

Compressed Sensing Environmental Mapping by an Autonomous Robot

Mitchell Horning
Harvey Mudd College

Simon Zou
University of California, Los Angeles

Matthew Lin
Harvey Mudd College

Matt Haberland
University of California, Los Angeles

Siddarth Srinivasan
Harvey Mudd College

Ke Yin
University of California, Los Angeles

Andrea L. Bertozzi^{*}
University of California, Los Angeles

ABSTRACT

This paper introduces the use of compressed sensing for autonomous robots performing environmental mapping in order to reduce data collection, storage, and transmission requirements. A prototype robot sends data collected over adaptively updated straight-line paths to a server, which reconstructs an image of the environment variable using Split-Bregman iteration. The amount of data collected is only 10% of the amount of data in the final map, yet the relative error is only 20%.

Categories and Subject Descriptors

I.2.9 [Robotics]: Autonomous Vehicles, Sensors ; C.4 [Performance of Systems]: Measurement Techniques, Modeling Techniques

General Terms

Algorithm, Measurement, Design, Experimentation

Keywords

Compressed sensing, autonomous control, robots

1. INTRODUCTION

Environmental mapping is essential to making truly autonomous robots and is critical for robotic exploration and discovery underwater, outdoors, within buildings, and in space [11]. However, the necessarily limited capabilities of mobile robots and inherent communication challenges of certain environments conspire to complicate the problem of mapping, even when the location of the robot(s) is known.

^{*}Corresponding author. email: bertozzi@math.ucla.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

RSN '15 Seattle, Washington USA

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

For instance, autonomous underwater vehicles would be immensely useful for mapping natural resources and scientific data collection in the ocean [2], but returning the information to a land-based server is difficult because radio communication is impractical underwater [1]. Acoustic communication is possible, but has greater bandwidth constraints and power requirements [10], hence there is a need to minimize data transmission. One approach would be to collect as much data as necessary and compress it prior to transmission, but this has significant drawbacks, as cost increases with the number of sensor nodes in a network and mobile robots are limited in their physical speed and data storage capability. A more promising solution is compressed sensing, a technique in which a relatively small amount of data is collected in order to reconstruct a higher-resolution signal or map. [7] considers the use of Random Access Compressed Sensing to create an energy-efficient network of static sensor nodes for oceanographic data collection, but in certain scenarios, such as exploration of uncharted territory, autonomous robots are clearly preferable to a network of fixed sensors. In this paper, we pioneer the use of mobile robots for compressed sensing environmental mapping. Section 2 describes the testbed hardware: a small wheeled robot tracked by an overhead camera measures the reflectivity of the testbed surface and sends the data to a central server. Section 3 documents the algorithm used by the server to reconstruct a complete map of the surface from the incomplete data collected by the robot. In Section 4 we describe both simulated and physical simulation results, and we conclude in Section 5.

2. TESTBED CONFIGURATIONS AND COMMUNICATIONS

The experimental setup consists of four main components: a testbed surface, a mobile robot, two Overhead Imaging Source DMK 21F04 1/4 Monochrome CCD cameras, and a server hosted on a windows computer.

The testbed surface is a 1.5m x 2.0m sheet of black asphalt felt paper with white paper patterns to be mapped. The robot itself consists of 5 main components:

1. **Arduino Uno:** The Arduino Uno, powered by a 9V battery, has very limited processing power (16 MHz) and memory (32 KB Flash and 2 KB SRAM). An

Adafruit WiFi Shield allows the Arduino to connect to a wireless network and thereby communicate with the server.

2. **Reflectance Sensor:** The reflectance sensor yields an analog voltage corresponding with the reflectance of the surface, which is then binarized: sensor readings above an experimentally determined threshold are read as white and those below are read as black.
3. **Chassis:** The chassis has four wheels, each directly connected to a gearmotor powered by five onboard AA rechargeable batteries.
4. **Motor controller:** The motor controller amplifies the 5V, logic-level output of the Arduino to power the motors with the AA batteries. The two motors on the left and the two motors on the right are connected in parallel to each other, enabling simple tracked-vehicle (tank) steering.
5. **Tag (Visual ID):** A unique white rectangular tag with a black-tape border and a black header strip allows the overhead cameras to visually identify the robot and determine the position of the reflectance sensor.

The two 640x480 pixel cameras send 30 frames per second to a processing computer via firewire. A python script uses OpenCV to recognize the robot identification tag and return the robot's position and orientation to the server over the serial port. Finally, the server reads the robot's location and orientation provided by the cameras, records data collected by the robot, and sends instructions to the robot. The logic and flow of data is illustrated in Figure 1.

3. ALGORITHM

3.1 Models and Assumptions

Without loss of generality, we assume the area to be explored is a rectangle denoted by Ω , and the interested environment variable $u(x, y)$ is a piece-wise constant function defined on Ω . See Figure 2 for an illustration. Assume the robot travels through n different paths, which are denoted by C_1, \dots, C_n . Along each path, the robot sensor reads the value of u at each point, and an integration of u is performed on-board by adding up the readings. Therefore, the robot collects the integral of u along C_k , which is denoted by b_k . The robot sends b_k to the server once it completes C_k . Formally, b_k is expressed as

$$b_k = \int_{C_k} u(x, y) d\Gamma, \quad k = 1, \dots, n. \quad (1)$$

For computational purpose, the whole domain is discretized into rectangular pixels (See Figure 2). Each path C_k defines a weight $a_{k,ij}$ for each pixel (i, j) . If C_k does not intersect with pixel (i, j) , then $a_{k,ij} = 0$, otherwise $a_{k,ij}$ is defined to be the length of the part of C_k that falls within pixel (i, j) . It is also assumed that the value of u is a constant within each pixel (i, j) , which is denoted by u_{ij} . After discretization, Eq (1) becomes

$$b_k = \sum_i \sum_j a_{k,ij} u_{ij}. \quad (2)$$




Figure 2: An illustration of the paths of robots. The shaded regions denote the area of interest (unknown), where the values of environment variable u is significantly different from the surroundings. The size of the unit pixel (each small rectangle) is determined by the accuracy of the positioning system.

Let $b = (b_1, \dots, b_n)^t$, $u = (u_{ij})$, then b and u have a linear relation, which is written as

$$Au = b, \quad (3)$$

where the linear operator A is specified in Eq (2). Eq (3) is the model equation, which poses an inverse problem. In this equation, A is determined from the paths that the robot travelled in - it travels in straight line paths, hence we only need the start and end locations to determine its trajectory. Each path contributes one row to A . b is obtained from the experiment, and is the vector of the integral of sensor values as the robot travelled along each of its paths. u is the variable that we want to solve. Of course, if the paths form a complete raster scan over the whole domain, then theoretically Eq (3) has a unique solution. For economical reasons, it is desirable to use much fewer paths and still being able to reconstruct a solution for u . The problem is well suited for compressed sensing based image reconstruction techniques, which provide a way to solve these underdetermined systems.

3.2 Solving the Inverse Problem

Based on the observation that the image u to be reconstructed is piecewise constant, its directional gradients, $\nabla_x u$ and $\nabla_y u$, are sparse over the domain. Also we assume that the interesting feature of the image is relatively small compared to the whole domain to be explored. Therefore, we postulate that the solution u minimize the energy functional

$$E_0(u) = \alpha |u|_0 + \beta (|\nabla_x u|_0 + |\nabla_y u|_0), \quad (4)$$

subject to the data fitting constraint (3). Here $|\cdot|_0$ denotes the ℓ^0 norm. α and β are parameters balancing the weights of the sparsity of the gradients and the feature. We note that minimizing ℓ^0 norm is generally considered intractable, we replace it with ℓ^p ($p > 0$) norm, which still enforces sparsity and much easier to solve:

$$E_p(u) = \alpha |u|_p + \beta (|\nabla_x u|_p + |\nabla_y u|_p). \quad (5)$$

Prior work in the field of compressed sensing and sparse image reconstruction normally chooses $p = 1$ [4], and some theoretical work has demonstrated its efficiency [3]. One popular algorithm for solving ℓ^1 minimization problems is based on Bregman iteration, which is carefully described in [8]. And we follow this method to solve our problem.

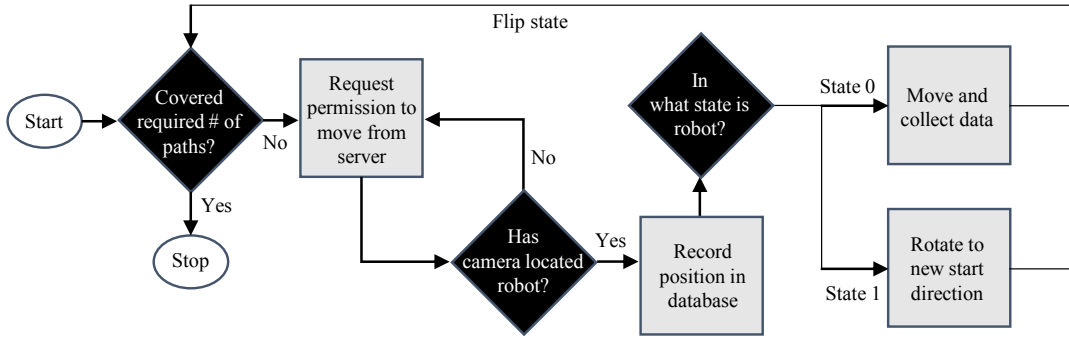


Figure 1: The robot operates as a finite state machine. During state 1, the robot attempts to rotate to a start orientation previously commanded by the server, then requests for the server to record its actual position from the video cameras. The robot switches to state 0 upon confirmation, integrates sensor data along a straight path for a duration commanded by the server, and requests for the server to record the actual final position.

Although it is a nonconvex optimization problem in the case $0 < p < 1$, some work such as [5, 6] has shown that it can also be solved using the Bregman-type method, and that the solution converges to the global minimum of the associated ℓ^0 formulation.

3.3 Path Planning Algorithm

The server program generates the next path for the robot every time the robot is in State 1 (see Figure 1). Without any prior information about the distribution of features in the domain to be explored, it is natural to choose random paths that are uniformly distributed over the whole domain. Although we would demonstrate later that it should work, it is desirable to adopt an adaptive path planning strategy based on the data already collected by the robot as it moves around.

Initially, the robot collects data along some constant number of uniformly distributed random paths. Then all the data that have already been collected are processed off-board on the server, and a coarse resolution image is reconstructed. Based on this image, the server program would segment interesting areas that contain more features than elsewhere. And it generate a new set of paths, still random but with higher probability passing through those interesting areas. As the robot travels more paths and collects more data, the resolution of the reconstructed image would improve, and the new paths based on that would be more likely directed into areas rich in features.

The segmentation of interesting areas consists of two steps. First a simple thresholding is applied to zero out small pixel values. Then it finds connected components containing more than n pixels and put them into the set of interesting areas. Every new path passes through the current location and the next location. The next location is chosen randomly from all viable pixels, with probability p_1 of being one inside interesting areas, and p_0 for non-interesting areas. It is required that $p_0 < p_1$ and the probabilities add up to 1. The pseudocode for adaptive path planning is in Algorithm 1.

With the adaptive path planning, the workflow of our algorithm is in Algorithm 2.

4. EXPERIMENT

4.1 Simulated Results

Algorithm 1 Adaptive path selection

```

1: function ADAPTIVEPATHS( $u, \Omega, p_0, p_1, \text{max\_num}, n$ )
2:   for  $i = 1$  to  $\text{max\_num}$  do
3:      $u_1 \leftarrow \text{THRESHOLDING}(u, \text{threshold})$ 
4:      $\Omega_1 \leftarrow \text{CONNECTEDCOMPONENTS}(\Omega, u_1, n)$ 
5:      $P_i \leftarrow \text{NEWPOINT}(\Omega, \Omega_1, p_0, p_1)$ 
6:      $C_i \leftarrow \text{NEWPATH}(\Omega, P_{i-1}, P_i)$ 
7:      $C_{\text{new}} \leftarrow C_{\text{new}} \cup \{C_i\}$ 
8:   end for
9:   return new paths  $C_{\text{new}}$ 
10: end function

```

Algorithm 2 Environment Mapping Algorithm

```

1: procedure ENVIRONMENTMAPPING
2:   initialize  $u, A, b$  with zeros
3:   while  $\|u_{\text{old}} - u_{\text{new}}\| < \text{tol}$  do
4:      $C_{\text{new}} \leftarrow \text{ADAPTIVEPATHS}(u_{\text{new}}, \Omega, p_0, p_1, \text{num}, n)$ 
5:      $b_{\text{new}} \leftarrow \text{COLLECTNEWDATA}(C_{\text{new}})$   $\triangleright$  on-board
6:      $A, b \leftarrow \text{UPDATESYSTEM}(\Omega, A, C_{\text{new}}, b, b_{\text{new}})$ 
7:      $u_{\text{old}} \leftarrow u_{\text{new}}$ 
8:      $u_{\text{new}} \leftarrow \text{SPLITBREGMAN}(A, b, \alpha, \beta)$ 
9:   end while
10:   $u \leftarrow u_{\text{new}}$ 
11:  return  $u$ 
12: end procedure

```

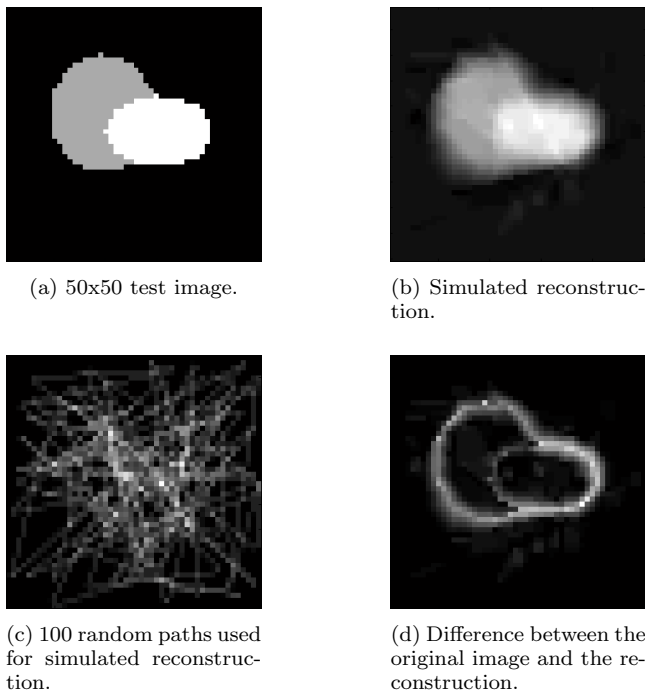


Figure 3: Simulated reconstruction using random paths.

Before actually using the robot to collect data, we run simulated reconstructions. Figure 3a is a 50x50 test image for simulations. We note that the test image has multiple gray scales. In our experiment the robot can only distinguish black and white colors since it is equipped with a binary sensor. With more advanced sensor our method should work as well.

Figure 3b is a reconstruction of a 50x50 image from 100 random paths, with data to unknowns ratio of 1:25. The shape and the intensity of the region of interest are successfully identified. Figure 3c illustrates the paths for the robot, with each one generated by connecting current position and randomly selected target position while ensuring the robot stays within the testbed. In the figure, brighter pixels represent points that are visited more times. Figure 3d shows the difference between the original image and the reconstructed one. It shows that most of the error in the reconstruction is around the edges of the region of interest. The relative error of the reconstruction is 0.19715.

4.2 Comparison with Other Approaches

In order to judge the effectiveness of our approach, we examine other methods of collecting data and reconstructing an image. An intuitive one is to sample individual pixels and try to reconstruct the image from these samples. One reconstruction approach is to use the same Split Bregman iteration that we are using for reconstruction from path integrals by assuming that each path is degenerated to a sample point. Alternatively, we can take the value found at each sampled pixel and assign the same value to nearby pixels, a method that we refer to as pixel expansion.

Figures 4 and 5 present reconstructions of the previously shown 50x50 test image, Figure 3a, using these approaches, both with randomly sampled points and points sampled on a grid. In all cases, 100 pixels are sampled, and the sampled

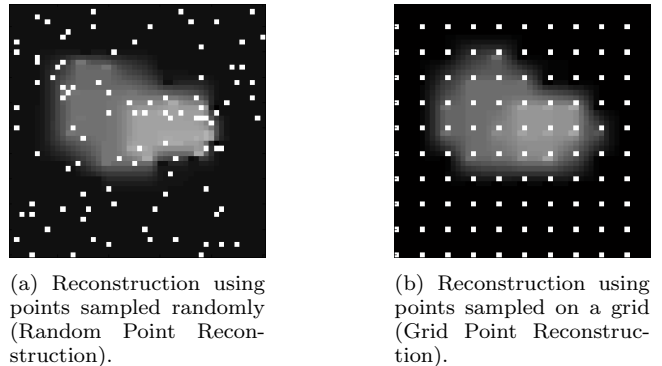


Figure 4: Values are found at sampled points, and Split Bregman iteration is used to reconstruct the image.

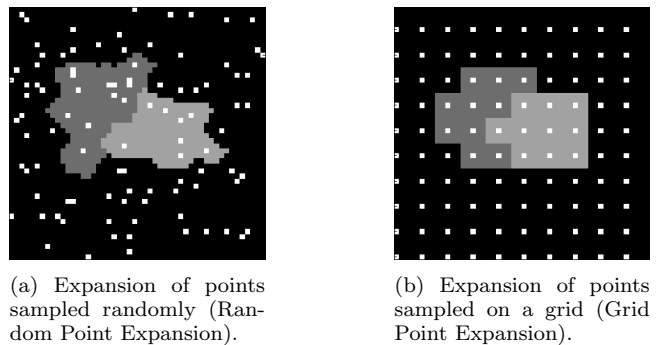


Figure 5: Values are found at sampled points, and those values fill the area surrounding the sampled points.

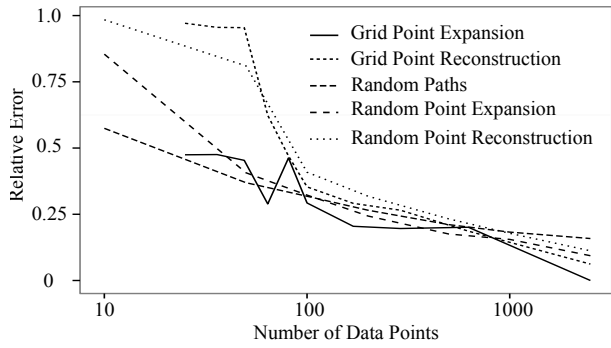


Figure 6: Plot of error for the various methods of reconstruction with respect to the number of data samples used.

points are colored white.

In Figure 4, Split Bregman iteration was used with the collected data for reconstruction. By comparing Figures 4a and 4b, it seems that the reconstruction is more accurate when the points are sampled on a grid. This is reasonable because by sampling on a grid, the spread of sampled points is very even and well distributed. When the points are sampled randomly, certain areas will be less thoroughly investigated than others, hindering accurate reconstruction.

Figure 5 illustrates pixel expansion. This process creates expanded pixels of a single color around each sampled pixel. When the pixels are sampled on a grid, the even distribution causes the expanded pixels to simply be squares, essentially creating a low resolution version of the image, as in Figure 5b. Figure 5a is a result of pixel expansion with randomly sampled pixels. The unpredictable spread of sampled pixels causes the expanded pixels to vary widely in shape and size, and the reconstruction does a poor job of determining the shape of the region of interest.

To quantitatively compare these methods, Figure 6 plots the errors of reconstruction, with respect to the number of data points collected. The reconstruction from randomly generated path integrals achieve lower error than the other methods for a small number of data points. As the number of data points grows, all methods converge to similar error; this is expected because enough data would turn the problem from under-determined to fully-determined. Worth noting is that while pixel expansion from samples on a grid appears to do better than random path integrals for some of the low data situations, the effectiveness of grid point expansion is highly variable, and as such, its overall usefulness is limited. The accuracy of reconstruction is very dependent on how well the grid aligns with the given image, and choosing the best grid size for a given environment would require prior knowledge of the environment. Another flaw with pixel expansion is the dependence on accurate sampling; pixel expansion will be very heavily affected by noise. If a pixel value is read inaccurately, the entire expanded pixel associated with that sample will be inaccurately reconstructed. Using path integrals mitigates the impact of inaccurate reading of individual pixels.


4.3 Experimental Results

We proceeded to carry out experimental testing with the actual robot. Figure 7a shows a testbed that we used. The testbed is modeled as a 70x90 image.

As a means of comparison, we first carried out the recon-




(a) Image of the testbed, sized at 70x90.



(b) Simulated reconstruction from the 178 paths that the robot travelled.



(c) Reconstruction from the data collected by the robot.



(d) Reconstruction from the data collected by the robot using tuned parameters.

Figure 7: Reconstruction using the data collected by the robot.

struction as a simulation, shown in Figure 7b. This simulated reconstruction uses the same set of 178 paths that the robot travelled along. Thus, this reconstruction is an ideal reconstruction.

The actual reconstruction from the robot collected data is shown in Figure 7c. Clearly, this reconstruction is more noisy, giving black spots in the middle of the region and creating white spots where it should be black. This is due to the fact that the data collection by the robot is not as perfect as the data in simulation. To make the model more robust against the noise, we retry the reconstruction with a lowered data fidelity parameter. This yields Figure 7d. It does a reasonably good job of showing the region of interest.

4.4 Adaptive Pathing


We first run simulations on a test image as in Figure 8a. As a comparison, we run the reconstruction using uniformly chosen random paths as well as adaptive paths as described in Sec 3.3. Each method uses the same Split-Bregman algorithm with the same number of paths.

Figure 8c shows the 200 uniformly random paths for the simulation. The reconstruction obtained with these paths is shown in Figure 9a. This reconstruction is able to recover successfully generally the region of interest and the size of the area, but is not able to ascertain some of the more detailed information about the shape and boundary of the area accurately. Figure 8d shows the difference between the the reconstructed image and the ground truth. The error is primarily along the edges but more prominent along the bottom edge. The relative error of the reconstruction is 0.5023.




(a) 100 x 100 pixel test image with features in bottom right.

(b) Simulated reconstruction from the 200 paths




(c) The 200 randomly generated paths.



(d) Difference between the original image and the reconstruction

Figure 8: Simulated reconstruction using uniformly distributed random paths.



(a) Simulated reconstruction from 200 adaptive paths.

(b) The 200 adaptive paths.


(c) The difference between the reconstruction and original image.

Figure 9: Reconstruction using the adaptive paths.



(a) Image of the testbed, sized at 70x90

(b) The 224 adaptive paths traveled by the robot



(c) Simulated reconstruction for the adaptive paths traveled by the robot



(d) Image reconstruction using robot collected data with tuned parameters

Figure 10: Experimental results of adaptive pathing using robot collected data

We compare these results against those displayed in Figure 9, which are generated by the adaptive pathing algorithm on the same test image in 8a. Figure 9b shows the 200 adaptive paths taken by the simulation. There are visibly clustered near the bottom of the image, where it identified early on had the area of interest. Figure 9a then shows the results of the Split Bregman reconstruction with the data collected from the adaptive paths. Because the paths tend to cluster in the region of interest, the reconstruction is able to more accurately capture the curve along the borders and thus the shape of the area itself. This is illustrated in Figure 9c, which shows the difference between this reconstruction and the original test image. The relative error of the reconstruction is .19614, which is nearly half as that for uniformly random paths.

We implement adaptive path planning for the robot and the results are shown in Figure 10. Figure 10a is an image of the original testbed, sized to 70x90, which contains multiple areas of interest that are away from the center of the test bed and have sharply defined edges. Figure 9b shows the 224 straight line paths the robot vehicle traveled on the testbed and it can be observed that it did not travel often to black areas or corners where there were no areas of interest. We then generate Figure 10c, which is an ideal reconstruction assuming no error in data collection. If we run the same reconstruction with robot data, the resulting reconstruction is significantly affected by the noise, as shown in Figure 10d. It is close to the reconstruction with ideal data.

5. CONCLUSION

We have demonstrated a simple yet useful prototype au-

tonomous robot for mapping environment variables. It can be applied to existing robot platforms that are equipped with wireless communication and tracking systems. The required computing power of the robot is only onboard addition and a storage big enough to hold one data point. The data collection strategy based on compressed sensing enables minimal data communication. It is further improved by adopting adaptive path planning using already collected data. The image reconstruction algorithm minimizes the data fitting error plus the ℓ^1 norm of the gradient of the image, which captures piecewise constant features. Split-Bregman algorithm is efficient for solving the minimization problem. Future work includes improving data processing methods that incorporates statistical methods such as cross validation and outlier pursuit, so that the reconstruction is more robust against the noise. Another work that can be done is to use sensors that can distinguish multiple levels of intensities, or even trying to use multiple sensors on the same robot. We also note that our method is scalable to using multiple robots, as long as the path planning can make robots avoid collisions.

6. ACKNOWLEDGMENTS

This work was supported by NSF grants CMMI-1435709, DMS-1045536, and DMS-1118971 and UC Lab Fees Research grant 12-LR-236660. The authors would also like to thank Rick Chartrand and collaborators at Los Alamos National Laboratory for their technical recommendations.

7. REFERENCES

- [1] I. F. Akyildiz, D. Pompili, and T. Melodia. Challenges for efficient communication in underwater acoustic sensor networks. *ACM Sigbed Review*, 1(2):3–8, 2004.
- [2] I. F. Akyildiz, D. Pompili, and T. Melodia. Underwater acoustic sensor networks: research challenges. *Ad hoc networks*, 3(3):257–279, 2005.
- [3] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.
- [4] E. J. Candès and M. B. Wakin. An introduction to compressive sampling. *Signal Processing Magazine, IEEE*, 25(2):21–30, 2008.
- [5] R. Chartrand. Nonconvex compressive sensing and reconstruction of gradient-sparse images: random vs. tomographic fourier sampling. In *Image Processing, 2008. ICIIP 2008. 15th IEEE International Conference on*, pages 2624–2627. IEEE, 2008.
- [6] R. Chartrand. Fast algorithms for nonconvex compressive sensing: Mri reconstruction from very few data. In *Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on*, pages 262–265. IEEE, 2009.
- [7] F. Fazel, M. Fazel, and M. Stojanovic. Random access compressed sensing for energy-efficient underwater sensor networks. *Selected Areas in Communications, IEEE Journal on*, 29(8):1660–1670, 2011.
- [8] T. Goldstein and S. Osher. The split bregman method for l_1 -regularized problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343, 2009.
- [9] M. Gonzalez, X. Huang, D. S. H. Martinez, C. H. Hsieh, Y. R. Huang, B. Irvine, M. B. Short, and A. L. Bertozzi. A third generation micro-vehicle testbed for cooperative control and sensing strategies. In *ICINCO (2)*, pages 14–20, 2011.
- [10] J. Heidemann, W. Ye, J. Wills, A. Syed, and Y. Li. Research challenges and applications for underwater sensor networking. In *Wireless Communications and Networking Conference, 2006. WCNC 2006. IEEE*, volume 1, pages 228–235. IEEE, 2006.
- [11] S. Thrun et al. Robotic mapping: A survey. *Exploring artificial intelligence in the new millennium*, pages 1–35, 2002.