INTERNATIONAL JOURNAL OF PURE SCIENCE RESEARCH STUDIES (IJPSRS)
(Open Access, Double-Blind Peer Reviewed Journal)
ISSN Online:
EDUSCHOOL ACADEMIC RESEARCH & PUBLISHERS

Compressed Sensing: Sparse Recovery Theory, Algorithms and Practical Applications

Sandhya E

Associate Professor, Department of Basic Sciences & Humanities, Jyothi Engineering College, Thrissur, India

Article Information
Received: 23rd April 2026
Received in revised form: 25th May 2026
Accepted: 27th June 2026
Available online: 14th August 2026
Volume: 2
Issue: 3
DOI: https://doi.org/10.5281/zenodo.21915660

Abstract

Classical sampling theory fixes the rate at which a signal must be measured by its bandwidth alone. Compressed sensing replaces that criterion with one based on structure: a signal that is sparse in some known basis can be reconstructed exactly from a number of linear measurements proportional to its sparsity and only logarithmic in its ambient dimension. This paper reviews the conditions under which such recovery is possible, the algorithms that achieve it, and the settings in which the theory has been applied. Uniqueness requires that the measurement matrix act almost isometrically on sparse vectors, a property that random matrices possess with high probability but that is difficult to certify for any given matrix. Under that property the combinatorial search for the sparsest solution can be replaced by minimisation of the sum of absolute values, a convex program, and the reconstruction is stable when measurements are noisy and the signal is only approximately sparse. Two numerical experiments carried out for this review illustrate the sharp boundary between success and failure as sparsity and the number of measurements are varied. Applications in magnetic resonance imaging and radio astronomy are examined together with the limits of the approach.

Keywords: Compressed Sensing, Sparse Recovery, Restricted Isometry, Basis Pursuit, Orthogonal Matching pursuit.

.

1. INTRODUCTION

The sampling theorem associated with Nyquist and Shannon states that a signal with no frequency content above a given limit is determined by samples taken at twice that limit. The result is exact and, read as a requirement rather than a sufficient condition, misleading. It assumes nothing about the signal except its bandwidth, whereas signals of practical interest are far from arbitrary: photographs compress well in a wavelet basis, and audio compresses well in a local cosine basis, precisely because most of their transform coefficients are negligible.

Conventional practice acquires everything and then discards most of it during compression. Compressed sensing asks whether the discarded measurements need to have been taken at all. Two papers published in 2006 showed that they do not.[1,2] If a signal has at most s non-zero coefficients in a known basis, then it can be recovered exactly from a number of suitably chosen linear measurements proportional to s times the logarithm of the ratio of dimension to sparsity, which is far fewer than the dimension when the signal is genuinely sparse.

The result rests on two ingredients that are worth separating. The first is a geometric condition on the measurement operator ensuring that distinct sparse signals produce distinguishable measurements. The second is an algorithmic fact: the natural but intractable problem of finding the sparsest consistent solution can, under that condition, be replaced by a convex program that is solvable in polynomial time. The sections that follow treat each in turn, then examine algorithms, numerical behaviour and applications.

2. CONDITIONS FOR RECOVERY

Write the measurement process as a matrix with m rows and n columns acting on an unknown vector, with m much smaller than n. The system is underdetermined and has infinitely many solutions, so recovery is possible only if the solution set contains exactly one sufficiently sparse vector.

The elementary condition is stated in terms of the spark of the matrix, the smallest number of linearly dependent columns. Any vector with fewer than half that many non-zero entries is the unique sparsest solution. Computing the spark requires examining every subset of columns, so this criterion is of theoretical value only. A weaker but computable substitute is the mutual coherence, the largest inner product between distinct normalised columns; guarantees expressed through coherence are easy to check and pessimistic, since coherence cannot fall below a bound determined by the matrix dimensions.

The condition that produced the strongest results is the restricted isometry property introduced by Candès and Tao.[3] A matrix satisfies it at order s with constant delta if it changes the length of every s-sparse vector by a factor lying between one minus delta and one plus delta. In other words, the matrix acts almost like an isometry on every s-dimensional coordinate subspace at once. When the constant at order 2s is small enough, distinct s-sparse vectors cannot produce the same measurements, and the convex relaxation described below recovers the correct one.

Verifying the property for a specific matrix is computationally hard, which is the awkward point in the theory. What can be shown is that certain random constructions satisfy it with high probability. Matrices with independent Gaussian or subgaussian entries do so once the number of rows exceeds a constant multiple of s times the logarithm of the ratio of n to s , and matrices formed from randomly selected rows of a discrete Fourier transform do so with an additional logarithmic factor .[4,5] The Fourier case matters in practice, because many measurement systems can select frequencies but cannot apply an arbitrary dense matrix to a physical signal.

3. CONVEX RELAXATION AND ITS GUARANTEES

Seeking the vector with fewest non-zero entries subject to the measurement constraints is a combinatorial problem. Replacing the count of non-zero entries by the sum of absolute values yields a convex program, known as basis pursuit, that can be recast as a linear program .[6] The substitution is not a heuristic in this setting: under the restricted isometry condition the two problems have the same solution.

Geometry explains why the sum of absolute values promotes sparsity. Its unit ball is a cross-polytope whose extreme points lie on the coordinate axes, so expanding this ball until it first meets the affine solution set typically produces contact at a vertex, which is a sparse vector. The Euclidean ball has no such structure and yields a dense solution.

The guarantees extend beyond exactly sparse signals and noiseless measurements, which is essential for applications. If measurements are corrupted by noise of bounded energy and the signal is compressible rather than sparse, the solution of the corresponding constrained problem lies within a distance of the true signal controlled by the sum of the noise level and the error of the best s-term approximation .[7] The bound degrades gracefully, so nothing catastrophic happens when the assumptions hold only approximately, which is the usual situation. Comprehensive treatments of these results are available in book form .[8]

4. ALGORITHMS

Three families of solvers are in common use, and they trade accuracy against cost differently. Table 1 summarises them.

Convex programming approaches solve the relaxed problem directly. Interior point methods are accurate and scale poorly; first-order methods based on iterative shrinkage are the practical choice for large problems, and the accelerated variant converges at a rate proportional to the inverse square of the iteration count .[9]

Greedy methods build the support one element at a time. Orthogonal matching pursuit selects at each step the column most correlated with the current residual, projects the measurements onto the span of the selected columns and repeats .[10] It is simple and fast, and its guarantees are weaker than those of convex relaxation, though for many random ensembles the practical performance is comparable. Compressive sampling matching pursuit adds and prunes several elements per iteration and recovers guarantees of the same order as basis pursuit .[11]

Thresholding methods iterate a gradient step followed by retention of the largest entries. Iterative hard thresholding is the simplest of these and converges under a restricted isometry condition .[12] Approximate message passing adds a correction term derived from a statistical physics analysis, which removes the correlation between iterates and yields performance that matches the theoretical phase transition of convex relaxation at a fraction of the cost.[13]

Fig. 1 reports a numerical experiment carried out for this review, in which random Gaussian matrices were used to measure random sparse vectors of dimension 200 and orthogonal matching pursuit was applied to recover them, with twenty independent trials at each combination of sparsity and number of measurements. The transition from reliable to unreliable recovery is sharp and follows a boundary that is close to linear in the two parameters, matching the phase transition behaviour predicted for this class of problem.[14] Fig. 2 shows a single instance from the recoverable region: eighty measurements of a 256-dimensional vector with ten non-zero entries, recovered to a relative error at the level of machine precision.

The experiment also illustrates the practical reading of such diagrams. At a sparsity of ten, recovery was reliable once about sixty measurements were available, which is roughly six times the sparsity and less than a third of the ambient dimension. At a sparsity of thirty, reliability required close to one hundred and fifty measurements, so the ratio of measurements to non-zero entries falls as sparsity increases, in line with the logarithmic term in the theory. A design should sit well above the transition rather than on it, since the region where the success probability is intermediate is also the region where reconstruction is most sensitive to noise.

Fig 1: Empirical probability of exact recovery by orthogonal matching pursuit as a function of sparsity and number of measurements, computed for this review with twenty trials per parameter combination.

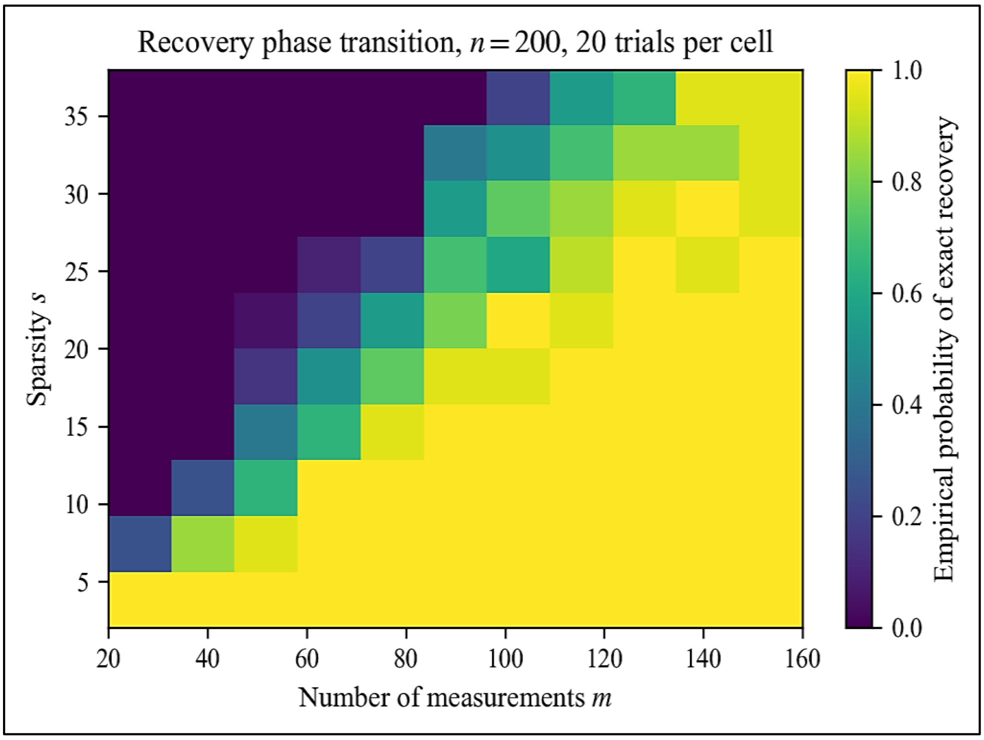

Table 1. Algorithms for sparse recovery

Algorithm	Type	Cost per iteration	Guarantee
Basis pursuit	Convex program	Depends on solver	Strongest, under restricted isometry [1,3]
Iterative shrinkage	First-order convex	One matrix product	Converges to the convex optimum [9]
Orthogonal matching pursuit	Greedy	One product and a least-squares fit	Weaker, competitive in practice [10]
Compressive sampling matching pursuit	Greedy	Several matrix products	Comparable order to basis pursuit [11]
Iterative hard thresholding	Thresholding	One matrix product	Requires small isometry constant [12]
Approximate message passing	Thresholding with correction	One matrix product	Attains the predicted transition [13]

Fig 2: A single recovery instance from the successful region of Fig. 1. Circles mark the true coefficients and crosses the values returned by orthogonal matching pursuit.

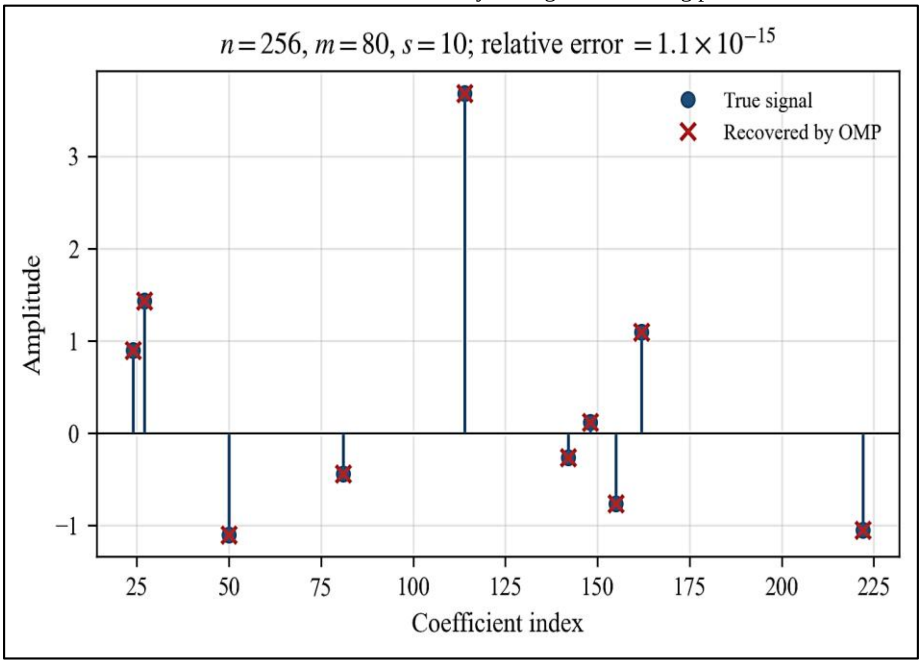

5. APPLICATIONS

Magnetic resonance imaging is the application where the idea has changed clinical practice. The scanner samples the Fourier transform of the image along trajectories in frequency space, and acquisition time is proportional to the number of samples taken. Medical images are compressible in wavelet bases, and randomised undersampling of frequency space produces artefacts that are incoherent and therefore removable by sparse reconstruction .[15] Paediatric imaging benefited first, since shorter scans reduce the need for sedation, and reconstruction of this type is now standard on commercial scanners.[16]

Optical imaging provided an early demonstration of a different kind. A single-pixel camera modulates the incoming image with a sequence of random patterns using a micromirror array and records only the total intensity for each pattern; the image is then reconstructed from far fewer measurements than pixels [17]. The device is of limited practical value at visible wavelengths, where detector arrays are cheap, but the principle matters at wavelengths where arrays are expensive.

Radio interferometry poses the same mathematical problem. An array samples the Fourier transform of the sky brightness at positions determined by the baselines between antennas, leaving large gaps. Reconstruction under sparsity assumptions is one of the methods used in the imaging of the black hole shadow reported by the Event Horizon Telescope, where several independent reconstruction pipelines were compared precisely because the inverse problem is underdetermined.[18]

The same circle of ideas extends to low-rank recovery. Replacing the count of non-zero entries by the rank of a matrix and the sum of absolute values by the sum of singular values gives conditions under which a matrix is recoverable from a small random sample of its entries, a result that underlies collaborative filtering and many problems in system identification.[19]

6. STRUCTURED SPARSITY AND THE DESIGN OF MEASUREMENTS

Plain sparsity ignores information that is usually available. The large wavelet coefficients of a natural image are not scattered at random; they cluster along edges and form connected trees across scales. Model-based compressive sensing restricts the search to signals whose support obeys such a structure, replacing the projection onto the set of sparse vectors by projection onto the structured model. The number of measurements needed then depends on the number of admissible supports rather than on all subsets of a given size, which removes the logarithmic factor and reduces the requirement in practice.[20] Similar arguments apply to block sparsity, where non-zero entries occur in groups, and to signals that are sparse in a redundant dictionary rather than an orthonormal basis.

The design of the measurement operator is equally consequential and is often outside the experimenter's control. Physical systems typically permit only structured randomness: partial Fourier sampling in magnetic resonance and radio imaging, random convolution in radar, and subsampled Hadamard patterns in optics. These ensembles satisfy the isometry condition with worse constants than fully random matrices, and the sampling pattern matters. Uniform random selection of Fourier frequencies performs poorly on images, whose energy is concentrated at low frequencies; variable density schemes that sample the centre of frequency space densely and the periphery sparsely perform much better, and this behaviour has been given a theoretical basis in terms of local coherence.[21] Practical sampling patterns in medical imaging follow this principle rather than uniform randomness.

A related point concerns the sparsifying transform. Guarantees are stated for a basis in which the signal is sparse, and the choice of that basis is a modelling decision that determines how many measurements are needed. Total variation, wavelets and learned dictionaries all serve, and the reconstruction quality obtained in an application usually depends more on this choice than on which of the algorithms in Table 1 is used to solve the resulting problem.

7. LIMITS AND COMMON MISUNDERSTANDINGS

Three points are frequently overstated. First, the theory assumes that measurements can be designed, and in many settings they cannot. If the physics dictates the measurement operator, and if that operator is coherent with the sparsity basis, the guarantees do not apply, however sparse the signal may be.

Second, sparsity is a modelling assumption rather than an observed fact. Natural signals are compressible, meaning their sorted coefficients decay, rather than exactly sparse, and the recovery error then contains a term proportional to the tail that no algorithm can remove. Reported reconstructions at very high undersampling factors often rely on additional structure, such as smoothness across coil channels or across time, and attributing the result to sparsity alone misstates what is doing the work.

Third, noise sets a floor. The stable recovery bounds are proportional to the noise level with a constant that grows as the isometry constant approaches its critical value, so a system operating near the recovery boundary amplifies noise substantially even though it does not fail outright.[7] Practical designs stay well inside the boundary shown in Fig. 1.

Reconstruction methods trained on data have displaced sparse optimisation in several imaging applications, since a learned prior represents the statistics of a class of images better than any fixed basis. What such methods retain from this theory is the sampling strategy: incoherent measurement remains the reason the reconstruction problem is solvable at all.

8. CONCLUSION

Compressed sensing established that the number of measurements needed to determine a signal depends on its information content rather than on its bandwidth. The mathematical statement is precise: under a restricted isometry condition satisfied by suitable random operators, minimising the sum of absolute values recovers any sufficiently sparse vector exactly, and recovers compressible vectors with an error controlled by the noise and by the approximation tail. The algorithmic landscape is mature, with convex, greedy and thresholding methods offering different balances of guarantee and cost, and the empirical phase transition is sharp enough to serve as a design rule. The practical influence has been greatest where measurement is expensive, in medical imaging and radio astronomy. The main open problems are the difficulty of certifying the isometry property for structured operators and the integration of learned priors with the guarantees that make the classical theory trustworthy.

REFERENCES

1. Candès EJ, Romberg J, Tao T. 2006. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. IEEE Trans Inf Theory. 52(2):489–509.
2. Donoho DL. 2006. Compressed sensing. IEEE Trans Inf Theory. 52(4):1289–306.
3. Candès EJ, Tao T. 2005. Decoding by linear programming. IEEE Trans Inf Theory. 51(12):4203–15.
4. Baraniuk R, Davenport M, DeVore R, Wakin M. 2008. A simple proof of the restricted isometry property for random matrices. Constr Approx. 28(3):253–63.
5. Rudelson M, Vershynin R. 2008. On sparse reconstruction from Fourier and Gaussian measurements. Commun Pure Appl Math. 61(8):1025–45.
6. Chen SS, Donoho DL, Saunders MA. 1998. Atomic decomposition by basis pursuit. SIAM J Sci Comput. 20(1):33–61.
7. Candès EJ, Romberg JK, Tao T. 2006. Stable signal recovery from incomplete and inaccurate measurements. Commun Pure Appl Math. 59(8):1207–23.
8. Foucart S, Rauhut H. 2013. A mathematical introduction to compressive sensing. New York (NY): Birkhäuser.
9. Beck A, Teboulle M. 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM J Imaging Sci. 2(1):183–202.
10. Tropp JA, Gilbert AC. 2007. Signal recovery from random measurements via orthogonal matching pursuit. IEEE Trans Inf Theory. 53(12):4655–66.
11. Needell D, Tropp JA. 2009. CoSaMP: iterative signal recovery from incomplete and inaccurate samples. Appl Comput Harmon Anal. 26(3):301–21.
12. Blumensath T, Davies ME. 2009. Iterative hard thresholding for compressed sensing. Appl Comput Harmon Anal. 27(3):265–74.
13. Donoho DL, Maleki A, Montanari A. 2009. Message-passing algorithms for compressed sensing. Proc Natl Acad Sci USA. 106(45):18914–9.

14. Donoho DL, Tanner J. 2010. Precise undersampling theorems. Proc IEEE. 98(6):913–24.
15. Lustig M, Donoho D, Pauly JM. 2007. Sparse MRI: the application of compressed sensing for rapid MR imaging. Magn Reson Med. 58(6):1182–95.
16. Vasanawala SS, Alley MT, Hargreaves BA, Barth RA, Pauly JM, Lustig M. 2010. Improved pediatric MR imaging with compressed sensing. Radiology. 256(2):607–16.
17. Duarte MF, Davenport MA, Takhar D, et al. 2008. Single-pixel imaging via compressive sampling. IEEE Signal Process Mag. 25(2):83–91.
18. Event Horizon Telescope Collaboration. 2019. First M87 Event Horizon Telescope results. IV. Imaging the central supermassive black hole. Astrophys J Lett. 875(1).
19. Candès EJ, Recht B. 2009. Exact matrix completion via convex optimization. Found Comput Math. 9(6):717–72.
20. Baraniuk RG, Cevher V, Duarte MF, Hegde C. 2010. Model-based compressive sensing. IEEE Trans Inf Theory. 56(4):1982–2001.
21. Krahmer F, Ward R. 2014. Stable and robust sampling strategies for compressive imaging. IEEE Trans Image Process. 23(2):612–22.