Bounding pixels in computational imaging

Keith Dillon* and Yeshaiahu Fainman

Department of Electrical and Computer Engineering, University of California, San Diego, 9500 Gilman Drive, La Jolla, California 92093-0407, USA

*Corresponding author: kdillon@ucsd.edu

Received 1 November 2012; revised 25 February 2013; accepted 7 March 2013; posted 12 March 2013 (Doc. ID 179122); published 22 March 2013

We consider computational imaging problems where we have an insufficient number of measurements to uniquely reconstruct the object, resulting in an ill-posed inverse problem. Rather than deal with this via the usual regularization approach, which presumes additional information which may be incorrect, we seek bounds on the pixel values of the reconstructed image. Formulating the inverse problem as an optimization problem, we find conditions for which a system's measurements can produce a bounded result for both the linear case and the non-negative case (e.g., intensity imaging). We also consider the problem of selecting measurements to yield the most bounded results. Finally we simulate examples of the application of bounded estimation to different two-dimensional multiview systems. © 2013 Optical Society of America

OCIS codes: 110.1758, 110.6955.

1. Introduction

We approach the computational imaging problem from an inverse systems perspective. We have a system such as in Fig. 1, where light from an unknown object is collected at one or more detectors. The points on the object are arranged (in arbitrary order) into the vector $\mathbf{x} = (x_1, x_2, \ldots, x_n)^T$. The collection of measurements on the detector is arranged into the vector $\mathbf{y} = (y_1, y_2, \ldots, y_m)^T$.

We denote vectors by bold lowercase letters and matrices by bold uppercase. The relationship between $\mathbf{x}$ and $\mathbf{y}$ is derived from fundamental physics considerations, and given by the forward model,

$$\mathbf{y} = \mathbf{f}(\mathbf{x}), \tag{1}$$

where $\mathbf{f}(\mathbf{x})$ is a vector-valued function of $\mathbf{x}$. We assume the form of $\mathbf{f}$ is known. As $\mathbf{y}$ is our known data and $\mathbf{x}$ is our unknown object, the problem of estimating $\mathbf{x}$ is an inverse problem. The common approach to solving such a problem is to minimize the following error function over all possible $\mathbf{x}$,

$$\phi(\mathbf{x}) = \|\mathbf{y} - \mathbf{f}(\mathbf{x})\|. \tag{2}$$

Here $\|\cdot\|$ is a norm function, for example the root mean-squared error. We focus on cases where $n > m$ and there are infinitely many potential solutions, owing to the fact that we have insufficient measurements. We neglect noise and as a result have an infinite set of potential solutions $\hat{\mathbf{x}}$ such that $\mathbf{f}(\hat{\mathbf{x}}) = \mathbf{y}$ and hence $\phi(\hat{\mathbf{x}}) = 0$. So minimizing $\phi(\mathbf{x})$ could result in a wide range of possible results. In this paper, we consider this range.

The typical approach in this situation is to employ regularization [1], which modifies the error function such that there is a unique minimum, in effect picking a single choice out of the infinite range of possibilities. For example, a common choice is to simultaneously minimize the norm of $\mathbf{x}$ itself, giving us the optimization problem

$$\min_{\mathbf{x}} \{\|\mathbf{y} - \mathbf{f}(\mathbf{x})\| + \mu\|\mathbf{x}\|\}, \tag{3}$$

where μ is a small regularization parameter, which trades off the effect of regularization and model error. Neglecting noise and model error, we can consider the closely related problem

1559-128X/13/100D55-09$15.00/0

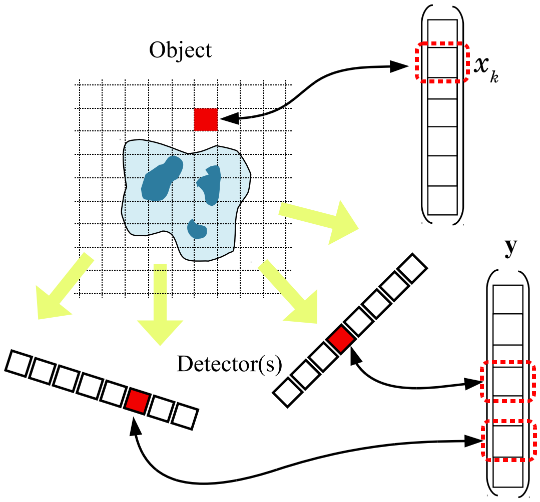

Fig. 1. (Color online) Input to forward model $\mathbf{x}$, a vector of pixels on the object, and output of model $\mathbf{y}$, a vector of detector samples. In this paper we are concerned with bounding the potential values for each object point x_k.

$$\min_{\mathbf{x}} \|\mathbf{x}\| \text{ Subject to } \mathbf{f}(\mathbf{x}) = \mathbf{y}, \tag{4}$$

where we use the model as a constraint. We follow the standard practice of writing constraints beneath the minimization objective, so that Eq. (4) is equivalent to Eq. (5):

$$\begin{aligned} &\min_{\mathbf{x}} \|\mathbf{x}\| \\ &\mathbf{f}(\mathbf{x}) = \mathbf{y}. \end{aligned} \tag{5}$$

Here we constrain $\mathbf{x}$ to be among the set of objects that would accurately reproduce our measured data, and find the object which has the least norm. This approach is less practical due to the reality of model error and noise but is a useful theoretical tool to make the meaning of the results more apparent.

Of course, the most studied situation is the system where $\mathbf{f}(\mathbf{x})$ is well approximated by a linear system $\mathbf{Ax} = \mathbf{y}$, where $\mathbf{A}$ is a $m \times n$ matrix. The linear system where $n > m$ is known as an underdetermined system of equations, in which case we have a problem of the form

$$\begin{aligned} &\min_{\mathbf{x}} \|\mathbf{x}\| \\ &\mathbf{Ax} = \mathbf{y}. \end{aligned} \tag{6}$$

When the 2-norm is used (i.e., mean-squared vector elements) this is analytically solvable and results in the least-length or pseudoinverse solution [2]. There is a great deal of recent interest in the case where the 1-norm, the sum of absolute values of elements, is used [3]. This results in a linear programming problem known as basis pursuit [4] and is important due to the fact that with the right conditions it can return a unique sparse solution. We should be careful to note, however, that though the system $\mathbf{Ax} = \mathbf{y}$ still has infinite solutions in general, it is only with the additional requirement that $\mathbf{x}$ is maximally sparse (while still reproducing the data) that the solution may potentially be unique.

A related situation occurs when the unknown, $\mathbf{x}$, is known to be non-negative. This condition is very important in imaging, for example dealing with any optical system which detects intensity, as well as systems which measure attenuation such as computed tomography. In such cases we have the constrained problem

$$\begin{aligned} &\min_{\mathbf{x}} \|\mathbf{x}\| \\ &\mathbf{Ax} = \mathbf{y} \\ &\mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{7}$$

where $\mathbf{x} \geq \mathbf{0}$ means $x_k \geq 0$ for $k = 1, \ldots, n$. $\mathbf{0}$ (a bold zero) denotes a vector of the appropriate size containing all zeros. If the 2-norm is used here, we have a version of the well-known non-negative least-squares problem [2]. For the problem of Eq. (7) there is also renewed recent interest [5–8], as it has been found that with the right conditions and with $\mathbf{x}$ sufficiently sparse (i.e., the number of nonzero elements is sufficiently small), the solution truly is unique, though the precise determination of what sparsity is sufficient may be very difficult to ever find, and in practice may never be known. At any rate, a key mathematical requirement for this case to occur, or necessary condition, is that the rowspace of matrix $\mathbf{A}$ intersects the positive orthant. This means that a linear combination of rows can be found which produces a positive vector (i.e., all elements are strictly greater than zero). Given this condition, a sufficiently sparse true vector $\mathbf{x}$ will mean that the set of $\mathbf{x}$ which fits the constraints above, i.e., $\{\mathbf{x} | \mathbf{Ax} = \mathbf{y}, \mathbf{x} \geq \mathbf{0}\}$, will contain only a single point, also known as a "singleton." Hence the above problem will find a unique result for any choice of norm. In imaging terms, if the object is known to be non-negative (e.g., we are forming an image of intensity), and if the matrix describing our imaging system has the above required property, and finally if the object itself happens to be sufficiently sparse, we can uniquely reconstruct the true object with far fewer measurements than would be required in a system that did not include the non-negativity property. Of course, if an object is not sufficiently sparse, we will not be able to uniquely reconstruct it using the very same imaging system. This is obviously a concern since the object is unknown and we cannot know whether the process will succeed.

In this paper, we take a different approach that sheds light on this problem and, where possible, reduces to the uniquely solvable situation. Rather than regularizing underdetermined problems to force a unique solution when one does not exist, or guessing that the object may be sparse, we compute bounds on the set of potential solutions to the inverse problem by directly estimating maximum and minimum values each pixel may take. In the next section, we start by considering the simplest linear case and formulate the bounds problem as an optimization problem which may be generalized to arbitrary optimization problems. Then we consider the non-negative case and find that the conditions for bounds are closely related to the conditions for finding a unique sparse solution, but that the bounds simply get looser when the object is not sufficiently sparse. Finally, we consider practical issues with estimating the bounds and provide simulated examples.

2. Theory

We start off by considering what is arguably the simplest possible optimization problem. We have the linear system as our forward model. For the objective, rather than a regularization function, we simply use the maximization or minimization of a single element of the unknown vector $\mathbf{x}$ (describing the object). This gives us the following pair of optimization problems for computing the bounds on a pixel:

$$x_k^{(\min)} = \min_{\mathbf{x}} x_k \qquad \mathbf{Ax} = \mathbf{y}, \tag{8}$$

$$x_k^{(\mathrm{mix})} = \min_{\mathbf{x}} x_k \qquad \mathbf{Ax} = \mathbf{y}. \tag{9}$$

In this way, we compute the bounds on the kth element of $\mathbf{x}$. For the entire object, we will have $2n$ of these optimization problems to solve, two per pixel. This allows us to explore that set of solutions to the forward model to some degree, to see what range a given pixel can take and still potentially fit our measured data. Equations (8) and (9) are forms of the equality-constrained linear program [9],

$$\max_{\mathbf{x}} \mathbf{c}^T\mathbf{x} \qquad \mathbf{Ax} = \mathbf{y}. \tag{10}$$

Here, $\mathbf{c}$ is $\pm\mathbf{e}_k$ and the vector $\mathbf{e}_k$ is a vector of zeros, with a one in the kth position, e.g., $\mathbf{e}_3 = (0, 0, 1, 0)^T$ where $n = 4$.

The condition for a bounded solution to exist for this optimization problem is $\mathbf{c}$ must be in the range of $\mathbf{A}^T$, the transpose of $\mathbf{A}$. In other words, a vector $\boldsymbol{\lambda}$ must exist where

$$\mathbf{A}^T\boldsymbol{\lambda} = \mathbf{e}_k. \tag{11}$$

If there is no solution to this system, it means $x_k^{(\min)}$ is $-\infty$ and $x_k^{(\max)}$ is ∞, i.e., this pixel can hold any possible value. If a solution does exist then there are finite solutions for both the max and min bound. And further (for this case) $x_k^{(\min)} = x_k^{(\max)}$, i.e., there is a unique solution for this pixel. So we have two possible situations for this problem: either a pixel is completely unbounded, or it is uniquely solvable.

As an example, we consider a system with two unknowns and one equation, given in Eq. (12):

$$y = (a_1, a_2)\begin{pmatrix} x_1 \\ x_2 \end{pmatrix}. \tag{12}$$

In this case, the single measurement, the scalar y, is a weighted sum of the two unknown pixel values, x_1 and x_2. The system matrix $\mathbf{A}$ consists of the single row, (a_1, a_2). So given the measurement y and matrix $\mathbf{A}$, we consider when we can find solutions to the linear programs of Eqs. (8) and (9). There are four of these problems in total, to compute the min and max for x_1 and min and max for x_2. For the max problem for x_1, we have the condition for boundedness,

$$\begin{pmatrix} a_1 \\ a_2 \end{pmatrix}\lambda = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \tag{13}$$

where λ is a scalar. Clearly this requires $a_2 = 0$, which gives us a rather trivial case but is the only way we can expect to bound the unknowns with this single measurement. Otherwise every real value is still possible for x_1.

We can view the uniquely solvable case as a kind of partial inverse of the matrix $\mathbf{A}^T$, just in terms of a single unknown. In the case where $\mathbf{A}$ is square and nonsingular and therefore we have a solution $\mathbf{A}^T\boldsymbol{\lambda}_k = \mathbf{e}_k$ for all $k = 1, \ldots, n$, we can combine the n conditions as

$$\mathbf{A}^T(\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots, \boldsymbol{\lambda}_n) = (\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_n) \qquad \mathbf{A}^T\Lambda = \mathbf{I}, \tag{14}$$

where Λ is the matrix with kth column $\boldsymbol{\lambda}_k$.

Generally, we might have some pixels for which unique solutions exist and some which have no bounds. Any estimate which finds an $\mathbf{x}$ that solves the system $\mathbf{Ax} = \mathbf{y}$ will immediately give us all the uniquely solvable elements. For example, we might compute the least-length (pseudoinverse) solution. The only question is which ones are truly bounded and which are only finite in our estimate due to the regularization inherent in the least-length solution. We can answer that by testing the system $\mathbf{A}^T\boldsymbol{\lambda}_k = \mathbf{e}_k$ for consistency for all $k = 1, \ldots, n$.

The equality-constrained case we just discussed can be addressed analytically, for example using the singular-value decomposition. But the advantage of formulating the bounds as an optimization problem is we may generalize it to much more complicated problems, such as by adding the constraint that the solution $\mathbf{x}$ be non-negative, in which case, our bounds problems become

$$x_k^{(\min)} = \min_{\mathbf{x}} x_k \qquad \mathbf{Ax} = \mathbf{y} \qquad \mathbf{x} \geq \mathbf{0}, \tag{15}$$

$$x_k^{(\mathrm{mix})} = \max_{\mathbf{x}} x_k \mathbf{A}\mathbf{x} = \mathbf{y}\mathbf{x} \geq \mathbf{0}. \tag{16}$$

These are cases of the standard form linear program [10],

$$\begin{aligned} &\max_{\mathbf{x}} \mathbf{c}^T\mathbf{x} \\ &\mathbf{A}\mathbf{x} = \mathbf{y} \\ &\mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{17}$$

where $\mathbf{c}$ is $\pm\mathbf{e}_k$ as before. Now, in addition to cases like before, where a pixel can be uniquely solvable or unbounded, we may also potentially have cases where the pixel has a finite range of possible values.

If we consider the two-dimensional case again, we can get an idea of when bounds exist. Figure 2 demonstrates a bounded case geometrically. The solid line represents the solutions to Eq. (12), which we can solve for analytically by solving for x_2 as a function of x_1. The range of solutions for which the variables are in the non-negative orthant give the bounds depicted on the axes in the Figure:

$$x_2 = -\frac{a_1}{a_2}x_1 + \frac{1}{a_2}y. \tag{18}$$

We can see that we will have finite upper bounds as long as both a_1 and a_2 are positive, in which case the upper bound for x_1 comes when x_2 is zero and vice versa.

Figure 3 demonstrates what happens when a_2 is negative. Now the solution set remains in the non-negative quadrant for an infinite length and we have no upper bounds. We do get a lower bound above zero in this case, from the x-intercept. From this figure we can also see that when a_1 or a_2 is zero, our solution set would be vertical or horizontal. In such a case one of the unknown pixels would be uniquely given with upper and lower bounds that are equal.

The above 1-by-2 case gives us some intuition for the property of $\mathbf{A}$ necessary for pixels to be bounded. We can imagine that if $\mathbf{A}$ has only positive elements,

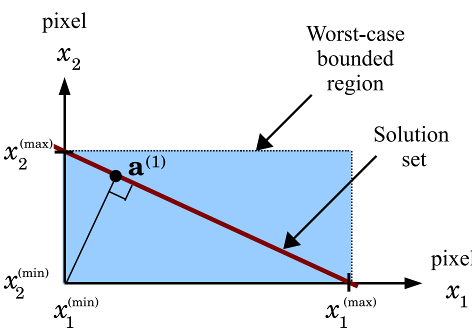

Fig. 2. (Color online) Two-dimensional case (two unknown pixels) with a single measurement. The solid line represents the solution to $\mathbf{A}\mathbf{x} = \mathbf{y}$ and the shaded region represents the values of $\mathbf{x}$ within the bounds. $\mathbf{a}^{(1)}$ is the first (and only) row of $\mathbf{A}$. In this case both pixels are bounded as the solution set is only in the non-negative quadrant for a finite range.

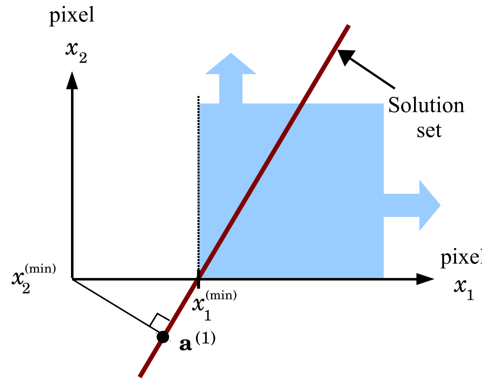

Fig. 3. (Color online) Two-dimensional unbounded case. Lower bounds exist but not upper bounds, as the solution set to $\mathbf{A}\mathbf{x} = \mathbf{y}$ continues to infinity in the non-negative quadrant.

we would have a similar situation in any number of dimensions. In fact, the requirement for finite pixel bounds is slightly looser than this and we will now derive it formally using duality theory.

The duals of the linear programs in Eqs. (16) and (15) are

$$\begin{aligned} \tilde{x}_k^{(\min)} &= \min_{\lambda} \mathbf{y}^T\boldsymbol{\lambda} \\ &\mathbf{A}^T\boldsymbol{\lambda} \geq -\mathbf{e}_k, \end{aligned} \tag{19}$$

$$\begin{aligned} \tilde{x}_k^{(\max)} &= \min_{\lambda} \mathbf{y}^T\boldsymbol{\lambda} \\ &\mathbf{A}^T\boldsymbol{\lambda} \geq -\mathbf{e}_k. \end{aligned} \tag{20}$$

Duality theory tells us that the primal problems in Eqs. (16) and (15) are feasible and bounded if the corresponding dual problems in Eqs. (20) and (19) are feasible and bounded. We assumed from the start that our system model $\mathbf{A}\mathbf{x} = \mathbf{y}$ was noise free and therefore feasible, which means the dual is bounded. So the only question is whether the duals are feasible, and if so it means the primals are bounded.

For the dual problem in Eq. (19) to be feasible, there must exist a solution to $\mathbf{A}^T\boldsymbol{\lambda} \geq -\mathbf{e}_k$. Obviously, this is trivially possible by simply choosing $\boldsymbol{\lambda} = \mathbf{0}$. This makes sense as it is also obvious that the primal problem for the min bound must be bounded, since we are minimizing x_k, which is constrained to be non-negative. As a worst case it has a lower bound of zero, which corresponds to the case where $\boldsymbol{\lambda} = \mathbf{0}$.

So the interesting question is the existence of the max bound. For the dual problem in Eq. (20) to be feasible, there must exist a solution to $\mathbf{A}^T\boldsymbol{\lambda} \geq \mathbf{e}_k$. If such a solution exists, say $\hat{\boldsymbol{\lambda}}$, then we also can say a solution exists for $\mathbf{A}^T\boldsymbol{\lambda} \geq \alpha\mathbf{e}_k$ for any scalar α, by forming $\alpha\hat{\boldsymbol{\lambda}}$. Therefore we may rewrite the requirement as: there exists a solution $\boldsymbol{\lambda}$ to $\mathbf{A}^T\boldsymbol{\lambda} = \boldsymbol{\beta}$ where $\boldsymbol{\beta} \geq \mathbf{0}$ and $\beta_k > 0$. In words, the rowspace of $\mathbf{A}$ must contain a non-negative vector which is positive in the kth element. The conditions for the max and min bounds are summarized in Table 1.

Table 1. Summary of Optimization Problems to Find Bounds on kth Pixel, the Corresponding Dual Problems, and the Resulting Conditions for Existence of Bounds[a]

Goal	Primal Problem	Dual Problem	Condition
Equality constrained pixel minimization	$\min_{\mathbf{x}} x_k\ \mathbf{Ax} = \mathbf{y}$	$\max_\lambda \mathbf{y}^T\boldsymbol{\lambda}\ \mathbf{A}^T\boldsymbol{\lambda} = -\mathbf{e}_k$	$\mathbf{A}^T\lambda = -\mathbf{e}_k$
Equality constrained pixel maximization	$\max_{\mathbf{x}} x_k\ \mathbf{Ax} = \mathbf{y}$	$\max_\lambda \mathbf{y}^T\boldsymbol{\lambda}\ \mathbf{A}^T\boldsymbol{\lambda} = \mathbf{e}_k$	$\mathbf{A}^T\lambda = \mathbf{e}_k$
Non-negativity constrained pixel minimization	$\min_{\mathbf{x}} x_k\ \mathbf{Ax} = \mathbf{y}\ \mathbf{x} \geq \mathbf{0}$	$\max_\lambda \mathbf{y}^T\boldsymbol{\lambda}\ \mathbf{A}^T\boldsymbol{\lambda} \geq -\mathbf{e}_k$	$\mathbf{A}^T\lambda \geq -\mathbf{e}_k$
Non-negativity constrained pixel maximization	$\max_{\mathbf{x}} x_k\ \mathbf{Ax} = \mathbf{y}\ \mathbf{x} \geq \mathbf{0}$	$\max_\lambda \mathbf{y}^T\boldsymbol{\lambda}\ \mathbf{A}^T\boldsymbol{\lambda} \geq \mathbf{e}_k$	$\mathbf{A}^T\lambda \geq \mathbf{e}_k$

[a] If a solution $\boldsymbol{\lambda}$ can be found to the condition, the corresponding max or min is finite. Note that the condition for the two equality-constrained bounds are mathematically equivalent.

The condition for the max bound for the non-negative case is a generalization of the necessary condition for a singleton solution to the non-negative linear system discussed in the literature (see e.g., [8]). In our case, instead of a fully positive vector, we only need one element to be positive with the rest non-negative. Further, we have the singleton case that if all pixels are bounded, then we have (for each k) a solution $\boldsymbol{\lambda}^{(k)}$ to $\mathbf{A}^T\boldsymbol{\lambda}^{(k)} = \boldsymbol{\beta}^{(k)}$ where $\boldsymbol{\beta}^{(k)} \geq \mathbf{0}$ and $\beta_k^{(k)} > 0$. By forming the vector $\boldsymbol{\lambda} = \boldsymbol{\lambda}^{(1)} + \boldsymbol{\lambda}^{(2)} + \cdots + \boldsymbol{\lambda}^{(n)}$, we have $\mathbf{A}^T\boldsymbol{\lambda} = \boldsymbol{\beta}$ where $\boldsymbol{\beta} > \mathbf{0}$. As the $\boldsymbol{\beta}^{(k)}$ resulting from each $\boldsymbol{\lambda}^{(k)}$ is positive in the kth element and non-negative in every other element, the result from the sum of all k of them is positive in every element.

Returning to our general requirement, to know if a given pixel (i.e., the kth pixel) is bounded in our system, we must determine if there exists a solution to $\mathbf{A}^T\boldsymbol{\lambda} = \boldsymbol{\beta}$ where $\boldsymbol{\beta} \geq \mathbf{0}$ and $\boldsymbol{\beta}_k > 0$. As above, we can combine all the solutions to the systems for the different pixels and form a vector $\boldsymbol{\beta}$ which is positive for all of the bounded elements. In other words, since we do not care whether $\beta_i > 0$ or $\beta_i = 0$ for $i \neq k$, we can perform all the tests for different pixels simultaneously by seeking a combined $\boldsymbol{\beta}$ that is positive in the maximum number of elements possible while still keeping other elements non-negative. We can write this as the following optimization problem:

$$K = \max_{\lambda,\beta} \|\boldsymbol{\beta}\|_0$$
$$\mathbf{A}^T\boldsymbol{\lambda} = \boldsymbol{\beta}$$
$$\boldsymbol{\beta} \geq \mathbf{0}, \tag{21}$$

where we have used the so-called zero norm, $\|\cdot\|_0$, which returns the number of nonzero elements in a vector. Normally, the zero norm is problematic as it does not fulfil the conditions of a true vector norm and so the more common goal of minimizing it leads to nonconvex optimization problems. But here we are maximizing it, plus we are constraining the vector to be non-negative. And in the non-negative orthant the zero norm (or any p-norm for small p, which may serve as an approximation) is concave so maximizing over it gives us a convex nonlinear optimization problem.

With the optimization problem of Eq. (21) we can find all the pixels which are bounded for a given system in a single optimization. For the case where the system matrix $\mathbf{A}$ is itself non-negative, the test becomes even simpler. We only need to find the columns which contain a positive element. This can be done by computing $\boldsymbol{\beta} = \mathbf{A}^T\mathbf{1}$, where $\mathbf{1}$ is a vector of ones of appropriate size. The elements for which $\boldsymbol{\beta}$ is nonzero (and therefore positive) are the elements for which the pixels of a measured object will always be bounded.

We can view the singleton solution as a special case of a bounded solution. A singleton further requires the additional condition that the unknown vector $\mathbf{x}$ describing the object be sufficiently sparse. One might imagine a continuum of problems where the difference between the max and min bounds drops from large to small to zero (singleton), as the true vector $\mathbf{x}$ becomes increasingly sparse. For imaging, it is much more useful to consider the effect of an increasingly large data collection with the same object. In this case, we would also expect the bounds to decrease until the object can be uniquely reconstructed.

As additional data is collected, bounds tighten monotonically. This we can demonstrate by simply noting that the feasible set (which we are finding bounds for the elements of), $F = \{\mathbf{x}|\mathbf{Ax} = \mathbf{y}, \mathbf{x} \geq \mathbf{0}\}$, can be viewed as the intersection of multiple sets. We may describe the new feasible set, formed by adding an additional measurement as a new row onto the matrix $\mathbf{A}$ and a new element onto the end of the vector $\mathbf{y}$, as

$$\begin{aligned} F' &= \left\{\mathbf{x}|\begin{pmatrix}\mathbf{A}\\ \mathbf{A}_1\end{pmatrix}\mathbf{x} = \begin{pmatrix}\mathbf{y}\\ \mathbf{y}_1\end{pmatrix}, \mathbf{x} \geq \mathbf{0}\right\} \\ &= \{\mathbf{x}|\mathbf{Ax} = \mathbf{y}, \mathbf{x} \geq \mathbf{0}\} \cap \{\mathbf{x}|\mathbf{A}_1\mathbf{x} = \mathbf{y}_1, \mathbf{x} \geq \mathbf{0}\} \\ &= F \cap F_1. \end{aligned} \tag{22}$$

So when we collect new data, the set of feasible solutions is the intersection of the original feasible set and the solutions which fulfil the new data. As an intersection can only reduce the size of a set or leave it unchanged, the bounds can only decrease or remain unchanged.

Now we consider the problem of choosing which additional measurement to perform, from multiple options. For example, we have a collection of measurements from a set of view angles and wish to decide the best view angle for the next measurement. This would be useful in an imaging scenario where performing actual measurements was costly but performing computations was cheap. As we do not have the data from the new measurement yet, we must make our decision purely from the change in the model.

Our approach here will be to select the new measurement which provides bounds on the most unbounded pixels. To do this, we do not need to know the result of the new measurement, just the forward model for it (e.g., $\mathbf{A}_1$ above).

For the equality-constrained system, our conditions for existence of bounds are given in Table 1. Note that the conditions for the minimum and maximum bound are equivalent; if a solution λ exists for the maximum bound condition, we can use its negative to prove the minimum bound condition is fulfilled. Hence we only need to consider one of the conditions which, with a new measurement, becomes

$$\begin{pmatrix} \mathbf{A} \\ \mathbf{A}_1 \end{pmatrix}^T \lambda = \mathbf{e}_k. \tag{23}$$

So in general for each potential measurement, we would check if a solution exists to Eq. (23) for every k, and then choose the measurement for which the most pixels have solutions. We do not need to check pixels for which we already know bounds exist without the new measurement.

One way to do these tests more efficiently would be to compute a basis for the nullspace using a singular value decomposition (SVD) of the augmented matrix of Eq. (23), then test whether each $\mathbf{e}_k$ is orthogonal to the entire basis simply by taking inner products with the basis vectors. The computational cost of this approach would be one SVD per potential new measurement.

For the non-negativity-constrained system, the conditions for existence of bounds are again given in Table 1. In this case, the condition for the minimum bound is trivially met as discussed earlier. Since x_k is constrained to be at least zero, we always have zero as a lower bound. As a result, when considering existence of bounds, we only need consider the condition for the maximum bound, which with new measurements is

$$\begin{pmatrix} \mathbf{A} \\ \mathbf{A}_1 \end{pmatrix}^T \lambda \geq \mathbf{e}_k. \tag{24}$$

Again, we would need to perform this test for each potential new measurement and for every pixel that is not already known to be bounded. However, we note that we can view this requirement as the system

$$\begin{pmatrix} \mathbf{A} \\ \mathbf{A}_1 \end{pmatrix}^T \lambda = \boldsymbol{\beta}, \tag{25}$$

where $\boldsymbol{\beta}$ is greater than or equal to zero and as dense as possible (i.e., has fewest zero elements). Similar to Eq. (21) we can write this as an optimization problem, yielding

$$K = \max_{\lambda,\boldsymbol{\beta}} \|\boldsymbol{\beta}\|_0 \;\; \begin{pmatrix} \mathbf{A} \\ \mathbf{A}_1 \end{pmatrix}^T \lambda = \boldsymbol{\beta} \;\; \boldsymbol{\beta} \geq \mathbf{0}. \tag{26}$$

The approach, therefore, is to compute K using this optimization program for each potential new measurement, and pick the new measurement which yields the highest K.

Finally, we address a couple practical issues in the use of bounds. The first issue we will consider is errors in the measurement vector $\mathbf{y}$, such as from noise and numerical precision. In practice we expect to get $\mathbf{y} + \Delta\mathbf{y}$, where $\Delta\mathbf{y}$ is some unknown error. While this might simply result in a variation in the resulting image, which may be quantified by classical linear algebra analysis, there is also a danger of the new system, $\mathbf{Ax} = \mathbf{y} + \Delta\mathbf{y}$, being incompatible. We address this by relaxing the constraint a limited amount, replacing the linear constraint $\mathbf{Ax} = \mathbf{y}$ with the nonlinear constraint $\|\mathbf{Ax} - \mathbf{y}\| < \delta$, where δ is a scalar meant to allow for the error in the measurement. Here, we will use a value of δ that is very small to primarily guard against numerical precision problems, without having a noticeable effect on the bounds. In a noisy system this can be adjusted to be larger.

The second practical issue we consider is the condition of the matrix $\mathbf{A}$. Just as a matrix inverse can be very sensitive to noise if it is poorly conditioned, we might expect bounds to be more sensitive to noise if $\mathbf{A}$ contains very small singular values (relative to the largest one). With regularization techniques, this is addressed by selecting a regularization parameter large enough such that small singular values are treated as if they were zero. We can get the same effect by performing a singular value decomposition of $\mathbf{A}$ and zeroing out the singular values below a chosen cutoff. This restricts the amount of change that errors may cause to the bounds.

3. Simulation

The system we simulate is a multiview imaging approach, one view of which is depicted in Fig. 4. In a multiview collection, we have the projection collected from multiple positions around the scene, giving an optical analog of computed tomography [11]. Here we will assume the object is sufficiently small and far away that the rays are approximately parallel when crossing it, to eliminate this variable from the system. Then the only parameter we need to describe each view is its axial direction. For example, in Fig. 4, an axial direction of 90° is used. The matrix $\mathbf{A}$ in this case is essentially the Radon transform (or some subset of its rows), which computes projections of the object at the collected view directions, with ones in the elements corresponding to visible pixels for a given measurement point. For the case where we introduce known occluding structures in the system, we simply zero out the elements for the hidden pixels at each view.

First, we demonstrate the application of estimating which pixels are bounded, via the optimization

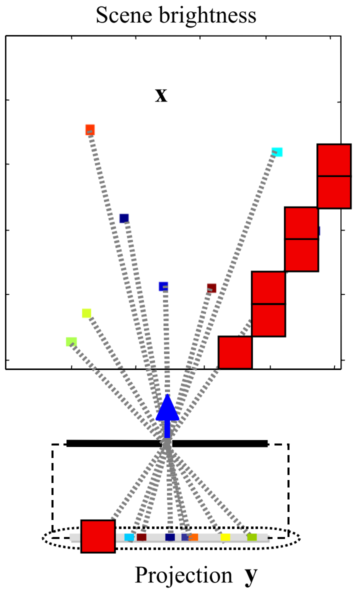

Fig. 4. (Color online) Pinhole camera in the high frequency limit. We assume there is no occlusion, so each point on the detector receives light from sources along a ray through the scene.

problems given in the last section. We consider two-dimensional systems with known occluders in the scene, restricting the view of the object. Two such systems of occluders we used are given in Figs. 5(a) and 5(b). For these systems, $n = 31 \times 31 = 961$. And $m = 31$ when we have a single view and $31 \times 2 = 62$ when we have two views. So a collected image is the

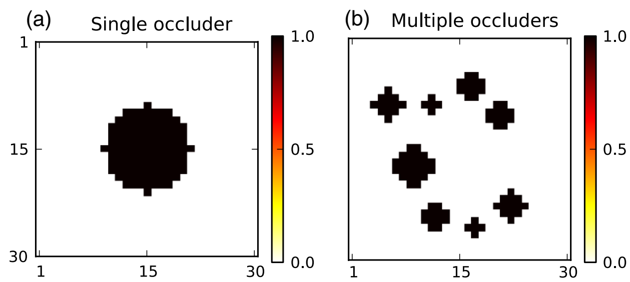

Fig. 5. (Color online) Simulated occluding structures.

same size as a cross section of the object with unit magnification. Note that Figs. 5(a) and 5(b) are not the unknown object to be imaged but rather structures blocking our view. The object would be values in the white region of these figures. The occluders are known and used to compute the system matrix $\mathbf{A}$, depending on which views are used in the data collection. For this simulation, we produced system matrices to demonstrate both Eqs. (21) and (26). For the first case, we computed the number of bounded pixels for each view in one-degree increments between 0 and 360. For the second case, we simulated two views, with one of the views fixed at 0 and the other view at each direction in one-degree increments between 0 and 360. The optimization problems were solved using CVX [12,13]. The numbers of bounded pixels, K in Eqs. (21) and (26), are plotted versus angle in Fig. 6. With a single view, we can see how there is nothing to be gained at different view angles with the single round occluder. However with multiple occluders, some views yield more bounded pixels than others. For the two-view case, we plotted the increase in bounded pixels over a single view, so in both cases we see a second view at 0° adds nothing as expected, and that there is a general trend toward getting more pixels as the second view is more opposite the first.

For the next simulation, we estimate the bounds for different simulated objects for a collection using an increasing range of angles (no occluding structures), again using two-dimensional objects. For these systems, $n = 21 \times 21 = 441$ and $m = 21 \times v$ where v is the number of views used. The objects are shown in Fig. 7, giving the brightness at each point in the imaged scene. The objects are normalized to a total mean-squared length of one. The first two objects are designed to be equally sparse, while the third object is simply random brightness in the imaged scene, and hence is as dense as possible.

A matrix which implements the collection of Fig. 4 was produced, which simulates the collection of multiple views of an object over a specified range of views in one degree increments. For example,

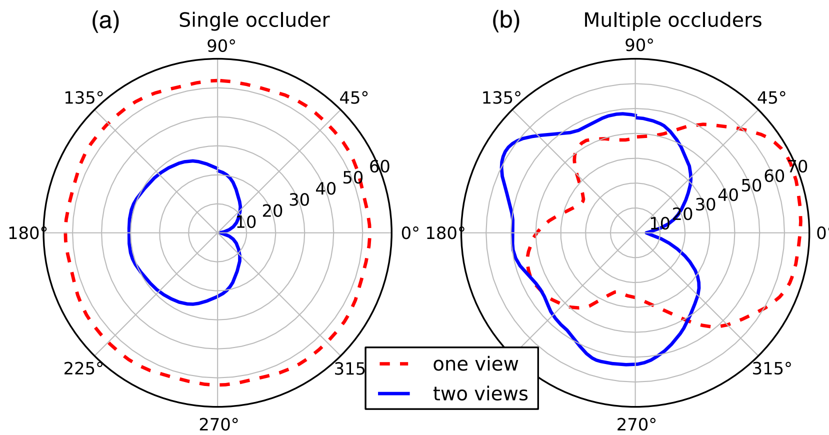

Fig. 6. (Color online) Polar plots of number of pixels seen given a single view for each view angle, and given two views, one at 0° and the other at the given range of view angles.

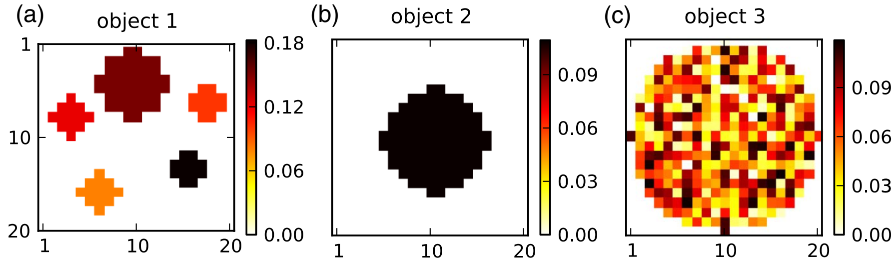

Fig. 7. (Color online) Two-dimensional simulated objects.

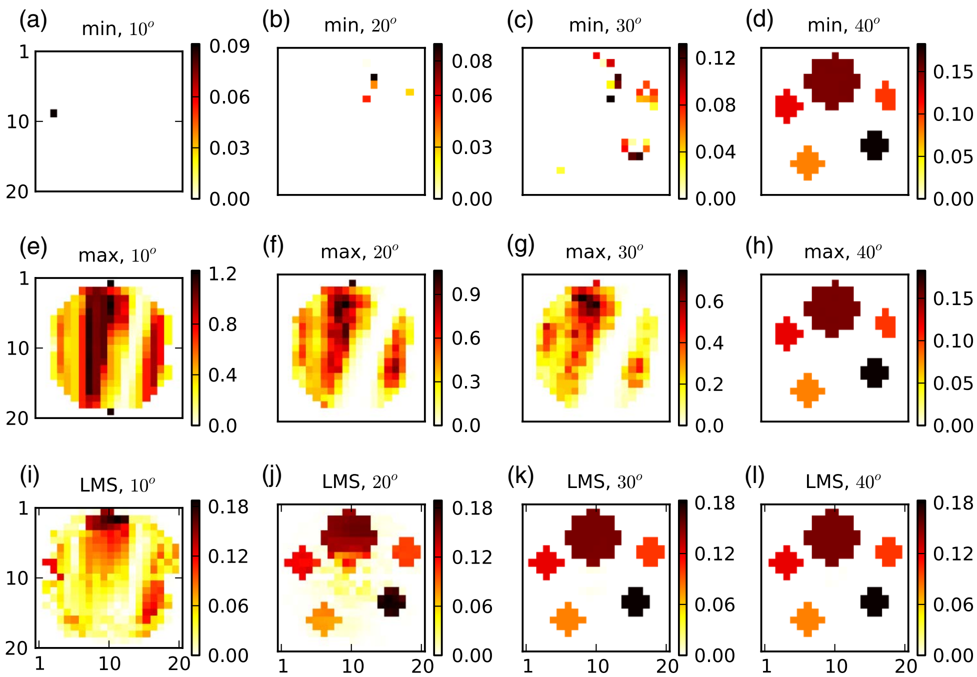

Fig. 8. (Color online) Max (a–d) and min (e–h) bounds for each pixel given the collected data for a system with the stated total collection size. Least-mean square estimate of image with the given collections shown for comparison (i–l).

the matrix for the 10° collection simulated ten equally spaced views a degree apart. The matrix was then conditioned by performing a singular-value decomposition and zeroing out singular values that were less than two orders of magnitude smaller then the maximum singular value. The bounds were estimated using the non-negative optimization problem with a δ value of 10^{-6} to protect against numerical precision issues. The optimization problems were again solved using CVX for finding the max and min bound on each pixel.

Examples of the bounds for the first object are shown in Fig. 8. We see that for increasing collection sizes, the max bounds decrease until they reach the true pixel values. Similarly, the min bounds increase for larger collections (where the true values are above zero). At around 40° collection size, the bounds are equal, demonstrating that it is here our sparser objects can be uniquely resolved, though the dense object still cannot.

In Fig. 9, we compute the bound range (maximum bound minus the minimum bound) averaged over the entire scene for each of the three objects for different ranges of collection sizes between 0 and 90°. We find that bounds for the first two objects decrease at roughly the same rate with increasing collection size but the third object continues to have unequal max

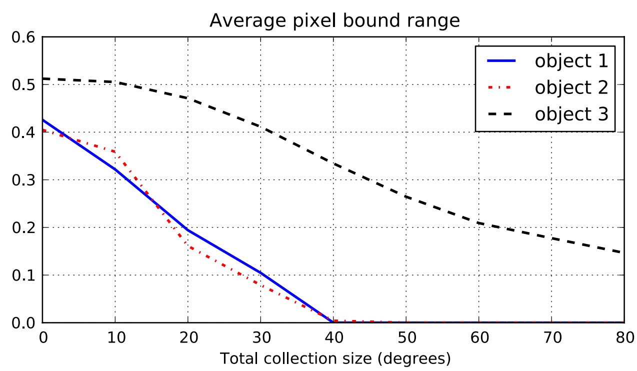

Fig. 9. (Color online) Average bounds for the three objects versus total collection size.

and min bounds for a longer range, as would be expected due to its lower sparsity.

4. Discussion

In this paper we have demonstrated the use of pixel bounds for computational imaging, with both theoretical and practical applications. The equality-constrained case was considered first, which is useful as a connection between the more classical approaches to treating linear systems and the application to more difficult constrained cases. In the equality-constrained case we also demonstrate how some pixel values may still be available in an ill-posed system. The case of a partially solvable image, while very unlikely for a random matrix, may occur quite easily in computational imaging as the collection geometry is often very structured.

The inequality-constrained case is more interesting in that we may have bounds which are not tight, i.e., both max and min are finite but not equal to each other. Further, we found the conditions for this to occur. In optical imaging, the condition (rowspace intersecting positive orthant) is very common. With an intensity-imaging approach for, example, the matrix itself will also be non-negative and it is very easy to show that the conditions of Table 1 are satisfied. The simulated system was such a case, where we saw finite bounds for every pixel and for every collection.

An important issue requiring further study is the effect of noise. In this paper, we considered solely the linear case. Ideally, to address noise, one would return to a problem such as Eq. (4), which would likely require a significantly different analysis and a nonlinear problem. However, we can address limited levels of noise by relaxing the constraint as we did by using the nonlinear constraint $\|\mathbf{Ax} - \mathbf{y}\| < \delta$, which roughly raises the bounds by δ, scaled by a term depending on the matrix condition. We can further control the matrix condition by truncating the singular vectors limiting this effect. The result is a loosening of the bounds.

Bounds are also a useful tool to approach nonnegative systems where a unique sparse reconstruction is desired. In such a system, we must have the conditions hold for every pixel, plus the object must be sufficiently sparse. The exact (both necessary and sufficient) condition for unique reconstruction is very difficult to compute for a matrix. But we may instead compute the bounds, which immediately tell us if a given object may be uniquely estimated with the system, and further, give us information even when it cannot be uniquely determined.

We also considered a strategy for actively collecting measurements in such a way that bounds may be made finite as fast as possible. This is useful in cases where the cost of performing measurements is greater than performing the computational effort of choosing the best measurement, for example if ionizing radiation is to be used.

We gratefully acknowledge helpful advice and assistance from Prof. Philip Gill.

References

1. A. K. Louis, "A unified approach to regularization methods for linear ill-posed problems," Inverse Probl. **15**, 489–498 (1999).
2. C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems* (Society for Industrial and Applied Mathematics, 1987).
3. M. Fornasier and H. Rauhut, "Compressive sensing," in *Handbook of Mathematical Methods in Imaging*, O. Scherzer, ed. (Springer, 2011), pp. 187–228.
4. S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," SIAM Rev. **43**, 129–159 (2001).
5. D. L. Donoho and J. Tanner, "Sparse nonnegative solution of underdetermined linear equations by linear programming," Proc. Natl. Acad. Sci. USA **102**, 9446–9451 (2005).
6. A. M. Bruckstein, M. Elad, and M. Zibulevsky, "Sparse nonnegative solution of a linear system of equations is unique," *3rd International Symposium on Communications, Control and Signal Processing, 2008* (ISCCSP, 2008), pp. 762–767.
7. A. M. Bruckstein, M. Elad, and M. Zibulevsky, "On the uniqueness of nonnegative sparse solutions to underdetermined systems of equations," IEEE Trans. Inf. Theory **54**, 4813–4820 (2008).
8. M. Wang and A. Tang, "Conditions for a unique non-negative solution to an underdetermined system," in *47th Annual Allerton Conference on Communication, Control, and Computing, 2009* (IEEE, 2009), pp. 301–307.
9. P. E. Gill, W. Murray, and M. H. Wright, *Numerical Linear Algebra and Optimization* (Addison-Wesley, 1991).
10. G. Dantzig, *Linear Programming and Extensions* (Princeton University, 1998).
11. D. L. Marks and D. J. Brady, "Three-dimensional source reconstruction with a scanned pinhole camera," Opt. Lett. **23**, 820–822 (1998).
12. "CVX: matlab software for disciplined convex programming | CVX research, inc." http://cvxr.com/cvx/.
13. M. Grant and S. Boyd, "Graph implementations for nonsmooth convex programs," in *Lecture Notes in Control and Information Sciences*, V. Blondel, S. Boyd, and H. Kimura, eds. vol. **371** (Springer, 2008), pp. 95–110.