

Compressed Sensing with Cross Validation

Rachel Ward

December 7, 2008

Abstract

Compressed Sensing decoding algorithms can efficiently recover an N dimensional real-valued vector x to within a factor of its best k -term approximation by taking $m = 2k \log N/k$ measurements $y = \Phi x$. If the sparsity or approximate sparsity level of x were known, then this theoretical guarantee would imply quality assurance of the resulting compressed sensing estimate. However, because the underlying sparsity of the signal x is unknown, the quality of a compressed sensing estimate \hat{x} of the unknown signal x using m measurements is not assured. Nevertheless, in this paper, we demonstrate that sharp bounds on the error $\|x - \hat{x}\|_{l_2^N}$ can be achieved with almost no effort. More precisely, we assume that a maximum number of measurements m is pre-imposed; we reserve $4 \log p$ of the original m measurements and compute a sequence of possible estimates $(\hat{x}_j)_{j=1}^p$ to x from the $m - 4 \log p$ remaining measurements; the errors $\|x - \hat{x}_j\|_{l_2^N}$ for $j = 1, \dots, p$ can then be bounded with high probability. As a consequence, numerical upper and lower bounds on the error between x and the best k -term approximation to x can be estimated for p values of k with almost no cost. Our observation has applications outside of compressed sensing as well.

Key words: Compressed Sensing, cross validation, measurements, best k -term approximation, Johnson Lindenstrauss Lemma, encoding/decoding, error estimates

1 Introduction

Compressed Sensing (CS) is a fast developing area in applied mathematics, motivated by the reality that most data we store and transmit contains far less information than its dimension suggests. For example, a one-dimensional slice through the pixels in a typical grayscale image will contain segments of smoothly varying intensity, with sharp changes between adjacent pixels appearing only at edges in the image. Often this sparsity in information translates into a sparse (or approximately sparse) representation of the data with respect to some standard basis; for the image example, the basis would be a wavelet or curvelet basis. For such N dimensional data vectors that are well approximated by a k -sparse vector (or a vector that contains at most $k \ll N$ nonzero entries), it is common practice to temporarily store the entire vector, possibly with the intent to go back and replace this vector with a smaller dimensional vector encoding the location and magnitude of its k significant coefficients. In compressed sensing, one instead collects fewer fixed linear measurements of the data to start with, sufficient in number to recover the location and numerical value of the k nonzero coordinates at a later time. Finding "good" linear measurements, as well as fast, accurate, and simple algorithms for recovering the original data from these measurements, are the twofold goals of compressed sensing research today.

Review of basic CS setup. The data of interest is taken to be a real-valued vector $x \in \mathbb{R}^N$ that is *unknown*, but from which we are allowed up to $m < N$ linear measurements, in the form of inner products of x with m vectors $v_j \in \mathbb{R}^N$ of our choosing. Letting Φ denote the $m \times N$ matrix whose j th row is the vector v_j , this is equivalent to saying that we have the

freedom to choose and store an $m \times N$ matrix Φ , along with the m -dimensional measurement vector $y = \Phi x$. Of course, since Φ maps vectors in \mathbb{R}^N to vectors in a smaller dimensional space \mathbb{R}^m , the matrix Φ is not invertible, and we thus have no hope of being able to reconstruct an arbitrary N dimensional vector x from such measurements.

However, if the otherwise unknown vector x is specified to be k -sparse, and k is fairly small compared with N , then there do exist matrices Φ for which $y = \Phi x$ uniquely determines x , and allows recovery of x using fast and simple algorithms. It was the interpretation of this phenomenon given by Candes and Tao [1], [2], and Donoho [3], that gave rise to compressed sensing. In particular, these authors define classes of matrices that possess this property. One particularly elegant characterization of this class is via the *Restricted Isometry Property* (RIP) [2]. A matrix Φ with unit normed columns is said to be k -RIP if all singular values of any k column submatrix of Φ lie in the interval $[1 - \delta_k, 1 + \delta_k]$ for a given constant $\delta_k < 1$. With high probability, $2k$ -RIP is obtained, with

$$k = K(m, N) := 2m / \log(N/m), \text{ where } m \leq \frac{1}{2}m, \quad (1)$$

for $m \times N$ matrices Φ whose entries $\Phi_{i,j}$ are independent realizations of a Gaussian or Bernoulli random variable [4]. Also with high probability, an $m \times N$ matrix obtained by selecting m rows at random from the $N \times N$ discrete Fourier matrix satisfies $2k$ -RIP of the same order as (1) up to an additional $\log^3 N$ factor [28]. In fact, the order of k given by (1) is optimal given m and N , as shown in [5] using classical results on Gelfand widths of l_1^N unit balls in l_2^N . To date, there exist no deterministic constructions of RIP matrices of this order.

Recovering or approximating x . As shown in [21], the following approximation results hold for matrices Φ that satisfy $2k$ -RIP with constant $\delta_{2k} \leq \sqrt{2} - 1$:

1. If $x \in \mathbb{R}^N$ is k -sparse, then x can be reconstructed from Φ and the measurement vector $y = \Phi x$ as the solution to the following ℓ_1 minimization problem:

$$x = \mathcal{L}_1(\Phi, y) := \arg \min_{\Phi z = y} \|z\|_1. \quad (2)$$

2. If x is not k -sparse, the error between x and the approximation $\hat{x} = \mathcal{L}_1(\Phi, y)$ is still bounded by

$$\|x - \hat{x}\|_2 \leq \frac{c}{\sqrt{k}} \sigma_k(x)_{l_p^N}, \quad (3)$$

where $c = 2(1 + \delta_{2k}) / (1 - \delta_{2k})$, and $\sigma_k(x)_{l_p^N} := \inf_{|z| \leq k} \|x - z\|_{l_p^N}$ denotes the best possible approximation error in the metric of l_p^N between x and the set of k -sparse signals in \mathbb{R}^N . The approximation error $\sigma_k(x)_{l_p^N}$ is realized by the k -sparse vector $x_k \in \mathbb{R}^N$ that corresponds to vector x with all but the k largest entries set to zero, independent of the l_p^N norm in the approximation $\sigma_k(x)_{l_p^N}$.

This immediately suggests to use the ℓ_1 -minimizer \mathcal{L}_1 as a means to recover or approximate an unknown x with sparsity constraint. Several other decoding algorithms are used as alternatives to ℓ_1 minimization for recovering a sparse vector x from its image $y = \Phi x$, not because they offer better accuracy (ℓ_1 minimization gives optimal approximation bounds when Φ satisfies RIP), but because they can be faster and easier to implement. For a comprehensive survey on compressed sensing decoding algorithms, we refer the reader to [30].

Estimating the accuracy of CS estimates. According to the bound (3), the quality of

a compressed sensing estimate $\hat{x} = \Delta(\Phi, y)$ depends on how well x can be approximated by a k -sparse vector, where the value of k is determined by the number of rows m composing Φ . While k is assumed to be known and fixed in the compressed sensing literature, no such bound is guaranteed for real-world signal models such as vectors $x \in \mathbb{R}^N$ corresponding to wavelet coefficient sequences of discrete photograph-like images. Thus, the quality of a compressed sensing estimate \hat{x} in general is not guaranteed.

If the error $\|x - \hat{x}\|_2/\|x\|_2$ incurred by a particular approximation \hat{x} were *observed* to be large, then decoding could be repeated using a larger number of measurements, perhaps at increasing measurement levels $\{m_1, m_2, \dots, m_p\}$, until the error $\|x - \hat{x}_j\|_2/\|x\|_2$ corresponding to m_j measurements were observed to be sufficiently small. Of course, the errors $\|x - \hat{x}_j\|_2$ and $\|x - \hat{x}_j\|_2/\|x\|_2$ are typically not known, as x is unknown. Our main observation is that one can apply the Johnson-Lindenstrauss lemma [13] to the set of p points,

$$\{(x - \hat{x}_1), (x - \hat{x}_2), \dots, (x - \hat{x}_p)\}. \quad (4)$$

In particular, $r = O(\log p)$ measurements of x , provided by $y_\Psi = \Psi x$, when Ψ is, e.g. a Gaussian or Bernoulli random matrix, are sufficient to guarantee that with high probability,

$$4/5\|y_\Psi - \Psi \hat{x}_j\|_2 \leq \|x - \hat{x}_j\|_2 \leq 4/3\|y_\Psi - \Psi \hat{x}_j\|_2 \quad (5)$$

and

$$1/3 \frac{\|y_\Psi - \Psi \hat{x}_j\|_2}{\|y_\Psi\|_2} \leq \frac{\|x - \hat{x}_j\|_2}{\|x\|_2} \leq 2 \frac{\|y_\Psi - \Psi \hat{x}_j\|_2}{\|y_\Psi\|_2} \quad (6)$$

for any p compressed sensing estimates. The equivalences (5) and (6) allow the *measurable* quantities $\|y_\Psi - \Psi \hat{x}_j\|_2$ and $\|y_\Psi - \Psi \hat{x}_j\|_2/\|y_\Psi\|_2$ to function as proxies for the *unknown* quantities $\|x - \hat{x}_j\|_2$ and $\|x - \hat{x}_j\|_2/\|x\|_2$; these proxies can be used to

- (a) provide tight numerical upper and lower bounds on the errors $\|x - \hat{x}_j\|_2$ and $\|x - \hat{x}_j\|_2/\|x\|_2$ at up to p compressed sensing estimates \hat{x}_j ,
- (b) provide estimates of the underlying k -term approximations $\|x - x_k\|_2$ of x for up to p different values of k , and
- (c) return from among a sequence of estimates $(\hat{x}_1, \dots, \hat{x}_p)$ with different initialization parameters, an estimate $\hat{x}_{cv} = \arg \min_j \|y_\Psi - \Psi \hat{x}_j\|_2$ having error $\|x - \hat{x}_{cv}\|_2$ that does not exceed a small multiplicative factor of the best possible error in the metric of ℓ_2^N between x and an element from the sequence at hand.

More precisely, all CS decoding algorithms require as input a parameter m corresponding to the number of rows in Φ ; some compressed sensing decoding algorithms (such as greedy algorithms) require also a parameter k indicating the sparsity level of x , and other algorithms require as input a bound γ on the expected amount of energy in x outside of its significant coefficients. All CS decoding algorithms can be symbolically represented by functions of the form $\Delta(\Phi, y, k, \gamma)$, and we will give examples where each of the parameters m , k , and γ can be optimized over a sequence of estimates $(\hat{x}_1, \hat{x}_2, \dots, \hat{x}_p)$ parametrized by increasing hypotheses on each of the variables m , k , and γ .

The estimation procedure described above, although novel in its proposed application, is by no means new. *Cross validation* is a technique used in statistics and learning theory whereby a data set is separated into a training/estimation set and a test/cross validation set, and the test set is used to prevent overfitting on the training set by estimating underlying noise parameters. We will take a set of m measurements of x , and use $m - r$ of these measurements,

Φx , in a compressed sensing decoding algorithm to return a sequence $(\hat{x}_1, \hat{x}_2, \dots)$ of candidate approximations to x . The remaining r measurements, Ψx , are then used to identify from among this sequence a single approximation $\hat{x} = \hat{x}_j$, along with an estimate of the sparsity level of x . The application of cross validation to compressed sensing has been studied by Boufounos, Duarte, and Baraniuk in [7], but in a different context from the present paper, and without the mathematical justification of the Johnson Lindenstrauss lemma that we present below.

2 Preliminary Notation

Throughout the paper, we will be dealing with large dimensional vectors that have few nonzero coefficients. We use the notation

$$|x| = n \tag{7}$$

to indicate that a vector $x \in \mathbb{R}^N$ has exactly n nonzero coefficients.

We will sometimes use the notation $a \sim_\epsilon b$ as shorthand for the multiplicative relation

$$(1 - \epsilon)a \leq b \leq (1 + \epsilon)a, \tag{8}$$

that can be worded as “the quantity a approximates the quantity b to within a multiplicative factor of $(1 \pm \epsilon)$ ”. Note that the relation \sim_ϵ is not symmetric. Properties of the relation $a \sim_\epsilon b$ are listed below; we leave the proofs (which amount to a string of simple inequalities) as an exercise for the reader.

Lemma 2.1. *Fix $\epsilon \in (0, 1)$.*

1. *If $a, b \in \mathbb{R}^+$ satisfy $a \sim_\epsilon b$, then $b / [(1 + \epsilon)(1 - \epsilon)] \sim_\epsilon a$.*
2. *If $a, b, c, d \in \mathbb{R}^+$ satisfy $a \sim_\epsilon b$ and $c \sim_\epsilon d$, then $a/c \sim_\delta b/d$ for parameter $\delta = 2\epsilon / (1 - \epsilon)$.*
3. *If (a_1, a_2, \dots, a_p) and (b_1, b_2, \dots, b_p) are sequences in \mathbb{R}^+ , and $a_j \sim_\epsilon b_j$ for each $1 \leq j \leq p$, then $\min_j a_j \sim_\epsilon \min_j b_j$.*

3 Mathematical Foundations

The Johnson Lindenstrauss (JL) lemma, in its original form, states that any set of p points in high dimensional Euclidean space can be embedded into $\epsilon^{-2} \log(p)$ dimensions, without distorting the distance between any two points by more than a factor of $(1 \pm \epsilon)$ [13]. In the same paper, it was shown that a random orthogonal projection would provide such an embedding with positive probability. Following several simplifications to the original proof [15], [12], [14], it is now understood that Gaussian random matrices, among other purely random matrix constructions, can substitute for the random projection in the original proof of Johnson and Lindenstrauss. Of the several versions of the lemma now appearing in the literature, the following variant presented in Matousek [16] is most applicable to the current presentation.

Lemma 3.1 (Johnson-Lindenstrauss Lemma). *Fix an accuracy parameter $\epsilon \in (0, 1/2]$, a confidence parameter $\delta \in (0, 1)$, and an integer $r \geq r_0 = C\epsilon^{-2} \log \frac{1}{2\delta}$.*

Let \mathcal{M} be a random $r \times N$ matrix whose entries $\mathcal{M}_{i,j}$ are independent realizations of a random variable R that satisfies:

1. $\text{Var}(R) = 1/r$ (so that the columns of \mathcal{M} have expected ℓ_2 norm 1)
2. $E(R) = 0$,
3. For some fixed $a > 0$ and for all λ ,

$$\text{Prob}[|R| > \lambda] \leq 2e^{-a\lambda^2} \quad (9)$$

Then for a predetermined $x \in \mathbb{R}^N$,

$$(1 - \epsilon)\|x\|_{l_2^N} \leq \|\mathcal{M}x\|_{l_2^r} \leq (1 + \epsilon)\|x\|_{l_2^N} \quad (10)$$

is satisfied with probability exceeding $1 - \delta$.

The constant C bounding r_0 in Lemma (3.1) grows with the parameter a specific to the construction of \mathcal{M} (9). Gaussian and Bernoulli random variables R will satisfy the concentration inequality (9) for a relatively small parameter a (as can be verified directly), and for these matrices one can take $C = 8$ in Lemma (3.1).

The Johnson Lindenstrauss lemma can be made intuitive with a few observations. Since $\mathbf{E}[R] = 0$ and $\mathbf{Var}[R] = \frac{1}{r}$, the random variable $\|\mathcal{M}x\|_2^2$ equals $\|x\|_2^2$ in expected value; that is,

$$\mathbf{E}[\|\mathcal{M}x\|_2^2] = \|x\|_2^2. \quad (11)$$

Additionally, $\|\mathcal{M}x\|_2^2$ inherits from the random variable R a nice concentration inequality:

$$\text{Prob}[\|\mathcal{M}x\|_2^2 - \|x\|_2^2 > \epsilon\|x\|_2^2] \leq e^{-a(2\epsilon\sqrt{r})^2} \leq \delta/2. \quad (12)$$

The first inequality above is at the heart of the JL lemma; its proof can be found in [16]. The second inequality follows using that $r \geq (2a\epsilon^2)^{-1} \log(\frac{\delta}{2})$ and $\epsilon \leq 1/2$ by construction. A bound similar to (12) holds for $\text{Prob}[\|\mathcal{M}x\|_2^2 - \|x\|_2^2 < -\epsilon\|x\|_2^2]$ as well, and combining these two bounds gives desired result (10).

For fixed $x \in \mathbb{R}^N$, a random matrix \mathcal{M} constructed according to Lemma (3.1) fails to satisfy the concentration bound (10) with probability at most δ . Applying Boole's inequality, \mathcal{M} then fails to satisfy the stated concentration on any of p predetermined points $\{x_j\}_{j=1}^p$, $x_j \in \mathbb{R}^N$, with probability at most $\xi = p\delta$. In fact, a specific value of $\xi \in (0, 1)$ may be imposed for fixed p by setting $\delta = \xi/p$. These observations are summarized in the following corollary to Lemma (3.1).

Corollary 3.2. Fix an accuracy parameter $\epsilon \in (0, 1/2]$, a confidence parameter $\xi \in (0, 1)$, and fix a set of p points $\{x_j\}_{j=1}^p \subset \mathbb{R}^N$. Set $\delta = \xi/p$, and fix an integer $r \geq r_0 = C\epsilon^{-2} \log \frac{1}{2\delta} = C\epsilon^{-2} \log \frac{p}{2\xi}$. If \mathcal{M} is a $r \times N$ matrix constructed according to Lemma (3.1), then with probability $\geq 1 - \xi$, the bound

$$(1 - \epsilon)\|x_j\|_{l_2^N} \leq \|\mathcal{M}x_j\|_{l_2^r} \leq (1 + \epsilon)\|x_j\|_{l_2^N} \quad (13)$$

obtains for each $j = 1, 2, \dots, p$.

4 Cross Validation in Compressed Sensing

We return to the situation where we would like to approximate a vector $x \in \mathbb{R}^N$ with an assumed sparsity constraint using $m < N$ linear measurements $y = \mathcal{A}x$ where \mathcal{A} is an $m \times N$ matrix of our choosing. Continuing the discussion in Section 1, we will not reconstruct x in the standard

way by $\hat{x} = \Delta(\mathcal{A}, y, k, \gamma)$ for fixed values of the input parameters, but instead separate the $m \times N$ matrix \mathcal{A} into an $n \times N$ *implementation* matrix Φ and an $r \times N$ *cross validation* matrix Ψ , and separate the measurements y accordingly into y_Φ and y_Ψ . We use the implementation matrix Φ and corresponding measurements y_Φ as input into the decoding algorithm to obtain a sequence of possible estimates $(\hat{x}_1, \dots, \hat{x}_p)$ corresponding to increasing one of the input parameters m , k , or γ . We reserve the cross validation matrix Ψ and measurements y_Ψ to estimate each of the error terms $\|x - \hat{x}_j\|_2$ in terms of the computable $\|y_\Psi - \Psi\hat{x}_j\|_2$. Our main result, which follows from Corollary (3.2), details how the number of cross validation measurements r should be chosen in terms of the desired accuracy ϵ of estimation, confidence level ξ in the prediction, and number p of estimates \hat{x}_j to be measured:

Theorem 4.1. *For a given accuracy $\epsilon \in (0, 1/2]$, confidence $\xi \in (0, 1)$, and number p of estimates $\hat{x}_j \in \mathbb{R}^N$, it suffices to allocate $r = \lceil C\epsilon^{-2} \log \frac{p}{2\xi} \rceil$ rows to a cross validation matrix Ψ of Gaussian or Bernoulli type, normalized according to Lemma (3.2) and independent of the estimates \hat{x}_j , to obtain with probability greater than or equal to $1 - \xi$, and for each $j = 1, 2, \dots, p$, the bounds*

$$\frac{1}{1 + \epsilon} \leq \frac{\|x - \hat{x}_j\|_{l_2^N}}{\|\Psi(x - \hat{x}_j)\|_{l_2^r}} \leq \frac{1}{1 - \epsilon} \quad (14)$$

and

$$\frac{1 - 3\epsilon}{(1 + \epsilon)(1 - \epsilon)^2} \leq \frac{\|x - \hat{x}_j\|_{l_2^N} / \|x\|_2}{\|\Psi(x - \hat{x}_j)\|_{l_2^r} / \|\Psi x\|_2} \leq \frac{1}{(1 - \epsilon)^2} \quad (15)$$

and also

$$\frac{1}{1 + \epsilon} \leq \frac{\eta_{or}}{\widehat{\eta}_{cv}} \leq \frac{1}{1 - \epsilon} \quad (16)$$

where $\eta_{or} = \min_{1 \leq j \leq p} \|x - \hat{x}_j\|_{l_2^N}$ is the unknown oracle error corresponding to the best possible approximation to x in the metric of l_2^N from the sequence $(\hat{x}_1, \dots, \hat{x}_p)$, and $\widehat{\eta}_{cv} = \min_{1 \leq j \leq p} \|\Psi(x - \hat{x}_j)\|_{l_2^r}$ is the observable cross validation error.

Proof. .

- The bounds in (14) are obtained by application of Lemma (3.2) to the p points $u_j = x - \hat{x}_j$, and rearranging the resulting bounds according to Lemma (2.1) part (1). The bound (16) follows from the bounds (14) and part (3) of Lemma (2.1).
- The bounds in (15) are obtained by application of Lemma (3.2) to the $p + 1$ points $u_0 = x, u_j = x - \hat{x}_j$, and regrouping the resulting bounds according to part (2) of Lemma (2.1). □

Remark 4.2. The measurements making up the cross validation matrix Ψ must be independent of the measurements comprising the rows of the implementation matrix Φ . This comes from the requirement in Lemma (3.1) that the matrix Ψ be independent of the points $u_j = x - \hat{x}_j$. This requirement is crucial, as observed when \hat{x} solves the ℓ_1 minimization problem

$$\hat{x} = \arg \min_{z \in \mathbb{R}^N} \|z\|_1 \text{ subject to } \Phi z = \Phi x, \quad (17)$$

in which case the constraint $\Phi(\hat{x} - x) = 0$ clearly precludes the rows of Φ from giving any information about the error $\|\hat{x} - x\|_2$.

Remark 4.3. Theorem (4.1) should be applied with a different level of care depending on what information about the sequence $(x - x_1, x - x_2, \dots, x - x_p)$ is sought. If the minimizer $\hat{x} = \arg \min_{1 \leq j \leq p} \|\Psi(x - \hat{x}_j)\|_{l_2^r}$ is sufficient for one's purposes, then the precise normalization of Ψ in Theorem (4.1) is not important. The normalization doesn't matter either for estimating the normalized quantities $\|x - \hat{x}_j\|_2 / \|x\|_2$. On the other hand, if one is using cross validation to obtain estimates for the quantities $\|x - \hat{x}_j\|_2$, then normalization is absolutely crucial, and one must observe the normalization factor given by Lemma (3.2) that depends on the number of rows r allocated to the cross validation matrix Ψ .

5 Applications of cross validation to compressed sensing

5.1 Estimation of the best k -term approximation error

We have already seen that if the $m \times N$ matrix Φ satisfies $2k$ -RIP with parameter $\delta \leq \sqrt{2} - 1$, and $\hat{x} = \mathcal{L}_1(\Phi, \Phi x)$ is returned as the solution to the ℓ_1 minimization problem (2), then the error between x and the approximation \hat{x} is bounded by

$$\|x - \hat{x}\|_2 \leq \frac{c}{\sqrt{k}} \sigma_k(x)_{l_1^N}. \quad (18)$$

Several other decoding algorithms in addition to ℓ_1 minimization enjoy the reconstruction guarantee (18) under similar bounds on Φ , such as the Iteratively Reweighted Least Squares algorithm (IRLS) [29], and the greedy algorithms CoSAMP [30] and Subspace Pursuit [31]. It has recently been shown [18] [20] that if the bound (18) is obtained, and if $x - \hat{x}$ lies in the null space of Φ (as is the case for the decoding algorithms just mentioned), then if Φ is a Gaussian or a Bernoulli random matrix, the error $\|x - \hat{x}\|_2$ also satisfies a bound, with high probability on Φ , with respect to the ℓ_2^N residual, namely

$$\|x - \hat{x}\|_2 \leq c \sigma_k(x)_{l_2^N}, \quad (19)$$

for a reasonable constant c depending on the RIP constant δ_{2k} of Φ . In the event that (19) is obtained, a cross validation estimate $\|\Psi(x - \hat{x})\|_{l_2^r}$ can be used to lower bound the residual $\sigma_k(x)_{l_2^N}$, with high probability, according to

$$(1 - \epsilon) \|\Psi(x - \hat{x})\|_{l_2^r} \leq \|x - \hat{x}\|_{l_2^N} \leq c \sigma_k(x)_{l_2^N}, \quad (20)$$

with $O(\frac{1}{\epsilon^2})$ rows reserved for the matrix Ψ (4.1). At this point, we will use Corollary 3.2 of [8], where it is proved that if the bound (18) holds for \hat{x} with constant c , then the same bound will hold for

$$\hat{x}_k = \arg \min_{z: |z| \leq k} \|\hat{x} - z\|_{l_2^N}, \quad (21)$$

the best k -sparse approximation to \hat{x} , with constant $\tilde{c} = 3c$. Thus, we may assume without loss of generality that \hat{x} is k -sparse, in which case $\|\Psi(x - \hat{x})\|_{l_2^r}$ also provides an upper bound on the residual $\sigma_k(x)_{l_2^N}$ by

$$(1 + \epsilon) \|\Psi(x - \hat{x})\|_{l_2^r} \geq \|x - \hat{x}\|_{l_2^N} \geq \sigma_k(x)_{l_2^N}. \quad (22)$$

With almost no effort then, cross validation can be incorporated into many decoding algorithms to obtain tight upper and lower bounds on the unknown k -sparse approximation error $\sigma_k(x)_{l_2^N}$ of x . More generally, the allocation of $10 \log p$ measurements to the cross validation matrix Ψ is sufficient to estimate the errors $(\|x - x_{k_j}\|_2)_{j=1}^p$ or the normalized approximation errors $(\|x - x_{k_j}\|_2 / \|x\|_2)_{j=1}^p$ at p sparsity levels k_j by decoding p times, adding m_j measurements to the implementation matrix Φ_j at each repetition. Recall that the quantities k_j and m_j are related by $k_j = 2m_j / \log(N/m_j)$ according to (1).

5.2 Choice of the number of measurements m

Photograph-like images have wavelet or curvelet coefficient sequences $x \in \mathbb{R}^N$ that are *compressible* [23] [32], having entries that obey a power law decay

$$|x|_{(k)} \leq c_s k^{-s}, \quad (23)$$

where $x_{(k)}$ denotes the k th largest coefficient of x in absolute value, the parameter $s > 1$ indicates the level of compressibility of the underlying image, and c_s is a constant that depends only on s and the normalization of x . From the definition (23), compressible signals are immediately seen to satisfy

$$\|x - x_k\|_1 / \sqrt{k} \leq c'_s k^{-s+1/2}, \quad (24)$$

so that the solution $\hat{x}_m = \mathcal{L}_1(\Phi, \Phi x)$ to the ℓ_1 minimization problem (2) using an $m \times N$ matrix Φ of optimal RIP order $k = 2m / \log(N/m)$ satisfies

$$\|x - \hat{x}_m\|_2 \leq c_{s,\delta} k^{-s+1/2}. \quad (25)$$

The number of measurements m needed to obtain an estimate \hat{x}_m satisfying $\|x - \hat{x}_m\|_2 \leq \tau$ for a predetermined threshold τ will vary according to the compressibility of the image at hand. Armed with a total of m measurements, the following decoding method that *adaptively* chooses the number of measurements for a given signal x presents a more democratic alternative to standard compressed sensing decoding structure:

Table 1: *CS decoding structure with adaptive number of measurements*

1. *Input:* The m -dimensional vector $y = \Phi x$, the $m \times N$ matrix Φ , (in some algorithms) the sparsity level k , and (again, in some algorithms) a bound γ on the noise level of x , the number p of of row subsets of Φ , $(\Phi_1, \Phi_2, \dots, \Phi_p)$, corresponding to increasing number of rows $m_1 < m_2 < \dots < m_p < m$, and threshold $\tau > 0$.
2. *Initialize* the decoding algorithm at $j = 1$.
3. *Estimate* $\hat{x}_j = \Delta(\Phi_{m_j}, y_{m_j}, k, \gamma)$ with the decoder Δ at hand, using only the first m_j measurement rows of Φ . The previous estimate \hat{x}_{j-1} can be used for “warm initialization” of the algorithm, if applicable. The remaining $r_j = m - m_j$ measurement rows are allocated to a cross validation matrix Ψ_j that is used to estimate the resulting error $\|x - \hat{x}_j\|_2 / \|x\|_2$.
4. *Increment* j by 1, and iterate from step 3 if stopping rule is not satisfied.
5. *Stop:* at index $j = j^* < p$ if $\|x - x_{m_j}\|_2 / \|x\|_2 \leq \tau$ holds with near certainty, as indicated by

$$\frac{\sqrt{r_j} \|\Psi(x - x_{m_j})\|_2 / \|\Phi x\|_2}{\sqrt{r_j} - 3 \log p} \leq \tau \quad (26)$$

according to Theorem (4.1). If the maximal number of decoding measurements $m_p < m$ have been used at iteration p , and (26) indicates that $\|x - \hat{x}_{m_p}\|_2 / \|x\|_2 > \tau$ still, return $\hat{x}_m = \Delta(\Phi, y, k, \gamma)$ using all m measurements, but with a warning that the underlying image x is probably too dense, and its reconstruction is not trustworthy.

5.3 Choice of regularization parameter in homotopy-type algorithms

Certain compressed sensing decoding algorithms iterate through a sequence of intermediate estimates \hat{x}_j that could be potential optimal solutions to x under certain reconstruction parameter choices. This is the case for greedy and homotopy-continuation based algorithms. In this section,

we study the application of cross validation to the intermediate estimates of decoding algorithms of homotopy-continuation type.

LASSO is the name coined in (??) for the problem of minimizing of the following convex program:

$$\hat{x}^{[\tau]} = \arg \min_{z \in \mathbb{R}^N} \|\Phi x - \Phi z\|_{\ell_2} + \tau \|z\|_1 \quad (27)$$

The two terms in the LASSO optimization problem (27) enforce data fidelity and sparsity, respectively, as balanced by the regularization parameter τ . In general, choosing an appropriate value for τ in (27) is a hard problem; when Φ is an underdetermined matrix, as is the case in compressed sensing, the function $f(\tau) = \|x - \hat{x}^{[\tau]}\|_2$ is unknown to the user but is seen empirically to have a minimum at a value of τ in the interval $[0, \|\Phi x\|_\infty]$ that depends on the unknown noise level and/or and compressibility level of x .

The *homotopy continuation* algorithm [26], which can be viewed as the appropriate variant of LARS [26], is one of many algorithms for solving the LASSO problem (27) at a predetermined value of τ ; it proceeds by first initializing τ' to a value sufficiently large to ensure that the ℓ_1 penalization term in (27) completely dominates the minimization problem and $x^{[\tau']} = 0$ trivially. The homotopy continuation algorithm goes on to generate $x^{[\tau']} = 0$ for decreasing τ' until the desired level for τ is reached. If $\tau = 0$, then the homotopy method traces through the entire solution path $x^{[\tau']} \in \mathbb{R}^N$ for $\tau' \geq 0$ before reaching the final algorithm output $x^{[0]} = \mathcal{L}_1(\Phi, y)$ corresponding to the ℓ_1 minimizer (2).

From the non-smooth optimality conditions for the convex functional (27), it can be shown that the solution path $\hat{x}^{[\tau]} \in \mathbb{R}^N$ is a piecewise-affine function of τ [26], with “kinks” possible only at a finite number of points $\tau \in \{\tau_1, \tau_2, \dots\}$. Theorem (4.1) suggests a method whereby an appropriate value of τ^* can be chosen from among a subsequence of the kinks $(\tau_1, \tau_2, \dots, \tau_p)$ by solving the minimization problem $\hat{x}^{[\tau^*]} = \arg \min_{j \leq p} \|\Psi(x - \hat{x}^{[\tau_j]})\|_2$ for appropriate cross validation matrix Ψ . Moreover, since the solution $x - \hat{x}^{[\tau]}$ for $\tau_j \leq \tau \leq \tau_{j+1}$ is restricted to lie in the two-dimensional subspace spanned by $x - \hat{x}^{[\tau_j]}$ and $x - \hat{x}^{[\tau_{j+1}]}$, one can combine the Johnson Lindenstrauss Lemma with a covering argument analogous to that used to derive the RIP property for Gaussian and Bernoulli random matrices in [4], to cross validate the entire *continuum* of solutions $\hat{x}^{[\tau]}$ between $\tau_1 \leq \tau \leq \tau_p$. More precisely, the following bound holds under the conditions outlined in Theorem (4.1), with the exception that $2r$ (as opposed to r) measurements are reserved to Ψ :

$$\frac{1}{1 + \epsilon} \leq \frac{\min_{\tau_1 \leq \tau \leq \tau_p} \|x - \hat{x}^{[\tau]}\|_2}{\min_{\tau_1 \leq \tau \leq \tau_p} \|\Psi(x - \hat{x}^{[\tau]})\|_2} \leq \frac{1}{1 - \epsilon} \quad (28)$$

Unfortunately, the bound (28) is not strong enough to *provably* evaluate the entire solution path $\hat{x}^{[\tau]}$ for $\tau \geq 0$, because the best upper bound on the number of kinks on a generic LASSO solution path can be very large. One can prove that this number is bounded by 3^N , by observing that if $\hat{x}^{[\tau_1]}$ and $\hat{x}^{[\tau_2]}$ have the same sign pattern, then $\hat{x}^{[\tau]}$ also has the same sign pattern for $\tau_1 \leq \tau \leq \tau_2$. Applying Theorem (4.1) to $p = 3^N$ points $x - \hat{x}_j$, this suggests that $O(N)$ rows would need to be allocated to a cross validation matrix Ψ in order for Theorem (4.1) and the corollary (28) to apply to the entire solution path, which clearly defeats the compressed sensing purpose. However, whenever the matrix Φ is an $m \times N$ compressed sensing matrix of random Gaussian, Bernoulli, or partial Fourier construction, it is observed *empirically* that the number of kinks along a homotopy solution path is bounded by $3m$, independent of the underlying vector $x \in \mathbb{R}^N$ used to generate the path. This suggests, at least heuristically, that the allocation of $O(\log m)$ out of m compressed sensing measurements of this type suffices to

ensure that the error $\|x - \hat{x}^{[\tau]}\|_2$ for the solution $\hat{x}^{[\tau]} = \arg \min_{\tau \geq 0} \|\Psi(x - \hat{x}^{[\tau]})\|_2$ will be within a small multiplicative factor of the best possible error in the metric of ℓ_2^N obtainable by any approximant $\hat{x}^{[\tau]}$ along the solution curve $\tau \geq 0$. At the value of τ corresponding to $\hat{x}^{[\tau]}$, the LASSO solution (27) can be computed using all m measurements Φ as a final approximation to x .

The *Dantzig selector* (DS) [22] refers to a minimization problem that is similar in form to the LASSO problem:

$$\hat{x}_\tau = \arg \min_{z \in \mathbb{R}^N} \|\Phi x - \Phi z\|_{\ell_\infty} + \tau \|z\|_1 \quad (29)$$

The difference between the DS (29) and LASSO (27) is the choice of norm (ℓ_∞ versus ℓ_2) on the fidelity-promoting term. Homotopy-continuation based algorithms have also been developed to solve the minimization problem (29) by tracing through the solution path $\hat{x}_{\tau'}$ for $\tau' \geq \tau$. As the minimization problem (29) can be reformulated as a linear program, its solution path $\hat{x}_\tau \in \mathbb{R}^N$ is seen to be a piecewise *constant* function of τ , in contrast to the LASSO solution path. In practice, the total number of breakpoints (τ_1, τ_2, \dots) in the domain $0 \leq \tau$ is observed to be on the same order of magnitude as m when the $m \times N$ matrix Φ satisfies RIP [24]; thus, the procedure just described to cross validate the LASSO solution path can be adapted to cross validate the solution path of (29) as well.

Thus far we have not discussed the possibility of using cross validation as a stopping criterion for homotopy-type decoding algorithms. Along the LARS homotopy curve (27), most of the breakpoints (τ_1, τ_2, \dots) appear only near the end of the curve in a very small neighborhood of $\tau = 0$. These breakpoints incur only miniscule changes in the error $\|x - \hat{x}_{\tau_j}\|_2$ even though they account for most of the computational expense of the LARS decoding algorithm. Therefore, it would be interesting to adapt such algorithms, perhaps using cross validation, to stop once τ^* is reached for which the error $\|x - \hat{x}_{\tau^*}\|_2$ is sensed to be sufficiently small.

5.4 Choice of sparsity parameter in greedy-type algorithms

Greedy compressed sensing decoding algorithms also iterate through a sequence of intermediate estimates \hat{x}_j that could be potential optimal solutions to x under certain reconstruction parameter choices. *Orthogonal Matching Pursuit* (OMP), which can be viewed as the prototypical greedy algorithm in compressed sensing, picks columns from the implementation matrix Φ one at a time in a greedy fashion until, after k iterations, the k -sparse vector \hat{x}_k , a linear combination of the k columns of Φ chosen in the successive iteration steps, is returned as an approximation to x . The OMP algorithm is listed in Table 2. Although we will not describe the algorithm in full detail, a comprehensive study of OMP can be found in [6]. Note in particular that OMP requires as input a parameter k corresponding to the expected sparsity level for $x \in \mathbb{R}^N$. Such input is typical among greedy algorithms in compressive sensing (in particular, we refer the reader to [30], [29], and [31]). As shown in [6], OMP will recover with high probability a vector x having at most $k \leq m/\log(N)$ nonzero coordinates from its image Φx if Φ is a (known) $m \times N$ Gaussian or Bernoulli matrix with high probability. Over the more general class of vectors $x = x_d + \mathcal{N}$ that can be decomposed into a d -sparse vector x_d (with d presumably less than or equal to k) and additive noise vector \mathcal{N} , we might expect an intermediate estimate \hat{x}_s to be a better estimate to x than the final OMP output \hat{x}_k , at least when $d \ll k$. Assuming that the signal x admits a decomposition of the form $x = x_d + \mathcal{N}$, the sequence of intermediate estimates ($\hat{x}_1, \dots, \hat{x}_k$) of an OMP algorithm can be cross validated in order to estimate the noise level and recover a better approximation to x . We will study this particular application of cross validation in more detail below.

Table 2: *Orthogonal Matching Pursuit Basic Structure*

1. *Input*: The m -dimensional vector $y = \mathcal{B}x$, the $m \times N$ encoding matrix Φ whose j^{th} column is labeled ϕ_j , and the sparsity bound k .
2. *Initialize* the decoding algorithm at $j = 1$, the residual $r_0 = y$, and the index set $\Lambda_0 = \emptyset$.
3. *Estimate*
 - (a) Find an index λ_j that realizes the bound $(\Phi^T r_{j-1})_{\lambda_j} = \|\Phi^T r_{j-1}\|_\infty$.
 - (b) Update the index set $\Lambda_j = \Lambda_{j-1} \cup \lambda_j$ and the submatrix of contributing columns: $\Phi_j = [\Phi_{j-1}, \phi_{\lambda_j}]$
 - (c) Update the residual:
$$\begin{aligned} s_j &= \arg \min_x \|\Phi_j x - y\|_2 = (\Phi_j^T \Phi_j)^{-1} \Phi_j^T y, \\ a_j &= \Phi_j x_j \\ r_j &= r_{j-1} - a_j. \end{aligned}$$
 - (d) The estimate \hat{x}_j for the signal has nonzero indices at the components listed in Λ_j , and the value of the estimate \hat{x}_j in component λ_i equals the i th component of s_j .
4. *Increment* j by 1 and iterate from step 3, if $j < k$.
5. *Stop*: at $j = k$. Output $\hat{x}_{omp} = \hat{x}_k$ as approximation to x .

6 Orthogonal Matching Pursuit: A case study

As detailed in Table 2, a single index λ_j is added to a set Λ_j estimated as the j most significant coefficients of x at each iteration j of OMP; following the selection of Λ_j , an estimate \hat{x}_j to x is determined by the least squares solution,

$$\hat{x}_j = \arg \min_{\text{supp}(z) \in \Lambda_j} \|\Phi z - y\|_2, \quad (30)$$

among the subspace of vectors $z \in \mathbb{R}^N$ having nonzero coordinates in the index set Λ_j . OMP continues as such, adding a single index λ_j to the set Λ_j at iteration j , until $j = k$ at which point the algorithm terminates and returns the k -sparse vector $\hat{x}_{omp} = \hat{x}_k$ as approximation to x .

Suppose x has only d significant coordinates. If d could be specified beforehand, then the estimate \hat{x}_d at iteration $j = d$ of OMP would be returned as an approximation to x . However, the sparsity d is not known in advance, and k will instead be an upper bound on d . As the estimate \hat{x}_j in OMP can be then identified with the hypothesis that x has j significant coordinates, the application of cross-validation as described in the previous section applies in a very natural way to OMP. In particular, we expect \hat{x}_{or} and \hat{x}_{cv} of Theorem (4.1) to be close to the estimate \hat{x}_j at index $j = |x|$ corresponding to the true sparsity of x ; furthermore, in the case that $|x|$ is significantly less than k , we expect the cross validation estimate \hat{x}_{cv} to be a better approximation to x than the OMP-returned estimate \hat{x}_k . We will put this intuition to the test in the following numerical experiment.

6.1 Experimental setup

We initialize a signal x_0 of length $N = 3600$ and sparsity level $d = 100$ as

$$x_0(j) = \begin{cases} 1, & \text{for } j = 1 \dots 100 \\ 0, & \text{else.} \end{cases} \quad (31)$$

Noise is then added to $x_a = x_0 + \mathcal{N}_a$ in the form of a Gaussian random variable \mathcal{N}_a distributed according to

$$\mathcal{N}_a \sim N(0, .05), \quad (32)$$

and the resulting vector x_a is renormalized to satisfy $\|x_a\|_{l_2^N} = 1$. This yields an expected noise level of

$$E(\sigma_d(x_a)) \approx .284. \quad (33)$$

We fix the input $k = 200$ in Table 2, and assume we have a total number of compressed sensing measurements $m = 800$. A number r of these m measurements are allotted to cross validation, while the remaining $n = m - r$ measurements are allocated as input to the OMP algorithm in Table 2. This experiment aims to numerically verify Theorem (4.1); to this end, we specify a confidence $\xi = 1/100$, and solve for the accuracy ϵ according to the relation $r = \epsilon^{-2} \log(\frac{k}{2\xi})$; that is,

$$\epsilon(r) = \sqrt{\frac{\log(\frac{k}{2\xi})}{r}} \approx \frac{3}{\sqrt{r}}. \quad (34)$$

Note that the specification (34) corresponds to setting the constant $C = 1$ in Theorem (4.1). Although $C \geq 8$ is needed for the proof of the Johnson Lindenstrauss lemma at present, we find that in practice $C = 1$ already upper bounds the optimal constant needed for Theorem (4.1) for Gaussian and Bernoulli random ensembles.

A single (properly normalized) Gaussian $n \times N$ measurement matrix Φ is generated (recall that $n = m - r$), and this matrix and the measurements $y = \Phi x$ are provided as input to the OMP algorithm; the resulting sequence of estimates $(\hat{x}_1, \hat{x}_2, \dots, \hat{x}_k)$ is stored. The final estimate \hat{x}_k from this sequence is the returned OMP estimate \hat{x}_{omp} to x . The error $\eta_{omp} = \|\hat{x}_{omp} - x\|_2$ is greater than or equal to the oracle error of the sequence, $\eta_{or} = \min_{\hat{x}_j} \|x - \hat{x}_j\|_2$.

With the sequence $(\hat{x}_1, \hat{x}_2, \dots, \hat{x}_k)$ at hand, we consider 1000 realizations Ψ_q of an $r \times N$ cross validation matrix having the same componentwise distribution as Φ , but normalized to have variance $1/r$ according to Theorem (3.1). The cross validation error

$$\widehat{\eta}_{cv}(q) = \min_j \|\Psi_q(x - \hat{x}_j)\|_{l_2^r} \quad (35)$$

is measured at each realization Ψ_q ; we plot the average $\widehat{\eta}_{cv}$ of these 1000 values and intervals centered at $\widehat{\eta}_{cv}$ having length equal to twice the empirical standard deviation. Note that we are effectively testing 1000 trials of OMP-CV, the algorithm which modifies OMP to incorporate cross validation so that $(\hat{x}_{cv}, \widehat{\eta}_{cv})$ are output instead of $\hat{x}_{omp} = \hat{x}_k$.

At the specified value of ξ , Theorem (4.1) part (15) (with constant $C = 1$) implies that

$$(1 - \epsilon)\eta_{or} \leq \widehat{\eta}_{cv}(q) \leq (1 + \epsilon)\eta_{or} \quad (36)$$

should obtain on at least 990 of the 1000 estimates $\widehat{\eta}_{cv}(q)$; in other words, at least 990 of the 1000 discrepancies $|\eta_{or} - \widehat{\eta}_{cv}(q)|$ should be bounded by

$$0 \leq |\widehat{\eta}_{cv}(q) - \eta_{or}| \leq \epsilon \eta_{or}. \quad (37)$$

Using the relation (34) between ϵ and r , this bound becomes tighter as the number r of CV measurements increases; however, at the same time, the oracle error η_{or} increases with r for fixed m as fewer measurements $n = m - r$ are input to OMP. An ideal number r of CV measurements should not be too large or too small; Figure 1 suggests that setting aside just enough measurements r such that $\epsilon \leq .6$ is satisfied in (34) serves as a good heuristic to choose the number of cross validation measurements (in Figure 1, $\epsilon \leq .6$ is satisfied by taking only $r = 30$ measurements).

We indicate the theoretical bound (36) with dark gray in Figure 1, which is compared to the interval in light gray of the 990 values of $\eta_{cv}(q)$ that are closest to η_{or} in actuality.

This experiment is run for several values of r within the interval $[5, 90]$; the results are plotted in Figure 1(a), with the particular range $r \in [5, 30]$ blown up in Figure 1(b).

We have also carried out this experiment with a smaller noise variance; i.e. $x_b = x_0 + \mathcal{N}_b$ is subject to additive noise

$$\mathcal{N}_b \sim N(0, .02). \quad (38)$$

The signal x_b is again renormalized to satisfy $\|x_b\|_{l_2^N} = 1$; it now has an expected noise level of

$$E(\sigma_d(x_b)) \approx .116. \quad (39)$$

The results of this experiment are plotted in Figure 1(c).

6.2 Experimental Results

1. We remind the reader that the cross-validation estimates $\widehat{\eta}_{cv}$ are observable to the user, while the values of η_{omp} , η_{or} , along with the noise level $\sigma_d(x)$, are not available to the user. Nevertheless, $\widehat{\eta}_{cv}$ can serve as a proxy for η_{or} according to (36), and this is verified by the plots in Figure 1. $\widehat{\eta}_{cv}$ can also provide an upper bound on $\sigma_d(x)$, as is detailed in Section 5.1.
2. The theoretical bound (36) is seen to be tight, when compared with the observed concentration bounds in Figure 1.
3. With high probability, the estimate $\hat{x}_{cv}(15)$ using $r = 15$ out of the allotted $m = 800$ measurements will be a better estimate of x than the OMP estimate: $\|\hat{x}_{cv}(15) - x\|_2 \leq \|\hat{x}_{omp}(15) - x\|_2$. With *overwhelming probability*, the estimate $\hat{x}_{cv}(30)$ will result in error $\|\hat{x}_{cv}(30) - x\|_2 \leq \|\hat{x}_{omp}(30) - x\|_2$. We note that the estimates $\hat{x}_{cv}(15)$ and $\hat{x}_{cv}(30)$ correspond to accuracy parameters $\epsilon(15) = .8405$ and $\epsilon(30) = .5943$ in (34), indicating that $\epsilon \leq .6$ is a good heuristic to determine when enough CV measurements have been reserved.
4. The OMP-CV estimate \hat{x}_{cv} will have more pronounced improvement over the OMP estimate \hat{x}_{omp} when there is larger discrepancy between the true sparsity d of x_0 and the upper bound k used by OMP (in Figure (1), $d = 100$ and $k = 200$). In contrast, OMP-CV will not outperform OMP in approximation accuracy when d is close to k ; however, the multiplicative relation (36) guarantees that OMP-CV will not underperform OMP, either.

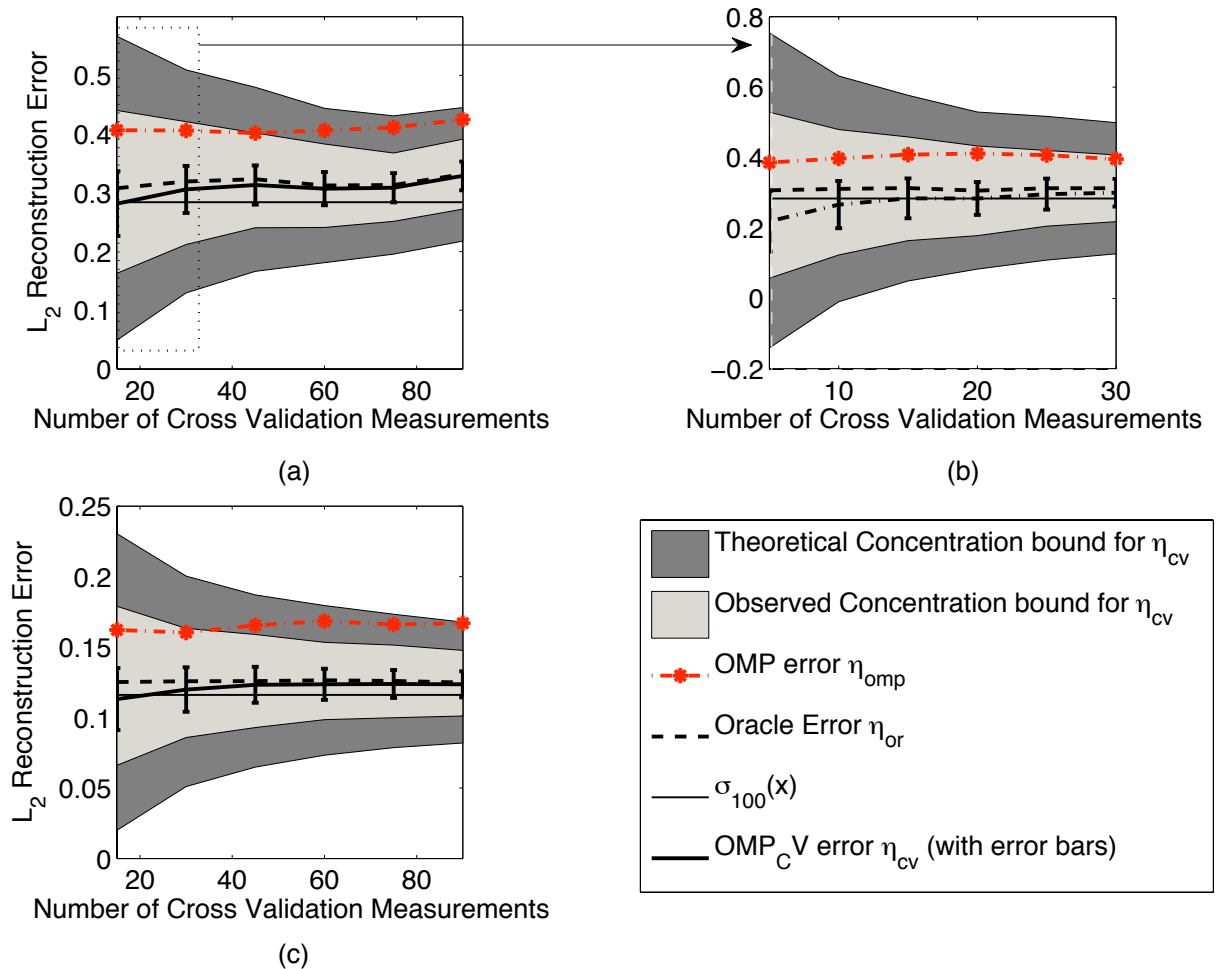


Figure 1: Comparison of the reconstruction algorithms OMP and OMP-CV. We fix the parameters $N = 3600$, $m = 800$, $k = 200$, and underlying sparsity $d = 100$, but vary the number r of the total m measurements reserved for cross validation from 5 to 90, using the remaining $n = m - r$ measurements for training. The underlying signal has residual $\sigma_{100}(x) \approx .284$ in Figures 1(a) and 1(b), and $\sigma_{100}(x) \approx .116$ in Figure 1(c), as shown for reference by the thin horizontal line. In both cases, the OMP-CV error η_{cv} (the solid black line with error bars; each point represents the average of 1000 trials) gives a better approximation to the residual error than does OMP (dot-dashed line) with very high probability, even when as few as 20 of the total 800 measurements are used for cross validation. Even though η_{cv} is guaranteed to provide a tighter bound for η_{or} as the number r of CV measurements increases, at the same time, the oracle error η_{or} becomes a worse indicator of the residual $\sigma_{100}(x)$ because fewer measurements $n = m - r$ are input to OMP.

7 Beyond Compressed Sensing

The Compressed Sensing setup can be viewed within the more general class of *underdetermined linear inverse problems*, in which $x \in \mathbb{R}^N$ is to be reconstructed from a known $m \times N$ underdetermined matrix \mathcal{A} and lower dimensional vector $y = \mathcal{A}x$ using a decoding algorithm $\Delta : \mathbb{R}^m \rightarrow \mathbb{R}^N$; in this broader context, \mathcal{A} is given to the user, but not necessarily *specified* by the user as in compressed sensing. In many cases, a prior assumption of sparsity is imposed on x , and an iterative decoding algorithm such as LASSO (27) will be used to reconstruct x from y [17]. If it is possible to take on the order of $r = \log p$ additional measurements of x by an $r \times N$ matrix Ψ satisfying the conditions of Lemma (3.1), then all of the analysis presented in this paper applies to this more general setting. In particular, the error $\|x - \hat{x}_j\|_{\ell_2^N}$ at up to $j \leq p$ successive approximations \hat{x}_j of the decoding algorithm Δ may be bounded from below and above using the quantities $\|\Psi(x - \hat{x}_j)\|_{\ell_2^r}$, and the final approximation \hat{x} to x can be chosen from among the entire sequence of estimates \hat{x}_j as outlined in Theorem (4.1); an earlier estimate \hat{x}_j may approximate x better than a final estimate \hat{x}_p which contains the artifacts of parameter overfitting occurring at later stages of iteration.

8 Extensions and Open Problems

We have presented an alternative approach to compressed sensing in which a certain number r of the m allowed measurements of a signal $x \in \mathbb{R}^N$ are reserved to track the error in decoding by the remaining $m - r$ measurements, allowing us to choose a best approximation to x in the metric of ℓ_2^N out of a sequence of p estimates $(\hat{x}_j)_{j=1}^p$, and estimate the error between x and its best approximation by a k -sparse vector, again with respect to the metric of ℓ_2^N . We detailed how the number r of such measurements should be chosen in terms of desired accuracy ϵ of estimation, confidence level ξ in the prediction, and number p of decoding iterations to be measured; in general, $r = O(\log(p))$ measurements suffice. Several important issues remain unresolved; we mention only a few below.

1. The cross validation technique promoted in this paper corresponds specifically to the technique of *holdout cross validation* in statistics, where a data set is partitioned into a single training and cross validation set (as a rule of thumb, the cross validation set is usually taken to be less than or equal to a third of the size of the training set; in the the current paper, we have shown that the Johnson Lindenstrauss lemma provides a theoretical justification of how many, or, more precisely, how few, cross validation measurements are needed in the context of compressed sensing). Other forms of cross validation, such as *repeated random subsampling* cross validation or *K-fold* cross validation, remain to be analyzed in the context of compressed sensing. The former technique corresponds to repeated application of holdout cross validation, with r cross validation measurements out of the total m measurements chosen by random selection at each application. The results are then averaged (or otherwise combined) to produce a single estimation. The latter technique, *K-fold* cross validation, also corresponds to repeated application of holdout cross validation. In this case, the m measurements are partitioned into K subsets of equal size r , and cross-validation is repeated exactly K times with each of the K subsets of measurements used exactly once as the validation set. The K results are again combined to produce a single estimation. Although Theorem (4.1) does not directly apply to these cross validation models, the experimental results of Section 6 suggest that, equipped with an $m \times N$ matrix satisfying the requirements of Lemma (3.1), the application of K fold cross validation to subsets of the measurements of size $r \ll m - r$ just large enough that $\epsilon > 0$ in Theorem (4.1) for fixed accuracy ξ and constant $C = 1$ can be combined to

accurately approximate the underlying signal with near certainty.

- It is not clear that the analysis in Theorem (4.1) can be extended to the noisy compressed sensing model,

$$y = \Phi x + \mathcal{N}, \quad (40)$$

where $\mathcal{N} \sim N(0, \sigma^2)$ is a Gaussian random variable that accounts for both noise and quantization error on the measurements Φx . Because measurement noise and quantization error are unavoidable in any real-world sensing device, any proposed compressed sensing technique should extend to the model (40). Indeed, cross validation is studied in [7] in this context as a stopping criterion for decoding algorithms of homotopy/greedy type, in the case that x is truly sparse and \mathcal{N} is Gaussian noise. The experimental results in [7] indicate that cross validation works well in this setting, but it remains to provide theoretical justification of these results.

- We have only considered cross validation over the metric of ℓ_2 . However, the error $\|x - \hat{x}\|_{\ell_2^N}$, or *root mean squared error*, is just one of several metrics used in image processing for analyzing the quality of a reconstruction $\hat{x} \in \mathbb{R}^N$ to a (known) image $x \in \mathbb{R}^N$. In fact, the ℓ_1 reconstruction error $\|x - \hat{x}\|_{\ell_1^N}$ has been argued to outperform the root mean squared error as an indicator of reconstruction quality [23]. Unfortunately, Theorem (4.1) cannot be extended to the metric of ℓ_1 , as there exists no ℓ_1 analog of the Johnson Lindenstrauss Lemma [27]. However, it remains to understand the extent to which cross validation in compressed sensing can be applied over a broader class of image reconstruction metrics, perhaps using more refined techniques than those considered in this paper.
- Many more compressed sensing matrices than just those satisfying the requirements of Lemma (3.1) are observed *empirically* to satisfy the Johnson Lindenstrauss Lemma; in particular, a properly normalized $r \times N$ matrix obtained by selecting r rows at random from the $N \times N$ discrete Fourier matrix is observed to satisfy (3.1) for number of measurements $r \geq \epsilon^{-2} \log \frac{1}{2\delta}$, and the empirical concentration bounds for these matrices appear to be *indistinguishable* from those of (properly normalized) random Gaussian and Bernoulli matrices of the same dimension. This suggests that cross validation should be considered as a general technique that can be applied to a set of m compressed sensing measurements, the theoretical justification of which is a very interesting problem that we hope to pursue in the future.

Acknowledgment

The author would like to thank Ingrid Daubechies and Albert Cohen for their insights and encouragement, without which this paper could not have been written.

References

- E. Candes, and T. Tao, *Decoding by Linear Programming*, IEEE Trans. Info. Theory, 51(12):4203-4215, Dec. 2005.
- E.. Candes, J. Romberg, and T. Tao. *Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information*. IEEE Transactions on Information Theory, vol. 52, no. 2, pp. 489 - 509, Feb. 2006.
- D. Donoho, and M. Elad, *Optimally sparse representation from overcomplete dictionaries via ℓ_1 norm minimization* Proc. Natl. Acad. Sci. USA, pp. 2197-2002, 2003.

- [4] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, *A simple proof of the restricted isometry property for random matrices*. Constructive Approximation, 2007.
- [5] A. Cohen, R. DeVore and W. Dahmen, *Compressed sensing and best k-term approximation*, J. Amer. Math. Soc. 22 (2009), 211-231.
- [6] J. Tropp and A. Gilbert. *Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit*. Information Theory, IEEE Transactions on Volume 53, Issue 12, Dec. 2007 Page(s):4655 - 4666.
- [7] P. Boufounos, M. Duarte, and R. Baraniuk, *Sparse Signal Reconstruction from Noisy Compressive Measurements Using Cross Validation*, IEEE Workshop on Statistical Signal Processing, 2007, Madison, WI, pp. 299-303.
- [8] D. Needell, R. Vershynin, *Signal Recovery from Inaccurate and Incomplete Measurements via Regularized Orthogonal Matching Pursuit*, Submitted, Dec. 2007.
- [9] D.L. Donoho. *Compressed Sensing*. IEEE Trans. Info. Theory, 52(4):1289-1306, Apr. 2006.
- [10] A. Nemirovski. *Advances in Convex Optimization: Conic Programming*. Volume I (Plenary Lectures) of International Congress of Mathematicians, Madrid 2006, European Mathematical Society, 2006.
- [11] D. Achlioptas, *Database-friendly random projections*, Journal of Computer and System Sciences, 2003.
- [12] P. Indyk, R. Motwani, *Approximate Nearest Neighbors: Towards Removing the Curse of Dimensionality*. Proc. 30th Symposium on Theory of Computing, 1998, pp. 604-613.
- [13] W. Johnson, J. Lindenstrauss, *Extensions of Lipschitz maps into a Hilbert space*. Contemp. Math. 26, 1984, pp. 189-206.
- [14] S. Dasgupta, A. Gupta, *An elementary proof of the Johnson-Lindenstrauss Lemma*. No. TR-99-006. (1999).
- [15] P. Frankl and H. Maehara. *The Johnson-Lindenstrauss Lemma and the Sphericity of Some Graphs*. Journal of Combinatorial Theory B, 44(1988):355-362.
- [16] J. Matousek. *On Variants of the Johnson-Lindenstrauss Lemma*. Random Struc. Algo., to appear.
- [17] I. Daubechies, M. Defrise, C. De Mol. *An iterative thresholding algorithm for linear inverse problems with a sparsity constraint*. Communications on Pure and Applied Mathematics 57, pp. 1413-1457, August 2004.
- [18] P. Wojtaszczyk. *Stability and instance optimality for Gaussian measurements in compressed sensing*. (Preprint, 2008)
- [19] E. Candes, *Compressive Sampling* Int. Congress of Mathematics, 3, pp. 1433-1452, Madrid, Spain, 2006.
- [20] R. DeVore, G. Petrova, and P. Wojtaszczyk, *Instance-optimality in probability with an ℓ_1 decoder*. (Preprint, 2008)
- [21] E. Candes. *The Restricted Isometry Property and its Implications for Compressed Sensing*. Comptes Rendus de l'Academie des Sciences, Paris, Series I, 346, pp. 589-594, 2008.

- [22] E. Candes and T. Tao, *The Dantzig Selector: Statistical estimation when p is much larger than n* Submitted, Annals of Statistics.
- [23] R. DeVore, B. Jawerth, and B. Lucier. *Image compression through wavelet transform coding.* IEEE Trans. Inform. Theory, 38(2):719-746, March 1992.
- [24] G. James, P. Radchenko, and J. Lv. *DASSO: Connections between the Dantzig Selector and Lasso.* Journal of the Royal Statistical Society, Series B, 71, 127-142, 2009.
- [25] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. *Least angle regression*, Annals of Statistics, vol. 32, no.2, pp. 407-499, 2004.
- [26] D. Malioutov, M. Cetin, A. Willsky. *Homotopy continuation for sparse signal representation* Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.
- [27] M. Charikar, A. Sahai. *Dimension Reduction in the ℓ_1 Norm.* Proceedings of the 43rd Symposium on Foundations of Computer Science, pages 551-560, 2002.
- [28] M. Rudelson, R. Vershynin, *Sparse reconstruction by convex relaxation: Fourier and Gaussian measurements*, CISS 2006 (40th Annual Conference on Information Sciences and Systems).
- [29] I. Daubechies, R. DeVore, M. Fornasier, and C. Sinan Gntkr, *Iteratively re-weighted least squares minimization for sparse recovery* (Preprint, 2008).
- [30] D. Needell and J. Tropp. *CoSaMP: Iterative signal recovery from incomplete and inaccurate samples.* Appl. Comp. Harmonic Anal., DOI: 10.1016/j.acha.2008.07.002.
- [31] W. Dai, O. Milenkovic. *Subspace Pursuit for Compressive Sensing: Closing the Gap Between Performance and Complexity.* Preprint, 2008.
- [32] E. Candes, L. Demanet, D. Donoho, and L. Ying. *Fast Discrete Curvelet Transforms*, Multiscale Model. Simul., 5 861-899.
- [33] R. Tibshirani. *Regression shrinkage and selection via the lasso.* J. Royal. Statist. Soc B., Vol. 58, No. 1, pages 267-288.