Sensor Distributions (27 Sensors)

1. Uniform Random

Based on the coherence identity shown in "The Central Connection", I hypothesized that the uniform random sensor distribution would provide the lowest mutual coherence and highest probabilities of OMP reconstruction success.

2. Normal Random

I chose to test this distribution in order to analyze the importance of the uniformity of a random distribution. In the aperture of size 1, the sensors had a standard deviation of about 0.16 in both the x- and y-directions from the mean point, (0.5, 0.5). Note the severe "clumping" of samples produced by this distribution.

3. VLA

The VLA has 4 different sensor configurations (A, B, C, and D), but they are essentially scaled versions of the same "Y" structure, whose arms are 120° apart. To model how the distances between adjacent sensors should change as I added more sensors, I interpolated from the best-fit polynomial that approximated the existing distances provided in [5].

4. VLA Randomized Along Arms (VLA RAA)

I used these last two sensor distributions to see whether and how introducing an element of randomness would affect the existing VLA configuration. In this distribution, sensors were randomly scattered along each of the arms of the Y. Somewhat similar to the normal random distribution, samples produced by VLA RAA tend to group together.

5. VLA Perpendicularly Perturbed 30% (VLA PP)

Sensors in the VLA distribution were perturbed by a random distance in the direction perpendicular to their respective arms. The maximum amount of random perturbation was measured as a percentage of the average distance between adjacent VLA sensors. It was found that, by 30% perturbation, the distribution already performed similarly to the uniform random distribution (see "Additional Data" Figures 1 and 2). Note that, compared to the original VLA distribution, this modification is very slight.

6. VLA Perpendicularly Perturbed 3 Outermost 30% (VLA PPO)

In the second phase of experimentation, I aimed to minimize the modification of VLA PP 30%, while still retaining the improvements it provided. In VLA PPO, only the 3 outermost sensors on each arm were perpendicularly perturbed.

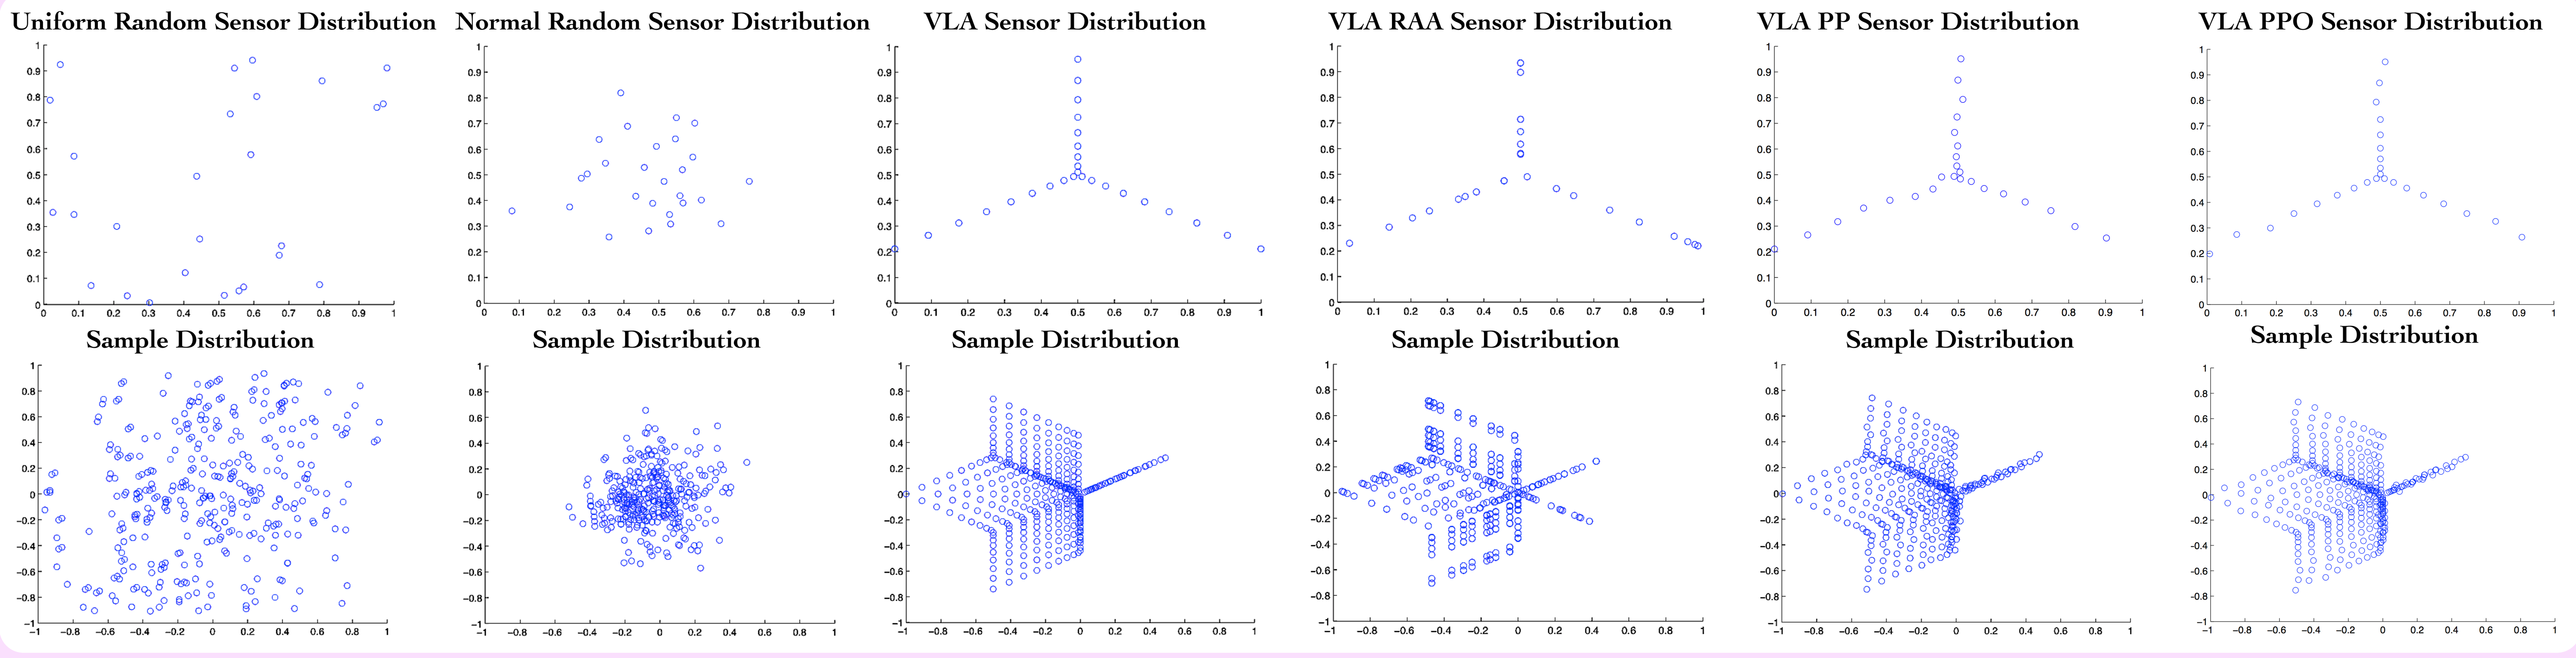

Experimental Methods

Mutual Coherence

As a theoretical indicator of OMP reconstruction success, I wrote MATLAB codes to find how the mutual coherence of sensing matrices was affected by different sample distributions. Low mutual coherence values would theoretically indicate high probabilities of OMP reconstruction success (see "Sensing Matrix Incoherence"). Because every sensor distribution except for the VLA had a degree of randomness, I took the average of the mutual coherence of 50 different sensing matrices for each distribution. Depending on how many sensors were used, the dimensions of the sensing matrices were between 36 x 3600 and 1275 x 3600 - all corresponding to severely underdetermined systems.

Probability of OMP Reconstruction Success

I wrote MATLAB codes for the OMP algorithm to simulate how sensor distributions affect the probability of image reconstruction success, as the sparsity of the signal increased from 1 to 300. For each distribution, I ran OMP reconstructions of 200 different objects for each sparsity level. A reconstruction was considered successful if the error relative to the original object was less than or equal to 1%. The number of successful trials for each sparsity was then divided by 2 to give the percent probability of OMP reconstruction success. To emphasize the efficiency of the compressed sensing approach, the number of samples in each reconstruction, 351, was **less than 10%** of the number of point sources of each signal, 3600. The reconstructions were run with 1% Gaussian-distributed random noise, to accommodate for error in the data.

The Well-Resolved Case: $A\ell/\lambda = 1$

Figure 13. **Mutual Coherence of Sensing Matrices by Number of Sensors**

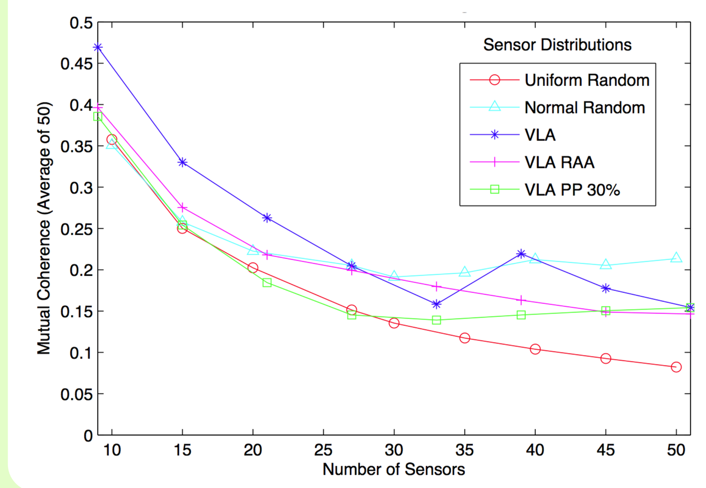

Figure 14. **Probability of OMP Reconstruction Success (1% Noise)**

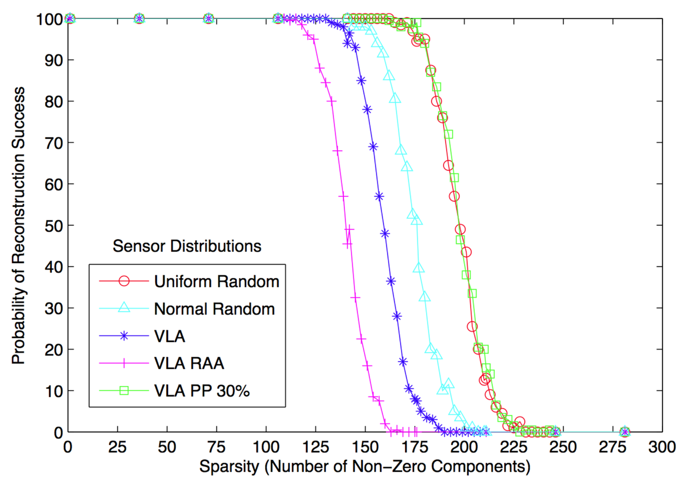

These results were simulated in the well-resolved case, where $A\ell/\lambda = 1$, providing the current resolution of the VLA. As predicted by the identity in "The Central Connection", in Figure 13 the mutual coherence given by the uniform random distribution continues to decrease as the number of sensors, n, increases. It also consistently provides the lowest mutual coherence of all the distributions, matched (for $n \leq 27$) only by VLA PP, one of the randomized modifications of VLA I designed. At 27 sensors, the number of sensors the VLA currently uses, VLA PP performs nearly identically to the uniform random distribution. This has very exciting implications - by only slightly perturbing the VLA distribution, we can significantly decrease its mutual coherence, which was among the worst of all the distributions. As predicted by the mutual coherence results, the reconstruction performances in Figure 14 of the uniform random distribution and VLA PP are nearly identical, and clearly superior to that of the VLA. Both distributions provide 100% reconstruction success for images of the greatest sparsity, about 200 non-zero components, whereas the VLA distribution can handle images of only about 150 non-zero components.

The other randomization modification of the VLA, VLA RAA, does not provide any noticeable improvement to the mutual coherence of the VLA at 27 sensors. Accordingly, in Figure 14 its reconstruction success is very poor - the worst of all the distributions, providing high probability of reconstruction success for images with a sparsity of only about 125. This may be because, like the normal random distribution (which also provided high mutual coherence), the samples generated by VLA RAA tend to clump together, instead of covering the aperture evenly (see "Sensor Distributions"). However, the mutual coherence provided by this distribution appears to continue decreasing as n increases, a favorable effect that the uniform random distribution is also proven to have. In comparison, the promising benefits of VLA PP seem to weaken as n increases. This is likely because the amount of perturbation (here, 30%) is measured as a percent of the average distance between sensors, and when n is large this average distance becomes very small. In light of the results in "Additional Data" Figures 1 and 2, increasing perturbation for larger numbers of sensors does not continue to better the performance.

The performance of the normal random distribution demonstrates the importance of uniformity in a random distribution. Its mutual coherence at 27 sensors is among the worst, and actually seems to increase slightly as n increases. Accordingly, its reconstruction results were considerably worse than the results of the uniform random distribution. Of course, as the standard deviation of the normal random distribution (about 0.16) increases and approaches that of the uniform random distribution (about 0.29), it is entirely plausible that its performance would improve and eventually mimic the uniform random distribution.

"Maybe we don't have to perturb *every* sensor…": VLA PPO

In the well-resolved case, the VLA PP distribution offered a distinct improvement to the VLA distribution - the results were nearly identical to those of the uniform random distribution.This inspired the next phase of experiments - can I reduce the degree of modification to the VLA even more? What is the minimum change needed to make the VLA distribution emulate a uniform random distribution?

Figure 17. **Mutual Coherence of Sensing Matrices (VLA PPI & VLA PPO)**

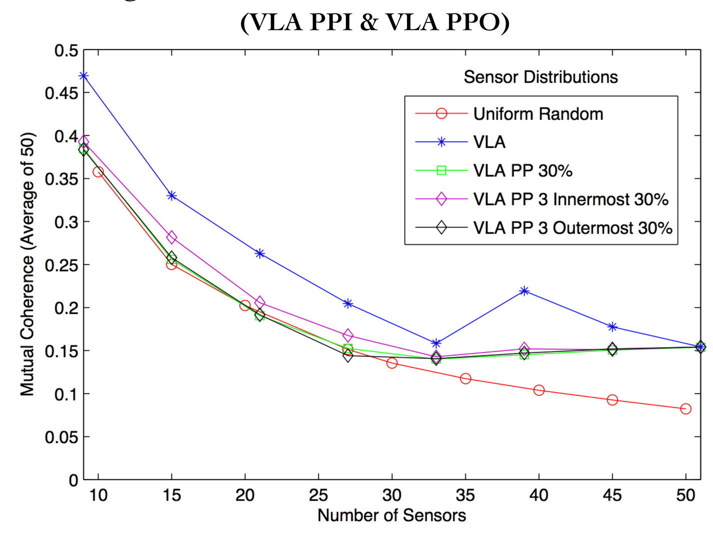

Figure 18. **Probability of OMP Reconstruction Success (VLA PPI & VLA PPO)**

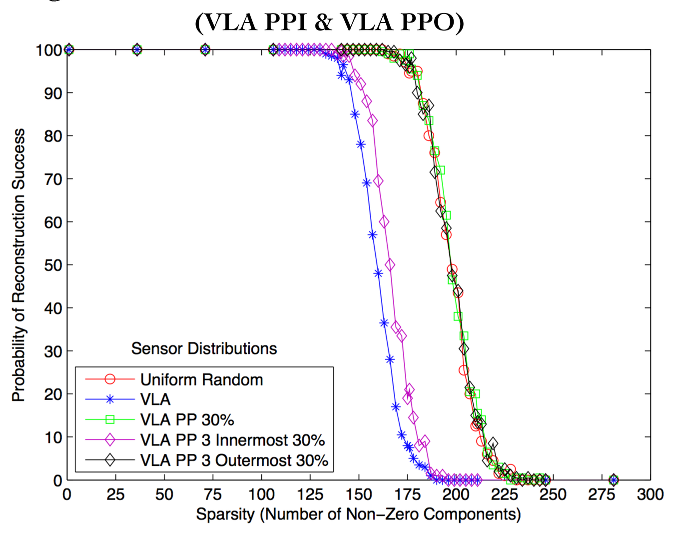

These results show the effects of perturbing only the 3 innermost sensors of each arm (VLA PPI), versus perturbing only the 3 outermost sensors of each arm (VLA PPO). Like VLA PP, when the number of sensors is not too large VLA PPO provides mutual coherence values that are nearly identical to those of the uniform random distribution. VLA PPI provides higher mutual coherence than either, though still considerably less than that of the VLA distribution. As in "Additional Data" Figures 1 and 2, it is apparent that less can be just as good as more - perturbing only the 3 outermost sensors on each arm provides the same benefits as perturbing all of the sensors. The results in Figure 18 correspond to the mutual coherence results, showing that VLA PPO is an extraordinarily simple modification to the VLA that allows it to emulate a uniform random distribution in reconstruction performance.

Uniform Randomness vs. Irregularity: Hammersley Points & GMRT

The capacities of the uniform random distribution in compressed sensing are clear. However, is the rigorous mathematical definition of uniform randomness essential - or will deterministic irregularity provide the same results? The **Hammersley point set** is a deterministically generated, low-discrepancy sampling pattern that is commonly used for **quasi-Monte-Carlo integration** to emulate aspects of a uniform random distribution. In two dimensions, it implements the **van der Corput** sequence, a function based on the dyadic expansions of numbers. Here I analyzed the comparability of this quasi-random distribution to a truly uniform random distribution.

The sensor distribution of the **Giant Metrewave Radio Telescope (GMRT)**, currently the world's largest radio interferometer, also includes an outer irregularized 'Y' array as well as a highly compact, central random array [6]. This set-up includes two key differences from a uniform random distribution - the non-random irregularity of the outer array, as well as the extreme normality of the central array.

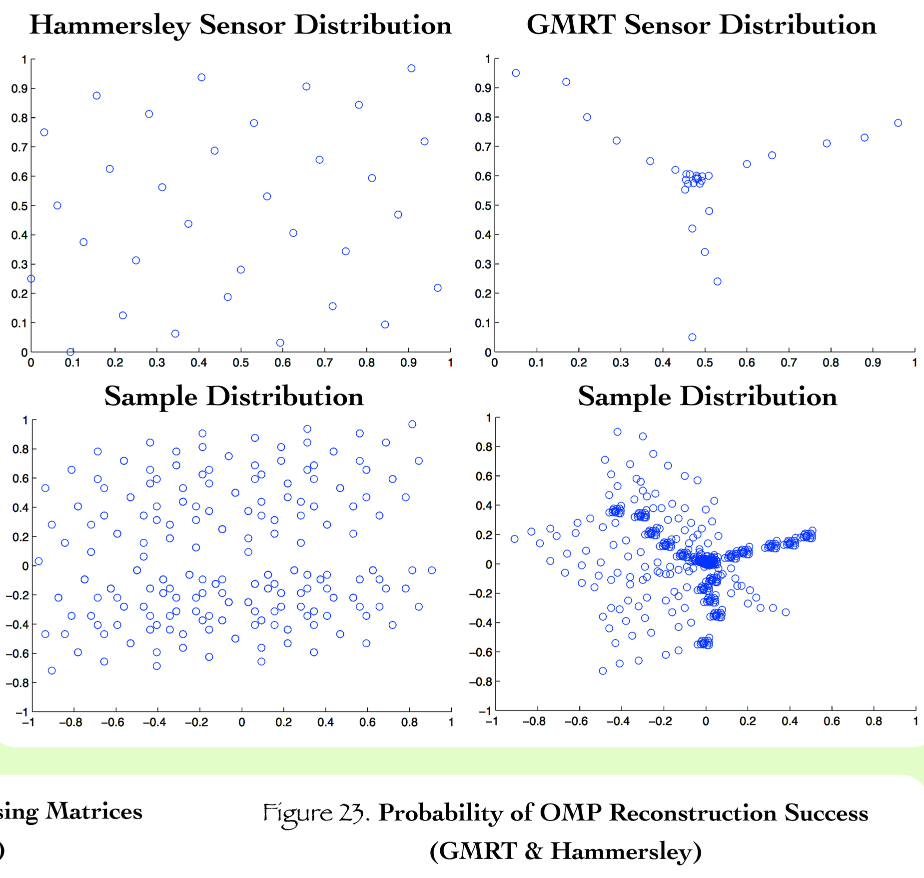

Figure 22. **Mutual Coherence of Sensing Matrices (GMRT & Hammersley)**

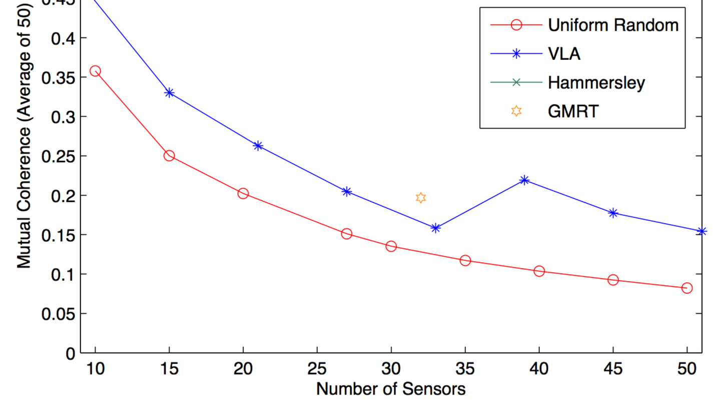

Figure 23. **Probability of OMP Reconstruction Success (GMRT & Hammersley)**

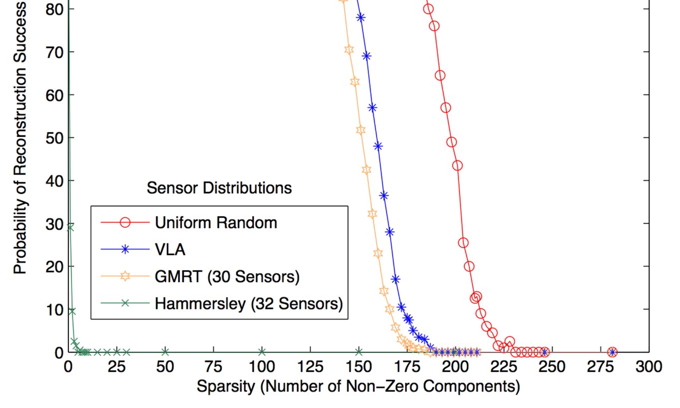

The mathematical specificity of uniform randomness is appears to be crucial in Figures 22 and 23. Though sensors in a Hammersley distribution have quasi-random qualities, the baseline distribution seems faintly formulaic, resulting in extraordinarily high coherence of the sensing matrix and the disastrous reconstruction results. (Thirty-two sensors were used because the formula provided in [7], which calculates an optimal shift of the point set, requires the number of points to be a power of two.) This points to the subtlety and inimitable qualities of rigorously defined uniform randomness.

Due to the outer array's non-random irregularity, I could not simulate increasing the number of sensors in the GMRT distribution. Regardless, the mutual coherence provided by the existing irregularized 30 sensors was not an improvement to the VLA, in contrast to the true randomization modification VLA PPO (see Figures 17 and 18). Along with the GMRT's subpar reconstruction performance, this is probably a result of both the outer array's non-random irregularity and the random central array's normality, rather than uniformity (compare the normal random distribution to the uniform random distribution in Figure 14).

The Under-Resolved Case: $A\ell/\lambda \leq 1$

As in many imaging processes, **super-resolution** - enhancing resolution beyond the diffraction limit - is a prize for physicists. In radio interferometry, super-resolution would allow us to reconstruct images with high resolution even when $A\ell/\lambda \leq 1$. In my experiments, I simulated this condition by assuming ℓ and λ to be constant, and decreasing the size of the aperture, A.

Figure 15. **Mutual Coherence of Sensing Matrices by Aperture Size**

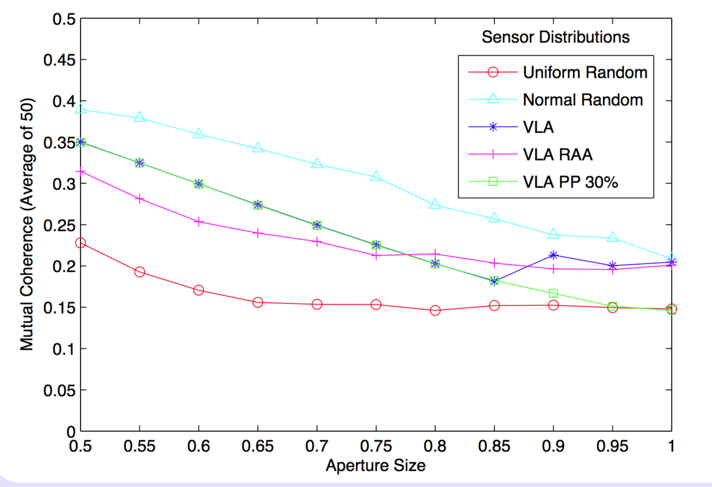

Figure 16. **Probability of OMP Reconstruction Success (0.5 Aperture)**

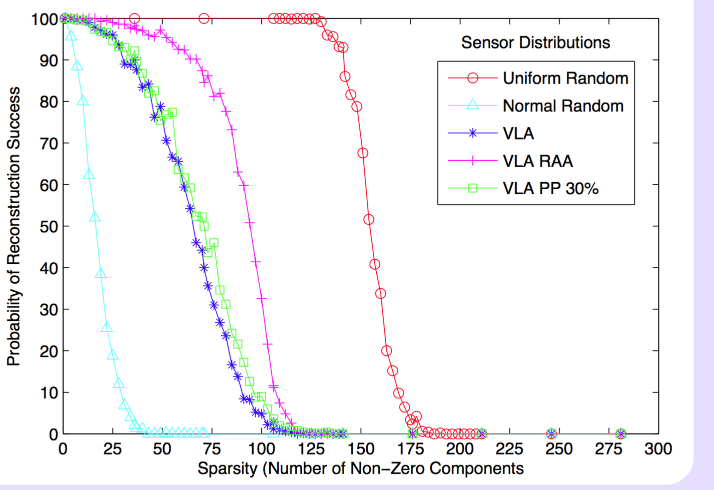

By about $A = 0.65$, the mutual coherence of the uniform random distribution appears to be as low as it is in the well-resolved case - no other distribution displays anything near this disposition toward super-resolution. In accordance with the mutual coherence results, Figure 16 clearly indicates the uniform random distribution's superiority in the under-resolved case, and again by an enormous degree - with $A = 0.5$, it provided 100% reconstruction success for images with a sparsity of about 125, whereas all of the other distributions began faltering in their performance beyond a sparsity of only about 25. In effect, implementing the compressed sensing approach with a uniform random distribution would allow interferometers not only to collect just a fraction of the samples for reconstruction, but also to **enhance image resolution beyond what the VLA is currently capable of.**

Unlike in the well-resolved case, VLA PP doesn't appear to provide noticeable improvement to VLA. The results in Figure 15 indicate near spot-on similarity in mutual coherence between the two until about $A = 0.9$. Generally, the results of VLA PP are probably lackluster because the amount of maximum perturbation is measured as a percent of the average distance between sensors, which becomes irrelevantly small at smaller apertures. However, following from the results in "Additional Data" Figures 1 and 2, simply increasing the percentage of perturbation is unlikely to improve the performance.

On the other hand, VLA RAA, which does not improve on the VLA results in the well-resolved case, provides noticeably lower mutual coherence than either VLA or VLA PP at relatively small apertures, and in Figure 16 gives the second-best reconstruction performance (though the uniform random distribution still out-performs it by a long run). The discrepancy of its results in the well-resolved and under-resolved cases points to the existence of qualities of a well-designed sensing matrix other than incoherence, which may provide an interesting arena for future research.

The normal random distribution's mutual coherence results correspond to those in the well-resolved case - they are among the highest of all the distributions. Accordingly, its reconstruction results in Figure 16 could be called a disaster, as the probability of reconstruction success rapidly deteriorates at a sparsity of just over 0. Again, these results highlight the importance of uniformity in a random sensor distribution.

Reconstructions of Galaxy 3C75

To expedite the process and provide general conclusions on sampling distribution, in the reconstruction probability experiments I used objects with randomly generated, randomly located non-zero components (see "Additional Data"). Here I simulated compressed sensing reconstructions of galaxy 3C75, which offered a more realistic context because the non-zero components are mostly adjacent to each other, a factor that can affect OMP's response. The number of samples used, 4851, was **only about 34%** the number of pixels, 14400. To emphasize the main results from the previous experiments, I reconstructed the image with the uniform random, VLA, and VLA PPO sensor distributions, in both well-resolved and under-resolved cases.

Figure 19. **Traditional Reconstruction**

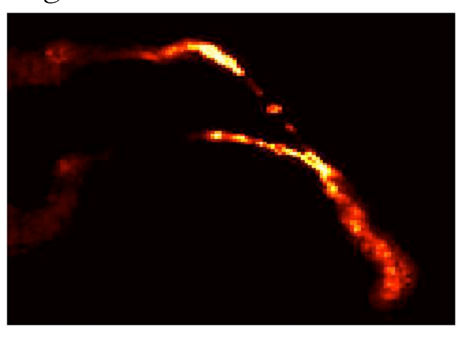

Figure 20. **Compressed Sensing Reconstructions (Well-Resolved Case)**

a. Uniform Random Distribution | **b. VLA Distribution** | **c. VLA PPO Distribution**

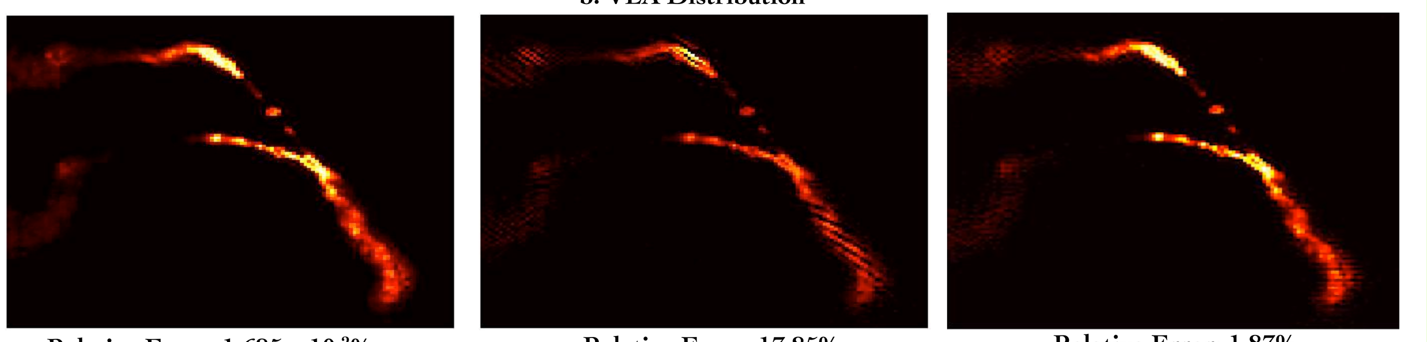

Relative Error: 1.685 x 10^{-2}% | **Relative Error: 17.85%** | **Relative Error: 1.87%**

The uniform random distribution gives an incredibly accurate reconstruction of the image (Figure 20.a), with a relative error of less than a tenth of a percent. With only very modest randomization modifications to the VLA, VLA PPO also provides a very accurate reconstruction (Figure 20.c). The VLA reconstruction (Figure 20.b), displays a distinct pattern of missing "streaks" of pixels, which is probably a direct consequence of relatively high mutual coherence. These streaks are mostly composed of adjacent point sources, which correspond to adjacent columns of the sensing matrix. Because adjacent point sources tend to be very similar, adjacent columns tend to be much more coherent than columns that are far apart. Therefore, in iteratively selecting new columns, the higher the mutual coherence is, the more OMP will tend to avoid columns that are adjacent or very near to previously selected columns. Missing adjacent columns leads to missing streaks of adjacent pixels, and the resulting visual patterns could easily give rise to faulty interpretations of the structure of the galaxy. The VLA PPO's reconstruction also tends to miss adjacent pixels in the "cloudy" ends of the galaxy, but the major parts of the galaxy remain unaffected, and the relative error is less than 2%.

Figure 21. **Compressed Sensing Reconstructions (Under-Resolved Case, $A\ell/\lambda = 0.75$)**

a. Uniform Random Distribution | **b. VLA Distribution** | **c. VLA PPO Distribution**

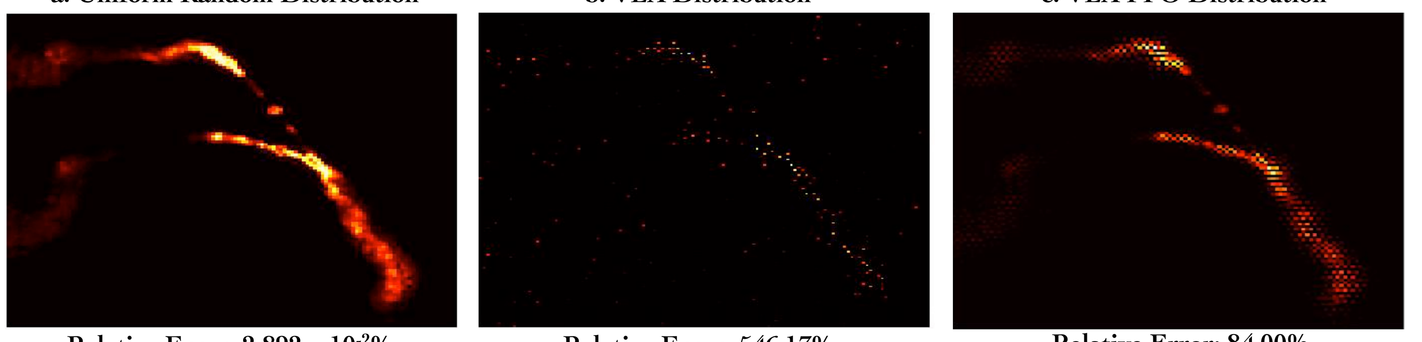

Relative Error: 2.892 x 10^{-2}% | **Relative Error: 546.17%** | **Relative Error: 84.00%**

In accordance with the impressive results in Figures 15 and 16, the accuracy of the uniform random distribution's reconstruction (Figure 21.a) appears unaffected in the under-resolved case, when the diffraction limit is broken. Maintaining a relative error of less than a tenth of a percent, the striking effects of super-resolution are clear. In contrast, the VLA distribution's reconstruction (Figure 21.b) had a massive relative error of over 500%; the reconstruction is hardly identifiable as a galaxy. Though VLA PPO provided a significantly improved reconstruction (Figure 21.c), it is still heavily affected by the missing-adjacent-pixels symptom that was apparent in the VLA distribution's well-resolved reconstruction.

Conclusions

Radio interferometers today collect massive amounts of interferometric data to reconstruct images - one sample for every point source of a radio signal. But how much of that data is actually necessary? Compressed sensing, a new counterintuitive approach that allows reconstruction with much fewer data, has been applied to many realms of signal processing since its research jumpstart several years ago. However, it has yet to be applied to radio interferometry in practice. My experiments reveal a key link between uniform randomization in radio sensor distribution, and high probability of image reconstruction success through compressed sensing - particularly with regards to how uniform randomization affects the mathematical quality of sensing matrices known as incoherence.

Figure 24. **VLA PPO Sensor Distribution**

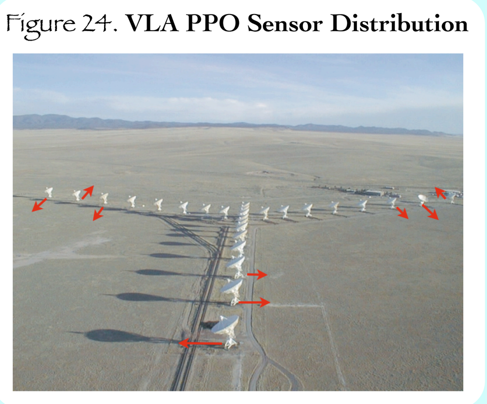

We can see the superiority of a uniform random distribution through algebraic proof (see "The Central Connection") and experimental demonstrations (Figures 13 and 15) of the high sensing matrix incoherence it provides, a theoretical indicator of compressed sensing success. Corresponding to these incoherence results, using just 351 samples (**a mere 9.75%** of the number of point sources!), the uniform random sensor distribution allows us to image radio sources far more successfully than with the VLA distribution (Figures 14). This enormous cutback of the samples needed in traditional image reconstruction - one for every point source - has great potential in improving the efficiency of sampling. The uniform random distribution has an especially acute advantage over the VLA distribution in the under-resolved case, where its super-resolution effects can allow an interferometer to enhance angular resolution, use a smaller aperture, or reconstruct signals of greater wavelengths (Figures 15 and 16). The subtlety of rigorously defined uniform randomness is also apparent - distributions with deterministic quasi-random or non-random irregularity are shown to be significantly outperformed (Figures 22 and 23).

These results also point to very simple, practical improvements to the current VLA distribution. In the well-resolved case, the VLA PPO distribution performs nearly identically to the uniform random distribution, and only requires a slight randomization modification to **just 9 sensors** in the existing VLA distribution (Figure 24).

As the uniform random distribution appears to be a theoretical and empirical benchmark, it is very plausible that adding amounts of uniform randomization to any deterministic sensor distribution would make interferometry more conducive to compressed sensing. Many similar interferometric devices (such as the **GMRT** - see "Uniform Randomness vs. Irregularity") and up-and-coming projects (such as the **Atacama Large Millimetre/submillimetre Array** and **Square Kilometer Array**) could be greatly enhanced with this new sampling perspective, allowing compressed sensing to make its first impact in astronomical studies.