Sequential Adaptive Compressed Sampling via Huffman Codes

Akram Aldroubi

Department of Mathematics, Vanderbilt University
Nashville, TN 37240
akram.aldroubi@vanderbilt.edu

Haichao Wang

Department of Mathematics, University of California
Davis, CA 95616
hchwang@ucdavis.edu

Kourosh Zarrinhalam

Department of Mathematics, Vanderbilt University
Nashville, TN 37240
Kourosh.Zarrinhalam@vanderbilt.edu

Abstract

In this paper we introduce an information theoretic approach to construct a sequence of binary sampling vectors to determine a sparse signal. Unlike the standard approaches, ours is adaptive in the sense that each sampling vector depends on the previous sample results. We prove that the expected total cost (number of measurements and reconstruction combined) we need for an s -sparse vector in $\mathbb{R}^n$ is no more than $s \log n + 2s$.

Key words and phrases : Sampling, Sparsity, Compressed Sampling, Huffman codes

2000 AMS Mathematics Subject Classification —94A12, 94A20

1 Introduction

Let x be a vector in $\mathbb{R}^n$ and assume that the vector x has at most $s \ll n$ nonzero components, denoted $\|x\|_0 \leq s$. In compressed sampling, the goal is to determine a set of linear functionals (*the sampling functions*) and associated reconstruction algorithms such that, if the set of functionals is applied to x , then the reconstruction algorithms will allow us to find x from the values of the functionals (measurements of x) in a computationally tractable way which is stable (i.e., that should produce a good approximation even when x is not

s -sparse) and robust to noise (i.e. that should produce a good approximation to x when the measurements (sampling) are corrupted by additive noise). There is a trade off between the number of sampling vectors we need to acquire x and the computational cost of the reconstruction algorithm which determines x from the samples.

Some of the earlier work uses an ℓ_1 minimization for finding x from a set of samples $\{y_i = \langle a_i, x \rangle, i = 1, \dots, m\}$ where the a_i s are vectors in $\mathbb{R}^n$, see e.g., [5, 6, 11, 12, 20, 30]. Letting $y = (y_1, \dots, y_m)^t$ and A be the $m \times n$ matrix whose rows are the vectors a_i , the ℓ_1 minimization approach finds the unknown s -sparse vector x (i.e., $\|x\|_0 \leq s$) by solving the constrained minimization problem

$$\min \|z\|_{\ell_1}, \quad Az = y. \quad (1)$$

The solution to the minimization problem produces x as its unique solution provided A is an appropriate sampling matrix, e.g., A satisfying the *Restricted Isometric Property* (RIP). The current methods for constructing matrices satisfying the needed RIP are probabilistic and produce matrices with $m = O(s \log(n/2s))$ rows (i.e., we need $m = O(s \log(n/2s))$ sampling vectors). The deterministic construction of a matrix A satisfying the needed RIP property for arbitrary n ($m \ll n$) is still open. Finding x via the ℓ_1 minimization problem involves linear programming with n variables and m constraints which can be computationally expensive.

Matching pursuits (OMP, ROMP, CoSaMP. etc., see [25, 26, 15] and the reference therein) form another class of sampling and reconstruction algorithms for finding an s -sparse signal x from the measurements $y = Ax$. These algorithms are iterative in nature and view the column of A as a dictionary for the reconstruction of x . At every step of the reconstruction algorithms, a certain number of columns (column indices) are chosen in order to minimize an equation of the form $\|r_n - Ax_n\|_2$, where $r_n = y - Ax_{n-1}$ is a residual from the previous step. The most performant of these type of reconstruction algorithms is relatively fast, stable, and robust to noise. However, similar to the ℓ_1 minimization algorithms, it requires a sampling matrix A satisfying the needed RIP. As before only probabilistic methods are known, and they produce matrices with $m = O(s \log(n/2s))$ rows that satisfy the RIP with high probability.

Other approaches for compressed sampling are combinatorial. A sampling matrix A is constructed using bipartite graphs, such as expander graphs, and the reconstruction finds an unknown s -sparse vector x using binary search methods, see e.g. [4, 17, 14, 29, 28, 18, 19, 31] and the references therein. Typically, the matrix A has binary entries. There exist fast algorithms for finding the solution x from the measurements (typically sublinear). However, the construction of A is still difficult to produce.

There are emerging new approaches for adaptive methods in compressed sampling. One approach uses a Bayesian method combined a gaussian model

for the measurements and a Laplacian model for the sparsity [22]. The sampling vectors are chosen by minimizing a differential entropy. Another approach uses a type of binary search algorithm that uses back projection and block adaptive sampling to focus on the possible nonzero component (see [21] the references therein).

On the other hand, information theoretic approaches have also been discussed. In particular, one can design sampling vectors in a signal-dependent, but not adaptive fashion (see [8]). In this method, a priori knowledge about the class of signals is given. The information is used to find a sensing matrix W that maximizes the mutual information between the signals and their projections. Once W is found, it is used to sample any signal from the class.

The new approach we present is also information theoretic and uses tools from the theory of Huffman codes to develop a deterministic construction of a sequence of binary sampling vectors a_Λ , i.e., the entries of a_Λ consist of 0 or 1. Similar to the informative sensing method described above, the design of sampling vector depends on knowledge about the class of signals. However, in our approach we do not design a fixed matrix W that can be applied to any signal of the class. Instead, our sampling is sequential, adaptive and depends on each signal, i.e., in our approach $W = W(x)$ is a function of the individual signal x . We assume that the signal $x \in \mathbb{R}^n$ is an instance of a vector random variable $X = (X_1, \dots, X_n)^t$ and we construct the i -th row a_i of A using the sample $y_{i-1} = \langle a_{i-1}, x \rangle$. Our goal is to make the average number of samples needed to determine a signal x as small as possible, and to make the reconstruction of x from those samples as fast as possible. We take advantage of the probability distribution of the random vector X to minimize the average number of samples needed to uniquely determine the unknown signal x . In our method, rather than constructing a fixed set of sampling vectors forming the rows of a single sampling matrix A for all possible signals, we construct s sequences of sampling vectors. Each sequence focuses on finding exactly one nonzero component of the signal x .

It is remarkable that the expected total cost of the combined sampling and reconstruction algorithms of an s -sparse vector x is no more than $s \log n + 2s$. If no information is available about the probability distribution of the random vector X , we can always assume that we have a uniform distribution, in which case the total cost of the combined sampling and reconstruction algorithms of an s -sparse vector x is equal to $s \log n + 2s$ even if the uniform assumption is erroneous.

This paper is organized as follows: In Section 2 we introduce the basic notation and definitions for sparse random vectors and trees. The new notion of Huffman tree and Huffman sampling vectors are introduced in Section 3.1 and 3.2. In Sections 3.3 and 3.4, we describe an information theoretic method for finding s -sparse vectors in $\mathbb{R}^n$. In Section 3.5 we describe variations for finding

s -sparse signals from noisy measurements and for finding compressible signals. Section 4 is devoted to examples, simulations, and testing of the algorithms on synthetic data.

2 Notation and Preliminaries

In this section we introduce the necessary notations and preliminaries needed in subsequent sections.

2.1 Sparse random vectors

1. We will use the notation $X = (X_1, \dots, X_n)^t$ to denote a vector of n random variables. An instance $x \in \mathbb{R}^n$ of X will be called a *signal*.
2. We will say that a signal $x \in \mathbb{R}^n$ has sparsity $s \leq n$ if x has at most s nonzero components ($\|x\|_0 \leq s$).
3. Let $\Omega = \{1, \dots, n\}$ be the set of all indices, and $\Lambda \subset \Omega$ be a subset of indices. Then $P_\Lambda = P(\Lambda)$ will denote the probability of X having nonzero components exactly at coordinates indexed by Λ , i.e., $P_\Lambda = Pr\{X_i \neq 0, i \in \Lambda; X_i = 0, i \in \Lambda^c\}$. Here Λ^c denotes the complement of Λ .
4. The pair (X, P) will be used to denote the random vector X together with the probability mass function P on the sample space $2^\Omega = 2^{\{1, \dots, n\}}$. Thus obviously $\sum_\Lambda P_\Lambda = 1$.
5. We will say that a random vector X is s sparse if $P_\Lambda = 0$ for all Λ with cardinality strictly larger than s , i.e., $\#(\Lambda) > s$ implies $P_\Lambda = 0$.
6. We will need the probabilities $q_\Lambda = q(\Lambda) = Pr(E_\Lambda)$ for the events $E_\Lambda = \{X_i \neq 0, \text{ for some } i \in \Lambda\}$, which is the probability that at least one of the components of X with index in Λ is nonzero. Note that q can be computed from P by

$$q_\Lambda = \sum_{\eta \cap \Lambda \neq \emptyset} P_\eta. \quad (2)$$

2.2 Trees

1. We consider finite *full binary trees* in which every node has zero or two children.
2. The *root* of the tree is the node with no parent nodes.
3. A *leaf* is a node with no children.

4. A *left (right) subtree* of a node v in a rooted binary tree is the tree whose root is the left (right) child of this node v .
5. The set of all nodes that can be reached from the root by a path of length L are said to be at *level* L .

2.3 Other notations

1. The notation χ_Λ will denote the characteristic function of a set Λ , i.e., $\chi_\Lambda(i) = 1$ for $i \in \Lambda$ and $\chi_\Lambda(i) = 0$ for $i \notin \Lambda$.
2. For a set Λ , $|\Lambda|$ will denote its cardinality.

3 Theory

In this section we describe our approach explicitly.

3.1 Huffman tree

Let (X, P) be a s -sparse random vector X in $\mathbb{R}^n$ together with the probability mass function P on the sample space $2^\Omega = 2^{\{1, \dots, n\}}$ the set of all subsets of $\Omega = \{1, \dots, n\}$. We define a Huffman tree to be a binary tree whose leaves are the sets $\{1\}, \dots, \{n\}$. We associate probabilities $q_{\{1\}}, \dots, q_{\{n\}}$ to these nodes respectively. The Huffman tree is constructed from the leaves to the root as follows: Suppose that the nodes at the l -th step are $\Lambda_1, \dots, \Lambda_m$. Let i and j be such that

$$\begin{aligned} q_{\Lambda_i} &= \min_{1 \leq \lambda \leq m} q_{\Lambda_\lambda}, \\ q_{\Lambda_j} &= \min_{1 \leq \lambda \leq m, \lambda \neq i} q_{\Lambda_\lambda}. \end{aligned}$$

Then, the nodes at the $(l+1)$ -th step are obtained by replacing Λ_i and Λ_j with $\Lambda_i \cup \Lambda_j$ in the list of the nodes at the l -th step and the probability associated to the node $\Lambda_i \cup \Lambda_j$, will be $q_{\Lambda_i \cup \Lambda_j}$, i.e., at each step the two nodes with smallest probabilities are combined. An illustrative example is shown below.

Example 3.1. Assume that $X \in \mathbb{R}^4$ is a 2-sparse random vector with probability mass function P defined by: $P_\emptyset = 0.02, P_{\{1\}} = 0.07, P_{\{2\}} = 0.05, P_{\{3\}} = 0.03, P_{\{4\}} = 0.1, P_{\{1,2\}} = 0.31, P_{\{1,3\}} = 0.2, P_{\{1,4\}} = 0.03, P_{\{2,3\}} = 0.06, P_{\{2,4\}} = 0.12, P_{\{3,4\}} = 0.01$. The nodes at the first step are: $\{1\}, \dots, \{4\}$. Using equation 2, we get that:

$$q_{\{1\}} = P_{\{1\}} + P_{\{1,2\}} + P_{\{1,3\}} + P_{\{1,4\}} = 0.61.$$

Similarly, $q_{\{2\}} = 0.54, q_{\{3\}} = 0.3$, and $q_{\{4\}} = 0.26$. Therefore the nodes at

Figure 1: Huffman Tree for the 2-sparse random vector $X \in \mathbb{R}^4$.

the second step are $\{1\}, \{2\}, \{3, 4\}$. Also $q_{\{3,4\}} = 0.55$ and hence the nodes at the third step are $\{1\}, \{2, 3, 4\}$. The root node is $\{1, 2, 3, 4\}$ with probability $q_{\{1,2,3,4\}} = 1$. This completes the Huffman tree (See Figure 1).

3.2 Huffman sampling vectors

In this section we introduce Huffman sampling vectors. Let (X, P) be as in the previous section. Our goal is to find a signal x which is an instance of X using, on average, a minimum number of samples. As mentioned before, we first focus on finding one nonzero component, if it exists. Other nonzero components are then iteratively found one at a time in a similar manner until all of the nonzero components are exhausted.

The average number of samples is measured by

$$L(X, P) = \sum_{\Lambda} P_{\Lambda} \ell_{\Lambda}, \quad (3)$$

where ℓ_{Λ} is the number of samples for finding x whose components $x_i \neq 0$ for $i \in \Lambda$ and $x_i = 0$ for $i \in \Lambda^c$.

We first construct a Huffman tree associated with the random vector X . Each node Λ which is not a leaf has two children Λ_1 and Λ_2 . Note that $\Lambda_1 \cap \Lambda_2 = \emptyset$ and $\Lambda_1 \cup \Lambda_2 = \Lambda$. We denote $\ell_{\Lambda_1} = q_{\Lambda_1}(\log |\Lambda_1| + 1) + (1 - q_{\Lambda_1})(\log |\Lambda_2| + 1)$ and $\ell_{\Lambda_2} = q_{\Lambda_2}(\log |\Lambda_2| + 1) + (1 - q_{\Lambda_2})(\log |\Lambda_1| + 1)$. We associate a sampling vector to such nodes Λ by:

$$a_{\Lambda} = \begin{cases} \chi_{\Lambda_1} & \text{if } \ell_{\Lambda_1} \leq \ell_{\Lambda_2}, \\ \chi_{\Lambda_2} & \text{if } \ell_{\Lambda_1} > \ell_{\Lambda_2}, \end{cases} \quad (4)$$

i.e., for $\ell_{\Lambda_1} \leq \ell_{\Lambda_2}$, we have $a_{\Lambda}(i) = 1$ for $i \in \Lambda_1$ and $a_{\Lambda}(i) = 0$ for $i \in \Lambda_1^c$, and for $\ell_{\Lambda_1} > \ell_{\Lambda_2}$ we have $a_{\Lambda}(i) = 1$ for $i \in \Lambda_2$ and $a_{\Lambda}(i) = 0$ for $i \in \Lambda_2^c$.

The choice of the sampling vector in (4) can be seen as follows: Since our goal is to find the nonzero component as quickly as possible, it seems that we would need $a_{\Lambda} = \chi_{\Lambda_i}$ for the set Λ_i with the highest probability q_{Λ_i} , $i = 1, 2$. However, the set Λ_i with the highest probability q_{Λ_i} may also have a large number of elements. Thus the choice should be a compromise between the size of the set and its probability. This particular choice of a_{Λ} will be apparent from the theorems and their proofs below.

3.3 Determination of a sparse vector x using Huffman sampling vectors.

Let x be an s -sparse signal in $\mathbb{R}^n$ which is an instance of (X, P) (we will write $x \sim (X, P)$). We make the additional assumption that the conditional probability $Pr(\sum_{i \in \Lambda} X_i \neq 0 | X_i \neq 0, i \in \Lambda) = 0$ holds for any $\Lambda \subset \Omega$ (recall that $\Omega = \{1, \dots, n\}$). This is a natural condition if the random variables X_i , $i = 1, \dots, n$, in the random vector X do not have a positive mass concentration except possibly at zero.

3.3.1 Finding a nonzero component

Algorithm 1 below is used to find the position and the corresponding value of one of the nonzero components of x (if any).

Algorithm 1.

1. *Initialization:* $\Lambda = \Omega$;
2. *Repeat until* $|\Lambda| = 1$
 - if* $\langle a_{\Lambda}, x \rangle \neq 0$, $\Lambda = \Lambda_1$
 - else* $\Lambda = \Lambda_2$*end repeat*
3. *Output the (only) element* $t_1 \in \Lambda$
4. *Output* $x_{t_1} = \langle \chi_{\{t_1\}}, x \rangle$

Remark 1.

- (i) *If the vector $x = 0$, then the algorithm will find an output $x_{t_1} = 0$, otherwise it will output the value x_{t_1} of one of the nonzero components of x and its index t_1 .*

- (ii) Note that the last inner product output in (4) is not always necessary since we have all the information needed to find x_{t_1} from the samples and the value of t_1 . However, this would involve solving a linear system of equations obtained from the sampling scheme. Thus this one extra sample can be considered as reconstruction step.
- (iii) Note that the sampling vectors depend on the instance x , i.e., the sampling vectors are adaptive.
- (iv) The number of possible sampling vectors is equal to the number of nodes in the Huffman tree, but only a subset of these vectors is used to determine a nonzero component of a given instance vector x .
- (v) If $P(X = 0) > 0$, then we choose the first sampling vector $a = (1, 1, \dots, 1)$. If $\langle a, x \rangle = 0$ we are done. Otherwise we proceed with Algorithm 1.

The first observation is that Algorithm 1 is optimal for 1-sparse vectors. This should not be a surprise since the algorithm was inspired by the theory of Huffman codes. We have

Theorem 3.2. *Given a 1-sparse vector $x \sim (X, P)$ in $\mathbb{R}^n$. Then the average number of samples $L_1(X, P)$ needed to find x using Algorithm 1 is less than or equal to the average number of samples $L_A(X, P)$ for finding x using any algorithm A with binary sampling vectors.*

Proof. Let $E = \{e_i\}_{i=1}^n$ be the canonical basis for $\mathbb{R}^n$. We first note that for any sampling algorithm with binary sampling vectors, the number of samples required to determine any nonzero multiple αe_i of e_i ($\alpha \in \mathbb{R}$) is equal to the number required to determine e_i . Hence for the remainder of this proof we will assume that x is binary, i.e., x is one of the canonical vectors e_i . Since any sampling vector a from an algorithm A is binary, the inner product $\langle a, x \rangle$ is either zero or one, i.e., binary. For each e_i , $i \in \Omega = \{1, \dots, n\}$, a binary algorithm uses a sequence of vectors $\{a_1^i, \dots, a_{\ell_i}^i\}$ where ℓ_i is the number of sampling vectors needed to determine e_i . We associate to each e_i the binary sequence $c^i = y_1^i y_2^i \dots y_{\ell_i}^i$ where $y_j^i = \langle a_j^i, e_i \rangle$. In this way each canonical vector e_i is associated to a unique sequence c^i , $c^i \neq c^j$ if $i \neq j$. Hence c^i is a binary code for the vectors of the canonical basis of $\mathbb{R}^n$. The code is a prefix (also known as instantaneous code, i.e., no code is a prefix of any other code [9]) since a binary algorithm terminates after finding the nonzero component (which correspond to the shorter code). Since x is 1 sparse, we have that $q_i = P_i$ which is the probability of component i being nonzero and all other components being zero. Hence for each algorithm A with binary sampling vectors, we associate a prefix code for $E = \{e_i\}_{i=1}^n$. Consequently, the average number of samplings required in this algorithm A is the same as the average length of the code:

$L_A(X, P) = \sum_{i=1}^n \ell_i P_i$. From the construction, Algorithm 1 is associated with the Huffman code whose average length is the shortest. Hence the average number of samples $L_1(X, P) \leq L_A(X, P)$ for any A . $\square$

The 1-sparse case is a special case. It is optimal because the sampling scheme can be associated exactly with the Huffman codes and we have that $q_\Lambda = P_\Lambda$ and $\sum_i q_{\{i\}} = \sum_i P_{\{i\}} = 1$. In fact, the choice of a_Λ in (4) can be chosen to be either χ_{Λ_1} or χ_{Λ_2} independently of the values of $\ell_{\Lambda_1}, \ell_{\Lambda_2}$. However, for the general s -sparse case, we do not have $q_\Lambda = P_\Lambda$ anymore, and the sampling vectors cannot be associated with Huffman codes directly. Thus for the s -sparse case, the average number of sampling vectors is not necessarily optimal and we need to estimate this number to have confidence in the algorithm.

From the construction of the Huffman tree in Section 3.1, it is not difficult to see that there is at most one node Λ (called *special node*) with children Λ_1 and Λ_2 such that $(\frac{1}{2} - q_{\Lambda_1})(\frac{1}{2} - q_{\Lambda_2}) < 0$, i.e., except for possibly the special node, all other nodes have the property that $q_{\Lambda_1}, q_{\Lambda_2}$ are either both larger than $1/2$ or both smaller than $1/2$. Thus, a Huffman tree can at most have one special node. We have the following lemmas:

Lemma 3.3. *For any fixed node Λ with children Λ_1 and Λ_2 , if Λ is not a special node, then $\min\{\ell_{\Lambda_1}, \ell_{\Lambda_2}\} \leq \log |\Lambda|$.*

Proof. Without loss of generality, we assume that $|\Lambda_1| \leq |\Lambda_2|$. Hence $\frac{|\Lambda_1|}{|\Lambda|} \leq \frac{1}{2}$ and $\frac{|\Lambda_2|}{|\Lambda|} \geq \frac{1}{2}$.

Consider the function $f_q(t) = t^q(1-t)^{1-q}$ for $t \in [0, 1]$ and $q \in [0, 1]$. Easy computations show that

$$f_q(t) \leq \frac{1}{2}, \text{ for } q \leq \frac{1}{2} \text{ and } t \geq \frac{1}{2}, \quad (5)$$

and

$$f_q(t) \leq \frac{1}{2}, \text{ for } q \geq \frac{1}{2} \text{ and } t \leq \frac{1}{2}. \quad (6)$$

Since Λ is not a special node, we have that $(q_{\Lambda_1} - \frac{1}{2})(q_{\Lambda_2} - \frac{1}{2}) \geq 0$.

If $q_{\Lambda_1} \leq \frac{1}{2}$ and $q_{\Lambda_2} \leq \frac{1}{2}$, then using the fact that $\frac{|\Lambda_2|}{|\Lambda|} \geq \frac{1}{2}$ and (5), we get $f_{q_{\Lambda_2}}(\frac{|\Lambda_2|}{|\Lambda|}) \leq \frac{1}{2}$, that is

$$(\frac{|\Lambda_2|}{|\Lambda|})^{q_{\Lambda_2}} (1 - \frac{|\Lambda_2|}{|\Lambda|})^{1-q_{\Lambda_2}} \leq \frac{1}{2},$$

which implies that

$$2|\Lambda_2|^{q_{\Lambda_2}} |\Lambda_1|^{1-q_{\Lambda_2}} \leq |\Lambda|,$$

after taking the log function on both sides, we have

$$\ell_{\Lambda_2} \leq \log |\Lambda|.$$

A similar calculation for the case $q_{\Lambda_1} \geq \frac{1}{2}$ and $q_{\Lambda_2} \geq \frac{1}{2}$ yields

$$\ell_{\Lambda_1} \leq \log |\Lambda|.$$

In either case, we have that $\min\{\ell_{\Lambda_1}, \ell_{\Lambda_2}\} \leq \log |\Lambda|$. $\square$

For general node, we have the following,

Lemma 3.4. *For any fixed node Λ with children Λ_1 and Λ_2 , we have that $\min\{\ell_{\Lambda_1}, \ell_{\Lambda_2}\} \leq \log |\Lambda| + 1$.*

Proof. For any $x \in [0, 1]$ and any $q \in [0, 1]$ we have that $f_q(x) = x^q(1-x)^{1-q} \leq 1$. We use this inequality for $x = \frac{|\Lambda_1|}{|\Lambda|}$ and $q = q_{\Lambda_1}$ or $x = \frac{|\Lambda_2|}{|\Lambda|}$ and $q = q_{\Lambda_2}$ to obtain the result. $\square$

Lemma 3.5. *Given a nonzero s -sparse vector $x \sim (X, P)$ in $\mathbb{R}^n$. If the Huffman tree associated with x has no special node, then the average number of samples L needed to find the position of one nonzero component of x using Algorithm 1 is at most $\log n$.*

Proof. We will use induction on n to prove this lemma. Suppose $\Lambda = \{1, \dots, n\}$ and Λ has children Λ_1 and Λ_2 .

For $n = 2$, $|\Lambda| = 2$, we only need one vector χ_{Λ_1} or χ_{Λ_2} to determine the position of one nonzero component of x . Hence the lemma holds trivially for this case.

Now assume the lemma is true for $n = k, k-1, \dots, 2$, we want to show it is also true for $n = k+1$.

If $|\Lambda| = k+1$, we must have $|\Lambda_1| \leq k$ and $|\Lambda_2| \leq k$. Without loss of generality, suppose $\ell_{\Lambda_1} \leq \ell_{\Lambda_2}$. By Algorithm 1, $a_\Lambda = \chi_{\Lambda_1}$. Then with probability q_{Λ_1} , we have $\langle a_\Lambda, x \rangle \neq 0$, in which case we need (on average) another L_{Λ_1} sampling vectors. With probability $1 - q_{\Lambda_1}$, we have $\langle a_\Lambda, x \rangle = 0$, and we need (on average) another L_{Λ_2} sampling vectors. By the induction hypothesis, we have $L_{\Lambda_1} \leq \log |\Lambda_1|$ and $L_{\Lambda_2} \leq \log |\Lambda_2|$.

Since by assumption the tree has no special node, using Lemma 3.3 we deduce that the average number of sampling vectors we need is

$$\begin{aligned} L &= q_{\Lambda_1}(1 + L_{\Lambda_1}) + (1 - q_{\Lambda_1})(1 + L_{\Lambda_2}) \\ &\leq q_{\Lambda_1}(1 + \log |\Lambda_1|) + (1 - q_{\Lambda_1})(1 + \log |\Lambda_2|) \\ &= \ell_{\Lambda_1} \leq \log |\Lambda| = \log k. \end{aligned}$$

$\square$

We are now ready to find an upper bound on the average number of sampling vectors needed for finding the position of one nonzero component in an s -sparse signal using Algorithm 1. Denoting by T_Λ the subtree with Λ as the root, we have

Theorem 3.6. *Given a nonzero s -sparse vector $x \sim (X, P)$ in $\mathbb{R}^n$ the average number of samples L needed to find the position of one nonzero component of x using Algorithm 1 is at most $\log n + 1$.*

Proof. We will use induction. The lemma holds trivially for $n = 2$.

Now assume the lemma is true for $n = k, k - 1, \dots, 2$, we want to show it is also true for $n = k + 1$.

If $|\Lambda| = k + 1$, we must have $|\Lambda_1| \leq k$ and $|\Lambda_2| \leq k$. Without loss of generality, suppose $\ell_{\Lambda_1} \leq \ell_{\Lambda_2}$. By Algorithm 1, the average number L of sampling vectors needed is

$$L = q_{\Lambda_1}(1 + L_{\Lambda_1}) + (1 - q_{\Lambda_1})(1 + L_{\Lambda_2}). \quad (7)$$

Since the Huffman tree can have at most one special node, we consider three cases:

Case(1): If the root of the tree $\Lambda = \{1, \dots, n\}$ is a special node and Λ_1, Λ_2 are its children. Then the subtrees T_{Λ_1} and T_{Λ_2} have no special nodes. Thus by Lemma 3.5, we have that $L_{\Lambda_1} \leq \log |\Lambda_1|$ and $L_{\Lambda_2} \leq \log |\Lambda_2|$.

From Lemma 3.4, $\ell_{\Lambda_1} \leq \log |\Lambda| + 1$. Thus we have that

$$\begin{aligned} L &= q_{\Lambda_1}(1 + L_{\Lambda_1}) + (1 - q_{\Lambda_1})(1 + L_{\Lambda_2}) \\ &\leq q_{\Lambda_1}(1 + \log |\Lambda_1|) + (1 - q_{\Lambda_1})(1 + \log |\Lambda_2|) \\ &= \ell_{\Lambda_1} \leq \log |\Lambda| + 1 = \log(k + 1) + 1. \end{aligned}$$

Case(2): If $\Lambda = \{1, \dots, n\}$ is not a special node and the subtree T_{Λ_1} has no special node, from Lemma 3.5, we have $L_{\Lambda_1} \leq \log |\Lambda_1|$. Since $|\Lambda_2| \leq k$, from the induction hypothesis, we have $L_{\Lambda_2} \leq \log |\Lambda_2| + 1$.

From Lemma 3.3, $\ell_{\Lambda_1} \leq \log |\Lambda|$. Thus we have that

$$\begin{aligned} L &= q_{\Lambda_1}(1 + L_{\Lambda_1}) + (1 - q_{\Lambda_1})(1 + L_{\Lambda_2}) \\ &\leq q_{\Lambda_1}(1 + \log |\Lambda_1|) + (1 - q_{\Lambda_1})(1 + \log |\Lambda_2| + 1) \\ &= \ell_{\Lambda_1} + 1 - q_{\Lambda_1} \leq \log |\Lambda| + 1 = \log(k + 1) + 1. \end{aligned}$$

Case(3): If $\Lambda = \{1, \dots, n\}$ is not a special node and the subtree T_{Λ_2} has no special node, then the same computation as in Case (2) gives $L \leq \log(k + 1) + 1$. $\square$

3.3.2 Iterative step for finding another nonzero component

For every subset $\omega \subset \Omega$ with $|\omega| < s$ we let $P_\Lambda^\omega = P^\omega(\Lambda)$ denote the conditional probabilities

$$Pr\{X_i \neq 0, i \in \Lambda \text{ and } X_i = 0, i \in (\Lambda \cup \omega)^c | X_i \neq 0, i \in \omega\},$$

for any $\Lambda \subset \omega^c$.

Similar to Section 3.1, we let $q_\Lambda^\omega = q^\omega(\Lambda) = \Pr(E_\Lambda^\omega)$ for the events $E_\Lambda^\omega = \{X_i \neq 0, \text{ for some } i \in \Lambda | X_i \neq 0, i \in \omega\}$, which is the conditional probability that at least one of the components of X with index in $\Lambda \subset \omega^c$ is nonzero given that $X_i \neq 0$ for $i \in \omega$. Using the same procedure as in Section 3.1, build a Huffman tree with leaves $\{i\}$, $i \in \omega^c$ with probabilities q_Λ^ω . Note that this tree has $n - |\omega|$ leaves.

As above (see (4)), we assign a sampling vector $a_\Lambda^\omega \in \mathbb{R}^n$ to every node Λ which is not a leaf. Note that from the construction of a_Λ^ω we have that $a_\Lambda^\omega(i) = 0$ for $i \in \omega$.

Let $k < s$ be the number of nonzero components of x that are found and let $\omega = \{t_1, \dots, t_k\}$ be the set of corresponding indices. The algorithm for finding the $(k+1)$ -th nonzero component of x (if any) is essentially the same as in Algorithm 1.

3.4 Algorithm for finding all nonzero components of x

The general algorithm for finding the s -sparse vector $x \sim (X, P)$ which is an instance of the s -sparse random vector (X, P) , can now be described as follows:

Algorithm 2.

Initialization: $s=1$; $\omega = \emptyset$;

Repeat until $k > s$ *or* $x_{t_k} = 0$

1. $\Lambda = \omega^c$;
2. *repeat until* $|\Lambda| = 1$
 - if* $\langle a_\Lambda^{\omega^c}, x \rangle \neq 0$, *if* $\ell_{\Lambda_1} \leq \ell_{\Lambda_2}$, $\Lambda = \Lambda_1$; *else* $\Lambda = \Lambda_2$;
 - if* $\langle a_\Lambda^{\omega^c}, x \rangle = 0$, *if* $\ell_{\Lambda_1} \leq \ell_{\Lambda_2}$, $\Lambda = \Lambda_2$; *else* $\Lambda = \Lambda_1$;
 - end repeat*
3. *Output the (only) element* $t_k \in \Lambda$;
4. *Output* $x_{t_k} = \langle \chi_{\{t_k\}}, x \rangle$;
5. $\omega = \omega \cup \{t_k\}$;
6. $k = k + 1$;

end repeat

Algorithm 2 repeats Algorithm 1 at most s times and adds one extra sample to determine the value of each nonzero component once its position is known. Thus, as a corollary of Theorem 3.6 we immediately get

Corollary 3.7. *Given a nonzero s -sparse vector $x \sim (X, P)$ in $\mathbb{R}^n$, the average number of sampling vectors L needed to find all nonzero components of x using Algorithm 2 is at most $s \log n + 2s$.*

Remark 2.

- (i) *Corollary 3.7 states that the upper bound on the expected total cost (number of measurements and reconstruction combined) that we need for an s -sparse vector in $\mathbb{R}^n$ using Algorithm 2 is no more than $s \log n + 2s$.*
- (ii) *If the probability distribution P is uniform then the combined cost of the measurements and reconstruction is exactly $s \log n + 2s$.*

3.5 Noisy measurements

In practice, the measurements $\{y_i\}$ maybe corrupted by noise. Typically the noise is modeled as additive and uncorrelated : $y_\Lambda = \langle x, a_\Lambda \rangle + \eta_\Lambda$ (see [7]). For this case the condition $Az = y$ in (1) for the ℓ_1 minimization technique is modified to $\|Az - y\|_2 \leq \epsilon$ where ϵ is of the same order as the standard deviation σ_η of the noise. With this modification, the ℓ_1 minimization technique yields a minimizer x^* satisfying $\|x^* - x\|_2 \leq C\epsilon$ where C is a constant independent of x . Similar modifications are made for the other techniques, e.g., ℓ_q minimization (see [16]).

Similarly, our algorithm needs to be modified accordingly to deal with the noisy measurements case. Algorithm 2 can be modified by changing the statement $\langle a_\Lambda^{\omega^c}, x \rangle \neq 0$ to the statement $|\langle a_\Lambda^{\omega^c}, x \rangle| > T$, where the threshold T is of the same order as the standard deviation of η .

Consider the model $Y = X + \eta$ where the signal $X \sim N(0, \sigma_X)$ and the noise $\eta \sim N(0, \sigma_\eta)$. Then $Y \sim N(0, \sqrt{\sigma_X^2 + \sigma_\eta^2})$. We set the threshold in Algorithm 2 to be $T = 3\sigma_\eta$, and consider a measure of error (for one sample) given by the probability

$$p(e) = P(|Y| < T \text{ and } |X| \geq T) + P(|Y| \geq T \text{ and } |X| < T)$$

After easy computation, we have that

$$p(e) = \operatorname{erf}\left(\frac{3\sigma_\eta}{\sqrt{2}\sigma_X}\right) + \operatorname{erf}\left(\frac{3\sigma_\eta}{\sqrt{2}\sigma_Y}\right) - \operatorname{erf}\left(\frac{3\sigma_\eta}{\sqrt{2}\sigma_X}\right)\operatorname{erf}\left(\frac{3\sigma_\eta}{\sqrt{2}\sigma_Y}\right),$$

where $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$ is the error function.

Using the Taylor series, we obtain

$$p(e) = \frac{12t}{\sqrt{2\pi}} + o(t),$$

where $t = \frac{\sigma_n}{\sigma_x}$ is the ratio of the standard deviations of the noise and the signal. Thus, for a relatively large signal to noise ratio, t will be small and we get that the probability of at least one error in the sampling-reconstruction for an s sparse vector is bounded above by the quantity

$$p_s(e) = 1 - (1 - p(e))^{s(\log n + 1)} = s(\log n + 1)p(e) + o(p(e)^2) \approx s(\log n + 1)\frac{12t}{\sqrt{2\pi}}.$$

Suppose that $\hat{x}$ is the reconstruction of x from Algorithm 2. From the algorithm, the maximum error measured by $\|\hat{x} - x\|_2$ is bounded by $\|x\|_2$. Thus, the expectation of the relative ℓ_2 error is bounded above by

$$E\left(\frac{\|\hat{x} - x\|_2}{\|x\|_2}\right) \leq p_s(e)\frac{\|x\|_2}{\|x\|_2} + (1 - p_s(e))0 = p_s(e).$$

It can be seen that the relative ℓ_2 error is essentially linear in the sparsity s , linear in $t = \frac{\sigma_n}{\sigma_x}$ and logarithmic in the dimension n , as can also be seen in the simulations below.

Note that the last sampling in the algorithm can be repeated several times. The values of the i^{th} component thus found can then be averaged reducing the effect of the noise on $\hat{x}$ while increasing the cost of the algorithm by a linear term proportional to s but independent of n .

3.6 Stability and compressible signals

One of the advantage of the standard compressed sensing methods is that they produce almost optimal results for signal that are not s -sparse. For example, if we let $\beta_s^1(x)$ denote the smallest possible error (in the ℓ^1 norm) that can be achieved by approximating a signal $x \in \mathbb{R}^n$ by an s -sparse vector z :

$$\beta_s(x) := \inf\{\|x - z\|_1, \|z\|_0 \leq s\},$$

then the vector x^* solution to the ℓ_1 reconstruction method (1) is quasi-optimal in the sense that $\|x - x^*\|_1 \leq C\beta_s(x)$ for some constant C independent of x . Since for a given $x \in \mathbb{R}^n$, the quantity $\beta_s^1(x)$ is the ℓ_1 norm of the smallest $n - s$ components of x , the previous result means that if x is not s -sparse, then x^* is close to the s -sparse vector x_s whose components are the s -largest components of x . In particular, if x is sparse, then $x^* = x$.

Clearly, our current approach cannot produce similar results since, in its current form, our method does not have any incentive for finding the largest nonzero components. However a possible modification of the method is to replace $\langle a_\Lambda^{\omega^c}, x \rangle \neq 0$ by $|\langle a_\Lambda^{\omega^c}, x \rangle| > \Delta_k(\Lambda)$ in Algorithm 2 where the threshold $\Delta_k(\Lambda)$ changes at each sampling vector and its value depends on the characteristics of the random variables X_i , and Λ . For example, when the probability distribution P is uniform then the Huffman tree for the signal is a binary tree. Suppose a

signal x has s “significant” components and $n - s$ “nonsignificant” components in the following sense: The magnitude for each “significant” component is $N(0, \sigma_X)$ and for each “nonsignificant” component is $N(0, \sigma_\eta)$ where $\frac{\sigma_\eta}{\sigma_X}$ is small. At the i th step, the sampling vector will have $\frac{n}{2^i}$ components with 1. The sum of the corresponding “nonsignificant” components $\eta_i \sim N(0, \sigma_{\eta_i})$ and $\sigma_{\eta_i} \approx \sqrt{\frac{n}{2^i}} \sigma_\eta$ and we will have the measurement $Y \sim N(0, \sigma_Y)$ where $\sigma_Y \geq \sqrt{\sigma_X^2 + \frac{n}{2^i} \sigma_\eta^2}$. We set the threshold in Algorithm 2 to be $T_i = 3\sqrt{\frac{n}{2^i}} \sigma_\eta$, and consider a measure of error (for one sample) given by the probability

$$p(e)_i \leq P(|Y| < T_i \text{ and } |X| \geq T_i) + P(|\eta_i| \geq T_i).$$

Note that $P(|\eta_i| \geq T_i) \leq 10^{-4}$ which can be ignored. From similar computations as above, we have that

$$p(e)_i \leq \frac{6t_i}{\sqrt{2\pi}} + o(t_i),$$

where $t_i = \frac{\sigma_\eta \sqrt{\frac{n}{2^i}}}{\sigma_X}$. Suppose that the ratio of the standard deviations of the “nonsignificant” components and the “significant” components $t = \frac{\sigma_\eta}{\sigma_X}$ is small. Since there are in total $\log n$ steps to find one nonzero component, the probability of at least one error in the sampling-reconstruction is

$$p(e) = 1 - \prod_{i=1}^{\log n} (1 - p(e)_i) \approx \frac{6t}{\sqrt{2\pi}(\sqrt{2} - 1)} \sqrt{n} \leq 6t\sqrt{n}.$$

Then we get that the probability of at least one error in the sampling-reconstruction for an s sparse vector is bounded above by the quantity

$$p_s(e) \approx 6st\sqrt{n}.$$

The expectation of the relative ℓ_2 error

$$E\left(\frac{\|\hat{x} - x\|_2}{\|x\|_2}\right) \leq p_s(e) \frac{\|x\|_2}{\|x\|_2} + (1 - p_s(e))0 = p_s(e).$$

It can be seen that the relative ℓ_2 error is essentially linear in the sparsity s , linear in t and linear in the square root of the dimension n , as can also be seen in the simulations below.

4 Examples and Simulations

In this section, we provide some examples and test our algorithm on synthetic data. All the experiments are done using Matlab 7.4 on a Macintosh MacBook Pro 2.16 GHz Intel Core Duo processor 1GB 667 MHz RAM.

4.1 Noiseless cases

In the first experiment we use an exponential distribution to generate the position of the nonzero components of the s -sparse vectors, and uniform distribution for their values. The signal x is generated by first generating an integer index $i \in [1, n]$ using the exponential pdf with a mean of 10, and then constructing the component $x(i) = A(rand - 0.5)$, where $rand$ is a random variable with uniform distribution in $[0, 1]$. In the second experiment, we use a uniform probability distribution for the signals. The signal x is generated by first generating an integer index $i \in [1, n]$ with a uniform distribution, and then constructing the component $x(i) = A(rand - 0.5)$. This process is repeated s times for s sparse signals.

Simulation 4.1. *Our first experiment is a sparse vector $x \sim (X, P)$ in a space of dimension $n = 2^{15}$ with an exponential pdf with mean 10 for the location of the nonzero components and a uniform distribution for the values of the components as described above. We have tested our algorithm with $s = 1, 3, 5, 7, 9, 11, 13$. The mean and variance of the number of sampling vectors needed for the various sparsity s (for the combined sampling and reconstruction) is shown in Table 1.*

s	1	3	5	7	9	11	13
$s \log n + s$	16	48	80	112	144	176	208
Mean	9.11	27.5	46.17	61.47	81.26	97.2	111.88
Var	10.15	16.6	25.5	25.88	31.88	30.94	37.37

Table 1: Mean and variance of number of sampling vectors needed to find x as functions of its sparsity s for $n = 2^{15} = 32768$.

Simulation 4.2. *Our second example is a sparse random vector X in a space of dimension $n = 1024$ with a uniform probability distribution. We have tested our algorithm on an example with $n = 1024$, $s = 1, 25, 50, 75, 100, 125, 150$. The time for finding the vector x for the various sparsity s is shown in Table 2. It is clear from the table that the methods is very performant.*

s	1	25	50	75	100	125	150
CPU time	0.0045	0.028	0.049	0.073	0.098	0.144	0.146

Table 2: CPU time as a function of sparsity s for $n = 1024$.

Simulation 4.3. *In this experiment, we use multi-modal distribution to generate the positions of the nonzero components of an s sparse vector in a space of dimension $n = 768 = 3 * 2^8$. We assume that the location of each nonzero component has probability A to be in $[1, 256]$ uniformly, probability B to be in $[257, 512]$ uniformly and probability C to be in $[513, 768]$ uniformly. We have tested our algorithm with $s = 1, 3, 5, 7, 9, 11, 13$ for $A = 40\%, B = 20\%, C = 40\%$ and $A = 20\%, B = 10\%, C = 70\%$. We construct a suboptimal Huffman tree, that is, we consider each nonzero component independently. We have repeated this experiment 100 times. The mean of the number of sampling vectors needed for the various sparsity s (for the combined sampling and reconstruction) is shown in Table 3.*

s	1	3	5	7	9	11	13
$s \log n + s$	10.6	31.8	53	74.2	95.4	116.6	137.8
Mean ($A=40\%, B=20\%, C=40\%$)	9.5	28.6	47.8	67.0	86.3	105.3	124.5
Mean ($A=20\%, B=10\%, C=70\%$)	9.3	27.9	46.3	65.2	83.6	102.2	120.8

Table 3: Mean number of sampling vectors needed to find x of different sparsity s for $n = 768 = 3 * 2^8$

Simulation 4.4. *In this experiment, we use another multi-modal distribution to generate the positions of the nonzero components of an s sparse vector in a space of dimension $n = 1024 = 2^{10}$. We assume that the location of each nonzero component has probability A to be in $[1, 256]$ uniformly, probability B to be in $[257, 768]$ uniformly and probability C to be in $[769, 1024]$ uniformly. We have tested our algorithm with $s = 1, 3, 5, 7, 9, 11, 13$ for $A = 40\%, B = 20\%, C = 40\%$ and $A = 20\%, B = 10\%, C = 70\%$. We construct a suboptimal Huffman tree, that is, we consider each nonzero component independently. We have repeated this experiment 100 times. The mean of the number of sampling vectors needed for the various sparsity s (for the combined sampling and reconstruction) is shown in Table 4.*

s	1	3	5	7	9	11	13
$s \log n + s$	11	33	55	77	99	121	143
Mean ($A=40\%, B=20\%, C=40\%$)	9.8	29.8	49.3	69.4	89.4	108.3	128.8
Mean ($A=20\%, B=10\%, C=70\%$)	9.4	28.1	47.1	65.7	84.4	103.1	122.3

Table 4: Mean number of sampling vectors needed to find x of different sparsity s for $n = 1024 = 2^{10}$

Figure 2: Relative ℓ^2 error of reconstruction from noisy measurements.

4.2 Noisy measurements

In all the experiments in this subsection, we use an exponential probability distribution for both signal and noise. The signal x is generated by first generating an integer index $i \in [1, n]$ with a uniform distribution, and then constructing the component $x(i)$ by an exponential distribution with mean 0 and standard deviation A . This process is repeated s times for s sparse signals. Each additive noise η_Λ , is generated by an exponential distribution with mean 0 and standard deviation N and added to the measurement y_Λ .

Simulation 4.5. *In this test we fix the following values: $n = 512$, $A = 20$, $N = 0.1$. For each value of s , we construct 100 s -sparse signals in $\mathbb{R}^n$. We test the effect of s on the ℓ^2 relative error $\frac{\|\hat{x} - x\|_2}{\|x\|_2}$ (in percent) as a function of the sparsity s . The results are displayed in Figure 2. The experiments shows that the relative error increases linearly with s .*

Simulation 4.6. *In this test we fix the following values: $n = 512$, $A = 20$, $s = 16$. We test the effect of noise on the ℓ^2 relative error $\frac{\|\hat{x} - x\|_2}{\|x\|_2}$ (in percent) as a function of the value N of the noise. The results are displayed in Figure 3. The experiments suggest that the the relative error increases linearly with N .*

Simulation 4.7. *In this test we fix the following values: $A = 20$, $N = 0.1$, $s = 8$. We test the effect of $n = 2^r$ on the ℓ^2 relative error $\frac{\|\hat{x} - x\|_2}{\|x\|_2}$ (in percent) as a function of the value r . The results are displayed in Figure 4.*

4.3 Compressible signals

In all the experiments in this subsection, we first use a uniform distribution to generate the position for one of the “significant” components for the signal x .

Figure 3: Relative ℓ^2 error of reconstruction from noisy measurements.

Figure 4: Relative ℓ^2 error of reconstruction from noisy measurements.

Figure 5: Relative ℓ^2 error of reconstruction for compressive signals.

The component $x(i)$ is constructed by an exponential distribution with mean 0 and standard deviation A . This process is repeated s times. Then each “nonsignificant” component is generated by an exponential distribution with mean 0 and standard deviation N .

Simulation 4.8. *In this test we fix the following values: $n = 512$, $A = 20$, $N = 0.1$. For each value of s , we construct 100 s -sparse compressive signals in $\mathbb{R}^n$. We test the effect of s on the ℓ^2 relative error $\frac{\|\hat{x}-x\|_2}{\|x\|_2}$ (in percent) as a function of the sparsity s . The results are displayed in Figure 5. The experiments shows that the relative error increases linearly with s .*

Simulation 4.9. *In this test we fix the following values: $n = 512$, $A = 20$, $s = 10$. We test the effect of standard deviation of the “compressive” signals on the ℓ^2 relative error $\frac{\|\hat{x}-x\|_2}{\|x\|_2}$ (in percent) as a function of the value N . The results are displayed in Figure 6. The experiments suggest that the the relative error increases linearly with N .*

Simulation 4.10. *In this test we fix the following values: $A = 20$, $N = 0.1$, $s = 8$. We test the effect of the dimension n on the ℓ^2 relative error $\frac{\|\hat{x}-x\|_2}{\|x\|_2}$ (in percent) as a function of the value $\sqrt{n}$. The results are displayed in Figure 7. The experiments suggest that the relative error increases linearly with $\sqrt{n}$.*

5 Conclusion

We have presented an information theoretic approach to compressed sampling of sparse signals. Using ideas similar to those in Huffman coding, we constructed an adaptive sampling scheme for sparse signals. In our scheme, the sampling

Figure 6: Relative ℓ^2 error of reconstruction for compressive signals.

Figure 7: Relative ℓ^2 error of reconstruction for compressive signals.

vectors are binary and their construction are deterministic and can be produced explicitly for each n . Without noise the reconstruction is exact, and the average cost for sampling and reconstruction combined for an s sparse vector is bounded by $s\log(n) + 2s$. We have also shown that the method is stable in noisy measurements and robust to compressible signals. We hope that the approach will stimulate further developments and interactions between the area of information theory and compressed sampling.

ACKNOWLEDGEMENT

The research of Akram Aldroubi is supported in part by NSF Grant DMS-0807464.

References

- [1] M. Aharon, M. Elad and A.M. Bruckstein, The k-svd: An algorithm for designing of overcomplete dictionaries for sparse representation, *IEEE Trans. Signal Processing*, **54**(11), 4311–4322, 2006.
- [2] M. Aharon, M. Elad and A.M. Bruckstein, On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them, *Linear Algebra Appl.*, **416**(1), 48–67, 2006.
- [3] R. Baraniuk, M. Davenport, R.A. DeVore and M. Wakin, A simple proof of the restricted isometry property for random matrices, *Constr. Approx.*, **28**(3), 253–263, 2008.
- [4] R. Berinde, A. C. Gilbert, P. Indyk, H. Karloff and M. J. Strauss, Combining geometry and combinatorics: A unified approach to sparse signal recovery, In *46th Annual Allerton Conference on Communication, Control, and Computing*, 798–805, 2009.
- [5] E. Candès and J. Romberg, Quantitative robust uncertainty principles and optimally sparse decompositions, *Foundations of Comput. Math.*, **6**, 227–254, 2006.
- [6] E. Candès, J. Romberg and T. Tao, Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information, *IEEE Trans. Information Theory*, **52**, 489–509, 2006..
- [7] E. Candès and M. Wakin, People hearing without listening: An introduction to compressive sampling, *IEEE Signal Processing Magazine*, **25** 21–30, 2008.
- [8] H. S. Chang, Y. Weiss, and W. T. Freeman, *Informative sensing*. arXiv:0901.4275 [cs:IT], Jan. 2009.

- [9] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley Interscience Publishing, New York, 1991.
- [10] E. Candès and T. Tao, Near optimal signal recovery from random projections: Universal encoding strategies, *IEEE Trans. Information Theory*, **52**, 5406–5425, 2006.
- [11] R.A. DeVore, Deterministic constructions of compressed sensing matrices, *Journal of Complexity* **23**, 918–925, 2007.
- [12] D.L. Donoho, Compressed sensing, *IEEE Trans. Information Theory*, **52**, 1289–1306, 2006.
- [13] P.L. Dragotti, M. Vetterli, and T. Blu, Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets strang-fix, *IEEE Trans. Signal Processing*, **55**, 1741–1757, 2007.
- [14] M. F. Duarte, M. B. Wakin, and R. G. Baraniuk. Fast reconstruction of piecewise smooth signals from random projections. In *Proc. SPARS05*, Rennes, France, Nov. 2005.
- [15] Y. C. Eldar, P. Kuppinger and H. Bolcskei, Compressed sensing of block-sparse signals: Uncertainty relations and efficient recovery, *IEEE Trans. Signal Processing*, **58**, 3042–3054, 2010.
- [16] S. Foucart and M.-J. Lai, Sparsest solutions of underdetermined linear systems via ℓ_q -minimization for $0 < q \leq 1$. *Applied and Computational Harmonic Analysis*, 26/3, 395–407, 2009.
- [17] A. C. Gilbert, Y. Kotidis, S. Muthukrishnan and M. Strauss, One-Pass Wavelet Decompositions of Data Streams *IEEE Trans. Knowl. Data Eng.*, **15** (3) (2003), 541–554.
- [18] A. C. Gilbert, M. J. Strauss, J. A. Tropp and R. Vershynin, Algorithmic linear dimension reduction in the ℓ_1 norm for sparse vectors, *Proc. 44th Allerton Conf. Communication, Control, and Computing*, Allerton, IL, Sep. 2006.
- [19] A. C. Gilbert, M. J. Strauss, J. A. Tropp and R. Vershynin, One sketch for all: fast algorithms for compressed sensing. In *ACM STOC 2007*, 237–246, 2007.
- [20] R. Gribonval and M. Nielsen, Sparse decompositions in unions of bases, *IEEE Trans. Information Theory*, **49**, 3320–3325, 2003.
- [21] J. Haupt, R. Castro and R. Nowak, Adaptive sensing for sparse signal recovery, *Proc. IEEE Digital Signal Processing Workshop and Workshop on Signal Processing Education*, Marco Island, FL, 2009, January 2009.

- [22] S. H. Ji, Y. Xue and L. Carin, Bayesian compressive sensing, *IEEE Trans. Signal Processing*, **56**(6), 2346 – 2356, 2008.
- [23] Y. Lu and M. N. Do, A theory for sampling signals from a union of subspaces, *IEEE Trans. Signal Processing*, **56**, 2334–2345, 2008.
- [24] I. Maravic and M. Vetterli, Sampling and reconstruction of signals with finite rate of innovation in the presence of noise, *IEEE Trans. Signal Processing*, **53**, 2788–2805, 2005.
- [25] D. Needell and J. A. Tropp, CoSaMP: Iterative signal recovery from incomplete and inaccurate samples, *Applied and Computational Harmonic Analysis*, **26**(3), 301–321, 2008.
- [26] D. Needell and R. Vershynin, Uniform Uncertainty Principle and signal recovery via Regularized Orthogonal Matching Pursuit, *Foundations of Computational Mathematics*, **9**(3), 317–334, 2009.
- [27] H. Rauhut, K. Schass and P. Vanderghelynst, Compressed sensing and redundant dictionaries, *IEEE Trans. Inform. Theory*, **54**, 2210–2219, 2008.
- [28] S. Sarvotham, D. Baron and R. G. Baraniuk, *Compressed sensing reconstruction via belief propagation*, Technical Report ECE-0601, Electrical and Computer Engineering Department, Rice University, 2006.
- [29] S. Sarvotham, D. Baron, and R. G. Baraniuk. Sudocodes - fast measurement and reconstruction of sparse signals, *IEEE International Symposium on Information Theory*, 2006.
- [30] J. A. Tropp, Greed is good: Algorithmic results for sparse approximation, *IEEE Trans. Informtion Theory*, **50**, 2231–2242, 2004.
- [31] W. Xu and B. Hassibi, Efficient compressive sensing with deterministic guarantees using expander graphs, *IEEE Information Theory Workshop*, 2007.