Research Article Vol. 34, No. 7 / July 2017 / Journal of the Optical Society of America B 1551

Journal of the Optical Society of America B

OPTICAL PHYSICS

Objective-trait-bias metaheuristics for design optimization of optical structures

Sacha Verweij[1,2,5] and Shanhui Fan[3,4]

[1]Department of Applied Physics, Stanford University, Stanford, California 94305, USA
[2]Institute for Computational & Mathematical Engineering, Stanford University, Stanford, California 94305, USA
[3]Ginzton Laboratory, Stanford University, Stanford, California 94305, USA
[4]e-mail: shanhui@stanford.edu
[5]e-mail: sacha@stanford.edu

Received 10 March 2017; revised 10 May 2017; accepted 18 May 2017; posted 19 May 2017 (Doc. ID 290297); published 29 June 2017

Search algorithms play a crucial role in systematic design optimization of optical structures. Though many sophisticated methods appear in the literature, physically motivated guiding principles for the design and enhancement of such methods are few. We introduce such a guiding principle—an *objective-trait-bias metaheuristic*—and demonstrate its value in practice. Specifically, we present a case study in which application of an instance of this metaheuristic—a *transmission-bias* metaheuristic—leads to significantly better performing variants of a simple stochastic local search algorithm—restarted iterative best improvement—on a challenging design optimization problem—combinatorial design optimization of a multi-spatial-mode photonic crystal waveguide bend that preserves modal content.

OCIS codes: (220.0220) Optical design and fabrication; (350.4600) Optical engineering; (230.0230) Optical devices; (000.5490) Probability theory, stochastic processes, and statistics; (350.4238) Nanophotonics and photonic crystals.

https://doi.org/10.1364/JOSAB.34.001551

1. INTRODUCTION

Systematic design optimization of optical structures is opening new and advancing existing frontiers in the control of light, from ultracompact wavelength- and mode-division multiplexing [1–13] to performance enhancement of photo- and thermophoto-voltaics [14–18]. To exploit systematic design optimization's full potential in optics, search algorithms capable of efficiently sifting large design spaces are critical [1–37].

The degree to which a search algorithm exploits problem structure determines its efficacy. Physically derived design optimization problems often contain rich structure that one can reason about on the basis of physical intuition; design optimization problems in optics are no exception [38]. General-purpose search algorithms exploit little if any of this rich structure, and thereby fail to achieve optimal performance on design optimization problems in optics. Augmenting such search algorithms with mechanisms that exploit this rich structure can significantly improve search efficacy.

Broadly applicable, physically motivated guiding principles for augmenting such search algorithms for design optimization problems in optics could substantially advance the field. The following *objective-trait-bias metaheuristic* exemplifies such guiding principles: Design objectives often involve some set of *traits*—necessary but insufficient conditions for performance. Examples include transmission or reflection, strong or weak absorption, broad or narrow bandwidth, modal conversion or preservation, or satisfaction or violation of symmetries. Now consider searching a design space for structures satisfying a design objective. Biasing the search toward design space regions that exhibit the requisite traits, and thereby more likely satisfy the design objective, should improve search performance.

This paper illustrates how one might apply such guiding principles as well as their potential value. A case study, in which applying the above metaheuristic leads to significantly better performing variants of a simple search algorithm on a challenging design optimization problem, provides this illustration.

The case consists of design optimization of a multi-spatial-mode photonic crystal waveguide bend. We previously studied this case in [38] highlighting the value of search landscape analysis in optics. The present work builds upon that work by applying an instance of the above metaheuristic—a *transmission-bias* metaheuristic—to substantially improve search performance over that discussed in [38]. The case study realizes this transmission-bias metaheuristic concretely through mechanisms that bias the search away from structures that contain too many or inappropriately placed scatterers in expected transmission paths.

Section 2 introduces the model design optimization problem about which the case study revolves (and, simultaneously, necessary terminology). Section 3 describes the baseline search

algorithm, which we augment via the above metaheuristic. Sections 4–6 each describe a variant of the baseline search algorithm that realizes the above metaheuristic, and demonstrate that variant's superior performance. Section 7 composes the enhancement mechanisms underlying the search algorithm variants introduced in Sections 4–6, yielding a corresponding search algorithm with performance exceeding that of the preceding, single-mechanism variants. To close, Section 9 discusses extensions and generalizations of the present work.

2. MODEL DESIGN OPTIMIZATION PROBLEM

As noted above, the design optimization problem this work considers is quite similar to that studied in this work's baseline [38]. The following paragraphs recapitulate this problem's essential definitions, but for additional details see [38].

Our case study revolves around design optimization of a multi-spatial-mode photonic crystal waveguide bend that preserves modal content; such multi-spatial-mode photonic crystal devices, for example those designed in [10–12], hold promise for mode-division multiplexing in future photonic integrated circuits. Particularly, we ultimately seek a device like that schematized in Fig. 1(b) and described below.

A square lattice of silicon cylinders in air, which supports a broad transverse-magnetic (TM) mode bandgap, hosts the device. Cylinder-deletion defects in this photonic crystal constitute the device itself. Two adjacent, horizontal lines of cylinder deletions form a waveguide incident from the left. Similarly, two adjacent, vertical lines of cylinder deletions form a waveguide exiting at the bottom. Each waveguide supports two TM spatial modes over the upper half of the crystal's bandgap. A mirror-symmetric, multi-defect junction terminates these waveguides at their intersection. Together, the waveguides and junction compose a symmetric multi-spatial-mode bend. Figures 1(c) and 1(d) respectively show the out-of-plane electric field patterns resulting from excitation of the first and second ingoing TM spatial modes on the horizontal waveguide; insofar as one can discern visually, these modes respectively transmit strictly to the first and second outgoing TM spatial modes on the vertical waveguide, exemplifying near-perfect modal-content-preserving transmission. Figure 1's caption provides additional detail.

To design such a device, we begin from the prototype schematized in Fig. 1(a); this prototype is identical to the device described above but for a junction solely containing defective cells. We generate a *candidate structure* from this prototype by placing nondefective cells at some set of the structure's *independent lattice sites*—the twenty-seven lattice sites in the junction's upper-left half, diagonal inclusive; enforcement of the prototype's diagonal mirror symmetry then dictates placement of nondefective cells in the junction's lower-right half. The desired device described above and shown in Fig. 1(b) exemplifies a candidate structure. We can generate 2^{27} distinct candidate structures via this procedure; the set of all such candidate structures constitutes our *search space* $\mathcal{S}$.

This discrete, *combinatorial* search space contrasts with continuous search spaces admitting, for example, continuous geometric or material-parametric deformations between candidate structures; for more on such search spaces, see [38].

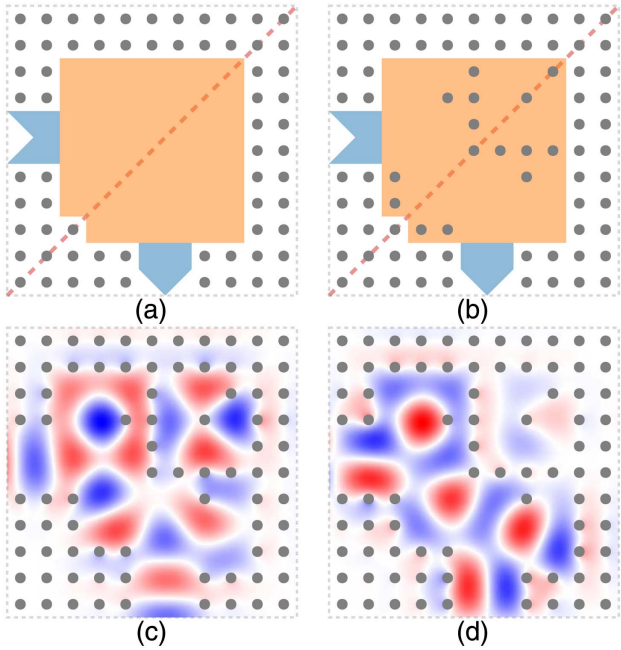

Fig. 1. (a) Design prototype for a multi-spatial-mode photonic crystal waveguide bend that preserves modal content. A square lattice of silicon cylinders (gray circles, $r/a = 0.2$, $n = 3.4$) in air (white backdrop), which supports a TM mode bandgap over frequencies $0.29 \cdot 2\pi c/a$ to $0.42 \cdot 2\pi c/a$, hosts the device. Cylinder-deletion defects in this photonic crystal constitute the device itself. Highlighted in blue, two adjacent horizontal lines of cylinder deletions form a waveguide incident from the left. Similarly highlighted in blue, two adjacent vertical lines of cylinder deletions form a waveguide exiting at the bottom. Each waveguide supports two TM spatial modes over frequencies $0.36 \cdot 2\pi c/a$ to $0.42 \cdot 2\pi c/a$. Highlighted in orange, a mirror-symmetric, multi-deletion junction terminates these waveguides at their intersection. Together, the waveguides and junction compose a symmetric multi-spatial-mode bend. The dashed red line indicates this prototype's diagonal mirror symmetry. We generate a *candidate structure* from this prototype, such as that shown in (b), by placing deletion cells at some of the structure's junction lattice sites consistent with the prototype's mirror symmetry. Parts (c) and (d) respectively show the out-of-plane electric field patterns resulting from excitation, at frequency $0.39 \cdot 2\pi c/a$, of the first and second ingoing TM spatial modes on the horizontal waveguide of this particular structure.

The metaheuristic investigated here, however, applies broadly, including to continuous design optimization problems in optics.

Next we define our *design objective*, which is to find a candidate structure that exhibits satisfactorily high modal-content-preserving transmission or, equivalently, satisfactorily low reflection and modal crosstalk in transmission. We capture our design objective quantitatively in an *objective function* $f : \mathcal{S} \to \mathcal{R}$. Though many possible objective functions reasonably correspond to our design objective, we focus on minimization of the ℓ^2-type objective function

$$f(s) = 2^{-1/2} \sqrt{\sum_{i,j} R^2_{i\to j}(s) + \sum_{i,j\neq i} T^2_{i\to j}(s)}, \tag{1}$$

where, under monochromatic excitation of ingoing TM spatial mode i on one port of candidate structure $s \in \mathcal{S}$, $R_{i\to j}(s)$

denotes the excitation-normalized power reflected into outgoing TM spatial mode j on the same port, and $T_{i\to j}(s)$ denotes the excitation-normalized power transmitted into outgoing TM spatial mode j on the other port. Accordingly, the first summation in this objective function captures reflection and the second modal crosstalk in transmission. Where i and j in the summations respectively run over all ingoing and outgoing spatial modes, the factor of $2^{-1/2}$ normalizes the objective function to the interval [0,1].

Finally, we form our *design optimization problem* from these elements: In principle, we seek a candidate structure s^* in search space $\mathcal{S}$ that minimizes objective function f, or $s^* = \text{argmin}_{s\in\mathcal{S}} f(s)$. In practice, however, we stop the search once we find an acceptable candidate structure s_a in search space $\mathcal{S}$ with objective (function) value below some acceptability threshold f_a, or $s_a \in \mathcal{S} : f(s_a) \leq f_a$.

Describing the algorithms and analysis below requires additional terminology and augmentation of our search space $\mathcal{S}$ with a sense of locality.

We measure the distance between candidate structures s and s' by the number of independent lattice sites at which they differ—a form of *Hamming distance* $\mathbf{H}(s, s')$. To illustrate, the candidate structures shown in Figs. 1(a) and 1(b) differ at nine lattice sites in the upper-left half of the junction, diagonal inclusive, and so by this metric the (Hamming) distance between them is nine.

With distance defined, we introduce a *neighborhood relation* $\mathcal{N} \subseteq \mathcal{S} \times \mathcal{S}$ that, intuitively, indicates which candidate structures are *neighbors* or *adjacent*: we say that s and s' are neighbors if and only if they differ at precisely one independent lattice site or, equivalently, their Hamming distance $\mathbf{H}(s, s')$ is one.

Our search space $\mathcal{S}$, neighborhood relation $\mathcal{N}$, and objective function f together constitute our *search landscape*.

3. BASELINE SEARCH ALGORITHM

Our baseline search algorithm is restarted iterative best improvement, a form of repeated local search where iterative best improvement is the subsidiary local search procedure.

In each step of *iterative best improvement*, we determine the objective values of all neighbors of an initially given *current* candidate structure. If any neighbor's objective value improves upon the current candidate structure's objective value, we declare the *best improving* neighbor the new current candidate and continue with the next search step. Otherwise, the current candidate structure—having no improving neighbors—is a *local minimum*, and iterative best improvement terminates.

Restarted iterative best improvement (RII) augments iterative best improvement with a rudimentary local minimum *escape mechanism*: In each search *super*step, we draw a candidate structure uniformly at random from our search space and, from that candidate structure, perform iterative best improvement. When iterative best improvement terminates at a local minimum, we *restart* the search by continuing with the next search superstep.

See [38] for extensive discussion and analysis of this algorithm's behavior on a similar problem. On the one hand, relative to approaches that exploit no search landscape structure, restarted iterative best improvement achieves good performance on our model design optimization problem. On the other hand, being a general-purpose search algorithm, restarted iterative best improvement makes no use of simple physical intuition regarding transmission.

Below we augment restarted iterative best improvement to exploit simple physical intuition regarding transmission, and show that such augmentation substantially improves search performance; we thereby illustrate the concrete application and impact of the transmission-bias metaheuristic described above.

Specifically, each of Sections 4–6 describes a kernel of intuition regarding transmission in our system, briefly provides empirical support for that intuition, introduces a variant of restarted iterative best improvement integrating that intuition to improve search performance, and then demonstrates that variant's superior performance. Presented in order of increasing sophistication, the first variant involves strictly excluding a search subspace from consideration during search, the second a biased restart mechanism, and the third a modified subsidiary local search procedure replacing iterative best improvement. These modifications being orthogonal, Section 7 discusses these variants' combinations and demonstrates those combinations' superior performance.

4. VARIANT I: SEARCH SUBSPACE EXCLUSION

Consider the two independent lattice sites at the input waveguide's insertion into the junction region [those containing black rather than gray inclusions in Figs. 2(a) and 2(b)]. Some candidate structures including scatterers at those lattice sites may exhibit satisfactory transmission. But intuition suggests that placing scatterers at those lattice sites, and thereby at least partially occluding the input waveguide's insertion into the junction region, will drive reflection into the input waveguide in most such structures. Illustrating this intuition, Fig. 2 shows the impact (namely, dramatically increased reflection) of adding scatterers at the identified independent lattice sites to the candidate structure shown in Figs. 1(b)–1(d). Additionally, in candidate structures that include scatterers at those lattice sites and nonetheless exhibit satisfactory transmission, resonant and hence narrowband or otherwise sensitive transmission is likely.

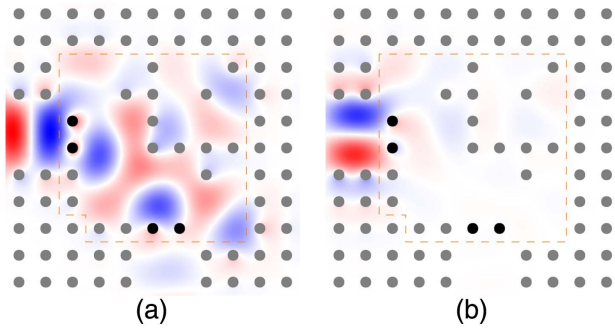

Fig. 2. Out-of-plane electric field patterns resulting from excitation, at frequency $0.39 \cdot 2\pi c/a$, of the (a) first and (b) second ingoing TM spatial modes on the shown structure's input waveguide. The structure is identical to that in Figs. 1(b)–1(d) except for scatterer-including cells at the two independent lattice sites identified in Section 4 (black inclusions at the input waveguide's insertion into the junction region).

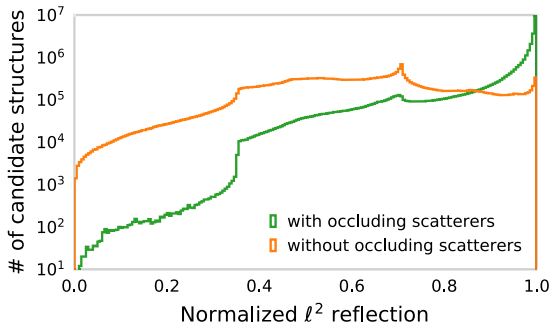

Fig. 3. Distribution of candidate structures by normalized ℓ^2 reflection $\left[g(s) = 2^{-1/2}\sqrt{\sum_{i,j} R^2_{i\to j}(s)}\right]$ for the search subspace consisting of candidate structures without (in orange) and with (in green) scatterers at both of the two independent lattice sites identified in Section 4 (at the input waveguide's insertion into the junction region).

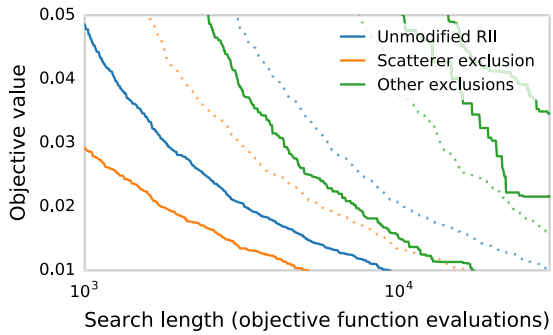

Fig. 4. Performance of unmodified restarted iterative best improvement in blue, the variant that avoids the search subspace including scatterers at the independent lattice sites identified in Section 4 in orange, and the variants that avoid search subspaces of the same size and freezing the same independent lattice sites as the preceding variant (but not the preceding specific search subspace) in green. The solid and dashed lines respectively reflect typical (0.5 quantile or median incumbent objective value) and near-worst-case (0.9 quantile incumbent objective value) performance. The objective function axis (vertical) terminates at the objective function's discretization-error noise floor ($f \approx 0.01$) and reflects operation at frequency $0.39 \cdot 2\pi c/a$.

So consider the search subspace consisting of candidate structures with scatterers at those independent lattice sites. Candidate structures drawn from that subspace seem significantly less likely to exhibit desirable, satisfactory transmission than candidate structures drawn from the complementary subspace. Figure 3 illustrates this intuition: The orange histogram shows the distribution, or normalized ℓ^2 reflection,

$$g(s) = 2^{-1/2}\sqrt{\sum_{i,j} R^2_{i\to j}(s)} \tag{2}$$

of candidate structures in the complementary search subspace (without scatterers at the identified independent lattice sites). For comparison, the green histogram shows the same for the search subspace consisting of candidate structures with scatterers at both the identified independent lattice sites. Clearly the without-scatterers search subspace contains many more high-transmission candidate structures than the with-scatterers search subspace.

This intuition leads us to avoid searching the with-scatterers subspace. Here we consider a restarted iterative best improvement variant in which we both disallow (re)start from candidate structures in that subspace and exclude such candidate structures from consideration during subsidiary local searches. (Alternatively, one can view this restarted iterative best improvement variant as unmodified restarted iterative best improvement operating in a reduced search space: by fixing as defective the two optimization variables corresponding to the identified independent lattice sites, one reduces the search space to a more-promising space of 2^{25} candidate structures.)

Figure 4 compares the statistical performance of unmodified restarted iterative best improvement and the above-described variant, specifically on the model design optimization problem outlined in Section 2. Restarted iterative best improvement is a *stochastic algorithm*. The behavior of an individual run of a stochastic algorithm is not particularly meaningful; rather, only the statistical behavior of a representative *ensemble* of runs necessarily has meaning. Figure 4 captures such statistical behavior as follows.

For each of 2000 independent runs of restarted iterative best improvement and the variant, we recorded the *incumbent objective value* at each search step, which is the objective value of the best candidate structure the search had found by that point. For each algorithm and search step, this procedure produced a sample of 2000 corresponding incumbent objective values, from which we calculated 0.5 and 0.9 quantile incumbent objective values. Figure 4's solid and dotted blue lines respectively show the calculated 0.5 and 0.9 quantile incumbent objective values at each search step for unmodified restarted iterative best improvement. The solid and dotted orange lines show the same for the variant.

Intuitively, our objective function captures the proportion of power lost to reflection and modal crosstalk in transmission. The solid and dotted lines thus respectively show the typical- and near-worst-case evolution of the best-encountered candidate structure's power loss as the search proceeds. Hence Fig. 4 shows that, in both typical and near-worst cases, the variant described above achieves an approximately 1.8× speedup over unmodified restarted best improvement in finding excellent candidate structures.

One could argue that this performance improvement arises from search space reduction in general, rather than from specific exclusion of scatterers at the waveguides' insertions into the junction region. As such Fig. 4 shows the performance of three additional variants in green. These variants are identical to the above-described variant, except instead of fixing the two identified independent lattice sites as defective, in each of these latter variants these two independent lattice sites contain some other cell combination—defective and nondefective, nondefective and defective, or both nondefective. Hence each of these variants excludes a search subspace of the same size and involving the same independent lattice sites as the earlier-described variant, but not the specific subspace excluded by the earlier-described variant. These latter variants exhibit markedly

worse typical- and near-worst-case performance than unmodified restarted iterative best improvement.

5. VARIANT II: RESTART BIASING

When scatterer density in the junction region exceeds some threshold, satisfactory transmission should become unlikely.

Consider a given spatial mode. Intuition suggests that structures allowing that spatial mode to flow through the junction region with little reflection should typically possess at least one relatively scatterer-devoid channel, where that channel's width matches or exceeds that of the mode in the input waveguide. This intuition lower-bounds the number of defective independent lattice sites a candidate structure should include to transmit that spatial mode; we denote that lower bound N.

Of course, satisfactory overall transmission requires that a candidate structure have such a channel for each spatial mode. Intuition suggests that typically those channels will not be the same for each spatial mode, nor will those channels be strictly disjoint. This intuition suggests that a satisfactorily transmitting candidate structure should have at least N defective independent lattice sites and as many as $2N$ or more.

As a conservative estimate of the number of defective independent lattice sites a single spatial mode's channel requires, we use the number of defective independent lattice sites for the bend formed from the simple intersection of the two waveguides—nine. The intuition above then suggests that a satisfactorily transmitting candidate structure should typically include at least nine and more typically toward eighteen defective independent lattice sites.

Figure 5 illustrates the above intuition, showing the distribution of candidate structures by number of defective independent lattice sites on the horizontal axis and normalized ℓ^2 reflection on the vertical axis: Focus on the lowermost row, corresponding to vanishing normalized ℓ^2 reflection. That row shows that, among candidate structures that exhibit such low reflection and hence high transmission, very few contain fewer than nine defective independent lattice sites, and most contain between thirteen and twenty defective independent lattice sites (with a peak around sixteen or seventeen).

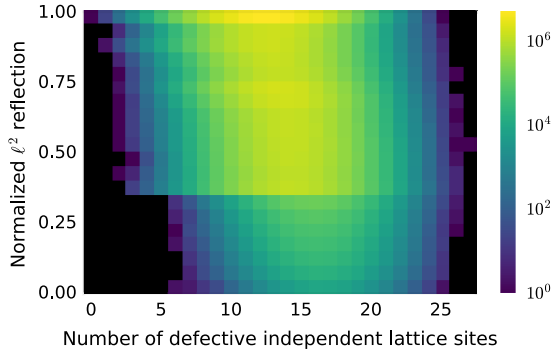

Fig. 5. Distribution of candidate structures by number of defective independent lattice sites on the horizontal axis and normalized ℓ^2 reflection $\left[g(s) = 2^{-1/2}\sqrt{\sum_{i,j} R^2_{i\to j}(s)}\right]$ on the vertical axis. Black cells represent an absence of candidate structures, (nonblack) dark cells a small number of candidate structures, and bright cells a large number of candidate structures. Intensity scaling is logarithmic.

This intuition leads us to bias search toward regions of the search space with an appropriate number of defective independent lattice sites. Here we consider a restarted iterative best improvement variant that accomplishes this biasing via the modified (re)start mechanism described below.

In unmodified restarted iterative best improvement, each search superstep begins with selection of a candidate structure uniformly at random from our search space. Given our search space's structure, this selection mechanism is equivalent to (1) selecting the number of defective independent lattice sites D according to the binomial distribution

$$\mathbf{Pr}(D = d) = \binom{N}{d} p^d (1-p)^{N-d}, \tag{3}$$

where $N = 27$ is the number of independent lattice sites and $p = \frac{1}{2}$ is the probability of a given independent lattice site being defective; and then (2) selecting a candidate structure uniformly at random from the subset of candidate structures with D defective independent sites. Figure 6 shows the preceding binomial distribution in blue.

To bias search toward search space regions containing an appropriate number of defective cells, in the variant we employ the same (re)start mechanism but with p chosen such that each subsidiary local search more probably begins from a candidate structure with an appropriate number of defective cells. We chose p *a priori* (specifically, prior to generation of Fig. 5) as follows.

The estimate above for the number of defective independent lattice sites a single mode's channel requires was conservative. Moreover, here we select the number of defective independent lattice sites for a *(re)start* candidate structure, and a (re)start candidate structure's defective cells are randomly located rather than appropriately placed for transmission. Hence in choosing p we err on the high side of the above range estimate, taking $p = \frac{2}{3}$, which places the corresponding distribution's mean precisely at eighteen defective independent lattice sites. Figure 6 shows the corresponding distribution in orange. Figure 7 shows typical (re)start candidate structures for unmodified restarted iterative best improvement (center), the above-described

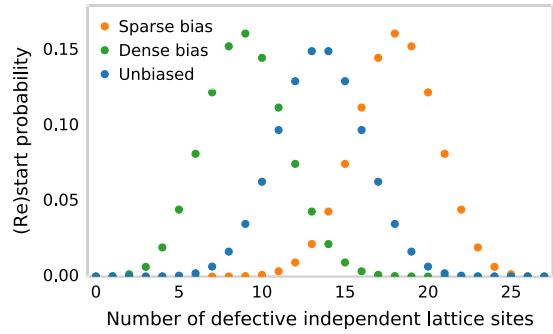

Fig. 6. Defect-count distributions chosen for (re)start candidate structures in unmodified restarted iterative best improvement (blue, $p = \frac{1}{2}$), the variant with (re)start bias toward appropriately sparse candidate structures (orange, $p = \frac{2}{3}$), and an identical variant but with (re)start bias toward equivalently dense candidate structures (green, $p = \frac{1}{3}$). Figure 7 shows typical restart structures for each of these distributions.

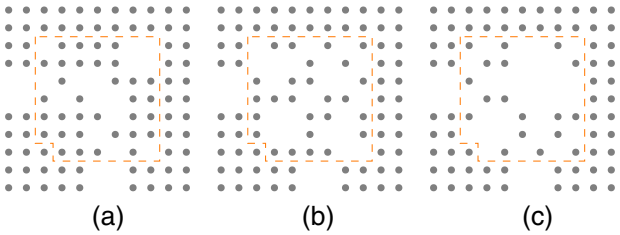

Fig. 7. Typical (re)start candidate structures for (b) unmodified restarted iterative best improvement (selected according to Fig. 6's blue distribution, yielding thirteen defective independent lattice sites), (c) the variant with (re)start bias toward appropriately sparse candidate structures (selected according to Fig. 6's orange distribution, yielding eighteen defective independent lattice sites), and (a) an identical variant but with (re)start bias toward equivalently dense candidate structures (selected in accord with Fig. 6's green distribution, yielding nine defective independent lattice sites).

variant (right), and another variant with $p = \frac{1}{3}$ (left) (corresponding to the green distribution in Fig. 6, which we introduce below for performance comparison).

Figure 8 compares the performance of unmodified restarted iterative improvement in blue and the above-described variant with (re)start bias toward appropriately sparse candidate structures in orange. In both typical and near-worst cases, the variant with (re)start bias toward appropriately sparse candidate structures achieves an approximately 1.8× speedup over unmodified restarted iterative best improvement in finding excellent candidate structures.

One could argue that this performance improvement arises from biasing in general rather than biasing specifically toward appropriately sparse candidate structures. As such Fig. 8 also shows the performance of a variant identical to that described above except with (re)start biasing instead toward equivalently *dense* candidate structures ($p = \frac{1}{3}$, see Fig. 6 in green and the left of Fig. 7) in green, which exhibits markedly worse

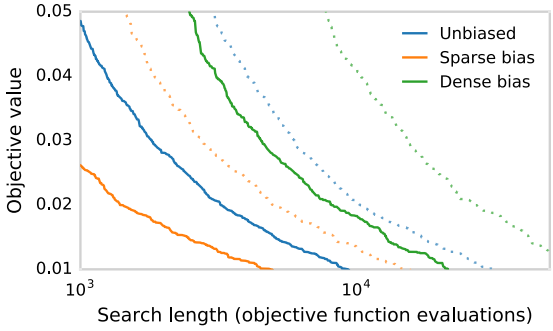

Fig. 8. Performance of unmodified restarted iterative best improvement in blue, the variant with (re)start bias toward appropriately sparse candidate structures in orange, and the variant with (re)start bias toward equivalently dense candidate structures in green, on the model design optimization problem outlined in Section 2. The solid and dashed lines respectively reflect typical- (0.5 quantile or median incumbent objective value) and near-worst-case (0.9 quantile incumbent objective value) performance. The objective function axis (vertical) terminates at the objective function's discretization-error noise floor ($f \approx 0.01$) and reflects operation at frequency $0.39 \cdot 2\pi c/a$.

typical- and near-worst-case performance than unmodified restarted iterative best improvement.

6. VARIANT III: LOCAL DESCENT BIASING

As an alternative to incorporating the intuition developed in Section 5 into restarted iterative best improvement's (re)start mechanism, here we incorporate that intuition into the subsidiary local search procedure. Specifically, we consider a variant in which we replace iterative best improvement with a tailored *variable neighborhood descent* (VND).

Recall that in each step of iterative best improvement we determine the objective values of all neighbors of the *current* candidate structure. If any neighbor's objective value improves upon the current candidate structure's objective value, we declare the *best improving neighbor* the new current candidate and continue with the next search step. Otherwise, the current candidate structure is a local minimum and iterative best improvement terminates.

Instead of simultaneously evaluating all neighbors of the current candidate structure as above, in each step of the tailored variable neighborhood descent considered here, we split the neighbors into two *subneighborhoods*: those neighbors that differ from the current candidate structure by removing a scatterer, and those neighbors that differ from the current candidate structure by introducing a scatterer. Initially we only determine the objective values of neighbors in the scatterer-removing subneighborhood. If any such scatterer-removing neighbor's objective value improves upon the current candidate structure's objective value, we declare the best improving scatterer-removing neighbor the new current candidate and continue with the next search step. If not, the current candidate structure is a local minimum within the scatterer-removing subneighborhood. In that case we expand the search neighborhood to include the scatterer-introducing subneighborhood, in hope of finding an improving neighbor in that expanded neighborhood. If any such scatterer-introducing neighbor's objective value improves upon the current candidate structure's objective value, we declare the best improving scatterer-introducing neighbor the new current candidate and continue with the next search step. Otherwise, the current candidate structure is a local minimum in both subneighborhoods, hence the complete neighborhood, and this variable neighborhood descent terminates.

By augmenting this variable neighborhood descent with the same restart mechanism with which we formed restarted iterative best improvement from iterative best improvement, we form a simple *restarted variable neighborhood descent*.

By preferentially considering scatterer removals over scatterer introductions, this variant biases search toward candidate structures including fewer scatterers and hence, following the intuition presented in Section 5, transmission. Unlike the variant in Section 5, however, this variant removes scatterers somewhat strategically; consequently it may prove an effective approach to biasing toward transmission even when scatterer density is low but the configuration lacks transmission-friendly structure, which may occur, for example, immediately following (re)start.

Figure 9 compares the performance of unmodified restarted iterative best improvement in blue and the above-described

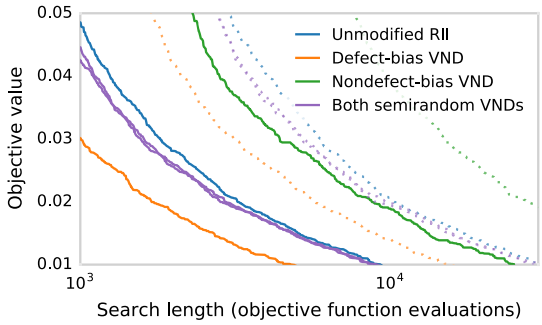

Fig. 9. Performance of unmodified restarted iterative best improvement in blue, the restarted variable neighborhood descent with scatterer-removal-biasing described first in Section 6 in orange, the similar restarted variable neighborhood descent with scatterer-introduction-biasing described later in Section 6 in green, and the two semirandom-neighborhood restarted variable neighborhood descents described in the last paragraph of Section 6 in purple. The solid and dashed lines respectively reflect typical- (0.5 quantile or median incumbent objective value) and near-worst-case (0.9 quantile incumbent objective value) performance. The objective function axis (vertical) terminates at the objective function's discretization-error noise floor ($f \approx 0.01$) and reflects operation at frequency $0.39 \cdot 2\pi c/a$.

restarted variable neighborhood descent in orange. In both typical and near-worst cases, this restarted variable neighborhood descent that biases toward appropriately sparse candidate structures achieves an approximately 1.8× speedup over unmodified restarted iterative best improvement in finding excellent candidate structures.

One could argue that this performance improvement arises from restarted variable neighborhood descent with neighborhood subsets approximately half the size of the complete neighborhood in general, rather than with neighborhoods specifically biasing toward appropriately sparse candidate structures. As such Fig. 9 shows the performance of a restarted variable neighborhood descent identical to that described above except with the neighborhood examination order reversed, thereby biasing toward scatterer introduction instead of scatterer removal, in green. This latter restarted variable neighborhood descent exhibits markedly worse typical- and near-worst-case performance than unmodified restarted iterative best improvement.

One could further argue that the immediately preceding choice of subneighborhoods and ordering (scatterer introduction followed by scatterer removal) is just particularly bad and hence fails to address the objection. Or one could argue that the performance difference between the two preceding restarted variable neighborhood descents arises from different characteristic or specific sizes of the subneighborhoods involved rather than biasing toward introduction or removal of scatterers. As such Fig. 9 shows the performance of another two similar restarted variable neighborhood descents in purple. The first of these restarted variable neighborhood descents considers subneighborhoods of, at each position in the search space, the same size as those for the first, scatterer-removal-biased variant described above, but with those subneighborhoods populated randomly (without replacement) from the complete neighborhood. The second of these restarted variable neighborhood descents is the analog for the scatterer-introduction-biased variant described in the preceding paragraph. In both typical and near-worst cases, these semirandom-neighborhood restarted variable neighborhood descents improve marginally at best upon unmodified restarted iterative best improvement.

7. COMBINATION

The transmission-biasing mechanisms of the search variants described above are largely independent, allowing construction of new search variants via combination of those mechanisms. The impact of those mechanisms on search performance, however, should not be strictly orthogonal; the performance improvements those mechanisms afford in isolation should not combine strictly geometrically.

Combining the search subspace exclusion (Section 4) and restart biasing (Section 5) mechanisms requires some care: Parameter selection for the restart biasing mechanism depends upon the search space's structure. And as noted in Section 4, one can view search subspace exclusion as a search space modification. As such the restart biasing mechanism's parameters should be adjusted when combining that mechanism with search subspace exclusion. That said, we leave the restart biasing parameter $p = \frac{2}{3}$ as is: On the one hand, we should slightly reduce p to account for effectively fixing as defective two additional independent lattice sites in search subspace exclusion. On the other hand, were we to make that reduction by following the estimation procedure in Section 5 to the paper, we would implicitly assume greater overlap between the two modes' channels. To account for that greater overlap, we would have to slightly increase the resulting p. In the interest of simplicity and consistency in comparison, we call the sum of these two opposing adjustments to p negligible and forgo adjusting p.

Figure 10 compares the performance of unmodified restarted iterative best improvement in blue, the single-mechanism transmission-biasing variants described in Sections 4–6 in purple,

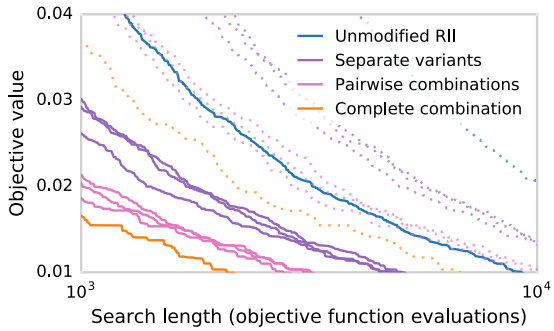

Fig. 10. Performance of unmodified restarted iterative best improvement in blue, the single-mechanism transmission-bias variants described in Sections 4–6 in purple, variants constituting all pairwise combinations of individual transmission-biasing mechanisms in pink, and a variant combining all transmission-biasing mechanisms in orange. The solid and dashed lines respectively reflect typical- (0.5 quantile or median incumbent objective value) and near-worst-case (0.9 quantile incumbent objective value) performance. The objective function axis (vertical) terminates at the objective function's discretization-error noise floor ($f \approx 0.01$) and reflects operation at frequency $0.39 \cdot 2\pi c/a$.

variants constituting all pairwise combinations of the individual transmission-biasing mechanisms in pink, and a variant combining all transmission-biasing mechanisms in orange. In both typical and near-worst cases, the pairwise-combination variants improve upon the performance of their constituents, and the complete-combination variant improves upon the performance of the pairwise combinations. As expected, combination yields subgeometric performance improvements. Nonetheless, the pairwise-combination variants provide an approximately 3× speedup over unmodified restarted iterative best improvement, and the complete-combination variant an approximately 4.5× speedup over the same.

8. DISCUSSION AND CONCLUSION

That each of the mechanisms realizing the transmission-bias metaheuristic independently yields an approximately 1.8× speedup warrants remark. This result perhaps reflects that these mechanisms exploit the same problem structure, in that they may derive roughly the same benefit from exploiting that structure even though they do so via different means.

One often has intuition for structural characteristics conducive to a trait. In such cases, biasing the search toward design space regions that exhibit those structural characteristics is one potential mechanism for biasing toward the trait. This paper illustrates such a case.

Other mechanisms, however, exist. For example, imagine augmenting restarted iterative best improvement with the following escape mechanism that—without injecting associated structural intuition—biases the search toward a trait: From a local minimum in the design objective function, perform an *escape* local search in an objective function that captures solely the trait. That escape local search terminates at a nearby local minimum in the trait objective function, which constitutes a trait-biased perturbation of the local minimum in the design objective function. That trait-biased perturbation then serves as a starting point for subsequent local search in the design objective function. (Of course, this escape mechanism necessitates additional machinery to avoid cycling, for example a *tabu* strategy [39].)

Additionally, penalty methods [39,40] provide an alternative means to achieve, and a lens through which to view, search biasing. For example, above we explicitly modified search restart and descent mechanisms to bias search toward transmission. Alternatively, however, one could achieve the same or similar effect by changing the figure(s) of merit used for candidate structure evaluation in those mechanisms. Specifically, one could modify the figure(s) of merit to penalize candidate structures not exhibiting, or unlikely to exhibit, satisfactory transmission (or, equivalently, to favor candidate structures exhibiting, or likely to exhibit, satisfactory transmission).

By virtue of its simplicity, restarted iterative best improvement served this work's end—illustrating the concrete application and potential value of broadly applicable, physically motivated guiding principles for design optimization in optics—better than more complex algorithms. Or more sophisticated and effective algorithms certainly exist. Where such algorithms do not exploit physically derived problem structure, such algorithms could of course also benefit from augmentation analogous to that presented in this work. More generally, algorithms that do not exploit physically derived problem structure in contexts beyond this work's combinatorial setting, for example general-purpose continuous optimization algorithms, could also benefit from analogous augmentation.

Broadly applicable, physically motivated guiding principles for design optimization in optics enable development of more effective search algorithms. In turn, these more effective search algorithms ultimately lead to better designs and new capabilities.

Funding. U.S. Department of Energy (DOE) Office of Science (SC) (DE-AC05-06OR23100); National Science Foundation (NSF); Air Force Office of Scientific Research (AFOSR) (FA9550-15-1-0335).

Acknowledgment. We thank an anonymous reviewer and J. E. Herriman for suggestions that considerably improved this paper. This research was supported in part by the Department of Energy Office of Science Graduate Fellowship Program (DOE SCGF), made possible in part by the American Recovery and Reinvestment Act of 2009, administered by ORISE-ORAU. This research was also supported in part by the National Science Foundation through XSEDE resources provided by the XSEDE Science Gateways program, and the United States Air Force Office of Scientific Research (USAFOSR).

REFERENCES

1. J. Jensen and O. Sigmund, "Systematic design of photonic crystal structures using topology optimization: low-loss waveguide bends," Appl. Phys. Lett. **84**, 2022–2024 (2004).
2. J. Jensen and O. Sigmund, "Topology optimization of photonic crystal structures: a high-bandwidth low-loss T-junction waveguide," J. Opt. Soc. Am. B **22**, 1191–1198 (2005).
3. N. Ikeda, Y. Sugimoto, Y. Watanabe, N. Ozaki, A. Mizutani, Y. Takata, J. Jensen, O. Sigmund, P. Borel, M. Kristensen, and K. Asakawa, "Topology optimised photonic crystal waveguide intersections with high-transmittance and low crosstalk," Electron. Lett. **42**, 1031–1032 (2006).
4. L. Frandsen, Y. Elesin, O. Sigmund, J. Jensen, and K. Yvind, "Wavelength selective 3D topology optimized photonic crystal devices," *CLEO: Science and Innovations*, June 9, 2013 (Optical Society of America), paper CTh4L-6.
5. Y. Elesin, B. Lazarov, J. Jensen, and O. Sigmund, "Time domain topology optimization of 3D nanophotonic devices," Photon. Nanostr. Fundam. Appl. **12**, 23–33 (2014).
6. A. Piggott, J. Lu, T. Babinec, K. Lagoudakis, J. Petykiewicz, and J. Vučković, "Inverse design and implementation of a wavelength demultiplexing grating coupler," Sci. Rep. **4**, 7210 (2014).
7. A. Piggott, J. Lu, K. Lagoudakis, J. Petykiewicz, T. Babinec, and J. Vučković, "Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer," Nat. Photonics **9**, 374–377 (2015).
8. J. Lu and J. Vučković, "Objective-first design of high-efficiency, small-footprint couplers between arbitrary nanophotonic waveguide modes," Opt. Express **20**, 7221–7236 (2012).
9. L. Frandsen, Y. Elesin, L. Frellsen, M. Mitrovic, Y. Ding, O. Sigmund, and K. Yvind, "Topology optimized mode conversion in a photonic crystal waveguide fabricated in silicon-on-insulator material," Opt. Express **22**, 8525–8532 (2014).
10. Y. Jiao, S. Fan, and D. Miller, "Systematic photonic crystal device design: global and local optimization and sensitivity analysis," IEEE J. Quantum Electron. **42**, 266–279 (2006).

11. V. Liu and S. Fan, "Compact bends for multi-mode photonic crystal waveguides with high transmission and suppressed modal crosstalk," Opt. Express **21**, 8069–8075 (2013).
12. S. Verweij, V. Liu, and S. Fan, "Accelerating simulation of ensembles of locally differing optical structures via a Schur complement domain decomposition," Opt. Lett. **39**, 6458–6461 (2014).
13. V. Liu, Y. Jiao, D. Miller, and S. Fan, "Design methodology for compact photonic-crystal-based wavelength division multiplexers," Opt. Lett. **36**, 591–593 (2011).
14. X. Sheng, S. Johnson, J. Michel, and L. Kimerling, "Optimization-based design of surface textures for thin-film Si solar cells," Opt. Express **19**, A841–A850 (2011).
15. C. Lin and M. Povinelli, "Optimal design of aperiodic, vertical silicon nanowire structures for photovoltaics," Opt. Express **19**, A1148–A1154 (2011).
16. V. Ganapati, O. Miller, and E. Yablonovitch, "Light trapping textures designed by electromagnetic optimization for subwavelength thick solar cells," IEEE J. Photovoltaics **4**, 175–182 (2014).
17. N. Sergeant, O. Pincon, M. Agrawal, and P. Peumans, "Design of wide-angle solar-selective absorbers using aperiodic metal-dielectric stacks," Opt. Express **17**, 22800–22812 (2009).
18. P. Bermel, M. Ghebrebrhan, W. Chan, Y. Yeng, M. Araghchini, R. Hamam, C. Marton, K. Jensen, M. Soljačić, J. Joannopoulos, S. Johnson, and I. Celanovic, "Design and global optimization of high-efficiency thermophotovoltaic systems," Opt. Express **18**, A314–A334 (2010).
19. P. Borel, A. Harpøth, L. Frandsen, M. Kristensen, P. Shi, J. Jensen, and O. Sigmund, "Topology optimization and fabrication of photonic crystal structures," Opt. Express **12**, 1996–2001 (2004).
20. J. Jensen and O. Sigmund, "Topology optimization for nano-photonics," Laser Photon. Rev. **5**, 308–321 (2011).
21. 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**, 235–244 (2005).
22. C. Lalau-Keraly, S. Bhargava, O. Miller, and E. Yablonovitch, "Adjoint shape optimization applied to electromagnetic design," Opt. Express **21**, 21693–21701 (2013).
23. J. Lu and J. Vučković, "Inverse design of nanophotonic structures using complementary convex optimization," Opt. Express **18**, 3793–3804 (2010).
24. J. Lu and J. Vučković, "Nanophotonic computational design," Opt. Express **21**, 13351–13367 (2013).
25. J. Andkjaer, V. Johansen, K. Friis, and O. Sigmund, "Inverse design of nanostructured surfaces for color effects," J. Opt. Soc. Am. B **31**, 164–174 (2014).
26. D. Liu, L. Gabrielli, M. Lipson, and S. Johnson, "Transformation inverse design," Opt. Express **21**, 14223–14243 (2013).
27. F. Poletti, V. Finazzi, T. Monro, N. Broderick, V. Tse, and D. Richardson, "Inverse design and fabrication tolerances of ultra-flattened dispersion holey fibers," Opt. Express **13**, 3728–3736 (2005).
28. A. Oskooi, A. Mutapcic, S. Noda, J. Joannopoulos, S. Boyd, and S. Johnson, "Robust optimization of adiabatic tapers for coupling to slow-light photonic-crystal waveguides," Opt. Express **20**, 21558–21575 (2012).
29. S. Bhargava and E. Yablonovitch, "Lowering HAMR near-field transducer temperature via inverse electromagnetic design," IEEE Trans. Magn. **51**, 3100407 (2015).
30. A. Gondarenko and M. Lipson, "Low modal volume dipole-like dielectric slab resonator," Opt. Express **16**, 17689–17694 (2008).
31. P. Seliger, M. Mahvash, C. Wang, and A. Levi, "Optimization of aperiodic dielectric structures," J. Appl. Phys. **100**, 034310 (2006).
32. W. Frei, D. Tortorelli, and H. Johnson, "Geometry projection method for optimizing photonic nanostructures," Opt. Lett. **32**, 77–79 (2007).
33. A. Hakansson, P. Sanchis, J. Sanchez-Dehesa, and J. Marti, "High-efficiency defect-based photonic-crystal tapers designed by a genetic algorithm," J. Lightwave Technol. **23**, 3881–3888 (2005).
34. E. Kerrinckx, L. Bigot, M. Douay, and Y. Quiquempois, "Photonic crystal fiber design by means of a genetic algorithm," Opt. Express **12**, 1990–1995 (2004).
35. V. Liu, D. Miller, and S. Fan, "Highly tailored computational electromagnetics methods for nanophotonic design and discovery," Proc. IEEE **101**, 484–493 (2013).
36. M. Kumar, S. Menabde, S. Yu, and N. Park, "Directional emission from photonic crystal waveguide terminations using particle swarm optimization," J. Opt. Soc. Am. B **27**, 343–349 (2010).
37. J. Jiang, J. Cai, G. Nordin, and L. Li, "Parallel microgenetic algorithm design for photonic crystal and waveguide structures," Opt. Lett. **28**, 2381–2383 (2003).
38. S. Verweij and S. Fan, "Understanding search behavior via search landscape analysis in design optimization of optical structures," J. Opt. Soc. Am. B **33**, 2457–2471 (2016).
39. H. Hoos and T. Stützle, *Stochastic Local Search: Foundations & Applications* (Morgan Kaufmann, 2004).
40. A. Smith and D. Coit, "Penalty functions," in *Evolutionary Computation 2*, T. Bäck, D. Fogel, and Z. Michalewicz, eds. (Taylor & Francis, 2000), pp. 41–48.