $z = A(\theta)^{-1}b(\theta)$, we can write problem (6) as the following optimization problem over $\theta \in \mathbf{R}^d$:

$$\begin{aligned} & \text{minimize} && f(A(\theta)^{-1}b(\theta)) \\ & \text{subject to} && \theta \in \Theta. \end{aligned} \quad (8)$$

When the function f is differentiable, we can easily compute its derivatives with respect to each component of θ :

$$\frac{\partial}{\partial \theta_i} f(A(\theta)^{-1}b(\theta)) = (\nabla f(z))^T A(\theta)^{-1} A_i z,$$

where $z = A(\theta)^{-1}b(\theta)$ is the solution to (4) for design parameters θ . We can write this in a slightly more compact form by defining $y = A(\theta)^{-T}(\nabla f(z))$, such that

$$\frac{\partial}{\partial \theta_i} f(A(\theta)^{-1}b(\theta)) = y^T A_i z. \quad (9)$$

Note that, to find z and y , we only need to solve two systems of linear equations, one over $A(\theta)$, and one over $A(\theta)^T$. This observation can be used to efficiently compute the gradient of the objective with respect to the design variables and is called the *adjoint method*, dating back to the control theory literature of the 1970s [45,46]. Other methods of computing the derivative when, for example, the matrix $A(\theta)$ is not invertible include automatic differentiation through the simulation (as in [47–49]), among many others.

Eliminating the design variables. In the case of diagonal design, and for some choices of the parameter constraint set Θ , it is also possible to eliminate the corresponding design variables [50]. For example, in the case where the parameter constraint set Θ is a box, $-1 \leq \theta_i \leq 1$, then, given a field z , there exist design parameters θ satisfying

$$(A_0 + \mathbf{diag}(\theta))z = b_0,$$

(*i.e.*, z satisfies the diagonal physics Eq. (5) with parameters θ) if, and only if the field z satisfies

$$|A_0 z - b_0| \leq |z|,$$

where the absolute value $|\cdot|$ is taken elementwise.

To see this, note that, from (5), we can write

$$A_0 z - b_0 = -\mathbf{diag}(\theta)z,$$

and, since we are free to choose any $-1 \leq \theta \leq 1$, we have that such a θ exists, if, and only if,

$$|A_0 z - b_0| \leq |z|.$$

This lets us write the diagonal physical design problem in the following equivalent way:

$$\begin{aligned} & \text{minimize} && f(z) \\ & \text{subject to} && |A_0 z - b_0| \leq |z|, \end{aligned} \quad (10)$$

where the only problem variable is the field $z \in \mathbf{R}^n$. We will call this formulation of the diagonal physical design problem, as in [50], the *absolute-upper-bound* formulation.

This rewriting also shows an interesting property of the physical design problem, whenever f is convex: if you know the signs of any optimal field, then the design problem becomes convex and therefore easy to solve globally. That is, if we know $s = \mathbf{sign}(z^*)$, where \mathbf{sign} is the signum function

$$\mathbf{sign}(z^*)_i = \begin{cases} 1 & z_i^* \geq 0 \\ -1 & z_i^* < 0, \end{cases}$$

and where z^* is optimal for (10), then any solution to the convex optimization problem

$$\begin{aligned} &\text{minimize} && f(z) \\ &\text{subject to} && |A_0 z - b_0| \leq \mathbf{diag}(s)z, \end{aligned} \tag{11}$$

with variable $z \in \mathbf{R}^n$ is a solution to (10) with the same optimal value as z^* . This follows from two basic facts. First, any z that is feasible for (11) is feasible for (10) with the same objective value because

$$|A_0 z - b_0| \leq \mathbf{diag}(s)z \leq |z|.$$

And, second, that z^* is feasible for (11) since

$$\mathbf{diag}(s)z^* = |z^*|,$$

and, by definition, z^* is feasible for (10) and so satisfies $|A_0 z^* - b_0| \leq |z^*|$. This observation also leads to an optimization algorithm which iteratively updates the signs and yields a sequence of feasible fields with decreasing objective value [50, §3], called sign-flip descent or SFD for short. We compare its performance against other basic solvers in §4.

A similar analysis holds when $\Theta = \{-1, 1\}^n$, *i.e.*, θ is constrained to be Boolean. In this case, we have that, given some field z , there exists a Boolean design $\theta \in \{-1, 1\}^n$ such that θ and z satisfy the diagonal physics equation if, and only if, z satisfies

$$|A_0 z - b_0| = |z|. \tag{12}$$

We note that, in photonics, the box-constrained formulation sometimes leads to designs that cannot be practically implemented because the resulting designs may have permittivities that lie along an interval (*i.e.*, θ_i may lie anywhere in $[-1, 1]$), while most fabrication methods only allow the use of two possible permittivities (*i.e.*, we must have $\theta_i \in \{-1, 1\}$). Despite this, the box-constrained formulation is often used as a good initialization for current inverse design algorithms, which then approximately solve the Boolean case [34].

Approximate solution methods. There are many practical methods for approximately solving (6). For example, many of the earliest solution methods approximately solve the physical design problem by applying zeroth order (or ‘derivative free’) optimization algorithms after eliminating the field variable, as shown in problem (8) [51–57]. Such methods are easy to implement in practice, as they only require the ability to perform a basic simulation; *i.e.*, to solve the physics Eq. (4) for a given $\theta \in \Theta$. Zeroth order optimization methods include hill-climbing, genetic algorithms [58], simulated annealing [59], Nelder-Mead [60], and adaptive coordinate descent [61], among many others [62]. While effective at finding designs with moderate to good performance, zeroth order optimization methods scale poorly and suffer from slow convergence when compared to higher-order methods. (See [63] and [64] for more information on zeroth order optimization methods.)

A second important family of optimization algorithms, which include the algorithms most used in practice, are the first order optimization algorithms, which are also almost always applied to problem (8). In these cases, such methods additionally make use of gradient information (9), leading to better computational performance and faster convergence times, at the expense of

higher implementation complexity, as these methods require additional information from the simulator. Examples of first order optimization algorithms used in practice include L-BFGS-B [65], proximal gradient methods [66], and the method of moving asymptotes [67], among many others. (See [68–70] for a comprehensive overview.)

We compare a few different methods in §4.1 in terms of computational performance and resulting design performance.

3. Performance limits

Any approximate optimization method for the physical design problem (6) can be used to generate approximately optimal designs. In other words, if we let p^* be the optimal value for (6), these procedures generate a design and field that satisfy the physics Eq. (1) whose objective value, say p , satisfies $p \geq p^*$. In general, because problem (6) is hard to solve, it is hard to know how far away our designs are from the true optimal value p^* . For example, once we have approximately optimized (6) and received some design with objective value p , it is not clear if this design is close to optimal (and no design can do significantly better) or if there are designs that have much better performance than the one we've found.

It is an old tradition in physics to then ask: what is the best possible value that we can hope to achieve? More specifically: is there some lower bound d such that we can guarantee that the optimal value of problem (6), p^* , is never smaller than this bound; *i.e.*, $p^* \geq d$? Such a bound can be interpreted in many ways. For example, it can be interpreted as an 'impossibility result', which states that no device that satisfies the physics Eq. (4) and the parameter constraints, $\theta \in \Theta$ can have objective value smaller than d . We can also interpret it as a 'certificate of optimality': given some design with objective value p , if p is close to d , then p must also be close to the optimal objective value, p^* , since $p \geq p^* \geq d$. Of course the best performance bound is p^* , but computing this is intractable, and we seek bounds that can be computed at reasonable cost.

Additionally, performance bounds can be very important in speeding up the design process. For example, it is often not clear how large a design needs to be in order to achieve reasonable performance. This often results in designers having to experiment with the total device size in order to find a design which has at least the desired performance. Lower bounds on these values, if they are efficiently computable, would give an indication of how large a design needs to be in order to achieve the designer's goals without additional (potentially very computationally expensive) experimentation.

Methods. Roughly speaking, there are two main approaches to the problem of finding lower bounds to the optimal objective value of problem (6). The first, and likely the earliest of the methods, is to make basic physical assumptions about the system (for example, that the system size is substantially smaller than the wavelength of the excitation [71,72]) and derive bounds on the corresponding quantities [73–77]. While these methods are historically important and yield good rule-of-thumb heuristics for design, many of the bounds derived in this way require assumptions that are not satisfied by the devices found by inverse design, or result in weak bounds. The second approach, which has become relatively popular recently (see, *e.g.*, [35,38,78–89]), essentially uses basic properties of the constraints and objective function of problem (6) to derive bounds on the best possible performance of the problem. Such approaches include algebraic manipulations of the physics Eq. (1) combined with the parameter constraints $\theta \in \Theta$, and applications of Lagrange duality to problem (6). The resulting bounds often do not have analytical forms, but can be numerically evaluated by an efficient algorithm and are therefore called *computational bounds*. We will discuss such bounds in this section.

3.1. Lagrange duality

The basic tool in a number of these bounds is the use of Lagrange duality. The idea is as follows. Given the optimization problem

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && h(x) \leq 0, \end{aligned} \tag{13}$$

for some objective function $f : \mathbf{R}^n \rightarrow \mathbf{R}$, constraint function $h : \mathbf{R}^n \rightarrow \mathbf{R}^m$, and optimization variable $x \in \mathbf{R}^n$, we form the *Lagrangian*:

$$L(x, \lambda) = f(x) + \lambda^T h(x),$$

where $\lambda \in \mathbf{R}^m$, with $\lambda \geq 0$ is a Lagrange multiplier or dual variable. This lets us define the *Lagrange dual function* $g : \mathbf{R}^m \rightarrow \mathbf{R}$, given by

$$g(\lambda) = \inf_x L(x, \lambda) = \inf_x (f(x) + \lambda^T h(x)).$$

See, e.g., [90, §5] for more information on Lagrange dual functions.

Lower bound property. The function g has a few interesting properties. First, for any $\lambda \geq 0$, $g(\lambda)$ is always smaller than p^* , the optimal value of (13), and is therefore a performance bound. More specifically, we have

$$g(\lambda) = \inf_x (f(x) + \lambda^T h(x)) \leq \inf_{h(x) \leq 0} (f(x) + \lambda^T h(x)) \leq \inf_{h(x) \leq 0} f(x) = p^*.$$

The first inequality follows from the fact that the set of x which satisfy $h(x) \leq 0$ is no larger than the set of all $x \in \mathbf{R}^n$, and the second follows from the fact that, because $h(x) \leq 0$ and $\lambda \geq 0$, then $\lambda^T h(x) \leq 0$. Thus the Lagrange dual function gives us a performance bound, parametrized by λ . (Depending on the problem and choice of λ , it can give the trivial lower bound $-\infty$.)

Concavity. Since $g(\lambda)$ is a performance bound for any $\lambda \geq 0$, it is then natural to ask, what is the best possible performance bound? In other words, what is the largest possible value of $g(\lambda)$ over the possible values of λ ? This problem is called the *dual problem* and can be written as

$$\begin{aligned} & \text{maximize} && g(\lambda) \\ & \text{subject to} && \lambda \geq 0. \end{aligned} \tag{14}$$

In general, evaluating $g(\lambda)$ at some $\lambda \geq 0$ is at least as hard as solving the original problem (13). On the other hand, when it is possible to efficiently evaluate $g(\lambda)$, it is almost always possible to efficiently find the optimal value of the dual problem (14) because the function g is always a concave function, even when the objective function f and constraints h in the original problem are not convex [90, §5.1.2].

Initializations. A solution to the dual problem (14) often suggests a good initialization for heuristics which attempt to minimize problem (13). Given some dual variable λ^* that is optimal for (14), there exists some x^0 which minimizes the Lagrangian at this choice of dual variable

$$x^0 \in \underset{x}{\operatorname{argmin}} L(x, \lambda^*),$$

under some basic assumptions on the objective function f and constraints h . In practice, x^0 is generally close to a reasonable design (see, e.g., [35]) even if it is not feasible; i.e., it need not satisfy $h(x^0) \leq 0$, except in some special scenarios such as when the functions f and h are both convex, in which case x^0 is globally optimal [90, §5.2]. Because it is often true that x^0 is easy to evaluate whenever $g(\lambda^*)$ is easy to evaluate, this initialization can be seen as a by-product of finding a solution to (14).

3.2. Local power conservation

The first approach to constructing bounds for (6) was presented originally in some generality in [78] and then [86] and later extended and fully clarified in [83,89,91] and subsequently fully generalized in [84] and [88] in the case where the parameters are Boolean (*i.e.*, $\Theta = \{-1, 1\}^n$). We will present a further generalization to the case where the parameters are box-constrained ($\Theta = [-1, 1]^n$), by a slightly different proof that considers a relaxation of the absolute-upper-bound formulation (10). (The Boolean case follows from an identical argument by considering (12), instead.) Bounds of this form are essentially power conservation laws over a given subdomain, which, in photonics, are often included under the name ‘optical theorem’ [92].

Power inequalities. Starting with the absolute-upper-bound formulation given in (10), we can square both sides of the inequality constraint to receive an equivalent formulation,

$$\begin{aligned} & \text{minimize} && f(z) \\ & \text{subject to} && (a_i^T z - (b_0)_i)^2 \leq z_i^2, \quad i = 1, \dots, n. \end{aligned} \quad (15)$$

In other words, a field z is feasible (*i.e.*, there exists a design $\theta \in \Theta$ such that z and θ satisfy the diagonal physics Eq. (5)) if, and only if the following quadratic inequalities are all satisfied:

$$(a_i^T z - (b_0)_i)^2 \leq z_i^2, \quad i = 1, \dots, n. \quad (16)$$

We can, of course, multiply these inequalities by a nonnegative value $\lambda_i \geq 0$ for $i = 1, \dots, n$ and add any number of them together to get another valid, quadratic inequality,

$$\sum_{i=1}^n \lambda_i (a_i^T z - (b_0)_i)^2 \leq \sum_{i=1}^n \lambda_i z_i^2,$$

where a_i^T is the i th row of A_0 . This can be more compactly expressed as

$$(A_0 z - b_0)^T D (A_0 z - b_0) \leq z^T D z, \quad (17)$$

where $D = \mathbf{diag}(\lambda)$ and has nonnegative entries along the diagonal.

The family of inequalities (17), parametrized by the nonnegative diagonal matrices D is the same as the family given in [84,88], except in the case where Θ specifies a box constraint, instead of a Boolean one. Additionally, because z satisfies (16) if, and only if, it satisfies (17) for all diagonal matrices D with nonnegative diagonals, then it follows that the family of quadratic inequalities in [84,88] is tight, in the following sense: any field z which satisfies these power conservation laws for all nonnegative diagonal matrices D must also have a corresponding design $\theta \in \Theta$ such that z and θ simultaneously satisfy the diagonal physics Eq. (5). (The inequalities in (17) are of a slightly different form than those presented in [84,88]. We show their equivalence in appendix 8.)

Relaxed formulation. Because any feasible field z satisfies (17) for any nonnegative diagonal matrix D , we can relax the family of inequalities (17) from all nonnegative diagonal matrices D to a finite number of them, which we will write as D_j for $j = 1, \dots, N$. The following problem can then be seen as a relaxation of (15):

$$\begin{aligned} & \text{minimize} && f(z) \\ & \text{subject to} && (A_0 z - b_0)^T D_j (A_0 z - b_0) \leq z^T D_j z, \quad j = 1, \dots, N, \end{aligned} \quad (18)$$

with the field $z \in \mathbf{R}^n$ as the only variable and problem data $A_0 \in \mathbf{R}^{n \times n}$, the excitation $b_0 \in \mathbf{R}^n$, and the diagonal matrices $D_j \in \mathbf{R}^{n \times n}$ with nonnegative entries along the diagonal, for $j = 1, \dots, N$.

Note that, if $D_j = E_{jj}$ for $j = 1, \dots, n$, where E_{jj} is the matrix with a single nonzero in the j, j entry (which is one) and is zero elsewhere, we recover the original problem (15). Additionally, while this problem is a relaxation of the original, it is still likely to be computationally hard to solve.

Quadratic objective. In general, finding lower bounds to the relaxed formulation (18) also need not be easy, except in some important special cases. For example, when the objective function f is a quadratic,

$$f(z) = \frac{1}{2}z^T Pz + q^T z + r,$$

for some symmetric matrix $P \in \mathbf{S}^n$, vector $q \in \mathbf{R}^n$, and $r \in \mathbf{R}$, then problem (18) is a quadratically constrained quadratic program (QCQP, see [93]), and the dual function g corresponding to this problem has a closed form solution. In fact, the dual problem of (18) is, in general, a convex semidefinite program (SDP) which can be solved in practice for moderate values of n and N [90, §1], allowing us to find a lower bound to (18) and therefore to (6) efficiently.

Dual problem. To find the dual problem, we first formulate the Lagrangian of (18):

$$\begin{aligned} L(z, \lambda) &= \frac{1}{2}z^T Pz + q^T z + r + \frac{1}{2} \sum_{j=1}^N \lambda_j ((A_0 z - b_0)^T D_j (A_0 z - b_0) - z^T D_j z) \\ &= \frac{1}{2}z^T T(\lambda)z + v(\lambda)^T z + u(\lambda), \end{aligned}$$

where $\lambda \in \mathbf{R}_+^N$ and we have defined

$$T(\lambda) = P + \sum_{j=1}^N \lambda_j (A_0^T D_j A_0 - D_j), \quad v(\lambda) = q - \sum_{j=1}^N \lambda_j (A_0^T D_j b_0), \quad u(\lambda) = \frac{1}{2} \sum_{j=1}^N \lambda_j b_0^T D_j b_0 + r,$$

such that $T(\lambda) \in \mathbf{S}^n$ and $v(\lambda) \in \mathbf{R}^n$, for notational convenience. (This follows from expanding the expression and collecting the quadratic, linear, and constant terms in the variable z .)

It is not hard to show that the dual function is, for $\lambda \geq 0$,

$$g(\lambda) = \inf_z L(y, \lambda) = \begin{cases} u(\lambda) - \frac{1}{2}v(\lambda)^T T(\lambda)^+ v(\lambda) & T(\lambda) \geq 0, \quad v(\lambda) \in \mathcal{R}(T(\lambda)) \\ -\infty & \text{otherwise.} \end{cases}$$

(See, e.g., [93, §3.2].) Here $T(\lambda)^+$ is the Moore-Penrose pseudoinverse [94, §11.5] of $T(\lambda)$, while $T(\lambda) \geq 0$ means that $T(\lambda)$ is positive semidefinite, and $\mathcal{R}(T(\lambda))$ is the range of $T(\lambda)$. The corresponding problem of maximizing g over $\lambda \in \mathbf{R}_+^N$ can be written as a standard form SDP, which we will call the *power dual bound*:

$$\begin{aligned} &\text{maximize} && u(\lambda) - (1/2)t \\ &\text{subject to} && \begin{bmatrix} t & v(\lambda)^T \\ v(\lambda) & T(\lambda) \end{bmatrix} \geq 0 \\ &&& \lambda \geq 0, \end{aligned} \tag{19}$$

with variables $t \in \mathbf{R}$ and $\lambda \in \mathbf{R}^N$. This problem can be easily specified using domain-specific languages such as CVXPY or JuMP.jl and solved using convex solvers that support SDPs, such as Mosek [95], SCS [96], or COSMO.jl [97]. Because of the sparsity of A_0 , it is often the case that such problems have chordal structure [98], which can be exploited to efficiently solve problem (19) by some solvers such as COSMO.jl.

Optimal choice of D_j matrices. Given the dual problem (19) of problem (18), the question remains of how to best choose the diagonal matrices D_j for $j = 1, \dots, N$ such that the optimal

value of (19) is maximized. Historically, the initial bounds in [78,83,86,89,91] assumed a number of fixed diagonal matrices D_j . Later, [84] and [88] generalized the approach to include any diagonal matrix D and [88] proposed an iterative algorithm which, starting with some diagonal matrix D_1 (such as $D_1 = I$) would solve a problem similar to (19) at iteration k and propose a new diagonal matrix D_{k+1} that would be appended to the constraints of (18). This method is conceptually similar to the cutting-plane method described in [93, §3.5] applied to the dual problem (19) with $D_j = E_{jj}$ for $j = 1, \dots, n$. (More accurately, the procedure proposed in [88] solves an SDP relaxation of (18) instead of the dual problem (19), but it can be shown that both problems have the same optimal value by strong duality [93, §3.3].)

It is also reasonable to ask: what are the best possible choices of D_j such that the optimal value of (19) is maximized? It is not hard to show that a single (correctly chosen) diagonal matrix, $D = \mathbf{diag}(\lambda)$, suffices and that this matrix can be efficiently found. To see this, note that we can choose $D_j = E_{jj}$ for $j = 1, \dots, n$ such that,

$$D = \sum_{j=1}^n \lambda_j D_j = \mathbf{diag}(\lambda). \quad (20)$$

This would let us write

$$T(\lambda) = P + A^T \mathbf{diag}(\lambda)A - \mathbf{diag}(\lambda), \quad v(\lambda) = q - A^T \mathbf{diag}(\lambda)b, \quad u(\lambda) = \frac{1}{2}b^T \mathbf{diag}(\lambda)b + r,$$

We then note that picking λ^* optimal for (19), when it exists, gives a diagonal matrix $D^* = \mathbf{diag}(\lambda^*)$. Additionally, solving (18) with this choice of matrix D^* is a special case where the number of quadratic constraints, N , equals 1, which can be efficiently solved and has the same optimal value as (19). (See, e.g., appendix B of [90].)

3.3. Diagonal physics dual

Another approach to computing lower bounds for the diagonal physical design problem (6) is by a direct application of Lagrange duality, originally given in [35] and later extended in [82] and [99]. This approach gives a lower bound when the objective function is separable:

$$f(z) = \sum_{i=1}^n f_i(z_i),$$

and the constraints are of the following form:

$$\Theta = \{\theta \mid -\mathbf{1} \leq \theta \leq \mathbf{1}, \theta_i = \theta_j \text{ for } i, j \in S_k, k = 1, \dots, K\}, \quad (21)$$

where $S_k \subseteq \{1, \dots, d\}$ for $k = 1, \dots, K$ are disjoint sets, specifying which entries of θ are constrained to be equal. We additionally note that a lower bound for this constraint set also yields a lower bound for the Boolean case, since this constraint set contains the Boolean one.

Problem Lagrangian. The basic trick here is to rewrite problem (6) as the following (equivalent) problem:

$$\begin{aligned} &\text{minimize} && \sum_{i=1}^n f_i(z_i) + I(\theta) \\ &\text{subject to} && (A_0 + \mathbf{diag}(\theta))z = b_0, \end{aligned} \quad (22)$$

where we have pulled out the constraint $\theta \in \Theta$ into the indicator function $I : \mathbf{R}^n \rightarrow \mathbf{R} \cup \{+\infty\}$ for the set Θ ,

$$I(\theta) = \begin{cases} 0 & \theta \in \Theta \\ +\infty & \theta \notin \Theta. \end{cases}$$

We can then easily formulate the Lagrangian of this problem:

$$L(z, \theta, \nu) = \sum_{i=1}^n f_i(z_i) + I(\theta) + \nu^T (A_0 + \mathbf{diag}(\theta))z - \nu^T b_0,$$

which we note is separable in terms of z :

$$L(z, \theta, \nu) = \sum_{i=1}^n \left(f_i(z_i) + (a_i^T \nu)z_i + \nu_i \theta_i z_i \right) + I(\theta) - \nu^T b_0$$

Dual function. As before, the dual function is defined as

$$g(\nu) = \inf_{\theta} \inf_z L(z, \theta, \nu).$$

Computing the inner infimum is relatively simple, which gives:

$$\begin{aligned} \inf_z L(z, \theta, \nu) &= \sum_{i=1}^n \inf_{z_i} \left(f_i(z_i) + (a_i^T \nu)z_i + \nu_i \theta_i z_i \right) + I(\theta) - \nu^T b_0 \\ &= - \sum_{i=1}^n f_i^*(-a_i^T \nu - \nu_i \theta_i) + I(\theta) - \nu^T b_0. \end{aligned}$$

Here, $f_i^* : \mathbf{R} \rightarrow \mathbf{R} \cup \{+\infty\}$ is the *convex conjugate* of f_i (sometimes called the Fenchel conjugate function or simply the conjugate function) defined as

$$f_i^*(u) = \sup_x (ux - f_i(x)),$$

and is well known for a number of functions [90, §3.3]. Additionally, we will make use of the fact that f_i^* is always a convex function, even when f_i is not convex.

To compute the outer infimum, we note that we can write

$$\begin{aligned} g(\nu) &= \inf_{\theta} \left(- \sum_{i=1}^n f_i^*(-a_i^T \nu - \nu_i \theta_i) + I(\theta) - \nu^T b_0 \right) \\ &= \inf_{\theta} \left(- \sum_{k=1}^K \sum_{i \in S_k} f_i^*(-a_i^T \nu - \nu_i \theta_k) + I(\theta) \right) - \nu^T b_0 \\ &= - \sum_{k=1}^K \left(\sup_{-1 \leq \theta_k \leq 1} \sum_{i \in S_k} f_i^*(-a_i^T \nu - \nu_i \theta_k) \right) - \nu^T b_0 \\ &= - \sum_{k=1}^K \max \left\{ \sum_{i \in S_k} f_i^*(-a_i^T \nu + \nu_i), \sum_{i \in S_k} f_i^*(-a_i^T \nu - \nu_i) \right\} - \nu^T b_0, \end{aligned}$$

where we have used the fact that $-\inf_x h(x) = \sup_x -h(x)$ for any function h , and the fact that a scalar convex function achieves its maximum over an interval at the boundary of that interval.

Dual problem. Given the dual function g , we can find a lower bound to the original problem by evaluating g for any $\nu \in \mathbf{R}^n$. We can then ask what's the best possible dual bound, which gives the following dual problem, which we will call the *diagonal dual bound*:

$$\text{maximize} \quad - \sum_{k=1}^K \max \left\{ \sum_{i \in S_k} f_i^*(-a_i^T \nu + \nu_i), \sum_{i \in S_k} f_i^*(-a_i^T \nu - \nu_i) \right\} - \nu^T b_0. \quad (23)$$

This problem is a convex optimization problem with variable $\nu \in \mathbf{R}^n$, whose optimal value, d^* , can almost always be efficiently found whenever the function f_i^* can be efficiently evaluated.

Field bounds. In some cases, the bounds given by (23) can sometimes be very weak; *e.g.*, if $f_i = 0$ for all indices i except one. One way of improving the lower bounds is to add a redundant constraint to (22); *i.e.*, a constraint such that, if z satisfies $(A_0 + \mathbf{diag}(\theta))z = b_0$, then it also satisfies $h(z) \leq 0$ for some function $h : \mathbf{R}^n \rightarrow \mathbf{R}$, such that the dual function of the resulting problem is still simple to evaluate.

One useful example, originally presented in [82], is to note that, if z satisfies

$$(A_0 + \mathbf{diag}(\theta))z = b_0,$$

then it also satisfies,

$$\|A_0 z\|_2 - \|\mathbf{diag}(\theta)z\|_2 \leq \|(A_0 + \mathbf{diag}(\theta))z\|_2 = \|b_0\|_2,$$

where we have taken the norm of both sides of the expression, and the inequality follows from the triangle inequality. Using the fact that

$$\|A_0 z\|_2 \geq \sigma_1(A_0)\|z\|_2, \quad \text{and} \quad \|\mathbf{diag}(\theta)z\|_2 \leq \|z\|_2,$$

where $\sigma_1(A_0)$ is the smallest singular value of A_0 , we find

$$\sigma_1(A_0)\|z\|_2 - \|z\|_2 \leq \|b_0\|_2,$$

or, after some rearrangement:

$$\|z\|_2 \leq \frac{\|b_0\|_2}{\sigma_1(A_0) - 1},$$

whenever $\sigma_1(A_0) > 1$. Evaluating the dual for this new problem is a simple extension of the procedure given above as in [35, §6]. In fact, the same procedure can be extended to any norm $\|\cdot\|$ by replacing $\sigma_1(A_0)$ with $1/\|A_0^{-1}\|$, where $\|A_0^{-1}\|$ is the induced operator norm of A_0^{-1} :

$$\|A_0^{-1}\| = \sup_{\|x\|=1} \|A_0^{-1}x\|,$$

but the resulting dual function need not be easy to evaluate.

Mode volume. Another extension, presented originally in [99] is for an objective function f of the form

$$f(z) = \frac{\|z\|_2^2}{z_i^2},$$

whenever $z_i \neq 0$ and is $+\infty$ otherwise. Here, i is a fixed index and we will assume there is no excitation; *i.e.*, $b_0 = 0$. Note that this objective function is similar to, and can be easily extended to include, the cavity mode volume. Because $b_0 = 0$, then the physics equation and objective are 0-homogeneous in z , so any feasible point z with $z_i \neq 0$ can be scaled by a nonzero value $\eta \in \mathbf{R}$, such that ηz is also feasible with the same objective value, $f(\eta z) = f(z)$. We can then fix a normalization by setting $z_i = 1$ to get an equivalent problem:

$$\begin{aligned} & \text{minimize} && \|z\|_2^2 \\ & \text{subject to} && (A_0 + \mathbf{diag}(\theta))z = 0 \\ & && z_i = 1 \\ & && -\mathbf{1} \leq \theta \leq \mathbf{1}, \end{aligned} \tag{24}$$

with variables z and θ and some fixed index i . The dual for problem (24) can be computed using the same method presented in this section.

4. Numerical examples

In this section, we show a basic comparison between some of the heuristics presented in §2.3 and the performance bounds presented in §3 on a small problem of designing a Helmholtz resonator in one and two dimensions. In these cases, we can certify that the heuristics find a device whose performance is at most 2% above the global optimum, even though the original problem is likely hard. We then show examples in the literature of much larger devices for which no lower bound has been computed (and, indeed, cannot be computed with the current methods in a reasonable amount of time), but has resulted in practically useful devices whose performance is much better than that of traditional designs.

4.1. Small examples

In this section, we will compare the performance of three heuristics, sign-flip descent (SFD), IPOPT, and genetic algorithms (GA), and the lower bounds presented in §3. We will first compare their respective performance on a small, one-dimensional design, where both the power lower bounds and the dual lower bounds can be computed in a reasonable amount of time, and then compare their performance on a larger two-dimensional design. The times reported in this section are from a dual-core 2015 MacBook Pro laptop running at 2.9GHz with 8GB of RAM. At the time of writing, this machine is 5 years old and is roughly representative of the computational power available in more recent standard and lower-end consumer laptops.

Problem formulation. In both scenarios, our objective function f will be to best match a desired field \hat{z} ,

$$f(z) = \|z - \hat{z}\|_2^2,$$

where the desired field is a cosine wave with a Gaussian envelope on the left half of the domain, and is equal to zero on the right half. The excitation is a single delta function in the center of the domain. In this problem, the designer is then allowed to choose the speed of the wave at each point in space (via the design parameters θ), such that the objective function is minimized. (For more details and code, see the appendix 9.)

1D problem. We compare the objective performance and computational performance of the heuristics and bounds in Table 1. We also plot the corresponding fields and the desired field in Fig. 1. While GA gives rather poor solutions (which only somewhat match the largest features of the desired field), IPOPT and SFD give approximately-optimal fields that are essentially indistinguishable from the desired one. We find, at least in this small scenario, that IPOPT and SFD have objective values that are extremely close to that of the power bounds and diagonal dual bounds, which means that the designs found in this scenario can, for all intents and purposes, be considered globally optimal. We also note that GA, while simple to implement, does not find a good solution even with some amount of tuning, while also having the worst performance of the available heuristics in terms of total time taken. Additionally, while the power bound was slightly tighter than the diagonal dual lower bound, it took nearly 90 times longer to converge for this small problem. This is due to the fact that SDPs solution times scale approximately cubically with the problem dimension (*i.e.*, are $O(n^3)$), which quickly becomes an issue for larger problems.

Table 1. Performance results for small design.

Algorithm	Objective value	Time (s)
Genetic algorithm	2.54	6.40
IPOPT	.652	1.70
Sign-flip descent	.642	.592
Power dual bound	.639	125
Diagonal dual bound	.634	1.39




Fig. 1. Approximate designs and desired field for 1D problem.

In fact, we note that it was difficult to find a desired field \hat{z} where the bounds and the performance of designs found by SFD differed significantly. We encourage readers to search for some cases where this is true by trying a few different desired fields in the code available for this paper. For more details, see appendix B.1.

2D problem. In the 2D problem, we again test the performance of GA, IPOPT, and SFD, and the diagonal design lower bound, though we do not compute the power bounds. We note that naïvely attempting to compute the power bounds in (19) by framing the problem as an SDP (using an SDP solver that does not support chordal sparsity) results in a dense matrix variable of size $(251^2)^2/2 = 251^4/2$, which requires approximately 16GB of memory to store, nearly double the available memory (8GB), not counting the additional memory required to perform operations on this matrix.

The results of this comparison are available in Table 2 and the fields of the approximately optimized designs are shown in Fig. 2. We terminated any algorithm whose runtime was longer than 30 minutes on the current computer. We note that, again, SFD has surprisingly good performance, and, when combined with the lower bound, yields a design that is guaranteed to be no more than $(11.9/11.7 - 1) \approx 1.7\%$ suboptimal, relative to the globally optimal value (in fact, as shown in Fig. 2, the resulting field from SFD and the desired field, \hat{z} , are difficult to visually distinguish). Additionally, solving the diagonal dual bound was faster than getting an approximate design from any of the heuristics, though we note that this is likely not the case with much larger designs or better-optimized heuristics. For more details, see appendix B.2.




Fig. 2. Approximate designs and desired field for 2D problem.

Table 2. Performance results for larger design.

Algorithm	Objective value	Time (s)
Genetic algorithm	N/A	>1800
IPOPT	190.71	1360
Sign-flip descent	11.9	364
Diagonal dual bound	11.7	48.5

Discussion. The examples in this section show that modern heuristics have very good practical performance when compared to the available lower bounds. In fact, we suspect that this is true more generally: modern heuristics with reasonable initializations likely return designs that are very good, if not globally optimal, even when there exist no bounds that can certify this, or when the available bounds are very weak. Additionally, while we have made some basic performance optimizations to the available code, we opted for clarity as opposed to pure performance in the implementation of the algorithms presented, so the numbers above should only be interpreted as general guidelines. The code to generate the plots and tables is available at <https://github.com/cvxgrp/pd-heuristics-and-bounds>.

4.2. Practical examples

In this section, we show practical examples in the literature where the devices found by heuristics have been fabricated and experimentally verified. While the current bounds cannot be used to certify that the performance of these designs is close to globally optimal in a reasonable amount of time, the resulting designs have much better performance than that of traditional designs. For a comprehensive overview of the history of inverse design and its applications in practice, including older literature, we refer the reader to [20].

Splitters. Perhaps the most striking applications of physical design is in the design of compact *splitters*—devices which, given some input in a specific scenario (for example, given an input at a specific frequency) must direct as much of the input as possible into a desired output location. Different scenarios would direct the input to different locations; *i.e.*, the input is ‘split’ to different outputs depending on the scenario. Some examples of such devices and their performance can be found in, *e.g.*, [37,100,101].

Figure 3 shows an example of the splitter designed and fabricated in [102]. In this figure, three different inputs are represented by the three different colors (blue, red, and green), which are initially ‘mixed’ in the input on the left-hand-side of the domain. The device then separates the wavelengths into each of the three output channels.




Fig. 3. Scanning-electron microscope image of a three-way wavelength splitter with simulated fields overlaid. Figure and design from [102].

Particle accelerators. There are other types of devices which have been designed and manufactured in practice. A recent demonstration of a laser-driven particle accelerator small enough to be placed on a chip [103] required several components to be efficient enough that the input power did not destroy the material out of which the components were made. The necessary components included the devices which coupled the laser to the structure, along with the actual accelerator. The resulting designs are unintuitive enough that it is not clear a human could manually design a device of the same (or smaller) size that had at least the same efficiency. The accelerator from the final device of [103] is shown in Fig. 4.




Fig. 4. Scanning-electron microscope image of an inverse-designed laser-driven particle accelerator with simulated fields overlaid. Figure and design from [103].

Lenses. There have been a few other types of devices that have been designed by these methods and created in practice. Some examples include flat lenses with a large depth of field that have the same focusing efficiency as traditional lenses [39], and flat lenses with an ultra-wide field of view [104]. There has been some additional work in creating deformable lenses whose focal length can be controlled by stretching or contracting the material, with performances exceeding those of traditional lenses [16].

Fabrication constraints. We note that, while there have been many numerical demonstrations of inverse designed photonic devices, the actual fabrication of such devices has been relatively difficult until recently, when methods that could include fabrication constraints [36,105] and could accommodate a large enough number of field variables became available and were fast enough to be used in practice. Additionally, fabrication constraints can also be applied to mass-produced photonic circuits in a foundry [106], which makes it possible to fabricate inverse-designed photonics at scale.

5. Future directions

While physical design and, more specifically, inverse design, has led to some dramatic improvements in the performance of photonic devices, there are still many questions left unanswered and possible future avenues for research.

Standard benchmarks. There is a need for standardized benchmarks for both heuristic algorithms and bounds. Similar to the standard performance benchmarks for machine learning (such as MNIST [107], ImageNet [108], CIFAR-10 [109], *etc.*) and for optimization (such as the Maros–Mészáros test set [110], MIPLIB 2017 [111], *etc.*), which are used to compare the performance of different proposed algorithms, it is feasible to have a standard library of design specifications (objective functions and constraints) which either have known global solutions or a best known solution and best known lower bound that is updated as better ones are found. This would allow for researchers to have a concrete library which can be used to compare proposed algorithms against existing ones, in terms of both objective value and computational performance.

Like in machine learning or optimization, we suspect that such benchmarks would clarify the performance tradeoffs of different algorithms and bounds. These benchmarks would help identify what problems are ‘hard’ for current heuristics and possibly lead to better approaches.

Improved heuristics. While some of the methods presented here can scale to very large designs with millions to tens of millions of field and tens of thousands of design variables, many of the current methods do not make use of the specific structure of the physical design problem and take a long time (days or weeks) even with large computers or clusters. For example, one possible avenue is the use of a primal-dual algorithm, which can lead to drastically improved performance for some problems (see, *e.g.*, [112, §3.5]). Such methods could also lead to algorithms whose natural byproduct of computing an approximately optimal design is a bound, like in §3, guaranteeing that the resulting design is close to the global optimum without requiring additional computational time. There has been additional work with combining machine learning approaches to speeding up both simulations [113] and optimization [114–117], which trade off training time for runtime performance.

Improved initializations. All of the heuristics used in practice often require good initializations in order to have reasonable performance. In practice, initializations that are guided by intuition appear to work well for many scenarios, but sometimes fail to reach what are known to be better designs. (See, *e.g.*, [34, §5].) As previously discussed, the bounds presented yield good initial designs as a by-product of the optimization procedure; but this need not be true in general, and we suspect that there might exist better methods which, while expensive to run exactly, might yield good initial designs. For example, while sign-flip descent might be potentially very expensive to run for large designs (as it requires solving a general constrained convex problem at each iteration) one could perform a small number of iterations to get a reasonably good initial field and then feed this resulting field into an optimization algorithm with faster convergence. (This latter procedure is sometimes called ‘solution polishing’ in combinatorial optimization.) Another avenue of similar work is the idea of ‘objective-first’ optimization [118] where the physics equation is relaxed to a penalty with a small weight and the resulting nonconvex problem is solved to get an infeasible design with good objective performance and (hopefully) small physics residual. This initial, infeasible design is then passed to any of the heuristics above to attempt to find a feasible design and field with similar performance.

Improved bounds. The small examples we have discussed here are simple cases where the bounds and the heuristics are very close; but this need not be the case. For example, the diagonal dual bounds given here can give weak bounds when the objective function depends only on a few entries of the field variable z , even with the additional extensions presented. While we, in general, suspect the power bound yields tighter results than the diagonal dual lower bound in these cases, current off-the-shelf solvers are too slow to solve the resulting problem other than for a small number of field variables, $n \lesssim 10^3$. One future research avenue is to create faster solvers for problem (19), by exploiting the specific structure available in these problems. A second possibility is to find some connection between the diagonal dual and power bounds; in particular, it is not clear how one is related to the other, if at all. We conjecture that the power bounds (19) are always guaranteed to be at least as tight as the diagonal dual ones (23) whenever the function f is a separable convex quadratic, and where the constraint sets S_k are singletons (Eq. (21) with $S_k = \{k\}$ for $k = 1, \dots, n$) but have been unable to prove this. It is also of practical interest to create lower bounds which give approximate scaling rules for designs, which would help a designer choose appropriate device sizes and materials for a desired objective.

6. Conclusion

There has been tremendous progress in the area of photonic inverse design, including foundry-based fabrication and the use of design software in a wide range of academic and industry labs. Still, there is a lot to be improved, including, but not limited to: designing standard benchmarks

and comparisons for heuristics, overcoming computational bottlenecks in order to design larger devices than those which are currently possible, and the improvement of bounds (both theoretically and in terms of computational performance) which currently either do not apply, or are very difficult to compute, for many of the devices used in practice. Yet, despite the drawbacks of the current approaches, inverse design has yielded incredible practical benefits, and we expect that such methods will, in the near future, become a standard approach to creating practical, efficient devices that far exceed the performance of their traditionally-designed counterparts.

Appendix A. Bound equivalence

We show that the bounds presented in [84,88] are equivalent to the ones given in section 3.2.

Derivation. First, we derive the bounds presented using this notation. We define the displacement field $y_i = \gamma_i z_i$, for $i = 1, \dots, n$, where $\gamma_i = (1 + \theta_i)/2$. (Note that $\theta_i \in [-1, 1]$ whenever $\gamma_i \in [0, 1]$.) The diagonal physics equation can then be written in terms of the field z and the displacement field y :

$$(2A_0 - I)z + y = 2b_0, \quad y_i = \gamma_i z_i, \quad i = 1, \dots, n,$$

where $\gamma \in [0, 1]^n$. We will define $G = (2A_0 - I)^{-1}$ (this can be recognized, roughly speaking, as the Green's function) and $b' = 2Gb_0$ (which is sometimes called the 'incident field') for notational convenience. We can then rewrite the diagonal physics equation using G and b' :

$$z_i + g_i^T y = b'_i, \quad y_i = \gamma_i z_i, \quad i = 1, \dots, n, \quad (25)$$

where g_i^T denotes the i th row of G . Multiplying on the left by y_i , we find that y and z must satisfy

$$y_i(z_i + g_i^T y) = y_i b'_i, \quad i = 1, \dots, n.$$

Note that $y_i^2 = \gamma_i^2 z_i^2 \leq \gamma_i z_i^2 = y_i z_i$ because $\gamma_i^2 \leq \gamma_i$, since $\gamma_i \in [0, 1]$. Using this result, we find the following quadratic inequalities depending only on the displacement field y :

$$y_i(y_i + g_i^T y) \leq y_i b'_i, \quad i = 1, \dots, n.$$

Scaling each of the i inequalities by any nonnegative value $\lambda_i \geq 0$ and summing them together implies that y must satisfy:

$$y^T D y + y^T D G y \leq y^T D b', \quad (26)$$

where $D = \text{diag}(\lambda)$ is any matrix with nonnegative diagonal entries. These are the bounds presented in [84,88] in the case where Θ is a box rather than a Boolean constraint.

Tightness. As expected, these bounds are also tight in the sense that, if y satisfies (26) for all nonnegative diagonal matrices D , then there exists a feasible design θ and a field z such that z and θ satisfy the diagonal physical equation, and the displacement field y satisfies $y_i = (1 + \theta_i)z_i/2$ for $-1 \leq \theta_i \leq 1$, or, equivalently, that $y_i = \gamma_i z_i$, where $0 \leq \gamma_i \leq 1$ for $i = 1, \dots, n$.

To show this, we will consider (as before) only the diagonal matrices $D = e_i e_i^T$. The bound then implies that

$$y_i(y_i + g_i^T y) \leq y_i b'_i, \quad i = 1, \dots, n.$$

We can then choose $z_i = b'_i - g_i^T y$ and

$$\gamma_i = \begin{cases} y_i / (b'_i - g_i^T y) & b'_i - g_i^T y \neq 0 \\ 0 & \text{otherwise,} \end{cases}$$

for $i = 1, \dots, n$. Note that z_i obviously satisfies $z_i + g_i^T y = b'_i$ and whenever $b'_i - g_i^T y \neq 0$, we have $y_i = \gamma_i z_i$, by construction. On the other hand, if $b'_i - g_i^T y = 0$, then

$$y_i^2 = y_i(y_i + g_i^T y - b'_i) \leq 0,$$

so $y_i = 0 = \gamma_i z_i$, since $\gamma_i = 0$ whenever $b'_i - g_i^T y = 0$, by definition. In other words, given a displacement field y satisfying (26) for all nonnegative diagonal matrices, we have constructed a field z and a design $\theta = 2\gamma - 1$ such that y , z , and $-1 \leq \theta \leq 1$ all satisfy the physics Eq. (25), or, equivalently, z and θ satisfy the diagonal physics equation.

Equivalence. The equivalence between the bounds follows immediately from the fact that (26) holds for all nonnegative diagonal matrices D if, and only if, the physics Eq. (25) holds for this choice of y , z , and θ , which, in turn, holds if, and only if, the original family of power bounds over z and θ holds for all diagonal matrices.

Appendix B. Numerical examples

In this section, we focus on two basic numerical experiments for comparing both heuristics and lower bounds. Additionally, all Julia [119] code for the examples can be found at <https://github.com/cvxgrp/pd-heuristics-and-bounds>. The optimization problems used in computing the bounds were specified using the JuMP [120] domain-specific language in Julia, and solved using Mosek [95].

Appendix B.1. Small design

In this scenario, we are given a device whose field must satisfy the discretized Helmholtz equation in one dimension. At every point in the domain, the designer is allowed to choose the speed of the wave in the material within some specified range, such that the resulting field best matches some desired field.

Helmholtz's equation. In one dimension, we can write Helmholtz's equation in the interval $[-1, 1]$ as:

$$\frac{\partial^2 u(x)}{\partial x^2} + \frac{\omega^2}{c(x)^2} u(x) = v(x), \quad -1 \leq x \leq 1, \quad (27)$$

where $u : [-1, 1] \rightarrow \mathbf{R}$ is the amplitude of the wave, while $\omega \in \mathbf{R}$ is the angular frequency and $c : [-1, 1] \rightarrow \mathbf{R}_{++}$ is the speed of the wave in the material, and $v : [-1, 1] \rightarrow \mathbf{R}$ is the excitation. We will assume Dirichlet boundary conditions for simplicity; *i.e.*, that $u(-1) = u(1) = 0$.

Discretization. We can write a simple discretization of the above equation by subdividing the interval $[-1, 1]$ into n equidistant points $\{x_i\}$ for $i = 1, \dots, n$. We then approximate the second derivative using finite-differences such that

$$\frac{\partial^2 u(x_i)}{\partial x^2} \approx \frac{z_{i-1} - 2z_i + z_{i+1}}{h^2}, \quad i = 1, \dots, n,$$

where h is the separation $x_i - x_{i-1}$ for any $i = 2, \dots, n$ and is equal to $h = 2/(n - 1)$, while $z_i = u(x_i)$ for $i = 1, \dots, n$. Additionally, we have defined $z_0 = z_{n+1} = 0$ for convenience.

If we then define $\omega^2/c(x_i)^2 = \theta_i$, then we can approximate Eq. (27) as a diagonal physics equation of the form (5), where

$$A_0 = \frac{1}{h^2} \begin{bmatrix} -2 & 1 & 0 & \dots & 0 \\ 1 & -2 & 1 & \dots & 0 \\ 0 & 1 & -2 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & -2 \end{bmatrix}, \quad (28)$$

and $(b_0)_i = v(x_i)$ for $i = 1, \dots, n$.

Problem data. In this problem, we will seek to minimize the squared Euclidean distance between the field z and some desired field \hat{z} ; *i.e.*, the objective is

$$f(z) = \|z - \hat{z}\|_2^2 = \sum_{i=1}^n (z_i - \hat{z}_i)^2,$$

and we choose \hat{z} to be a cosine wave with a Gaussian envelope of width $\sigma/\sqrt{2}$, whenever $x_i < 0$ and 0 when $x_i \geq 0$:

$$\hat{z}_i = \begin{cases} \cos(\omega x_i) \exp(-x_i^2/\sigma^2) & x_i < 0 \\ 0 & x_i \geq 0. \end{cases} \quad (29)$$

Note that this function is discontinuous at 0. For the small numerical experiment, we will choose $\sigma = 1/2$, $\omega = 6\pi$, and $\Theta = [1, 1.5]^n$. We will also set $b_{(n+1)/2} = 2$ and zero otherwise as our excitation, and set n , the number of gridpoints, to be $n = 1001$. This means that there are 1001 design parameters and field variables, which, in practice, is a very small number when compared to current applications.

Appendix B.2. Larger design

In this example, we will show an example of a larger physical design problem with $n = d = 251^2 = 63001$, that is similar in spirit to the smaller design, but is large enough that sparsity has to be exploited in order to have reasonable run time performance.

2D Helmholtz equation. For this problem, we discretize the 2D Helmholtz's equation:

$$\frac{\partial^2 u(x, y)}{\partial x^2} + \frac{\partial^2 u(x, y)}{\partial y^2} + \frac{\omega^2}{c(x, y)^2} u(x, y) = v(x, y),$$

over the domain $(x, y) \in [-1, 1]^2$. Here, as before $u : [-1, 1]^2 \rightarrow \mathbf{R}$ is the amplitude of the wave, while $\omega \in \mathbf{R}$ is the angular frequency, $c : [-1, 1]^2 \rightarrow \mathbf{R}_{++}$ is the speed of the wave in the material, and $v : [-1, 1]^2 \rightarrow \mathbf{R}$ is the excitation. We assume Dirichlet boundary conditions; *i.e.*, that $u(x, y) = 0$ at the boundary of the domain.

Discretization. Assuming that we have $n = l^2$ equally spaced gridpoints $(x_i, y_i) \in [-1, 1]^2$ (*i.e.*, there are l points along a given axis), we will let z_i be the approximate value of $u(x_i, y_i)$. Writing the second-order difference matrix in (28) as $\Delta \in \mathbf{R}^{l \times l}$, we can approximate the sum of the partial derivatives as

$$\frac{\partial^2 u(x_i, y_i)}{\partial x^2} + \frac{\partial^2 u(x_i, y_i)}{\partial y^2} \approx ((I_l \otimes \Delta + \Delta \otimes I_l)z)_i, \quad i = 1, \dots, n,$$

where I_l is the $l \times l$ identity matrix, and \otimes is the Kronecker product. The resulting problem is then a diagonal physical design problem with

$$A_0 = I_l \otimes \Delta + \Delta \otimes I_l,$$

and $(b_0)_i = v(x_i, y_i)$ for $i = 1, \dots, n$. Note that the size of A_0 is $n \times n = l^2 \times l^2$; so the total number of entries of the matrix grows in the fourth power of the side length. On the other hand, the resulting matrix is very sparse because less than $5n$ of the entries are not zero. In other words, when $l = 100$, the percentage of nonzero entries is less than .05%.

Problem data. As before, the objective is to minimize the squared difference of the resulting field z and some desired field \hat{z} :

$$f(z) = \|z - \hat{z}\|_2^2.$$

Here \hat{z} is the two-dimensional analogue of (29),

$$\hat{z}_i = \begin{cases} \cos(\omega x_i) \cos(\omega y_i) \exp(-(x_i^2 + y_i^2)/\sigma^2) & x_i < 0 \\ 0 & x_i \geq 0, \end{cases}$$

for $i = 1, \dots, n$. We will again choose $\Theta = [1, 1.5]^n$, $\omega = 6\pi$ and $\sigma = 1/2$, and $b_{(n^2+1)/2} = 1$, along with $l = 251$ (and $n = l^2 = 65001$), which results in a problem with 65001 design parameters and field variables. We encourage the reader to try different choices parameters in the provided code and explore the resulting heuristic performance, bounds, and time taken by the algorithms.

Funding. Advanced Research Projects Agency - Energy (DE-AR0001212); National Science Foundation (DGE-1656518).

Acknowledgements. The authors would like to thank Rahul Trivedi for useful discussions regarding the current literature on bounds and Carlos G. Gonzalez for comments and corrections. The authors would also like to acknowledge financial support from ARPA-E with Agreement No. DE-AR0001212. Guillermo Angeris is supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1656518.

Disclosures. The authors declare that there are no conflicts of interest related to this article.

References

1. K. Y. Yang, J. Skarda, M. Cotrufo, A. Dutt, G. H. Ahn, M. Sawaby, D. Vercruyse, A. Arbabian, S. Fan, A. Alù, and J. Vučković, "Inverse-designed non-reciprocal pulse router for chip-based LiDAR," *Nat. Photonics* **14**(6), 369–374 (2020).
2. L. A. Kamensky and C. N. Liu, "Computer-Automated Design of Multifont Print Recognition Logic," *IBM J. Res. Dev.* **7**(1), 2–13 (1963).
3. A. Jameson, "Aerodynamic design via control theory," *J. Sci. Comput.* **3**(3), 233–260 (1988).
4. A. Jameson, L. Martinelli, and N. Pierce, "Optimum Aerodynamic Design Using the Navier-Stokes Equations," *Theor. Comput. Fluid Dyn.* **10**(1-4), 213–237 (1998).
5. A. Lee, M. Ahmadi, G. Jullien, W. Miller, and R. Lashkari, "Digital filter design using genetic algorithm," in *1998 IEEE Symposium on Advances in Digital Filtering and Signal Processing. Symposium Proceedings (Cat. No. 98EX185)*, (IEEE, Victoria, BC, Canada, 1998), pp. 34–38.
6. J. J. Reuther, A. Jameson, J. J. Alonso, M. J. Rimlinger, and D. Saunders, "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 2," *J. Aircr.* **36**(1), 61–74 (1999).
7. S. Nadarajah and A. Jameson, "A comparison of the continuous and discrete adjoint approach to automatic aerodynamic optimization," in *38th Aerospace Sciences Meeting and Exhibit*, (American Institute of Aeronautics and Astronautics, Reno, NV, U.S.A., 2000).
8. H. Guan, Y.-J. Chen, Y.-C. Loo, Y.-M. Xie, and G. Steven, "Bridge topology optimisation with stress, displacement and frequency constraints," *Comput. & Struct.* **81**(3), 131–145 (2003).
9. M. P. Bendsøe and O. Sigmund, *Topology Optimization* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2004).
10. C. Y. Kao, S. Osher, and E. Yablonovitch, "Maximizing band gaps in two-dimensional photonic crystals by using level set methods," *Appl. Phys. B* **81**(2-3), 235–244 (2005).
11. Y. Jiao, S. Fan, and D. A. B. Miller, "Demonstration of systematic photonic crystal device design and optimization by low-rank adjustments: An extremely compact mode separator," *Opt. Lett.* **30**(2), 141 (2005).
12. S. Ehara and Y. Kanno, "Topology design of tensegrity structures via mixed integer programming," *Int. J. Solids Struct.* **47**(5), 571–579 (2010).
13. C. M. Lalau-Keraly, S. Bhargava, O. D. Miller, and E. Yablonovitch, "Adjoint shape optimization applied to electromagnetic design," *Opt. Express* **21**(18), 21693 (2013).
14. V. Ganapati, O. D. Miller, and E. Yablonovitch, "Light Trapping Textures Designed by Electromagnetic Optimization for Subwavelength Thick Solar Cells," *IEEE J. Photovoltaics* **4**(1), 175–182 (2014).
15. H. Men, K. Y. K. Lee, R. M. Freund, J. Peraire, and S. G. Johnson, "Robust topology optimization of three-dimensional photonic-crystal band-gap structures," *Opt. Express* **22**(19), 22632 (2014).
16. F. Callewaert, V. Velev, S. Jiang, A. V. Sahakian, P. Kumar, and K. Aydin, "Inverse-designed stretchable metalens with tunable focal distance," *Appl. Phys. Lett.* **112**(9), 091102 (2018).
17. T. W. Hughes, M. Minkov, I. A. D. Williamson, and S. Fan, "Adjoint Method and Inverse Design for Nonlinear Nanophotonic Devices," *ACS Photonics* **5**(12), 4781–4787 (2018).

18. Y. Shi, W. Li, A. Raman, and S. Fan, "Optimization of Multilayer Optical Films with a Memetic Algorithm and Mixed Integer Programming," *ACS Photonics* **5**(3), 684–691 (2018).
19. H. E. Ahmed, B. Salman, A. Kherbeet, and M. Ahmed, "Optimization of thermal design of heat sinks: A review," *Int. J. Heat Mass Transfer* **118**, 129–153 (2018).
20. S. Molesky, Z. Lin, A. Y. Piggott, W. Jin, J. Vučković, and A. W. Rodriguez, "Inverse design in nanophotonics," *Nat. Photonics* **12**(11), 659–670 (2018).
21. J. Wang, Y. Shi, T. Hughes, Z. Zhao, and S. Fan, "Adjoint-based optimization of active nanophotonic devices," *Opt. Express* **26**(3), 3236 (2018).
22. R. E. Christiansen, J. Michon, M. Benzaouia, O. Sigmund, and S. G. Johnson, "Inverse design of nanoparticles for enhanced Raman scattering," *Opt. Express* **28**(4), 4444 (2020).
23. A. Michaels, M. C. Wu, and E. Yablonovitch, "Hierarchical Design and Optimization of Silicon Photonics," *IEEE J. Sel. Top. Quantum Electron.* **26**(2), 1–12 (2020).
24. Z. Lin, V. Liu, R. Pestourie, and S. G. Johnson, "Topology optimization of freeform large-area metasurfaces," *Opt. Express* **27**(11), 15765 (2019).
25. A. Thompson, "An evolved circuit, intrinsic in silicon, entwined with physics," in *Evolvable Systems: From Biology to Hardware*, vol. 1259 G. Goos, J. Hartmanis, J. Leeuwen, T. Higuchi, M. Iwata, and W. Liu, eds. (Springer Berlin Heidelberg, Berlin, Heidelberg, 1997), pp. 390–405.
26. J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic Crystals: Molding the Flow of Light* (Princeton University Press, Princeton, 2008), 2nd ed.
27. P. Amestoy, A. Buttari, J.-Y. L'Excellent, and T. Mary, "Performance and Scalability of the Block Low-Rank Multifrontal Factorization on Multicore Architectures," *ACM Trans. Math. Softw.* **45**(1), 1–26 (2019).
28. Y. Chen, T. Davis, W. Hager, and S. Rajamanickam, "Algorithm 887: CHOLMOD, Supernodal Sparse Cholesky Factorization and Update/Downdate," *ACM Trans. Math. Softw.* **35**(3), 1–14 (2008).
29. M. Hestenes and E. Stiefel, "Methods of conjugate gradients for solving linear systems," *J. Res. Natl. Bur. Stan.* **49**(6), 409 (1952).
30. Y. Saad and M. H. Schultz, "GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems," *SIAM J. Sci. and Stat. Comput.* **7**(3), 856–869 (1986).
31. B. F. Smith, "Domain Decomposition Methods for Partial Differential Equations," in *Parallel Numerical Algorithms*, vol. 4 M. D. Salas, D. E. Keyes, A. Sameh, and V. Venkatakrishnan, eds. (Springer Netherlands, Dordrecht, 1997), pp. 225–243.
32. K. Urban, *Wavelet Methods for Elliptic Partial Differential Equations* (Oxford University Press, 2008).
33. A. Gillman, "Fast direct solvers for elliptic partial differential equations," Ph.D. thesis, University of Colorado (2011).
34. L. Su, D. Vercruysee, J. Skarda, N. V. Saprà, J. A. Petykiewicz, and J. Vučković, "Nanophotonic inverse design with SPINS: Software architecture and practical considerations," *Appl. Phys. Rev.* **7**(1), 011407 (2020).
35. G. Angeris, J. Vučković, and S. Boyd, "Computational bounds for photonic design," *ACS Photonics* **6**(5), 1232–1239 (2019).
36. A. Piggott, J. Petykiewicz, L. Su, and J. Vučković, "Fabrication-constrained nanophotonic inverse design," *Sci. Rep.* **7**(1), 1786 (2017).
37. A. Y. Piggott, J. Lu, K. G. Lagoudakis, J. Petykiewicz, T. M. Babinec, and J. Vučković, "Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer," *Nat. Photonics* **9**(6), 374–377 (2015).
38. H. Chung and O. D. Miller, "High-NA achromatic metalenses by inverse design," *Opt. Express* **28**(5), 6945 (2020).
39. E. Bayati, R. Pestourie, S. Colburn, Z. Lin, S. G. Johnson, and A. Majumdar, "Inverse Designed Metalenses with Extended Depth of Focus," *ACS Photonics* **7**(4), 873–878 (2020).
40. J. Lu and J. Vučković, "Inverse design of nanophotonic structures using complementary convex optimization," *Opt. Express* **18**(4), 3793 (2010).
41. J. Lu and J. Vučković, "Nanophotonic computational design," *Opt. Express* **21**(11), 13351 (2013).
42. A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Math. Program.* **106**(1), 25–57 (2006).
43. R. Karp, "Reducibility among combinatorial problems," in *Complexity of Computer Computations*, R. Miller, J. Thatcher, and J. Bohlinger, eds. (Springer US, Boston, MA, 1972), pp. 85–103.
44. S. Aaronson, "Why Philosophers Should Care About Computational Complexity," arXiv:1108.1791 [quant-ph] (2011).
45. J. L. Lions, *Optimal Control of Systems Governed by Partial Differential Equations* (Springer, Berlin, 1971).
46. M. B. Giles and N. A. Pierce, "An Introduction to the Adjoint Approach to Design," *Flow, Turbul. Combust.* **65**(3/4), 393–415 (2000).
47. F. de Avila Belbute-Peres, K. Smith, K. Allen, J. Tenenbaum, and J. Z. Kolter, "End-to-end differentiable physics for learning and control," in *Advances in Neural Information Processing Systems*, (2018), pp. 7178–7189.
48. Y. Hu, L. Anderson, T.-M. Li, Q. Sun, N. Carr, J. Ragan-Kelley, and F. Durand, "DiffTaichi: Differentiable programming for physical simulation," in *International Conference on Learning Representations*, (2019).
49. M. Minkov, I. A. D. Williamson, L. C. Andreani, D. Gerace, B. Lou, A. Y. Song, T. W. Hughes, and S. Fan, "Inverse Design of Photonic Crystals through Automatic Differentiation," *ACS Photonics* **7**(7), 1729–1741 (2020).
50. G. Angeris, J. Vučković, and S. Boyd, "A New Heuristic for Physical Design," arXiv:2002.05644 [physics] (2020).

51. R. Haupt, "Comparison between genetic and gradient-based optimization algorithms for solving electromagnetics problems," *IEEE Trans. Magn.* **31**(3), 1932–1935 (1995).
52. J. Jervase and H. Bourdoucen, "Design of resonant-cavity-enhanced photodetectors using genetic algorithms," *IEEE J. Quantum Electron.* **36**(3), 325–332 (2000).
53. G. Cormier, R. Boudreau, and S. Thériault, "Real-coded genetic algorithm for Bragg grating parameter synthesis," *J. Opt. Soc. Am. B* **18**(12), 1771 (2001).
54. E. Kerrinckx, L. Bigot, M. Douay, and Y. Quiquempois, "Photonic crystal fiber design by means of a genetic algorithm," *Opt. Express* **12**(9), 1990 (2004).
55. L. Sanchis, A. Håkansson, D. López-Zanón, J. Bravo-Abad, and J. Sánchez-Dehesa, "Integrated optical devices design by genetic algorithm," *Appl. Phys. Lett.* **84**(22), 4460–4462 (2004).
56. A. Hakansson, J. Sanchez-Deh, and L. Sanchis, "Inverse design of photonic crystal devices," *IEEE J. Sel. Areas in Commun.* **23**(7), 1365–1371 (2005).
57. J. Goh, I. Fushman, D. Englund, and J. Vučković, "Genetic optimization of photonic bandgap structures," *Opt. Express* **15**(13), 8218 (2007).
58. G. Hornby, A. Globus, D. Linden, and J. Lohn, "Automated Antenna Design with Evolutionary Algorithms," in *Space 2006*, (American Institute of Aeronautics and Astronautics, San Jose, California, 2006).
59. S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by Simulated Annealing," *Science* **220**(4598), 671–680 (1983).
60. J. Nelder and R. Mead, "A Simplex Method for Function Minimization," *The Comput. J.* **7**(4), 308–313 (1965).
61. I. Loshchilov, M. Schoenauer, and M. Sebag, "Adaptive coordinate descent," in *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation - GECCO '11*, (ACM Press, Dublin, Ireland, 2011), p. 885.
62. L. M. Rios and N. Sahinidis, "Derivative-free optimization: A review of algorithms and comparison of software implementations," *J. Glob. Optim.* **56**(3), 1247–1293 (2013).
63. A. R. Conn, K. Scheinberg, and L. N. Vicente, *Introduction to Derivative-Free Optimization* (Society for Industrial and Applied Mathematics, 2009).
64. C. Audet and W. Hare, *Derivative-Free and Blackbox Optimization*, *Springer Series in Operations Research and Financial Engineering* (Springer International Publishing, Cham, 2017).
65. R. Byrd, P. Lu, J. Nocedal, and C. Zhu, "A Limited Memory Algorithm for Bound Constrained Optimization," *SIAM J. Sci. Comput.* **16**(5), 1190–1208 (1995).
66. N. Parikh, "Proximal Algorithms," *Foundations Trends Optim.* **1**(3), 127–239 (2014).
67. K. Svanberg, "The method of moving asymptotes—a new method for structural optimization," *Int. J. Numer. Meth. Engng.* **24**(2), 359–373 (1987).
68. D. Bertsekas, *Nonlinear Programming* (Athena Scientific, Belmont, Massachusetts, 2016), third edition ed.
69. H. A. Eiselt and C.-L. Sandblom, *Nonlinear Optimization: Methods and Applications*, vol. 282 of *International Series in Operations Research & Management Science* (Springer International Publishing, Cham, 2019).
70. M. Kochenderfer and T. Wheeler, *Algorithms for Optimization* (The MIT Press, Cambridge, Massachusetts, 2019).
71. H. Wheeler, "Fundamental Limitations of Small Antennas," *Proc. IRE* **35**(12), 1479–1484 (1947).
72. L. Chu, "Physical Limitations of Omni-Directional Antennas," *J. Appl. Phys.* **19**(12), 1163–1175 (1948).
73. B. McKellar, M. Box, and C. Bohren, "Sum rules for optical scattering amplitudes," *J. Opt. Soc. Am.* **72**(5), 535 (1982).
74. Z. Yu, A. Raman, and S. Fan, "Fundamental limit of nanophotonic light trapping in solar cells," *Proc. Natl. Acad. Sci.* **107**(41), 17491–17496 (2010).
75. Z. Yu, A. Raman, and S. Fan, "Thermodynamic Upper Bound on Broadband Light Coupling with Photonic Structures," *Phys. Rev. Lett.* **109**(17), 173901 (2012).
76. O. D. Miller and E. Yablonovitch, "Photon extraction: The key physics for approaching solar cell efficiency limits," in *SPIE NanoScience + Engineering*, G. S. Subramania and S. Foteinopoulou, eds. (San Diego, California, United States, 2013), p. 880807.
77. A. E. Miroshnichenko and M. I. Tribelsky, "Ultimate Absorption in Light Scattering by a Finite Obstacle," *Phys. Rev. Lett.* **120**(3), 033902 (2018).
78. O. Miller, A. Polimeridis, M. T. Homer Reid, C. W. Hsu, B. DeLacy, J. D. Joannopoulos, M. Soljačić, and S. G. Johnson, "Fundamental limits to optical response in absorptive systems," *Opt. Express* **24**(4), 3329 (2016).
79. O. D. Miller, O. Ilic, T. Christensen, M. T. H. Reid, H. A. Atwater, J. D. Joannopoulos, M. Soljačić, and S. G. Johnson, "Limits to the Optical Response of Graphene and Two-Dimensional Materials," *Nano Lett.* **17**(9), 5408–5415 (2017).
80. S. Molesky, W. Jin, P. Venkataram, and A. Rodriguez, "Bounds on absorption and thermal radiation for arbitrary objects," *Phys. Rev. Lett.* **123**(25), 257401 (2019).
81. H. Shim, L. Fan, S. Johnson, and O. Miller, "Fundamental Limits to Near-Field Optical Response over Any Bandwidth," *Phys. Rev. X* **9**(1), 011043 (2019).
82. R. Trivedi, G. Angeris, L. Su, S. Boyd, S. Fan, and J. Vučković, "Bounds for Scattering from Absorptionless Electromagnetic Structures," *Phys. Rev. Appl.* **14**(1), 014025 (2020).
83. S. Molesky, P. Chao, W. Jin, and A. W. Rodriguez, "Global T operator bounds on electromagnetic scattering: Upper bounds on far-field cross sections," *Phys. Rev. Res.* **2**(3), 033172 (2020).
84. S. Molesky, P. Chao, and A. W. Rodriguez, "Hierarchical Mean-Field T Operator Bounds on Electromagnetic Scattering: Upper Bounds on Near-Field Radiative Purcell Enhancement," arXiv:2008.08168 [physics] (2020).

85. S. Molesky, P. Chao, and A. Rodriguez, "T-Operator Limits on Electromagnetic Scattering: Bounds on Extinguished, Absorbed, and Scattered Power from Arbitrary Sources," arXiv:2001.11531 [physics] (2020).
86. S. Molesky, P. Venkataram, W. Jin, and A. Rodriguez, "Fundamental limits to radiative heat transfer: Theory," *Phys. Rev. B* **101**(3), 035408 (2020).
87. J. Michon, M. Benzaouia, W. Yao, O. D. Miller, and S. G. Johnson, "Limits to surface-enhanced Raman scattering near arbitrary-shape scatterers: Erratum," *Opt. Express* **28**(15), 22264 (2020).
88. Z. Kuang and O. Miller, "Computational bounds to light-matter interactions via local conservation laws," arXiv:2008.13325 [physics] (2020).
89. M. Gustafsson, K. Schab, L. Jelinek, and M. Capek, "Upper bounds on absorption and scattering," *New Journal of Physics* (2020).
90. S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, UK ; New York, 2004).
91. Z. Kuang, L. Zhang, and O. D. Miller, "Maximal single-frequency electromagnetic response," arXiv:2002.00521 [physics] (2020).
92. R. G. Newton, "Optical theorem and beyond," *Am. J. Phys.* **44**(7), 639–642 (1976).
93. J. Park and S. Boyd, "General Heuristics for Nonconvex Quadratically Constrained Quadratic Programming," arXiv:1703.07870 [math] (2017).
94. S. Boyd and L. Vandenberghe, *Introduction to Applied Linear Algebra: Vectors, Matrices, and Least Squares* (Cambridge University Press, 2018), 1st ed.
95. M. ApS, "MOSEK Optimizer API for Python 9.1.5," <https://docs.mosek.com/9.1/pythonapi/index.html> (2019).
96. B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd, "Conic Optimization via Operator Splitting and Homogeneous Self-Dual Embedding," *J. Optim. Theory Appl.* **169**(3), 1042–1068 (2016).
97. M. Garstka, M. Cannon, and P. Goulart, "COSMO: A conic operator splitting method for large convex problems," in *2019 18th European Control Conference (ECC)*, (IEEE, Naples, Italy, 2019), pp. 1951–1956.
98. L. Vandenberghe and M. Andersen, "Chordal Graphs and Semidefinite Optimization," *Foundations Trends Optim.* **1**(4), 241–433 (2015).
99. Q. Zhao, L. Zhang, and O. D. Miller, "Minimum Dielectric-Resonator Mode Volumes," arXiv:2008.13241 [physics] (2020).
100. B. Shen, P. Wang, R. Polson, and R. Menon, "An integrated-nanophotonics polarization beamsplitter with 2.4 by 2.4 μm^2 footprint," *Nat. Photonics* **9**(6), 378–382 (2015).
101. J. C. C. Mak, C. Sideris, J. Jeong, A. Hajimiri, and J. K. S. Poon, "Binary particle swarm optimized 2 by 2 power splitters in a standard foundry silicon photonic platform," *Opt. Lett.* **41**(16), 3868 (2016).
102. L. Su, A. Y. Piggott, N. V. Sapro, J. Petykiewicz, and J. Vučković, "Inverse Design and Demonstration of a Compact on-Chip Narrowband Three-Channel Wavelength Demultiplexer," *ACS Photonics* **5**(2), 301–305 (2018).
103. N. V. Sapro, K. Y. Yang, D. Vercruysee, K. J. Leadle, D. S. Black, R. J. England, L. Su, R. Trivedi, Y. Miao, O. Solgaard, R. L. Byer, and J. Vučković, "On-chip integrated laser-driven particle accelerator," *Science* **367**(6473), 79–83 (2020).
104. S. Banerji, M. Meem, A. Majumder, J. C. Garcia, P. Hon, C. Pies, T. Oberbiermann, B. S. Rodriguez, and R. Menon, "Inverse Designed Flat Optics with Diffractive Lenses," in *Imaging and Applied Optics Congress*, (OSA, Washington, DC, 2020), p. ITh5E.3.
105. D. Vercruysee, N. V. Sapro, L. Su, R. Trivedi, and J. Vučković, "Analytical level set fabrication constraints for inverse design," *Sci. Rep.* **9**(1), 8999 (2019).
106. A. Y. Piggott, L. Su, N. V. Sapro, A. M. Netherton, A. S. P. Khope, J. E. Bowers, and J. Vučković, "Foundry-fabricated Inverse Designed Photonics," in *Frontiers in Optics + Laser Science APS/DLS*, (OSA, Washington, DC, 2019), p. JW3A.20.
107. Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE* **86**(11), 2278–2324 (1998).
108. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, (IEEE, Miami, FL, 2009), pp. 248–255.
109. A. Krizhevsky, "Learning Multiple Layers of Features from Tiny Images," (2009).
110. I. Maros and C. Mészáros, "A repository of convex quadratic programming problems," *Optim. Methods Softw.* **11**(1-4), 671–681 (1999).
111. A. Gleixner, G. Hendel, G. Gamrath, T. Achterberg, M. Bastubbe, T. Berthold, P. Christophel, K. Jarck, T. Koch, J. Linderoth, M. Lubbecke, H. D. Mittelmann, D. Ozyurt, T. K. Ralphs, D. Salvagnin, and Y. Shinano, "MIPLIB 2017: Data-Driven Compilation of the 6th Mixed-Integer Programming Library,"
112. D. O'Connor and L. Vandenberghe, "Primal-Dual Decomposition by Operator Splitting and Applications to Image Deblurring," *SIAM J. Imaging Sci.* **7**(3), 1724–1754 (2014).
113. R. Trivedi, L. Su, J. Lu, M. F. Schubert, and J. Vuckovic, "Data-driven acceleration of photonic simulations," *Sci. Rep.* **9**(1), 19728 (2019).
114. D. Liu, Y. Tan, E. Khoram, and Z. Yu, "Training Deep Neural Networks for the Inverse Design of Nanophotonic Structures," *ACS Photonics* **5**(4), 1365–1369 (2018).
115. J. Jiang and J. A. Fan, "Global Optimization of Dielectric Metasurfaces Using a Physics-Driven Neural Network," *Nano Lett.* **19**(8), 5366–5372 (2019).

116. S. So, T. Badloe, J. Noh, J. Bravo-Abad, and J. Rho, “Deep learning enabled inverse design in nanophotonics,” *Nanophotonics* **9**(5), 1041–1057 (2020).
117. C.-T. Chen and G. X. Gu, “Generative Deep Neural Networks for Inverse Materials Design Using Backpropagation and Active Learning,” *Adv. Sci.* **7**(5), 1902607 (2020).
118. J. Lu and J. Vučković, “Objective-first design of high-efficiency, small-footprint couplers between arbitrary nanophotonic waveguide modes,” *Opt. Express* **20**(7), 7221 (2012).
119. J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, “Julia: A fresh approach to numerical computing,” *SIAM Rev.* **59**(1), 65–98 (2017).
120. I. Dunning, J. Huchette, and M. Lubin, “JuMP: A modeling language for mathematical optimization,” *SIAM Rev.* **59**(2), 295–320 (2017).