

LOW-RANK SPECTRAL OPTIMIZATION VIA GAUGE DUALITY*

MICHAEL P. FRIEDLANDER[†] AND IVES MACÊDO[†]

Abstract. Various applications in signal processing and machine learning give rise to highly structured spectral optimization problems characterized by low-rank solutions. Two important examples that motivate this work are optimization problems from phase retrieval and from blind deconvolution, which are designed to yield rank-1 solutions. An algorithm is described that is based on solving a certain constrained eigenvalue optimization problem that corresponds to the gauge dual which, unlike the more typical Lagrange dual, has an especially simple constraint. The dominant cost at each iteration is the computation of rightmost eigenpairs of a Hermitian operator. A range of numerical examples illustrate the scalability of the approach.

Key words. convex optimization, gauge duality, semidefinite optimization, sparse optimization, low-rank solutions, phase retrieval

AMS subject classifications. 90C15, 90C25

DOI. 10.1137/15M1034283

1. Introduction. There are a number of applications in signal processing and machine learning that give rise to highly structured spectral optimization problems. We are particularly interested in the class of problems characterized by having solutions that are very low rank and by involving linear operators that are best treated by matrix-free approaches. This class of problems is sufficiently restrictive such that it allows us to design specialized algorithms that scale well and lend themselves to practical applications, but it is still sufficiently rich to include interesting problems. Two examples include the nuclear-norm minimization for problems such as blind deconvolution (Ahmed, Recht, and Romberg, 2014) and the PhaseLift formulation of the celebrated phase retrieval problem (Candès, Strohmer, and Voroninski, 2013). The problems can be cast generically in the semidefinite programming (SDP) framework, for which a variety of algorithms are available. However, typical applications can give rise to enormous optimization problems that challenge the very best workhorse algorithms.

Denote the set of complex-valued $n \times n$ Hermitian matrices by \mathcal{H}^n . The algorithm that we propose is designed to solve the problems

$$(1a) \quad \underset{X \in \mathcal{H}^n}{\text{minimize}} \quad \text{trace } X \quad \text{subject to} \quad \|b - \mathcal{A}X\| \leq \epsilon, \quad X \succeq 0,$$

$$(1b) \quad \underset{X \in \mathbb{C}^{n_1 \times n_2}}{\text{minimize}} \quad \|X\|_1 := \sum_i \sigma_i(X) \quad \text{subject to} \quad \|b - \mathcal{A}X\| \leq \epsilon,$$

where the parameter ϵ controls the admissible deviations between the linear model $\mathcal{A}X$ and the vector of observations b . (The particular properties of the vectors b and of the linear operators \mathcal{A} are detailed in section 1.2.) Our approach for both problems

*Submitted to the journal's Methods and Algorithms for Scientific Computing section August 7, 2015; accepted for publication (in revised form) March 16, 2016; published electronically June 1, 2016.

<http://www.siam.org/journals/sisc/38-3/M103428.html>

[†]Department of Computer Science, University of British Columbia, Vancouver, BC, V6T 1Z4, Canada (mpf@cs.ubc.ca, ijamj@cs.ubc.ca). The first author's research was supported by ONR award N00014-16-1-2242 and by an AWS in Education Research grant. The second author's research was supported by NSERC Discovery grant 312104 and NSERC Collaborative Research and Development grant 375142-08.

is based on first solving a related Hermitian eigenvalue optimization problem over a very simple constraint and then using that solution to recover a solution of the original problem. This eigenvalue problem is highly structured, and because the constraint is easily handled, we are free to apply a projected first-order method with inexpensive per-iteration costs that scales well to very large problems.

The key to the approach is to recognize that the problems (1) are members of the family of gauge optimization problems, which admit a duality concept different from the Lagrange duality prevalent in convex optimization. Gauges are nonnegative, positively homogeneous convex functions that vanish at the origin. They significantly generalize the familiar notion of a norm, which is a symmetric gauge function. The class of gauge optimization problems, as defined by Freund’s seminal 1987 work, can be stated simply: find the element of a convex set that is minimal with respect to a gauge. These conceptually simple problems appear in a remarkable array of applications and include an important cross-section of convex optimization. For example, all of conic optimization can be phrased within the class of gauge optimization; see Friedlander, Macêdo, and Pong (2014, Example 1.3) and section 2 below.

Problem (1a) is not explicitly stated as a gauge problem because the objective is not nonnegative everywhere on its domain, as required in order for it to be a gauge function. It is, however, nonnegative on the feasible set, and the problem can easily be cast in the gauge framework simply by changing the objective function to

$$(2) \quad \text{trace } X + \delta(X \mid \cdot \succeq 0), \quad \text{where } \delta(X \mid \cdot \succeq 0) = \begin{cases} 0 & \text{if } X \succeq 0, \\ +\infty & \text{otherwise.} \end{cases}$$

This substitution yields an equivalent problem, and the resulting convex function is nonnegative and positively homogeneous—and therefore a gauge function. More generally, it is evident that any function of the form $\gamma + \delta(\cdot \mid \mathcal{K})$ is a gauge, in which γ is a gauge and $\delta(\cdot \mid \mathcal{K})$ is the indicator of a convex cone \mathcal{K} .

The method that we develop applies to the much broader class of semidefinite optimization problems with nonnegative objective values, as described in section 6. We pay special attention to the low-rank spectral problems just mentioned because they have a special structure that can be exploited both theoretically and computationally.

1.1. Notation. To emphasize the role of the vector of singular values $\sigma(X)$, we adopt the Schatten p -norm notation for the matrix-norms referenced in this paper, i.e., $\|X\|_p := \|\sigma(X)\|_p$. Thus, the nuclear, Frobenius, and spectral norms of a matrix X are denoted by $\|X\|_1$, $\|X\|_2$, and $\|X\|_\infty$, respectively. The notation for complex-valued quantities, particularly in the SDP context, is not entirely standard. Here we define some objects we use frequently. Define the complex inner product $\langle X, Y \rangle := \text{trace } XY^*$, where Y^* is the conjugate transpose of a complex matrix Y , i.e., $Y^* = \bar{Y}^T$. The set of $n \times n$ Hermitian matrices is denoted by \mathcal{H}^n , and $X \succeq 0$ (resp., $X \succ 0$) indicates that the matrix X is both Hermitian and positive semidefinite (resp., definite). Let $\lambda(A)$ be the vector of ordered eigenvalues of $A \in \mathcal{H}^n$, i.e., $\lambda_1(A) \geq \lambda_2(A) \geq \dots \geq \lambda_n(A)$. (An analogous ordering is assumed for the vector of singular values.) For $B \succ 0$, let $\lambda(A, B)$ denote the vector of generalized eigenvalues of the pencil (A, B) , i.e., $\lambda(A, B) = \lambda(B^{-1/2}AB^{-1/2})$. Let $\Re(\cdot)$ and $\Im(\cdot)$ denote the real and imaginary parts of their arguments. The norm dual to $\|\cdot\| : \mathbb{C}^m \rightarrow \mathbb{R}_+$ is defined by

$$(3) \quad \|x\|_* := \sup_{\|z\| \leq 1} \Re\langle z, x \rangle.$$

The positive part of a scalar is denoted by $[\cdot]_+ = \max\{0, \cdot\}$.

When we make reference to *one-* and *two-dimensional* signals, our intent is to differentiate between problems that involve discretized functions of *one* and *two* variables, respectively, rather than to describe the dimension of the ambient space. Hence the terms *two-dimensional signals* and *two-dimensional images* are used interchangeably.

Generally we assume that the problems are feasible, although we highlight in section 2 how to detect infeasible problems. We also assume that $0 \leq \epsilon < \|b\|$, which ensures that the origin is not a trivial solution. In practice, the choice of the norm that defines the feasible set will greatly influence the computational difficulty of the problem. Our implementation is based on the 2-norm, which often appears in many practical applications. Our theoretical developments, however, allow for any norm.

1.2. Problem formulations. Below we describe two applications, in phase retrieval and blind deconvolution, that motivate our work. There are other relevant examples, such as matrix completion (Recht, Fazel, and Parrilo, 2010), but these two applications require optimization problems that exemplify properties we exploit in our approach.

1.2.1. Phase retrieval. The phase retrieval problem is concerned with recovery of the phase information of a signal—e.g., an image—from magnitude-only measurements. One important application is X-ray crystallography, which generates images of the molecular structure of crystals (Harrison, 1993). Other applications are described by Candès, Strohmer, and Voroninski (2013) and Waldspurger, d’Aspremont, and Mallat (2015). They describe the following recovery approach, based on convex optimization.

Magnitude-only measurements of the signal $x \in \mathbb{C}^n$ can be described as quadratic measurements of the form

$$b_k = |\langle x, a_k \rangle|^2$$

for some vectors a_k that encode the waveforms used to illuminate the signal. These quadratic measurements of x can be understood as linear measurements

$$b_k = \langle xx^*, a_k a_k^* \rangle = \langle X, A_k \rangle$$

of the lifted signal $X := xx^*$, where $A_k := a_k a_k^*$ is the k th lifted rank-1 measurement matrix.

In the matrix space, the trace of the unknown lifted signal X acts as a surrogate for the rank function. This is analogous to the 1-norm, which stands as a convex surrogate for counting the number of nonzeros in a vector. This leads us to an optimization problem of the form (1a), where $\mathcal{A} : \mathcal{H}^n \rightarrow \mathbb{R}^m$ is defined by $(\mathcal{A}X)_k := \langle X, A_k \rangle$. The parameter ϵ anticipates noise in the measurements. Candès, Strohmer, and Voroninski (2013) call this the *PhaseLift* formulation. In section 5.1, we give numerical examples for recovering one- and two-dimensional signals with and without noise.

1.2.2. Biconvex compressed sensing and blind deconvolution. The biconvex compressed sensing problem (Ling and Strohmer, 2015) aims to recover two signals from a number of sesquilinear measurements of the form

$$b_k = \langle x_1, a_{1k} \rangle \overline{\langle x_2, a_{2k} \rangle},$$

where $x_1 \in \mathbb{C}^{n_1}$ and $x_2 \in \mathbb{C}^{n_2}$. In the context of blind deconvolution, x_1 and x_2 correspond to coefficients of the signals in some bases. The lifting approached used in

the phase retrieval formulation can again be used, and the measurements of x_1 and x_2 can be understood as coming from linear measurements

$$b_k = \langle x_1 x_2^*, a_{1k} a_{2k}^* \rangle = \langle X, A_k \rangle$$

of the lifted signal $X = x_1 x_2^*$, where $A_k := a_{1k} a_{2k}^*$ is the lifted asymmetric rank-1 measurement matrix. Ahmed, Recht, and Romberg (2014) study conditions on the structure and the number of measurements that guarantee that the original vectors (up to a phase) may be recovered by minimizing the sum of singular values of X subject to the linear measurements. This leads to an optimization problem of the form (1b), where $\mathcal{A} : \mathbb{C}^{n_1 \times n_2} \rightarrow \mathbb{C}^m$ is defined by $(\mathcal{A}X)_k := \langle X, A_k \rangle$.

In section 5.2, we describe a two-dimensional blind deconvolution application from motion deblurring, and there we provide further details on the structure of the measurement operators A_k and report on numerical experiments.

1.3. Reduction to Hermitian SDP. It is convenient, for both our theoretical and algorithmic development, to embed the nuclear-norm minimization problem (1b) within the symmetric SDP (1a). The resulting theory is no less general, and it permits us to solve both problems with what is essentially a single software implementation. The reduction to the Hermitian trace-minimization problem (1a) takes the form

$$\begin{aligned}
 (4) \quad & \underset{\substack{U \in \mathcal{H}^{n_1}, V \in \mathcal{H}^{n_2} \\ X \in \mathbb{C}^{n_1 \times n_2} \\ r_1, r_2 \in \mathbb{R}^m}}{\text{minimize}} & \quad \frac{1}{2} \left\langle \begin{pmatrix} I & 0 \\ 0 & I \end{pmatrix}, \begin{pmatrix} U & X \\ X^* & V \end{pmatrix} \right\rangle \\
 & \text{subject to} & \quad \frac{1}{2} \left\langle \begin{pmatrix} 0 & A_k \\ A_k^* & 0 \end{pmatrix}, \begin{pmatrix} U & X \\ X^* & V \end{pmatrix} \right\rangle + r_{1k} = \Re b_k, \\
 & & \quad \frac{i}{2} \left\langle \begin{pmatrix} 0 & A_k \\ -A_k^* & 0 \end{pmatrix}, \begin{pmatrix} U & X \\ X^* & V \end{pmatrix} \right\rangle + r_{2k} = \Im b_k, \\
 & & \quad \|r_1 + ir_2\| \leq \epsilon, \quad \begin{pmatrix} U & X \\ X^* & V \end{pmatrix} \succeq 0, \quad k = 1, \dots, m.
 \end{aligned}$$

The residual variables $r_1, r_2 \in \mathbb{R}^m$ merely serve to allow the compact presentation above, as they can be eliminated using the equality constraints. The additional variables U and V , at the minimizer, correspond to $(XX^*)^{1/2}$ and $(X^*X)^{1/2}$, respectively; the variable X retains its original meaning. This reduction is based on a well-known reformulation of the nuclear norm as the optimal value of an SDP; see Fazel (2002, Lemma 2) or Recht, Fazel, and Parrilo (2010, Proof of Proposition 2.1).

Although this reduction is convenient for our presentation, it is not strictly necessary, as will be clear from the results in section 6. In fact, the resulting increase in problem size might affect a solver’s performance, as would likely be noticeable if dense linear solvers are employed. Our focus, however, is on large problems that require matrix-free operators and exhibit low-rank solutions. Throughout this paper, we focus entirely on the SDP formulation (1a) without loss of generality.

1.4. Approach. Our strategy for these low-rank spectral optimization problems is based on solving the constrained eigenvalue optimization problem

$$(5) \quad \underset{y \in \mathbb{R}^m}{\text{minimize}} \quad \lambda_1(\mathcal{A}^*y) \quad \text{subject to} \quad \langle b, y \rangle - \epsilon \|y\|_* \geq 1$$

resulting from applying gauge duality (Friedlander, Macêdo, and Pong, 2014; Freund, 1987) to a suitable reformulation of (1a). This is outlined in section 2. The dimension

of the variable y in the eigenvalue optimization problem corresponds to the number of measurements. In the context of phase retrieval and blind deconvolution, Candès and Li (2014) and Ahmed, Recht, and Romberg (2014) show that the number of measurements needed to recover with high probability the underlying signals is within a logarithmic factor of the signal length. The crucial implication is that the dimension of the dual problem grows slowly as compared to the dimension of the primal problem, which grows as the square of the signal length.

In our implementation, we apply a simple first-order projected subgradient method to solve the eigenvalue problem. The dominant cost at each iteration of our algorithm is the computation of rightmost eigenpairs of the $n \times n$ Hermitian linear operator \mathcal{A}^*y , which are used to construct descent directions for (5). The structure of the measurement operators allows us to use Krylov-based eigensolvers, such as ARPACK (Lehoucq, Sorensen, and Yang, 1998), for obtaining these leading eigenpairs. Primal solution estimates X are recovered via a relatively small constrained least-squares problem, described in section 4.

An analogous approach based on the classical Lagrangian duality also leads to a dual optimization problem in the same space as our dual eigenvalue problem:

$$(6) \quad \underset{y \in \mathbb{R}^m}{\text{maximize}} \quad \langle b, y \rangle - \epsilon \|y\|_* \quad \text{subject to} \quad \mathcal{A}^*y \preceq I.$$

Note that the Lagrange dual possesses a rather simple objective and a difficult linear matrix inequality of order n as a constraint. Precisely the reverse situation holds for the gauge dual (5), which has a relatively simple constraint.

It is well known that SDPs with a constant-trace property—i.e., $\mathcal{A}X = b$ implies $\text{trace}(X)$ is constant—have Lagrange dual problems that can be formulated as unconstrained eigenvalue problems. This approach is used by Helmberg and Rendl (2000) to develop a spectral bundle method. The applications that we consider, however, do not necessarily have this property.

1.5. Reproducible research. The data files and MATLAB scripts used to generate the numerical results presented in section 5 can be obtained at the following URL: <http://www.cs.ubc.ca/~mpf/low-rank-opt>.

1.6. Related work. Other researchers have recognized the need for algorithms with low per-iteration costs that scale well for large-scale, low-rank spectral optimization problems. Notable efforts include Hazan (2008), Laue (2012), and Freund, Grigas, and Mazumder (2015), who advocate variations of the Frank–Wolfe (FW) method to solve some version of the problem

$$(7) \quad \underset{X}{\text{minimize}} \quad f(X) \quad \text{subject to} \quad \text{trace}(X) \leq \tau, X \succeq 0,$$

where f is a differentiable function. For example, the choice $f(X) = \frac{1}{2} \|\mathcal{A}X - b\|_2^2$ yields a problem related to (1a). The asymmetric version of the problem with $\|X\|_1 \leq \tau$ is easily accommodated by simply replacing the above constraints. For simplicity, here we focus on the symmetric case, though our approach applies equally to the asymmetric case. The main benefit of using FW for this problem is that each iteration requires only a rightmost eigenvalue of the gradient $\nabla f(X)$, and therefore has the same per-iteration cost of the method that we consider, which requires a rightmost eigenvalue of the same-sized matrix \mathcal{A}^*y . The same Krylov-based eigensolvers apply in both cases.

There are at least two issues that need to be addressed when comparing the FW algorithm to the approach we take here. First, as Freund, Grigas, and Mazumder

(2015) make clear, even in cases where low-rank solutions are expected, it is not possible to anticipate the rank of early iterates X_k generated by the FW method. In particular, they observe that the rank of X_k quickly increases during early iterations and only slowly starts to decrease as the solution is approached. This motivates their development of algorithmic devices that attenuate rank growth in intermediate iterates. Any implementation, however, must be prepared to increase storage for the factors of X_k during intermediate iterates. In contrast, a subgradient method applied to the gauge dual problem can be implemented with constant storage. Second, although in principle there exists a parameter τ that causes the optimization problems (7) and (1a) to share the same solution, this parameter is not generally known in advance. One way around this is to solve a sequence of problems (7) for varying parameters τ_k using, for example, a level-set procedure described by Aravkin et al. (2016).

As an alternative to applying the FW algorithm to (7), we might instead consider applying a variation of the FW method directly to the gauge dual problem (5). Because the gauge dual objective is not differentiable and the feasible set is not compact if $\epsilon = 0$, some modification to the standard FW method is required. Argyriou, Signoretto, and Suykens (2014), Bach (2015), and Nesterov (2015) propose variations of FW that involve smoothing the objective. These smoothing approaches are typically based on infimal convolution with a smooth kernel, which may lead to a function whose gradient is expensive to compute. For example, the “soft-max” smooth approximation of $\lambda_1(\cdot)$ is the function $\mu \log \sum_{i=1, \dots, n} \exp(\lambda_i(\cdot)/\mu)$. Forming the gradient of this smooth function requires computing all eigenvalues of an n -by- n Hermitian matrix.

Laue (2012) proposes a hybrid algorithm that interleaves a nonconvex subproblem within the FW iterations. If the local minimum of the nonconvex subproblem improves the objective value, it is used to replace the current FW iterate. This approach is similar in spirit to the primal-dual refinement that we describe in section 4.3, but because Laue’s method is entirely primal, it has the benefit of not requiring a procedure to feed the improved primal sequence back to the dual sequence.

2. Spectral gauge optimization and duality. The derivation of the eigenvalue optimization problem (5) as a dual to (1a) follows from a more general theory of duality for gauge optimization. Here we provide some minimal background for our derivations related to spectral optimization; see Freund (1987) and Friedlander, Macêdo, and Pong (2014) for fuller descriptions. We begin with a general description of the problem class.

Let $\kappa : \mathcal{X} \mapsto \mathbb{R} \cup \{+\infty\}$ and $\rho : \mathcal{Y} \mapsto \mathbb{R} \cup \{+\infty\}$ be gauge functions, where $A : \mathcal{X} \mapsto \mathcal{Y}$ is a linear operator that maps between the finite-dimensional real inner-product spaces \mathcal{X} and \mathcal{Y} . The polar

$$f^\circ(y) := \inf \{ \mu > 0 \mid \langle x, y \rangle \leq \mu f(x) \ \forall x \}$$

of a gauge f plays a key role in the duality of gauge problems. The problems

$$(8a) \quad \underset{x \in \mathcal{X}}{\text{minimize}} \quad \kappa(x) \quad \text{subject to} \quad \rho(b - Ax) \leq \epsilon,$$

$$(8b) \quad \underset{y \in \mathcal{Y}}{\text{minimize}} \quad \kappa^\circ(A^*y) \quad \text{subject to} \quad \langle b, y \rangle - \epsilon \rho^\circ(y) \geq 1$$

are dual to each other in the following sense: all primal-dual feasible pairs (x, y) satisfy the weak-duality relationship

$$(9) \quad 1 \leq \kappa(x) \kappa^\circ(A^*y).$$

Moreover, a primal-dual feasible pair is optimal if this holds with equality. This strong-duality relationship provides a certificate of optimality.

The SDP problem (1a) can be cast into the mold of the canonical gauge formulation (8a) by using the redefined objective (2) and making the identifications

$$\kappa(X) = \text{trace } X + \delta(X \mid \cdot \succeq 0) \quad \text{and} \quad \rho(r) = \|r\|.$$

We use the polar calculus described by Friedlander, Macêdo, and Pong (2014, section 7.2.1) together with the definition of the dual norm to obtain the corresponding polar functions:

$$\kappa^\circ(Y) = [\lambda_1(Y)]_+ \quad \text{and} \quad \rho^\circ(y) = \|y\|_*.$$

It then follows from (8) that the following are a dual gauge pair:

$$(10a) \quad \underset{X \in \mathcal{H}^n}{\text{minimize}} \quad \text{trace } X + \delta(X \mid \cdot \succeq 0) \quad \text{subject to} \quad \|b - \mathcal{A}X\| \leq \epsilon,$$

$$(10b) \quad \underset{y \in \mathbb{R}^m}{\text{minimize}} \quad [\lambda_1(\mathcal{A}^*y)]_+ \quad \text{subject to} \quad \langle b, y \rangle - \epsilon \|y\|_* \geq 1.$$

The derivation of gauge dual problems relies on the polarity operation applied to gauges. When applied to a norm, for example, the polar is simply the dual norm. In contrast, Lagrange duality is intimately tied to conjugacy, which is what gives rise to the dual problem (6). Of course, the two operations are closely related. For any gauge function κ , for example, $\kappa^*(y) = \delta_{\kappa^\circ(\cdot) \leq 1}(y)$. These relationships are described in detail by Rockafellar (1970, section 15) and Friedlander, Macêdo, and Pong (2014, section 2.3).

We can simplify the dual objective and safely eliminate the positive-part operator: because $\kappa(X)$ is necessarily strictly positive for all nonzero X , and is additionally finite over the feasible set of the original problem (1a), it follows from (9) that $\kappa^\circ(\mathcal{A}^*y)$ is positive for all dual feasible points. In other words,

$$(11) \quad 0 < [\lambda_1(\mathcal{A}^*y)]_+ = \lambda_1(\mathcal{A}^*y)$$

for all dual feasible points y . Hence we obtain the equivalent dual problem (5).

In practice, we need to be prepared to detect whether the primal problem (10a) is infeasible. The failure of condition (11) in fact furnishes a certificate of infeasibility for (1a): if $\lambda(\mathcal{A}^*y) = 0$ for some dual-feasible vector y , it follows from (9) that $\kappa(X)$ is necessarily infinite over the feasible set of (10a)—i.e., $X \not\succeq 0$ for all X feasible for (10a). Thus, (1a) is infeasible.

3. Derivation of the approach. There are two key theoretical pieces needed for our approach. The first is the derivation of the eigenvalue optimization problem (5), as shown in section 2. The second piece is the derivation of a subproblem that allows recovery of a primal solution X from a solution of the eigenvalue problem (5).

3.1. Recovering a primal solution. Our derivation of a subproblem for primal recovery proceeds in two stages. The first stage develops necessary and sufficient optimality conditions for the primal-dual gauge pair (10a) and (10b). The second stage uses these to derive a subproblem that can be used to recover a primal solution from a dual solution.

3.1.1. Strong duality and optimality conditions. The weak duality condition (9) holds for all primal-feasible pairs (X, y) . The following result asserts that if the pair is optimal, then that inequality must necessarily hold tightly.

PROPOSITION 1 (strong duality). *If (1a) is feasible and $0 \leq \epsilon < \|b\|$, then*

$$(12) \quad \left[\min_{\substack{X \in \mathcal{H}^n \\ \|b - \mathcal{A}X\| \leq \epsilon}} \text{trace } X + \delta(X \mid \cdot \succeq 0) \right] \cdot \left[\inf_{\substack{y \in \mathbb{R}^m \\ \langle b, y \rangle - \epsilon \|y\|_* \geq 1}} [\lambda_1(\mathcal{A}^*y)]_+ \right] = 1.$$

Proof. We proceed by reasoning about the Lagrangian-dual pair (1a) and (6). We then translate these results to the corresponding gauge-dual pair (10a) and (10b).

The primal problem (1a) is feasible by assumption. Because its Lagrange dual problem (6) admits strictly feasible points (e.g., $y = 0$), it follows from Rockafellar (1970, Theorems 28.2 and 28.4) that the primal problem attains its positive minimum value and that there is zero duality gap between the Lagrange-dual pair.

Moreover, because the primal problem (1a) attains its positive minimum value for some \widehat{X} , and there is zero duality gap, there exists a sequence $\{y_j\}$ such that $[\lambda_1(\mathcal{A}^*y_j)]_+ \leq 1$ and $\langle y_j, b \rangle - \epsilon \|y_j\|_* \nearrow \text{trace } \widehat{X}$. Because $\text{trace } \widehat{X} > 0$, we can take a subsequence $\{y_{j_k}\}$ for which $\langle y_{j_k}, b \rangle - \epsilon \|y_{j_k}\|_*$ is uniformly bounded above zero. Define the sequence $\{\widehat{y}_k\}$ by $\widehat{y}_k := y_{j_k} (\langle y_{j_k}, b \rangle - \epsilon \|y_{j_k}\|_*)^{-1}$. Then $\langle \widehat{y}_k, b \rangle - \epsilon \|\widehat{y}_k\|_* = 1$ for all k , which is a feasible sequence for the gauge dual problem (10b). Weak gauge duality (9) and the definition of \widehat{y}_k then imply that

$$(\text{trace } \widehat{X})^{-1} \leq [\lambda_1(\mathcal{A}^*\widehat{y}_k)]_+ \leq (\langle y_{j_k}, b \rangle - \epsilon \|y_{j_k}\|_*)^{-1} \searrow (\text{trace } \widehat{X})^{-1}.$$

Multiply the series of inequalities by $\text{trace } \widehat{X}$ to obtain (12). □

Note the lack of symmetry in the statement of Proposition 1: the primal problem is stated with a “min” but the dual problem is stated with an “inf”. This is because the dual Slater condition—i.e., strict feasibility of the corresponding Lagrange-dual problem (6)—allows us to assert that a primal optimal solution necessarily exists. However, we cannot assert in general that a dual optimal solution exists because the corresponding primal feasible set does not necessarily satisfy the Slater condition.

Although in this work we do not attempt to delineate conditions under which dual attainment holds, a practical case in which it always does is when the primal objective is a norm and the measurement operator is surjective. In that case, the dual gauge objective $\|\mathcal{A}^* \cdot\|_*$ defines a norm in \mathbb{R}^m , which has compact level sets. Hence a dual solution always exists. We comment further on this theoretical question in section 7.

The following result characterizes gauge primal-dual optimal pairs. It relies on von Neumann’s trace inequality: for Hermitian matrices A and B ,

$$\langle A, B \rangle \leq \langle \lambda(A), \lambda(B) \rangle,$$

and equality holds if and only if A and B admit a simultaneous ordered eigendecomposition, i.e., $A = U \text{Diag}[\lambda(A)]U^*$ and $B = U \text{Diag}[\lambda(B)]U^*$ for some unitary matrix U ; see Lewis (1996).

PROPOSITION 2 (optimality conditions). *If (1a) is feasible and $0 \leq \epsilon < \|b\|$, then $(X, y) \in \mathcal{H}^n \times \mathbb{R}^m$ is primal-dual optimal for the gauge dual pair (10a) and (10b) if and only if the following conditions hold:*

1. $X \succeq 0$ and $\|b - \mathcal{A}X\| = \epsilon$;
2. $\langle y, b \rangle - \epsilon \|y\|_* = 1$;
3. $\langle y, b - \mathcal{A}X \rangle = \|y\|_* \|b - \mathcal{A}X\|$;
4. $\lambda_i(X) \cdot (\lambda_1(\mathcal{A}^*y) - \lambda_i(\mathcal{A}^*y)) = 0, i = 1, \dots, n$;
5. X and \mathcal{A}^*y admit a simultaneous ordered eigendecomposition.

Proof. By strong duality (Proposition 1), the pair $(X, y) \in \mathcal{H}^m \times \mathbb{R}^m$ is primal-dual optimal if and only if they are primal-dual feasible and the product of their corresponding objective values is equal to one. In this case,

$$\begin{aligned}
 \text{(strong duality)} \quad & 1 = [\text{trace } X + \delta(X \mid \cdot \succeq 0)] \cdot [\lambda_1(\mathcal{A}^*y)]_+ \\
 & = \langle e, \lambda(X) \rangle \cdot \lambda_1(\mathcal{A}^*y) \\
 & = \langle \lambda_1(\mathcal{A}^*y) \cdot e, \lambda(X) \rangle \\
 (\lambda_1(\mathcal{A}^*y) \geq \lambda_i(\mathcal{A}^*y) \text{ and } X \succeq 0) \quad & \geq \langle \lambda(\mathcal{A}^*y), \lambda(X) \rangle \\
 \text{(von Neumann's trace inequality)} \quad & \geq \langle \mathcal{A}^*y, X \rangle \\
 & = \langle y, \mathcal{A}X \rangle \\
 & = \langle y, b \rangle - \langle y, b - \mathcal{A}X \rangle \\
 \text{(Cauchy-Schwarz inequality)} \quad & \geq \langle y, b \rangle - \|y\|_* \|b - \mathcal{A}X\| \\
 \text{(primal feasibility)} \quad & \geq \langle y, b \rangle - \epsilon \|y\|_* \\
 \text{(dual feasibility)} \quad & \geq 1.
 \end{aligned}$$

Thus all of the above inequalities hold with equality. This proves conditions 1–4. Condition 5 follows from again invoking von Neumann’s trace inequality and noting its implication that X and \mathcal{A}^*y share a simultaneous ordered eigenvalue decomposition. Sufficiency of those conditions can be verified by simply following the reverse chain of reasoning and again noticing that the inequalities can be replaced by equalities. \square

3.2. Primal recovery subproblem. The optimality conditions that are stated in Proposition 2 furnish the means for deriving a subproblem that can be used to recover a primal solution from a dual solution. The next result establishes an explicit relationship between primal solutions X and \mathcal{A}^*y for an arbitrary optimal dual solution y .

COROLLARY 3. *Suppose that the conditions of Proposition 2 hold. Let $y \in \mathbb{R}^m$ be an arbitrary optimal solution for the dual gauge program (10b), $r_1 \in \{1, \dots, n\}$ be the multiplicity of $\lambda_1(\mathcal{A}^*y)$, and $U_1 \in \mathbb{C}^{n \times r_1}$ be the matrix formed by the first r_1 eigenvectors of \mathcal{A}^*y . Then a matrix $X \in \mathcal{H}^n$ is a solution for the primal problem (10a) if and only if there exists an $r_1 \times r_1$ matrix $S \succeq 0$ such that*

$$(13) \quad X = U_1 S U_1^* \quad \text{and} \quad (b - \mathcal{A}X) \in \epsilon \partial \|\cdot\|_*(y).$$

Proof. The assumptions imply that the optimal dual value is positive. If $y \in \mathbb{R}^m$ is an optimal solution to (10b), the positive homogeneity of its objective and constraint, and the positivity of the optimal value, allow us to deduce that the dual constraint must be active, i.e., $\langle y, b \rangle - \epsilon \|y\|_* = 1$. Thus condition 2 of Proposition 2 holds.

The construction of X in (13) guarantees that it shares a simultaneous ordered eigendecomposition with \mathcal{A}^*y and that it has rank of r_1 at most. Thus, conditions 4 and 5 of Proposition 2 hold.

We now show that conditions 1 and 3 of the proposition hold. The subdifferential $\partial \|\cdot\|_*$ corresponds to the set of maximizers of the linear function that defines the

dual ball; see (3). Then because $(b - \mathcal{A}X) \in \epsilon \partial \|\cdot\|_*(y)$, it holds that $\|b - \mathcal{A}X\| \leq \epsilon$ and $\epsilon \|y\|_* = \langle y, b - \mathcal{A}X \rangle \leq \|y\|_* \|b - \mathcal{A}X\| \leq \epsilon \|y\|_*$, implying that $\|b - \mathcal{A}X\| = \epsilon$ and $\langle y, b - \mathcal{A}X \rangle = \|y\|_* \|b - \mathcal{A}X\|$. This way, conditions 1 and 3 of Proposition 2 are also satisfied. Hence, all the conditions of the proposition are satisfied, and the pair $(X, y) \in \mathcal{H}^n \times \mathbb{R}^m$ is optimal.

Suppose now that $X \in \mathcal{H}^n$ is optimal for (10a). We can invoke Proposition 2 on the pair $(X, y) \in \mathcal{H}^n \times \mathbb{R}^m$. Condition 4 implies that any eigenvector of \mathcal{A}^*y associated to an eigenvalue $\lambda_i(\mathcal{A}^*y)$ with $i > r_1$ is in the nullspace of X ; therefore, there is an $r_1 \times r_1$ matrix $S \succeq 0$ such that $X = U_1 S U_1^*$. Conditions 1 and 3 imply that $\|b - \mathcal{A}X\| \leq \epsilon$ and $\langle y, b - \mathcal{A}X \rangle = \epsilon \|y\|_*$, thus verifying that $(b - \mathcal{A}X) \in \epsilon \partial \|\cdot\|_*(y)$, as required. \square

Corollary 3 thus provides us with a way to recover a solution to our model problem (1a) after computing a solution to the gauge dual problem (10a). When the residual in (1a) is measured in the 2-norm, condition (13) simplifies and implies that the matrix S that defines $X = U S U^*$ can be obtained by solving

$$(14) \quad \underset{S \succeq 0}{\text{minimize}} \quad \|\mathcal{A}(U_1 S U_1^*) - b_\epsilon\|^2, \quad \text{with } b_\epsilon := b - \epsilon y / \|y\|.$$

When the multiplicity r_1 of the eigenvalue $\lambda_1(\mathcal{A}^*y)$ is much smaller than n , this optimization problem is relatively inexpensive. In particular, if $r_1 = 1$ —which may be expected in some applications such as PhaseLift—the optimization problem is over a scalar s that can be obtained immediately as

$$s = [\langle \mathcal{A}(u_1 u_1^*), b_\epsilon \rangle]_+ / \|\mathcal{A}(u_1 u_1^*)\|^2,$$

where u_1 is the rightmost eigenvalue of \mathcal{A}^*y . This approach exploits the complementarity relation on eigenvalues in condition 4 of Proposition 2 to reduce the dimensionality of the primal solution recovery. Its computational difficulty effectively depends on finding a dual solution y at which the rightmost eigenvalue has low multiplicity r_1 .

4. Implementation. The success of our approach hinges on efficiently solving the constrained eigenvalue optimization problem (6) in order to generate solution estimates y and rightmost eigenvector estimates U_1 of \mathcal{A}^*y that we can feed to (14). The two main properties of this problem that drive our approach are that it has a nonsmooth objective and that projections on the feasible set are inexpensive. Our implementation is based on a basic projected-subgradient descent method, although certainly other choices are available. For example, Nesterov (2009) and Richtárik (2011) propose specialized algorithms for minimizing positively homogeneous functions with affine constraints; some modification of this approach could possibly apply to (6). Another possible choice is the spectral bundle method of Helmberg and Rendl (2000). For simplicity, and because it has proven sufficient for our needs, we use a standard projected subgradient method, described below.

4.1. Dual descent. The generic subgradient method is based on the iteration

$$(15) \quad y_+ = \mathcal{P}(y - \alpha g),$$

where g is a subgradient of the objective at the current iterate y , α is a positive steplength, and the operator $\mathcal{P} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ gives the Euclidean projection onto the feasible set. For the objective function $f(y) = \lambda_1(\mathcal{A}^*y)$ of (5), the subdifferential has the form

$$(16) \quad \partial f(y) = \{ \mathcal{A}(U_1 T U_1^*) \mid T \succeq 0, \text{ trace } T = 1 \},$$

where U_1 is the $n \times r_1$ matrix of rightmost eigenvectors of \mathcal{A}^*y ; see Overton (1992, Theorem 3). A Krylov-based eigenvalue solver can be used to evaluate $f(y)$ and a subgradient $g \in \partial f(y)$. Such methods require products of the form $(\mathcal{A}^*y)v$ for arbitrary vectors v . In many cases, these products can be computed without explicitly forming the matrix \mathcal{A}^*y . In particular, for the applications described in section 1.2, these products can be computed entirely using fast operators such as the FFT. Similar efficiencies can be used to compute a subgradient g from the forward map $\mathcal{A}(U_1TU_1^*)$.

For large problems, further efficiencies can be obtained simply by computing a single eigenvector u_1 , i.e., any unit-norm vector in the range of U_1 . In our implementation, we typically request at least *two* rightmost eigenpairs: this gives us an opportunity to detect whether the leading eigenpair is isolated. If it is, then the sub-differential contains only a single element, which implies that f is differentiable at that point.

Any sequence of steplengths $\{\alpha_k\}$ that satisfies the generic conditions

$$\lim_{k \rightarrow \infty} \alpha_k = 0, \quad \sum_{k=0}^{\infty} \alpha_k = \infty$$

is sufficient to guarantee that the value of the objective at y_k converges to the optimal value (Bertsekas, 2015, Proposition 3.2.6). A typical choice is $\alpha_k = \mathcal{O}(1/k)$. Our implementation defaults to a Barzilai–Borwein steplength (Barzilai and Borwein, 1988) with a nonmonotonic linesearch (Zhang and Hager, 2004) if it is detected that a sequence of iterates is differentiable (by observing separation of the leading eigenpair); and otherwise it falls back to a decreasing step size.

The projection operator \mathcal{P} onto the dual-feasible set (5) is inexpensive when the residual is measured in the 2-norm. In particular, if $\epsilon = 0$, the dual-feasible set is a halfspace, and the projection can be accomplished in linear time. When ϵ is positive, the projection requires computing the roots of a one-dimensional degree-4 polynomial, which in practice requires little additional time.

4.2. Primal recovery. At each iteration of the descent method (15) for the eigenvalue optimization problem (5), we compute a corresponding primal estimate

$$(17) \quad X_+ = U_1 S_+ U_1^*$$

maintained in factored form. The matrix U_1 has already been computed in the evaluation of the objective and its subgradient; see (16). The positive semidefinite matrix S_+ is the solution of the primal-recovery problem (14).

A byproduct of the primal-recovery problem is that it provides a suitable stopping criterion for the overall algorithm. Because the iterations y_k are dual feasible, it follows from Corollary 3 that if (14) has a zero residual, then the dual iterate y_k and the corresponding primal iterate X_k are optimal. Thus, we use the size of the residual to determine a stopping test for approximate optimality.

4.3. Primal-dual refinement. The primal-recovery procedure outlined in section 4.2 is used only as a stopping criterion and does not directly affect the sequence of dual iterates from (15). In our numerical experiments, we find that significant gains can be had by refining the primal estimate (17) and feeding it back into the dual sequence. We use the following procedure, which involves two auxiliary subproblems that add relatively little to the overall cost.

The first step is to refine the primal estimate obtained via (14) by using its solution to determine the starting point $Z_0 = U_1 S_+^{1/2}$ for the smooth unconstrained nonconvex

problem

$$(18) \quad \underset{Z \in \mathbb{C}^{n \times r}}{\text{minimize}} \quad h(Z) := \frac{1}{4} \|\mathcal{A}(ZZ^*) - b_\epsilon\|^2.$$

In effect, we continue to minimize (14), where additionally U_1 is allowed to vary. Several options are available for solving this smooth unconstrained problem. Our implementation has the option of using a steepest-descent iteration with a spectral steplength and nonmonotone linesearch (Zhang and Hager, 2004), or a limited-memory BFGS method (Nocedal and Wright, 2006, section 7.2). The main cost at each iteration is the evaluation of the gradient

$$(19) \quad \nabla h(Z) = \mathcal{A}^*(\mathcal{A}(ZZ^*) - b_\epsilon)Z.$$

We thus obtain a candidate improved primal estimate $\widehat{X} = \widehat{Z}\widehat{Z}^*$, where \widehat{Z} is a solution of (18). When $\epsilon = 0$, this nonconvex problem coincides with the problem used by Candès, Li, and Soltanolkotabi (2015). They use the initialization $Z_0 = \gamma u_1$, where u_1 is a leading eigenvector of \mathcal{A}^*b , and $\gamma = n \sum_i b_i / \sum_i \|a_i\|^2$. Our initialization, on the other hand, is based on a solution of the primal-recovery problem (14).

The second step of the refinement procedure is to construct a candidate dual estimate \widehat{y} from a solution of the constrained linear-least-squares problem

$$(20) \quad \underset{y \in \mathbb{R}^m}{\text{minimize}} \quad \frac{1}{2} \|(\mathcal{A}^*y)\widehat{Z} - \widehat{\lambda}\widehat{Z}\|^2 \quad \text{subject to} \quad \langle b, y \rangle - \epsilon \|y\|_* \geq 1,$$

where $\widehat{\lambda} := 1/\text{trace } \widehat{X} \equiv 1/\|\widehat{Z}\|_F^2$ is the reciprocal of the primal objective value associated with \widehat{X} . This constrained linear-least-squares problem attempts to construct a vector \widehat{y} such that the columns of \widehat{Z} correspond to eigenvectors of $\mathcal{A}^*\widehat{y}$ associated with $\widehat{\lambda}$. If $f(\widehat{y}) < f(y_+)$, then \widehat{y} improves on the current dual iterate y_+ obtained by the descent method (15), and we are free to use \widehat{y} in its place. This improved estimate, which is exogenous to the dual descent method, can be considered a “spacer” iterate, as described by Bertsekas (1999, Proposition 1.2.6). Importantly, it does not interfere with the convergence of the underlying descent method. The projected-descent method used to solve the dual sequence can also be applied to (20), though in this case the objective is guaranteed to be differentiable.

4.4. Algorithm summary. The steps at the end of this section summarize one iteration of the dual-descent algorithm: y is the current dual iterate, and y^+ is the updated iterate. The primal iterate X is maintained in factored form. Steps 5–7 implement the primal-dual refinement strategy described in section 4.3 and constitute a heuristic that may improve the performance of the dual descent algorithm without sacrificing convergence guarantees.

In step 1, the rightmost eigenpair (λ_1, U_1) of \mathcal{A}^*y is computed. The eigenvectors in the matrix U_1 are used in step 2 to compute a subgradient g for the dual objective. Any PSD matrix T that has trace equal to 1 can be used in step 2. For example, the case where only a single rightmost eigenvector u_1 can be afforded corresponds to setting T so that $U_1 T U_1^* = u_1 u_1^*$. Step 3 is a projected subgradient iteration with steplength α . Step 4 solves the primal-recovery problem to determine the matrix S_+ used to define a primal estimate $X_+ = U_1 S_+ U_1^*$; cf. (13). We use the factorization $Z_0 := U_1 S_+^{1/2}$ to initialize the algorithm in the next step. Step 5 applies an algorithm to the nonlinear least-squares problem (18) to obtain a stationary point \widehat{Z} used to define the dual-refinement problem used in the next step. Step 6 computes a candidate dual solution \widehat{y} that—if it improves the dual objective—is used to replace the latest dual estimate y_+ .

1	$(\lambda_1, U_1) \leftarrow \lambda_1(\mathcal{A}^*y)$	[eigenvalue computation]
2	$g \leftarrow \mathcal{A}(U_1 T U_1^*)$	[gradient of dual objective; cf. (16)]
3	$y^+ \leftarrow \mathcal{P}(y - \alpha g)$	[projected subgradient step; cf. (15)]
4	$S_+ \leftarrow$ solution of (14)	[primal recovery subproblem]
5	$\widehat{Z} \leftarrow$ solution of (18) initialized with S_+	[primal refinement]
6	$\widehat{y} \leftarrow$ solution of (20) initialized with \widehat{Z}	[dual refinement]
	if $\lambda_1(\mathcal{A}^*\widehat{y}) < \lambda_1$ then	
7	$y^+ \leftarrow \widehat{y}$	[spacer step]
	end	

5. Numerical experiments. This section reports on a set of numerical experiments for solving instances of the phase retrieval and blind deconvolution problems described in section 1.2. The various algorithmic pieces described in section 4 have been implemented as a MATLAB software package. The implementation uses the MATLAB `eigs` routine for the eigenvalue computations described in section 4.1. We implemented a projected gradient-descent method, which is used for solving (15), (18), and (20).

5.1. Phase recovery. We conduct three experiments for phase retrieval via the PhaseLift formulation. The first experiment is for a large collection of small one-dimensional random signals, and is meant to contrast the approach against a general-purpose convex optimization algorithm and a specialized nonconvex approach. The second experiment tests problems where the vector of observations b is contaminated by noise, hence testing the case where $\epsilon > 0$. The third experiment tests the scalability of the approach on a large two-dimensional natural image.

Our phase retrieval experiments follow the approach outlined in Candès, Li, and Soltanolkotabi (2015). The diagonal matrices $C_k \in \mathbb{C}^{n \times n}$ encode diffraction patterns that correspond to the k th “mask” ($k = 1, \dots, L$) through which a signal $x_0 \in \mathbb{C}^n$ is measured. The measurements are given by

$$b = \mathcal{A}(x_0 x_0^*) := \text{diag} \left[\begin{pmatrix} FC_1 \\ \vdots \\ FC_L \end{pmatrix} (x_0 x_0^*) \begin{pmatrix} FC_1 \\ \vdots \\ FC_L \end{pmatrix}^* \right],$$

where F is the unitary discrete Fourier transform (DFT). The adjoint of the associated linear map \mathcal{A} is then

$$\mathcal{A}^*y := \sum_{k=1}^L C_k^* F^* \text{Diag}(y_k) F C_k,$$

where $y = (y_1, \dots, y_L)$ and $\text{Diag}(y_k)$ is the diagonal matrix formed from the vector y_k . The main cost in the evaluation of the forward map $\mathcal{A}(VV^*)$ involves L applications of the DFT for each column of V . Each evaluation of the adjoint map applied to a vector v —i.e., $(\mathcal{A}^*y)v$ —requires L applications of both the DFT and its inverse. In the experimental results reported below, the columns labeled “nDFT” indicate the total number of DFT evaluations used over the course of a run. The costs of these

TABLE 1

Phase retrieval comparisons for random complex Gaussian signals of size $n = 128$ measured using random complex Gaussian masks. Numbers of the form n_{-e} are shorthand for $n \cdot 10^{-e}$.

L	GAUGE		GAUGE-feas		TFOCS			WFLOW		
	nDFT	xErr	nDFT	xErr	nDFT	xErr	%	nDFT	xErr	%
12	18,330	1.6 ₋₆	3,528	1.3 ₋₆	2,341,800	3.6 ₋₃	100	5,232	1.2 ₋₅	100
11	19,256	1.5 ₋₆	3,344	1.4 ₋₆	2,427,546	4.3 ₋₃	100	4,906	1.6 ₋₅	100
10	19,045	1.4 ₋₆	3,120	1.6 ₋₆	2,857,650	5.5 ₋₃	100	4,620	2.1 ₋₅	100
9	21,933	1.6 ₋₆	2,889	1.4 ₋₆	$1.2 \cdot 10^7$	7.5 ₋₃	89	4,374	2.5 ₋₅	100
8	23,144	2.1 ₋₆	2,688	1.9 ₋₆	$1.1 \cdot 10^7$	1.2 ₋₂	22	4,080	3.3 ₋₅	100
7	25,781	1.8 ₋₆	2,492	2.0 ₋₆	6,853,245	2.4 ₋₂	0	3,836	5.2 ₋₅	95
6	34,689	3.0 ₋₆	2,424	2.5 ₋₆	2,664,126	6.4 ₋₂	0	3,954	9.5 ₋₅	62

DFT evaluations are invariant across the different algorithms and dominate the overall computation.

5.1.1. Random Gaussian signals. In this section we consider a set of experiments for different numbers of masks. For each value of $L = 6, 7, \dots, 12$, we generate a fixed set of 100 random complex Gaussian vectors x_0 of length $n = 128$ and a set of L random complex Gaussian masks C_k .

Table 1 summarizes the results of applying four different solvers to each set of 100 problems. The solver **GAUGE** is our implementation of the approach summarized in section 4.4; **TFOCS** (Becker, Candès, and Grant, 2011) is a first-order conic solver applied to the primal problem (1a). The version used here was modified to avoid explicitly forming the matrix \mathcal{A}^*y (Strohmer, 2013). The algorithm **WFLOW** (Candès, Li, and Soltanolkotabi, 2015) is a nonconvex approach that attempts to recover the original signal directly from the feasibility problem (18), with $\epsilon = 0$. To make sensible performance comparisons to **WFLOW**, we add to its implementation a stopping test based on the norm of the gradient (19); the default algorithm otherwise uses a fixed number of iterations.

We also show the results of applying the **GAUGE** code in a “feasibility” mode that exits as soon as the primal-refinement subproblem (see step 7 of the algorithm summary in section 4.4) obtains a solution with a small residual. This resulting solver is labeled **GAUGE-feas**. This variant of **GAUGE** is in some respects akin to **WFLOW**, with the main difference that **GAUGE-feas** uses starting points generated by the dual-descent estimates and generates search directions and step-lengths for the feasibility problem from a spectral gradient algorithm. The columns labeled “xErr” report the median relative error $\|x_0 x_0^* - \widehat{x} \widehat{x}^*\|_F / \|x_0\|_2^2$ of the 100 runs, where \widehat{x} is the solution returned by the corresponding solver. The columns labeled “%” give the percentage of problems solved to within a relative error of 10^{-2} . At least on this set of artificial experiments, the **GAUGE** solver (and its feasibility variant **GAUGE-feas**) appear to be most efficient. Table 2 provides an additional comparison of **GAUGE** with the variation **GAUGE-nodfp**, which ignores the “spacer” iterate computed by (20). There seems to be significant practical benefit in using the refined primal estimate to improve the dual sequence. The columns labeled “%” are excluded for all versions of **GAUGE** because these solvers obtained the prescribed accuracy for all problems in each test set.

Note that the relative accuracy “xErr” is often slightly better for **GAUGE-feas** than for **GAUGE**. These small discrepancies are explained by the different stopping criteria between the two versions of the solver. In particular, **GAUGE** will continue iterating

TABLE 2
Additional comparisons for the random examples of Table 1.

L	GAUGE		GAUGE-nodfp	
	nDFT	xErr	nDFT	xErr
12	18,330	1.6 ₋₆	277,722	1.6 ₋₆
11	19,256	1.5 ₋₆	314,820	1.6 ₋₆
10	19,045	1.4 ₋₆	374,190	2.0 ₋₆
9	21,933	1.6 ₋₆	485,658	1.9 ₋₆
8	23,144	2.1 ₋₆	808,792	1.9 ₋₆
7	25,781	1.8 ₋₆	2,236,885	2.3 ₋₆
6	34,689	3.0 ₋₆	14,368,437	2.9 ₋₆

past the point at which **GAUGE-feas** normally terminates because it is searching for a dual certificate that corresponds to the recovered primal estimate. This slightly changes the computed subspaces U_1 , which influence subsequent primal estimates. Similar behavior is exhibited in the noisy cases that we consider in the next section.

As the number of measurements (L) decreases, we expect the problem to be more difficult. Indeed, we can observe that the total amount of work, as measured by the number of operator evaluations (i.e., the ratio between **nDFT** and L), increases monotonically for all variations of **GAUGE**.

5.1.2. Random problems with noise. In this set of experiments, we assess the effectiveness of the SDP solver to problems with $\epsilon > 0$, which could be useful in recovering signals with noise. For this purpose, it is convenient to generate problem instances with noise and known primal-dual solutions, which we can do by using Corollary 3. Each instance is generated by first sampling octanary masks C_k —as described by Candès, Li, and Soltanolkotabi (2015)—and real Gaussian vectors $y \in \mathbb{R}^m$; a solution $x_0 \in \mathbb{C}^n$ is then chosen as a unit-norm rightmost eigenvector of \mathcal{A}^*y , and the measurements are computed as $b := \mathcal{A}(x_0x_0^*) + \epsilon y/\|y\|$, where ϵ is chosen as $\epsilon := \|b - \mathcal{A}(x_0x_0^*)\| = \eta\|b\|$ for a given noise-level parameter $\eta \in (0, 1)$.

For these experiments, we generate 100 instances with $n = 128$ for each pairwise combination (L, η) with $L \in \{6, 9, 12\}$ and $\eta \in \{0.1\%, 0.5\%, 1\%, 5\%, 10\%, 50\%\}$. Table 3 summarizes the results of applying the three variations of **GAUGE**, and the **WFLOW** solver, to these problems. It is not clear that **GAUGE-feas** and **WFLOW** are relevant for this experiment, but for interest we include them in the results. As with the experiments in section 5.1.1, a solve is “successful” if it recovers the true solution with a relative error of 10^{-2} . The median relative error for all solvers is comparable, and hence we omit the column “xErr.” **GAUGE-nodfp** is generally successful in recovering the rank-1 minimizer for most problems—even for cases with significant noise, though in these cases the overall cost increases considerably. On the other hand, **GAUGE** is less successful: it appears that although the primal-dual refinement procedure can help to reduce the cost of successful recovery in low-noise settings, in high-noise settings it may obtain primal solutions that are not necessarily close to the true signal. For noise levels over 5%, **GAUGE-feas** and **WFLOW** are unable to recover a solution within the prescribed accuracy, which points to the benefits of the additional cost of obtaining a primal-dual optimal point, rather than just a primal feasible point.

5.1.3. Two-dimensional signal. We conduct a second experiment on a stylized application in order to assess the scalability of the approach to larger problem sizes. In this case the measured signal x_0 is a two-dimensional image of size 1600×1350

TABLE 3

Phase retrieval comparisons for problems with noise, i.e., $\epsilon > 0$. Numbers of the form n_{-e} are shorthand for $n \cdot 10^{-e}$.

L	η	GAUGE		GAUGE-nodfp		GAUGE-feas		WFLOW	
		nDFT	%	nDFT	%	nDFT	%	nDFT	%
12	0.1%	4,584	100	29,988	100	936	100	14,856	100
9	0.1%	3,222	100	36,292	100	774	100	11,511	100
6	0.1%	2,232	100	50,235	100	612	100	8,922	98
12	0.5%	3,768	100	27,252	100	936	100	14,808	100
9	0.5%	2,934	100	31,032	100	774	100	11,430	100
6	0.5%	2,148	100	38,766	100	606	100	8,790	98
12	1.0%	3,744	100	22,620	97	936	100	14,712	100
9	1.0%	2,934	100	24,813	96	774	100	11,331	99
6	1.0%	$1 \cdot 10^7$	97	$2 \cdot 10^5$	98	600	53	8,634	8
12	5.0%	95,952	26	$9 \cdot 10^5$	90	936	0	14,148	0
9	5.0%	$2 \cdot 10^6$	89	$8 \cdot 10^5$	90	774	0	10,701	0
6	5.0%	$2 \cdot 10^6$	82	$5 \cdot 10^5$	98	600	0	7,728	0
12	10.0%	$1 \cdot 10^5$	17	$1 \cdot 10^6$	89	912	0	13,548	0
9	10.0%	$8 \cdot 10^5$	78	$7 \cdot 10^5$	91	765	0	10,125	0
6	10.0%	$7 \cdot 10^5$	90	$5 \cdot 10^5$	100	588	0	7,098	0
12	50.0%	$2 \cdot 10^5$	53	$5 \cdot 10^5$	94	888	0	11,424	0
9	50.0%	$1 \cdot 10^5$	42	$3 \cdot 10^5$	99	738	0	8,586	0
6	50.0%	$1 \cdot 10^5$	24	$2 \cdot 10^5$	95	588	0	7,176	0



FIG. 1. Image used for phase retrieval experiment; size 1600×1350 pixels (7.5MB).

pixels, shown in Figure 1, which corresponds to $n = 2.2 \cdot 10^6$. The size of the lifted formulation is on the order of $n^2 \approx 10^{12}$, which makes it clear that the resulting SDP is enormous and must be handled by a specialized solver. We have excluded TFOCS from the list of candidate solvers because it cannot make progress on this example. We generate 10 and 15 octonary masks. Table 4 summarizes the results. The column headers carry the same meaning as Table 1.

5.2. Blind deconvolution. In this blind deconvolution experiment, the convolution of two signals $s_1 \in \mathbb{C}^m$ and $s_2 \in \mathbb{C}^m$ is measured. Let $B_1 \in \mathbb{C}^{m \times n_1}$ and $B_2 \in \mathbb{C}^{m \times n_2}$ be two bases. The circular convolution of the signals can be described

TABLE 4
Phase retrieval comparisons on a two-dimensional image.

L	GAUGE		GAUGE-feas		WFLOW	
	nDFT	xErr	nDFT	xErr	nDFT	xErr
15	200,835	$2.1 \cdot 10^{-6}$	5,700	$2.1 \cdot 10^{-6}$	8,100	$4.1 \cdot 10^{-6}$
10	195,210	$5.8 \cdot 10^{-7}$	12,280	$9.1 \cdot 10^{-7}$	12,340	$2.1 \cdot 10^{-5}$

by

$$\begin{aligned} b &= s_1 * s_2 = (B_1 x_1) * (B_2 x_2) \\ &= F^{-1} \text{diag}((FB_1 x_1)(FB_2 x_2)^T) \\ &= F^{-1} \text{diag}((FB_1)(x_1 \bar{x}_2^*)(\overline{FB_2})^*) =: \mathcal{A}(x_1 \bar{x}_2^*), \end{aligned}$$

where \mathcal{A} is the corresponding asymmetric linear map with the adjoint

$$\mathcal{A}^* y := (FB_1)^* \text{Diag}(Fy)(\overline{FB_2}).$$

Because F is unitary, it is possible to work instead with measurements

$$\hat{b} \equiv Fb = \text{diag}((FB_1)(x_1 \bar{x}_2^*)(\overline{FB_2})^*)$$

in the Fourier domain. For the experiments that we run, we choose to work with the former real-valued measurements b because they do not require accounting for the imaginary parts, and thus the number of constraints in (4) that would be required otherwise is reduced by half.

We follow the experimental setup outlined by Ahmed, Recht, and Romberg (2014) and use the data and code that they provide. In that setup, B_1 is a subset of the Haar wavelet basis, and B_2 is a mask corresponding to a subset of the identity basis. The top row of Figure 2 shows the original image, the blurring kernel, and the observed blurred image. The second row of the figure shows the image reconstructed using the augmented Lagrangian code provided by Ahmed, Recht, and Romberg (2014), GAUGE, and GAUGE-feas. Table 5 summarizes the results of applying the three solvers. The columns headed “nDFT” and “nDWT” count the number of discrete Fourier and wavelet transforms required by each solver; the columns headed “xErr1” and “xErr2” report the relative errors $\|x_i - \hat{x}_i\|_2 / \|x_i\|_2$, $i = 1, 2$, where \hat{x}_i are the recovered solutions; the column headed “rErr” reports the relative residual error $\|b - \mathcal{A}(\hat{x}_1 \bar{\hat{x}}_2^*)\|_2 / \|b\|_2$. Although the nonconvex augmented Lagrangian approach yields visibly better recovery of the image, the table reveals that the solutions are numerically similar and are recovered with far less work.

6. Extensions. The problems (1) that we have considered so far are stated in their simplest form. General semidefinite optimization problems with nonnegative optimal value, and reweighted formulations for rank minimization, as introduced by Mohan and Fazel (2010) and Candès et al. (2013), are also useful and can be accommodated by our approach.

In the context of rank minimization over the PSD cone, an approximate minimum-rank solution \hat{X} (e.g., computed via trace minimization) might be used to obtain an even better approximation by using the weighted objective $\langle C, X \rangle$, where $C := (\delta I + \hat{X})^{-1}$ and δ is a small positive parameter. We might reasonably expect that

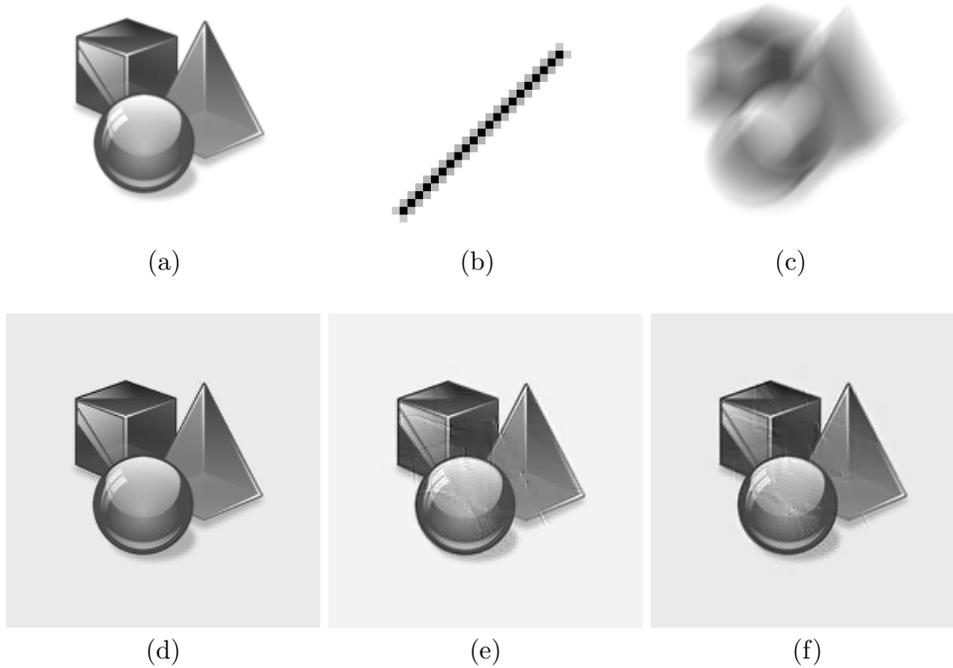


FIG. 2. Images used for the blind deconvolution experiments: (a) original image; (b) zoom of the motion-blurring kernel; (c) blurred image; (d) image recovered by the augmented Lagrangian approach; (e) image recovered by GAUGE; (f) image recovered by GAUGE-feas. The shaded background on the recovered images is an artifact of the uniform scaling used to highlight any error between the original and recovered signals.

TABLE 5
Blind deconvolution comparisons.

Solver	nDFT	nDWT	xErr1	xErr2	rErr
aug Lagrangian	92,224	76,872	$7.9 \cdot 10^{-2}$	$5.0 \cdot 10^{-1}$	$1.4 \cdot 10^{-4}$
GAUGE	17,432	8,374	$1.9 \cdot 10^{-2}$	$5.4 \cdot 10^{-1}$	$3.8 \cdot 10^{-4}$
GAUGE-feas	4,128	2,062	$8.4 \cdot 10^{-2}$	$5.5 \cdot 10^{-1}$	$4.0 \cdot 10^{-4}$

the objective value at a minimizer of this objective more closely matches the rank function. Candès et al. (2013) show that such an iteratively reweighted sequence of trace minimization problems can improve the range of signals recoverable using PhaseLift. Each problem in that sequence uses the previous solution \hat{X} to derive a weighting matrix $C = (\delta I + \hat{Z}\hat{Z}^*)^{-1}$ for the next problem. The inverse of the matrix C is a low-rank update $\hat{Z}\hat{Z}^* \approx \hat{X}$ to a small regularizing multiple δ of the identity matrix. This idea generalizes that of reweighting the 1-norm for cardinality minimization problems in compressed sensing, where the number of nonzero entries of a vector x is approximated by $\sum_i |x_i| / (|\hat{x}_i| + \delta)$ for small δ and an available approximation \hat{x} (e.g., computed via 1-norm minimization).

In the next sections we derive the corresponding gauge duals for the weighted formulations of both trace minimization in the PSD cone and nuclear-norm mini-

mization.

6.1. Nonnegative semidefinite optimization. Consider the semidefinite optimization problem

$$(21) \quad \underset{X \in \mathcal{H}^n}{\text{minimize}} \quad \langle C, X \rangle \quad \text{subject to} \quad \|b - \mathcal{A}X\| \leq \epsilon, \quad X \succeq 0,$$

where $C \succ 0$. Define the maps

$$\mathcal{C}(\cdot) := C^{-\frac{1}{2}}(\cdot)C^{-\frac{1}{2}} \quad \text{and} \quad \mathcal{A}_C := \mathcal{A} \circ \mathcal{C}^{-1}.$$

It is evident that $X \succeq 0$ if and only if $\mathcal{C}(X) \succeq 0$, and so the SDP problem can be stated equivalently as

$$\underset{X \in \mathcal{H}^n}{\text{minimize}} \quad \text{trace } \mathcal{C}(X) \quad \text{subject to} \quad \|b - \mathcal{A}_C(\mathcal{C}(X))\| \leq \epsilon, \quad \mathcal{C}(X) \succeq 0.$$

Because \mathcal{C} is a bijection, we can optimize over $\widehat{X} := \mathcal{C}(X)$ instead of X :

$$(22) \quad \underset{\widehat{X} \in \mathcal{H}^n}{\text{minimize}} \quad \text{trace } \widehat{X} \quad \text{subject to} \quad \|b - \mathcal{A}_C \widehat{X}\| \leq \epsilon, \quad \widehat{X} \succeq 0.$$

This clearly falls within the structure of (1a) and has the corresponding gauge dual

$$(23) \quad \underset{y \in \mathbb{R}^m}{\text{minimize}} \quad [\lambda_1(\mathcal{A}_C^* y)]_+ \quad \text{subject to} \quad \langle b, y \rangle - \epsilon \|y\|_* \geq 1.$$

Observe that $\lambda_1(\mathcal{A}_C^* y) = \lambda_1(C^{-\frac{1}{2}}(\mathcal{A}^* y)C^{-\frac{1}{2}}) = \lambda_1(\mathcal{A}^* y, C)$. Then

$$(24) \quad \underset{y \in \mathbb{R}^m}{\text{minimize}} \quad [\lambda_1(\mathcal{A}^* y, C)]_+ \quad \text{subject to} \quad \langle b, y \rangle - \epsilon \|y\|_* \geq 1.$$

This shows that the introduction of a weighting matrix C that is not a simple multiple of the identity leads to a dual gauge problem involving the minimization of the rightmost *generalized* eigenvalue of $\mathcal{A}^* y$ with respect to that weight. Now that we have a formulation for the gauge dual problem, we focus on how a primal solution to the original weighted trace minimization can be computed given a dual minimizer. This extends Corollary 3.

COROLLARY 4. *Suppose that problem (21) is feasible and $0 \leq \epsilon < \|b\|$. Let $y \in \mathbb{R}^m$ be an arbitrary optimal solution for the dual gauge (24), $r_1 \in \{1, \dots, n\}$ be the multiplicity of $\lambda_1(\mathcal{A}^* y, C)$, and $U_1 \in \mathbb{C}^{n \times r_1}$ be the matrix formed by the first r_1 generalized eigenvectors of $\mathcal{A}^* y$ with respect to C . Then $X \in \mathcal{H}^n$ is a solution for the primal problem (21) if and only if there exists $S \succeq 0$ such that*

$$X = U_1 S U_1^* \quad \text{and} \quad (b - \mathcal{A}X) \in \epsilon \partial \|\cdot\|_*(y).$$

Proof. A solution for (24) is clearly a solution for (23). We may thus invoke Corollary 3 and assert that $\widehat{X} \in \mathcal{H}^n$ is a solution for (22) if and only if there is $S \succeq 0$ such that $\widehat{X} = \widehat{U}_1 S \widehat{U}_1^*$ and $(b - \mathcal{A}_C \widehat{X}) \in \epsilon \partial \|\cdot\|_*(y)$, where $\widehat{U}_1 \in \mathbb{C}^{n \times r_1}$ is a matrix formed by the first r_1 eigenvectors of $\mathcal{A}_C^* y = C^{-\frac{1}{2}}(\mathcal{A}^* y)C^{-\frac{1}{2}}$. From the structure of \mathcal{C} , we have that X is a solution to (21) if and only if $X = \mathcal{C}(\widehat{X})$. Thus, $X = C^{-\frac{1}{2}} \widehat{U}_1 S \widehat{U}_1^* C^{-\frac{1}{2}} = U_1 S U_1^*$, where $U_1 := C^{-\frac{1}{2}} \widehat{U}_1$ corresponds to the first r_1 generalized eigenvectors of $\mathcal{A}^* y$ with respect to C . \square

Once again, this provides us with a way to recover a solution to the weighted trace minimization problem by computing a solution to the gauge dual problem (now involving the rightmost generalized eigenvalue) and then solving a problem of potentially much reduced dimensionality.

6.2. Weighted affine nuclear-norm optimization. We can similarly extend the reweighted extension to the asymmetric case (1b). Let $C_1 \in \mathcal{H}^{n_1}$ and $C_2 \in \mathcal{H}^{n_2}$ be invertible. The weighted nuclear-norm minimization problem becomes

$$(25) \quad \underset{X \in \mathbb{C}^{n_1 \times n_2}}{\text{minimize}} \quad \|C_1 X C_2^*\|_1 \quad \text{subject to} \quad \|b - \mathcal{A}X\| \leq \epsilon.$$

Define the weighted quantities

$$\mathcal{C}(\cdot) = C_1^{-1}(\cdot)C_2^{-*} : \mathbb{C}^{n_1 \times n_2} \rightarrow \mathbb{C}^{n_1 \times n_2}, \quad \mathcal{A}_C = \mathcal{A} \circ \mathcal{C}, \quad \text{and} \quad \widehat{X} := \mathcal{C}(X).$$

The weighted problem can then be stated equivalently as

$$\underset{\widehat{X} \in \mathbb{C}^{n_1 \times n_2}}{\text{minimize}} \quad \|\widehat{X}\|_1 \quad \text{subject to} \quad \|b - \mathcal{A}_C \widehat{X}\| \leq \epsilon,$$

which, following the approach introduced in Fazel (2002), can be embedded in a symmetric problem:

$$(26) \quad \underset{\substack{\widehat{U} \in \mathcal{H}^{n_1} \\ \widehat{V} \in \mathcal{H}^{n_2} \\ \widehat{X} \in \mathbb{C}^{n_1 \times n_2}}}{\text{minimize}} \quad \left\langle \frac{1}{2}I, \begin{pmatrix} \widehat{U} & \widehat{X} \\ \widehat{X}^* & \widehat{V} \end{pmatrix} \right\rangle$$

$$\text{subject to} \quad \begin{pmatrix} \widehat{U} & \widehat{X} \\ \widehat{X}^* & \widehat{V} \end{pmatrix} \succeq 0 \quad \text{and} \quad \|b - \mathcal{A}_C \widehat{X}\| \leq \epsilon.$$

Define the measurement operator from $\mathcal{M} : \mathcal{H}^{n_1+n_2} \rightarrow \mathbb{C}^m$ by the map

$$\begin{pmatrix} \widehat{U} & \widehat{X} \\ \widehat{X}^* & \widehat{V} \end{pmatrix} \mapsto \mathcal{A}_C \widehat{X},$$

and identify \mathbb{C}^m with \mathbb{R}^{2m} as a real inner-product space. The adjoint of the measurement operator is then given by

$$\mathcal{M}^* y = \begin{pmatrix} 0 & \mathcal{A}_C^* y \\ (\mathcal{A}_C^* y)^* & 0 \end{pmatrix},$$

where $\mathcal{A}_C^* y = \frac{1}{2} \sum_{i=1}^m C_1^{-1} A_i C_2^{-*} y_i$. We can now state the gauge dual problem:

$$(27) \quad \underset{y \in \mathbb{C}^m}{\text{minimize}} \quad \left[\lambda_1(\mathcal{M}^* y, \frac{1}{2}I) \right]_+ \quad \text{subject to} \quad \Re \langle b, y \rangle - \epsilon \|y\|_* \geq 1.$$

Observe the identity

$$\begin{aligned} \lambda_1(\mathcal{M}^* y, \frac{1}{2}I) &= \lambda_1(2\mathcal{M}^* y) \\ &= \left[\lambda_1 \left(\begin{pmatrix} 0 & \sum_{i=1}^m C_1^{-1} A_i C_2^{-*} y_i \\ (\sum_{i=1}^m C_1^{-1} A_i C_2^{-*} y_i)^* & 0 \end{pmatrix} \right) \right]_+ \\ &= \left[\|C_1^{-1}(\mathcal{A}^* y)C_2^{-*}\|_\infty \right]_+ = \|C_1^{-1}(\mathcal{A}^* y)C_2^{-*}\|_\infty. \end{aligned}$$

We can now deduce the simplified form for the gauge dual problem:

$$(28) \quad \underset{y \in \mathbb{C}^m}{\text{minimize}} \quad \|C_1^{-1}(\mathcal{A}^* y)C_2^{-*}\|_\infty \quad \text{subject to} \quad \Re \langle b, y \rangle - \epsilon \|y\|_* \geq 1.$$

This weighted gauge dual problem can be derived from first principles using the tools from section 2 by observing that the primal problem is already in standard gauge form. We chose this approach, however, to make explicit the close connection between the (weighted) nuclear-norm minimization problem and the (weighted) trace-minimization problem described in section 6.1.

The following result provides a way to characterize solutions of the nuclear-norm minimization problem when a solution to the dual gauge problem is available.

COROLLARY 5. *Suppose that problem (25) is feasible and $0 \leq \epsilon < \|b\|$. Let $y \in \mathbb{C}^m$ be an arbitrary optimal solution for the dual gauge problem (28), $r_1 \in \{1, \dots, n\}$ be the multiplicity of $\sigma_1(C_1^{-1}(\mathcal{A}^*y)C_2^{-*})$, $U_1 \in \mathbb{C}^{n_1 \times r_1}$ and $V_1 \in \mathbb{C}^{n_2 \times r_1}$ be the matrices formed by the first r_1 left and right singular vectors of $C_1^{-1}(\mathcal{A}^*y)C_2^{-*}$, respectively. Then $X \in \mathbb{C}^{n_1 \times n_2}$ is a solution for the primal problem (25) if and only if there exists $S \succeq 0$ such that $X = (C_1^{-1}U_1)S(C_2^{-1}V_1)^*$ and $(b - \mathcal{A}X) \in \epsilon \partial \|\cdot\|_*(y)$.*

Proof. A solution for (28) is clearly a solution for (27); in this way we invoke Corollary 3 and have that $(\widehat{U}, \widehat{V}, \widehat{X}) \in \mathcal{H}^{n_1} \times \mathcal{H}^{n_2} \times \mathbb{C}^{n_1 \times n_2}$ induce a solution for (26) if and only if there is $S \succeq 0$ such that $\widehat{X} = \widehat{U}_1 S \widehat{V}_1^*$ and $(b - \mathcal{A}_C \widehat{X}) \in \epsilon \partial \|\cdot\|_*(y)$, where $\widehat{U}_1 \in \mathbb{C}^{n_1 \times r_1}$ and $\widehat{V}_1 \in \mathbb{C}^{n_2 \times r_1}$ are matrices formed by the first r_1 left and right singular vectors of $\mathcal{A}_C^* y = C_1^{-1}(\mathcal{A}^*y)C_2^{-*}$. From the structure of \mathcal{C} , we have that X is a solution to (25) if and only if $X = \mathcal{C}(\widehat{X})$. This way, $X = (C_1^{-1}\widehat{U}_1)S(C_2^{-1}\widehat{V}_1)^*$. \square

7. Conclusions. The phase retrieval and blind deconvolution applications are examples of convex relaxations of nonconvex problems that give rise to large spectral optimization problems with strong statistical guarantees for correctly reconstructing certain signals. One of the criticisms that has been leveled at these relaxation approaches is that they lead to problems that are too difficult to be useful in practice. This has led to work on nonconvex recovery algorithms that may not have as-strong statistical recovery guarantees but are nonetheless effective in practice; see Netrapalli, Jain, and Sanghavi (2013); Candès, Li, and Soltanolkotabi (2015); White, Sanghavi, and Ward (2015). Our motivation is to determine whether it is possible to develop convex optimization algorithms that are as efficient as nonconvex approaches. The numerical experiments on these problems suggest that the gauge-dual approach may prove effective. Indeed, other convex optimization algorithms may be possible, and clearly the key to their success will be to leverage the special structure of these problems.

A theoretical question we have not addressed is to delineate conditions under which dual attainment will hold. In particular, the conclusion (12) of Theorem 1 is asymmetric: we can assert that a primal solution exists that attains the primal optimal value (because the Lagrange dual is strictly feasible), but we cannot assert that a dual solution exists that attains the dual optimal value. A related theoretical question is to understand the relationship between the quality of suboptimal dual solutions, and the quality of the primal estimate obtained by the primal recovery procedure.

In our experiments, we have observed that the rightmost eigenvalue of \mathcal{A}^*y remains fairly well separated from the others across iterations. This seems to contribute to the overall effectiveness of the dual-descent method. Is there a special property of these problems or of the algorithm that encourages this separation property? It seems likely that there are solutions y at which the objective is not differentiable, and in that case, we wonder whether there are algorithmic devices that could be used to avoid such points.

The dual-descent method we use to solve the dual subproblem (cf. section 4.1) is only one possible algorithm among many. Other more specialized methods, such as the spectral bundle method of Helmberg and Rendl (2000), its second-order variant (Helmberg, Overton, and Rendl, 2014), or the stochastic-gradient method of d’Aspremont and El Karoui (2014), may prove effective alternatives.

We have found it convenient to embed the nuclear-norm minimization problem (1b) in the SDP formulation (1a) because it allows us to use the same solver for both problems. Further efficiencies, however, may be gained by implementing a solver that applies directly to the corresponding gauge dual

$$\underset{y \in \mathbb{C}^m}{\text{minimize}} \quad \|\mathcal{A}^* y\|_\infty \quad \text{subject to} \quad \Re\langle b, y \rangle - \epsilon \|y\|_* \geq 1.$$

This would require an iterative solver for evaluating leading singular values and singular vectors of the asymmetric operator $\mathcal{A}^* y$, such as PROPACK (Larsen, 2001).

Acknowledgments. We extend sincere thanks to our colleague Nathan Krislock, who was involved in an earlier incarnation of this project, and to our colleague Ting Kei Pong, who was our partner in establishing the crucial gauge duality theory in Friedlander, Macêdo, and Pong (2014). We are also grateful to Xiaodong Li and Mahdi Soltanolkotabi for help with using their WFLOW code. Finally, we wish to thank the two anonymous referees who provided a careful list of comments and suggestions that helped to clarify our presentation.

REFERENCES

- A. AHMED, B. RECHT, AND J. ROMBERG (2014), *Blind deconvolution using convex programming*, IEEE Trans. Inform. Theory, 60, pp. 1711–1732.
- A. Y. ARAVKIN, J. V. BURKE, D. DRUSVYATSKIY, M. P. FRIEDLANDER, AND S. ROY (2016), *Level-Set Methods for Convex Optimization*, arXiv:1602.01506 [math.OC].
- A. ARGYRIOU, M. SIGNORETTO, AND J. SUYKENS (2014), *Hybrid conditional gradient-smoothing algorithms with applications to sparse and low rank regularization*, in Regularization, Optimization, Kernels, and Support Vector Machines, CRC Press, Boca Raton, FL, pp. 53–82.
- F. BACH (2015), *Duality between subgradient and conditional gradient methods*, SIAM J. Optim., 25, pp. 115–129, <http://dx.doi.org/10.1137/130941961>.
- J. BARZILAI AND J. M. BORWEIN (1988), *Two-point step size gradient methods*, IMA J. Numer. Anal., 8, pp. 141–148.
- S. BECKER, E. J. CANDÈS, AND M. GRANT (2011), *Templates for convex cone problems with applications to sparse signal recovery*, Math. Program. Comput., 3, pp. 165–218.
- D. P. BERTSEKAS (1999), *Nonlinear Programming*, Athena Scientific, 2nd ed., Belmont, MA.
- D. P. BERTSEKAS (2015), *Convex Optimization Algorithms*, Athena Scientific, Belmont, MA.
- E. J. CANDÈS, Y. C. ELДАР, T. STROHMER, AND V. VORONINSKI (2013), *Phase retrieval via matrix completion*, SIAM J. Imaging Sci., 6, pp. 199–225, <http://dx.doi.org/10.1137/110848074>.
- E. J. CANDÈS AND X. LI (2014), *Solving quadratic equations via PhaseLift when there are about as many equations as unknowns*, Found. Comput. Math., 14, pp. 1017–1026.
- E. J. CANDÈS, X. LI, AND M. SOLTANOLKOTABI (2015), *Phase retrieval via Wirtinger flow: Theory and algorithms*, IEEE Trans. Inform. Theory, 61, pp. 1985–2007.
- E. J. CANDÈS, T. STROHMER, AND V. VORONINSKI (2013), *PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming*, Comm. Pure. Appl. Math. 66, pp. 1241–1274.
- A. D’ASPREMONT AND N. EL KAROUI (2014), *A stochastic smoothing algorithm for semidefinite programming*, SIAM J. Optim., 24, pp. 1138–1177, <http://dx.doi.org/10.1137/12088728X>.
- M. FAZEL (2002), *Matrix Rank Minimization with Applications*, Ph.D. thesis, Electrical Engineering Department, Stanford University, Stanford, CA.
- R. M. FREUND (1987), *Dual gauge programs, with applications to quadratic programming and the minimum-norm problem*, Math. Program., 38, pp. 47–67.
- R. M. FREUND, P. GRIGAS, AND R. MAZUMDER (2015), *An Extended Frank-Wolfe Method with “In-Face” Directions, and its Application to Low-Rank Matrix Completion*, arXiv:1511.02204

- [math.OC].
- M. P. FRIEDLANDER, I. MACÊDO, AND T. K. PONG (2014), *Gauge optimization and duality*, SIAM J. Optim., 24, pp. 1999–2022, <http://dx.doi.org/10.1137/130940785>.
- R. W. HARRISON (1993), *Phase problem in crystallography*, J. Opt. Soc. Am. A, 10, pp. 1046–1055.
- E. HAZAN (2008), *Sparse approximate solutions to semidefinite programs*, in LATIN 2008: Theoretical Informatics, Springer, Berlin, pp. 306–316.
- C. HELMBERG AND F. RENDL (2000), *A spectral bundle method for semidefinite programming*, SIAM J. Optim., 10, pp. 673–696, <http://dx.doi.org/10.1137/S1052623497328987>.
- C. HELMBERG, M. L. OVERTON, AND F. RENDL (2014), *The spectral bundle method with second-order information*, Optim. Methods Softw., 29, pp. 855–876.
- R. M. LARSEN (2001), *Combining Implicit Restart and Partial Reorthogonalization in Lanczos Bidiagonalization*, <http://sun.stanford.edu/~rmunk/PROPACK/>.
- S. LAUE (2012), *A hybrid algorithm for convex semidefinite optimization*, in Proceedings of the 29th International Conference on Machine Learning (ICML-12).
- R. B. LEHOUCQ, D. C. SORENSEN, AND C. YANG (1998), *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*, Software Environ. Tools 6, SIAM, Philadelphia, <http://dx.doi.org/10.1137/1.9780898719628>.
- A. S. LEWIS (1996), *Convex analysis on the Hermitian matrices*, SIAM J. Optim., 6, pp. 164–177, <http://dx.doi.org/10.1137/0806009>.
- S. LING AND T. STROHMER (2015), *Self-calibration and biconvex compressive sensing*, Inverse Problems, 31, 115002.
- K. MOHAN AND M. FAZEL (2010), *Reweighted nuclear norm minimization with application to system identification*, in Proceedings of the American Control Conference (ACC), pp. 2953–2959.
- Y. NESTEROV (2009), *Unconstrained convex minimization in relative scale*, Math. Oper. Res., 34, pp. 180–193.
- Y. NESTEROV (2015), *Complexity Bounds for Primal-Dual Methods Minimizing the Model of Objective Function*, Tech. rep., Center for Operations Research and Econometrics, Louvain-la-Neuve, Belgium.
- P. NETRAPALLI, P. JAIN, AND S. SANGHAVI (2013), *Phase retrieval using alternating minimization*, in Advances in Neural Information Processing Systems 26, C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, eds., The MIT Press, Cambridge, MA, pp. 2796–2804.
- J. NOCEDAL AND S. J. WRIGHT (2006), *Numerical Optimization*, 2nd ed., Springer, New York.
- M. L. OVERTON (1992), *Large-scale optimization of eigenvalues*, SIAM J. Optim., 2, pp. 88–120, <http://dx.doi.org/10.1137/0802007>.
- B. RECHT, M. FAZEL, AND P. A. PARRILO (2010), *Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization*, SIAM Rev., 52, pp. 471–501, <http://dx.doi.org/10.1137/070697835>.
- P. RICHTÁRIK (2011), *Improved algorithms for convex minimization in relative scale*, SIAM J. Optim., 21, pp. 1141–1167, <http://dx.doi.org/10.1137/090747142>.
- R. T. ROCKAFELLAR (1970), *Convex Analysis*, Princeton University Press, Princeton, NJ.
- T. STROHMER (2013), *private communication*.
- I. WALDSPURGER, A. D'ASPREMONT, AND S. MALLAT (2015), *Phase recovery, maxcut and complex semidefinite programming*, Math. Program., 149, pp. 47–81.
- C. D. WHITE, S. SANGHAVI, AND R. WARD (2015), *The Local Convexity of Solving Systems of Quadratic Equations*, arXiv:1506.07868 [math.NA].
- H. ZHANG AND W. W. HAGER (2004), *A nonmonotone line search technique and its application to unconstrained optimization*, SIAM J. Optim., 14, pp. 1043–1056, <http://dx.doi.org/10.1137/S1052623403428208>.