Proceedings of the ASME 2019
Dynamic Systems and Control Conference
DSCC2019
October 8-11, 2019, Park City, Utah, USA

DSCC2019-9201

COMPRESSIVE SENSING-BASED RECONSTRUCTION OF LISSAJOUS-LIKE NODDING LIDAR DATA

Michael T. Benson
Dept of Mechanical Engineering
Villanova University
Villanova, PA 19085

Harish Sathishchandra
Dept. of Electrical and Computer Engineering
Boston University
Boston, MA 02215

Garrett M. Clayton
Dept of Mechanical Engineering
Villanova University
Villanova, PA 19085
Email: garrett.clayton@villanova.edu

Sean B. Andersson
Dept. of Mechanical Engineering
Div. of Systems Engineering
Boston University
Boston, MA 02215
Email: sanderss@bu.edu

ABSTRACT

In this article, a compressive sensing (CS) reconstruction algorithm is applied to data acquired from a nodding multi-beam Lidar system following a Lissajous-like trajectory. Multi-beam Lidar systems provide 3D depth information of the environment for applications in robotics, but the vertical resolution of these devices may be insufficient to identify objects, especially when the object is small and/or far from the robot. In order to overcome this issue, the Lidar can be nodded in order to obtain higher vertical resolution with the side-effect of increased scan time, especially when raster scan patterns are used. Such systems, especially when combined with nodding, also yield large volumes of data which may be difficult to store and mange on resource constrained systems. Using Lissajous-like nodding trajectories allows for the trade-off between scan time and horizontal and vertical resolutions through the choice of scan parameters. These patterns also naturally sub-sample the imaged area and the data can be further reduced by simply not collecting each data point along the trajectory. The final depth image must then be reconstructed from the sub-sampled data. In this article, a CS reconstruction algorithm is applied to data collected during a fast and therefore low-resolution Lissajous-like scan. Experiments and simulations show the feasibility of this method and compare its results to images produced from simple nearest-neighbor interpolation.

INTRODUCTION

For autonomous robots, having range information about the environment and objects within it is extremely useful, if not necessary, for the robot to navigate and react to its surroundings. The data can be used for localization and/or mapping in the forms of visual odometry [1–3], SLAM [4–6], and more [7–9]. In recent years, Lidar has perhaps become the most widespread way to collect this information, in large part due to increasing availability and falling prices as a results of their use in self-driving automobiles. While Lidar configurations come in many different forms, recent systems tend to favor multi-beam arrangements, such as those from Velodyne [10] or Ouster [11] which rotate multiple beams to provide a more complete 3D view of the scene.

While these rotating multi-beam configurations do provide 3D data, due to limitations on how closely the beams can be

spaced, the vertical resolution can still be insufficient for specific applications; due to these limitations, objects can "disappear" as they pass between visibility of one beam to another. For example, one of Velodyne's most affordable models, the VLP-16 "Puck", cannot consistently see objects one meter tall at distances over 30 meters. Note that at these ranges the Lidar could, but is not guaranteed to, see the object with one beam at most. Even if seen, it is unlikely that detection by a single beam would provide sufficient information to successfully classify the object. Newer, higher resolution models face the same problem, but with smaller heights at larger ranges.

Before multi-beam Lidar systems became more widely available, single-beam planar Lidar units were (and still are, in some applications) used to collect 3D depth information by placing them on a nodding or spinning platform, allowing the plane in which the Lidar scans to be rotated [12–17]. Typically, these configurations are driven through a raster scan pattern in which the assembly is rotated in a stair-step pattern between scans of the Lidar. This method, results in high resolution in the dimension added due to rotation, but at the expense of significantly increased scan times. In prior work of two of the authors [18], the applicability of raster scan patterns for multi-beam Lidar to help solve the "disappearing" problem, and how other scan trajectories, such as Lissajous-like scan patterns, can be used to provide a trade-off between scan period and resolution was explored. While Lissajous-like scanning could replicate depth images from the raster scan to some degree, the resulting decrease in data gathered resulted in smaller objects being more difficult to identify.

While it is often useful to trade-off some resolution for higher speed imaging, it is typically still desirable to have the highest-resolution scan possible to support detection, mapping, and SLAM. Inspired by recent results of one of the authors on increasing the imaging rate in atomic force microscopy through sub-sampling, combined with image reconstruction algorithms to recover high resolution detail [19–21], in this paper we explore the use of reconstruction algorithms combined with the Lissajous sampling pattern in an attempt to get both high speed and high resolution. The essential idea is to view a high-resolution raster-scan image as a "ground truth" image that has been sub-sampled through the Lissajous-like scan. The final image is formed from the data using a recent optimization algorithm for sparse reconstruction of 3-D depth images [22]. This approach differs from common Lidar reconstruction methods, which infer some information about the detected objects from segmentation or classification to determine how to perform reconstruction on the Lidar data, such as in [23–25]. Instead, in this work, reconstruction is performed directly on the raw Lidar data using compressive sensing techniques, without performing any other processing on the data.

In this article, experimental results are presented to establish the feasibility of this method. Specifically, Lissajous-like scans that sample less than 10% of the full raster scan data can reproduce the full raster data set with high fidelity when the CS algorithm is applied. To further show the advantages of this method, simulated scans at much lower resolution are also used to reconstruct the full raster data with good results.

LISSAJOUS-LIKE PATTERNS IN NODDING LIDAR

Much of the work relevant to this section was presented in [18], and as a result will be presented in minimal depth here. The nodding Lidar system used for this effort is shown in Figure 1a. The apparatus features a direct drive system for nodding the Lidar platform and is designed so the distance separation between the Lidar's optical center and the axis of rotation is small. The Lidar used is the Velodyne VLP-16 Puck, which rotates its beams at a rate of 20 Hz, with a horizontal resolution of 900 points per scan over a full 360° range, and a vertical resolution of 16 equally spaced beams over a 30° field of view, which in total produces a maximum output of 14,400 points per scan in single return modes.

The relevant axes and geometry of the assembly is shown in Figure 1b. Note that for this work, the angular velocity f_l which dictates the angle of the beams inside the Lidar θ_l and the incident angle of the beams θ_i are considered fixed; only trajectory design of the nodding angle θ_n is discussed. In this work, the spinning of the beam assembly θ_l rotates according to:

$$\theta_l(t) = -\pi + 2\pi(f_l t \bmod 1) \tag{1}$$

where $f_l = 20Hz$ is the fixed rotational velocity of the Lidar, and the incident angles of the 16 beams θ_i are fixed:

$$\theta_i \in \{-15^\circ, -13^\circ, \ldots -1^\circ, 1^\circ, \ldots 15^\circ\}. \tag{2}$$

Because the parameters of the Lidar are fixed, the distribution of the scan readings is determined by the choice of trajectory for the nodding platform. For a raster scan, the trajectory is a stair-step pattern, with the Lidar nodding in small angle increments between scans of the Lidar, with a trajectory spanning from some arbitrary angle $-A$ to A. For Lissajous-like scanning, the Lidar is nodded with a triangular trajectory between $-A$ to A with frequency ω_n (we call this "Lissajous-like" because Lissajous trajectories are typically sinusoidal in both axes). The path of one beam of the Lidar for each of these scans trajectories, as well as the stationary Lidar, is shown in Figure 2. In each plot, the x-axis represents the spinning of the Lidar beams θ_l, while the y-axis represents the angle due to the nodding motion applied to the Lidar θ_n.

In this work, an indoor and an outdoor space were imaged by the Lidar. The Lidar followed a nodding trajectory of $\pm 20^\circ$,

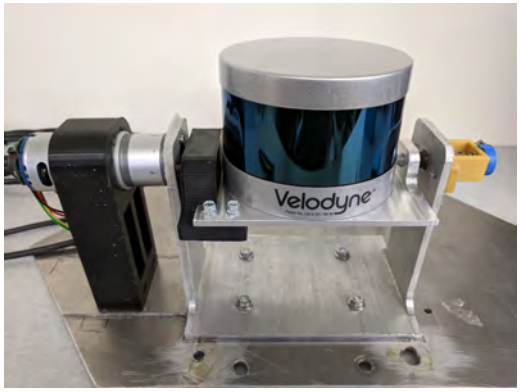

(a)

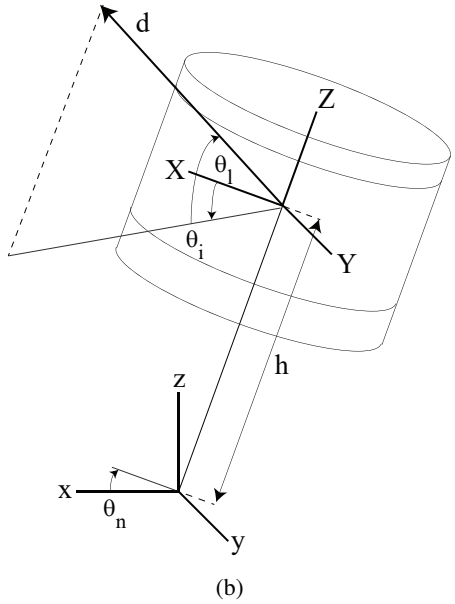

(b)

Figure 1: Figures showing a) the Lidar/cradle apparatus used in this research and b) the relevant axes and notation.

which provided a sufficient field of view to visualize the relevant objects in the space (but not things considered unnecessary, such as the ceiling of the indoor space). The spaces are first scanned by the Lidar undergoing a raster scan trajectory with a period of 10 s. Then, the spaces are scanned by the Lidar with Lissajous-like trajectories with periods of 1 s, 0.5 s, and 0.25 s. These Lissajous-like scans can be treated as "downsampled" versions of the raster scan, which is treated as the ground truth, with downsampling percentages of 10%, 5%, and 2.5%, respectively.

Figure 3 shows the indoor space imaged by a color camera and the nodding Lidar. A panoramic color image is presented in Figure 3a to familiarize the reader with the space. The raster scan is shown in Figure 3b, and the 1 s and 0.25 s Lissajous-like scans are shown in Figures 3c and 3d respectively. The data from the Lidar scans are displayed as depth images, with closer objects

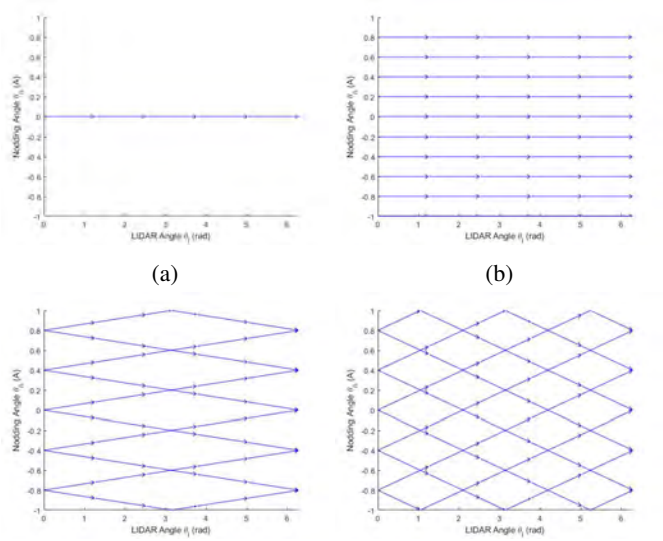

(a) (b) (c) (d)

Figure 2: Trajectories of one beam of the Lidar for a) a single, stationary scan of the Lidar with a period of 0.05s, b) a raster scan with a period of 1s, and (c-d) Lissajous-like scans with periods of c) 0.5s and d) 0.167s. The nodding angle θ_n is on the y-axis, and the Lidar's rotation angle θ_l is on the x-axis.

being whiter, and black where there is no data from the scan. Note that, because of how the grid is generated for the depth images, the downsampling percentages of the Lissajous-like scans do not correspond to the percentage of populated pixels in the depth images. This is because of cases where multiple points are associated with the same pixel, commonly along the sides of the scan where measurements are more clustered.

RECONSTRUCTION OF LIDAR IMAGES FROM SUB-SAMPLED DATA

The goal of reconstruction is to fill in the missing pixels in the depth image (namely, those that were not sampled) with values that depend on the measured data. Perhaps the simplest way to do this is nearest neighbor interpolation. Such an approach can be effective if there is sufficient data, but it does not take into account any structure in the image and may therefore produce less than ideal results. An alternative approach is to leverage results from compressive sensing (CS). CS is a joint measurement and signal processing technique which can produce good (or even exact) reconstructions of signals from significantly fewer measurements than the Nyquist-Shannon sampling theorem requires [26]. At the heart of CS is the assumption of compressibility (or true sparsity) of the signal of interest, that is, when described in an appropriate basis, most of the coefficients are negligible (compressibility) or exactly zero (sparsity).

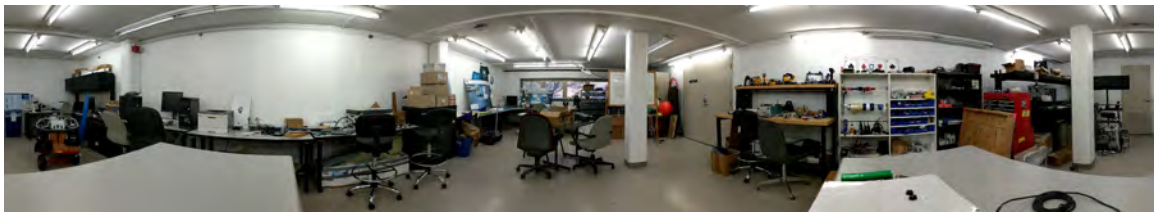

(a) Indoor color panorama

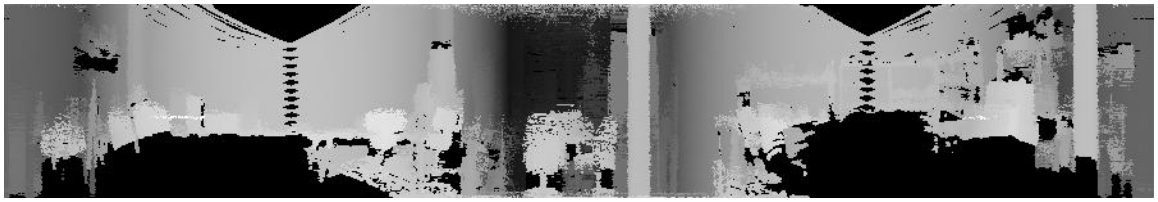

(b) Raster scan

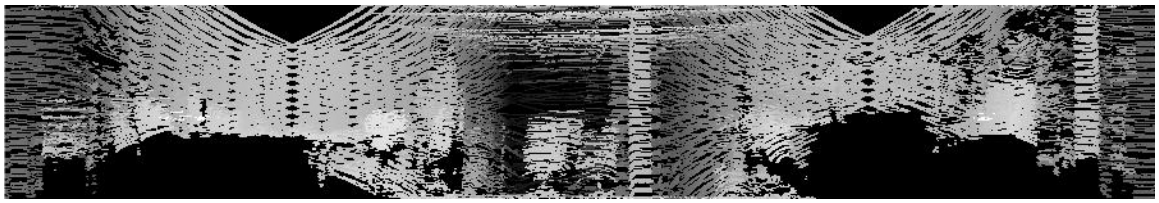

(c) 1Hz Lissajous-like scan

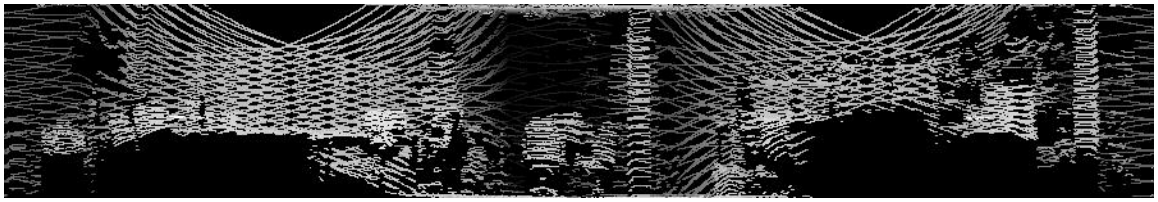

(d) 4Hz Lissajous-like scan

Figure 3: a) A color image from an RGB camera, and (b-d) depth images of the same space sampled by the nodding Lidar. The depth images are derived from point clouds generated by b) a raster scan with a scan period of 10s, and (c-d) Lissajous-like scans with periods of c) 1/10th (10%) and d) 1/40th (2.5%) of that of the raster scan. In the depth images (b-d), closer points are whiter, and the image is black where no depth information is available. Note that the camera and depth images suffer from different distortion effects, so they are not perfectly aligned.

In the general setting, a signal is reconstructed from measurements by solving an ℓ_1 minimization problem,

$$\min \|\eta\|_1 \quad \text{subject to} \quad y = \Phi x = \Phi\Psi\eta \triangleq A\eta \tag{3}$$

where $x \in \mathbb{R}^n$ is the true signal, $y \in \mathbb{R}^m$ is the vector of measurements, Φ is an $m \times n$ measurement matrix, Ψ is an $n \times n$ orthonormal basis for $\mathbb{R}^n$ (often referred to as the *sparsity basis*), and η is the sparse representation of the true signal x in the domain of Ψ. We define the product $\Phi\Psi$ as the $m \times n$ matrix A.

In the context of reconstruction of the depth image from the nodding Lidar scan, measurements are restricted to values of the pixels of the image that lie on the scanned path. As a result, each row in Φ consists of a single one (indicating which pixel is measured) with all other entries being zero. Thus, the measurement matrix Φ is given by extracting rows from an $n \times n$ identity matrix. This has implications for a property know as the *mutual coherence* [27], generally implying that reconstructions from data sampled by Φ from solving (3) will not be exact and in general may have a large error.

As it turns out, in structured environments such as one finds in indoor environments, the depth profile is often well described as consisting of many planar regions connected by a few "edges". The sparsity of the data, then, is not with respect to a particular sparsity basis, but rather with respect to these edges. This case was first considered in [22] where a reconstruction algorithm was

created using a cosparsity model. We give a brief description here.

To reflect the fact that we are after a depth profile, let the underlying signal now be denoted as $Z \in \mathbb{R}^{n\times n}$, viewed as a matrix. Since Z is assumed to be formed from (not too many) planar regions, the goal is to find a profile that fits the data while minimizing the number of depth changes in the horizontal, vertical, and diagonal directions. These changes can be measured using three different convolutional filters with kernels

$$D_{xx} = \begin{bmatrix} 0 & 0 & 0 \\ 1 & -2 & 1 \\ 0 & 0 & 0 \end{bmatrix}, D_{yy} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & -2 & 0 \\ 0 & 1 & 0 \end{bmatrix}, D_{xy} = \begin{bmatrix} -1 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & -1 \end{bmatrix}. \quad (4)$$

Reconstruction of the full depth profile from a collection of measurements y obtained by a measurement matrix Φ comes from solving the optimization problem

$$\begin{aligned} &\min_{Z} \|\mathrm{vec}\,(Z \circledast D_{xx}) + \mathrm{vec}\,(Z \circledast D_{yy}) + \mathrm{vec}\,(Z \circledast D_{xy})\,\|_1 \\ &\quad \text{subject to} \|\Phi \mathrm{vec}\,(Z) - y\|_\infty \leq \varepsilon, \end{aligned} \quad (5)$$

where $\mathrm{vec}\,(\cdot)$ is the vectorization operator which transforms a matrix into a vector by stacking the columns, $\circledast$ is the convolution operator, and ε is a user-defined parameter that accounts for noise in the measurements.

EXPERIMENTAL IMAGING AND RECONSTRUCTION

In the following, an experimental investigation of the Lidar reconstruction algorithm shown above is presented. The goal of these experiments is to show the efficacy of compressive sensing reconstruction for Lissajous nodding Lidar scans for different environments.

EXPERIMENTAL SETUP

Lidar scans of both an indoor and outdoor space were obtained. In both scenarios, the Lidar assembly was controlled to follow a nodding trajectory with $A = 20^\circ$. This provided a sufficient field of view to visualize the relevant objects in the space. In order to provide a ground-truth example for the reconstruction, raster Lidar scans of both spaces are performed. The scan trajectory had a period of 10 s. The spaces are then scanned by the Lidar with Lissajous-like trajectories with periods of 1 s, 0.5 s, and 0.25 s. These Lissajous-like scans can be treated as "downsampled" versions of the raster scan with downsampling percentages of 10%, 5%, and 2.5%, respectively. While these percentages are already quite low, they still represent a significant amount of data. To explore performance at extremely low sampling, the Lissajous scan data were further sub-sampled by discarding some of the data. This was done in two ways: first by keeping only every 10^{th} sample (thus regularly downsampling) and second by randomly selecting 10% of the sample points. This second approach is motivated by the fact that in general, CS-based reconstruction performs best when the sample locations exhibit some level of randomness. (In practice, of course, one would not discard data but rather choose to only acquire the limited set.) Both the regular and the random yield a total downsampling of 1%, 0.5%, and 0.25%. Two reconstructions were performed on the data, one a nearest neighbor interpolation, and the other the compressive sensing approach discussed in the previous section. Note that experimental results are only shown for the 2.5% Lissajous scan since it is the sparsest data and, thus, the most difficult to reconstruct.

EXPERIMENTAL RESULTS

Reconstructions of the indoor Lidar data from the 2.5% sampling scan in Fig. 3d are shown in Fig. 4. Both the nearest neighbor interpolation (top image) and the CS based reconstruction (bottom image) are visually quite good. While there are only minor differences between them, in general, edges are (slightly) crisper in the CS-based reconstruction.

These differences are accentuated at the very low sampling rates of 0.25% after either regular or randomly downsampling the Lissajous data. Reconstructions for the indoor data are shown in Fig. 5 for the regularly subsampled data and in Fig. 6 for the randomly subsampled set. In both cases, both the nearest neighbor and the CS reconstructions are remarkably good given the extremely low sampling, though both are clearly degraded relative to the higher sampling rate. The CS version, however, now has significantly better edges and a smoother overall appearance. In addition, the randomly subsampled image is visually superior to the regularly-spaced version despite the same level of subsampling.

As a comparison between all the indoor cases presented above, focus on the electrical bench and chair to the right of the column at the center of Fig. 3a. Qualitatively speaking, the appearance of the bench and chair in the 2.5% scans in Fig. 4 is smoother in the CS reconstruction when compared to the nearest neighbor reconstruction. In the 0.25% scans in Figs. 5 and 6, the same relation between the CS and nearest neighbor interpolations is observed. When comparing the 2.5% to the 0.25% scans, the electrical bench and chair is less blurry in the 2.5% scan as expected, with the electrical bench and chair being more clear in the random sampled 0.25% scan in Fig. 6.

Results from the outdoor Lidar scans are shown in Fig. 7. An image of the space is shown in Fig. 7(a), followed by a raster scan in Fig. 7(b), the 2.5% Lissajous image in Fig. 7(c), and the reconstruction results in Fig. 7(d) and (e). As before, the reconstructions are fairly faithful to the raster scan, though smaller features such as the windows in the building on the left which do

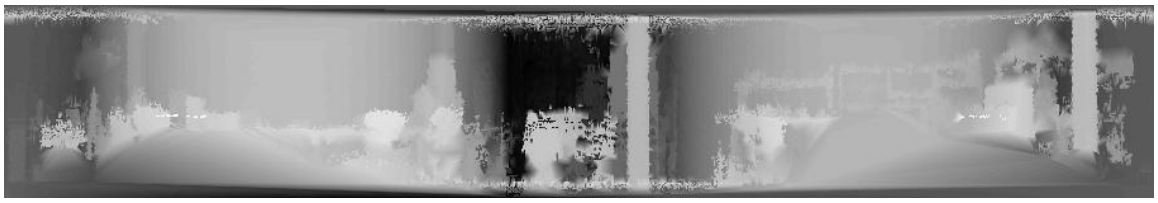

(a) Nearest neighbor interpolation

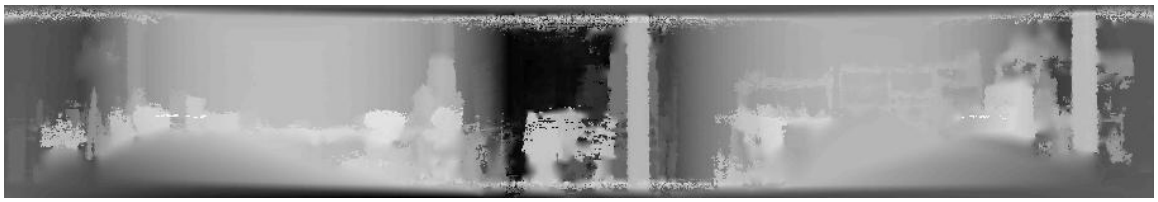

(b) CS reconstruction

Figure 4: Reconstruction results for the indoor lab space using the 2.5% sampling shown in Fig.3d for a) nearest neighbor interpolation and b) the compressive sensing reconstruction algorithm. These should be compared to the ground truth in 4b.

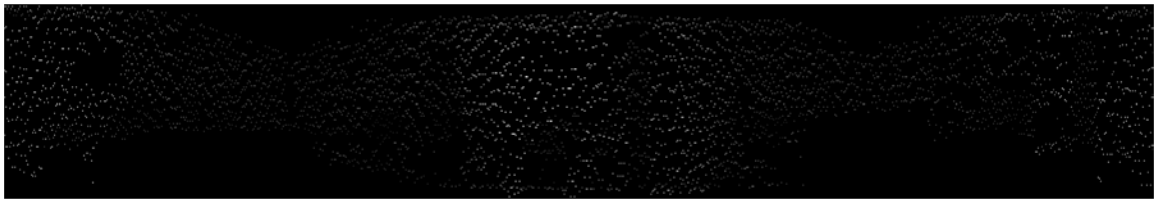

(a) Regularly subsampled point cloud

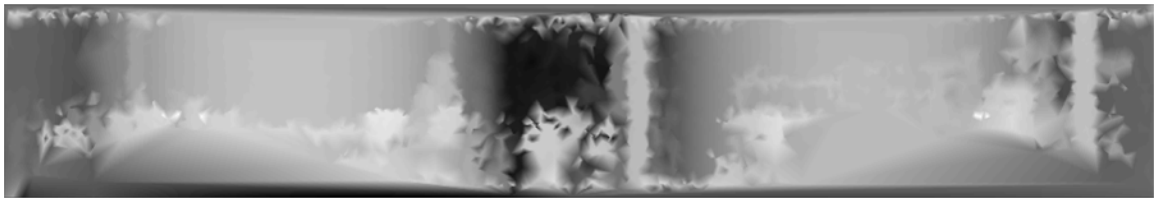

(b) Nearest neighbor interpolation

(c) CS reconstruction

Figure 5: Reconstruction results for the indoor lab space using 10% regularly spaced downsampling of the 2.5% Lissajous scan shown in Fig.3d (for a net downsampling of 0.25%). (a) Sub-sampled data used in the reconstructions. (b) Nearest neighbor interpolation. (c) CS reconstruction. These should be compared to the ground truth in 4b.

appear, at least roughly, in the raster scan in (b) are completely absent in the subsampled data.

Results from the 0.25% sampling data are shown in Fig. 8 for the regularly sub-sampled set and in Fig. 9 for the randomly sub-sampled data. Both results are similar and clearly much poorer than the reconstructions based on 2.5% data.

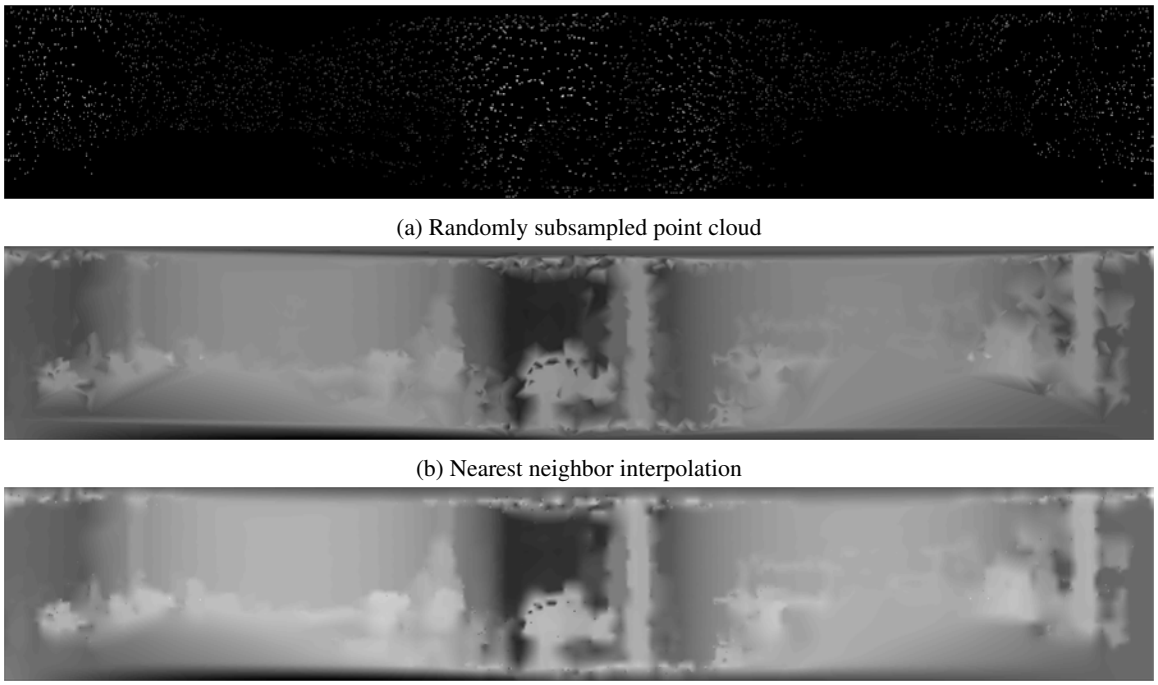

(a) Randomly subsampled point cloud

(b) Nearest neighbor interpolation

(c) CS reconstruction

Figure 6: Reconstruction results for the indoor lab space using 10% randomly selected downsampling of the 2.5% Lissajous scan shown in Fig.3d (for a net downsampling of 0.25%). (a) Sub-sampled data used in the reconstructions. (b) Nearest neighbor interpolation. (c) CS reconstruction. These should be compared to the ground truth in 4b.

DISCUSSION

These results, while admittedly qualitative, are nonetheless interesting in that reasonably good depth images of the indoor are generated even for very low levels of subsampling. While there is loss of detail as one goes from full raster to the 2.5% Lissajous scan, and again when going down to 0.25% sampling, major features are preserved. This could allow, for example, for adaptive sampling that uses a high speed, very low resolution scan to get initial information followed by detailed, regional scans to fill out details in areas of interest.

The results in the outdoor setting are of lower quality, especially at the lowest sub-sampling. This is driven not only by the low sampling but also by the fact that the outdoor environment does not meet the assumption that the depth image is built mainly from planes with only a few edges. Since the CS reconstruction algorithm relies on this, it is not surprising that the resulting images are of low quality.

CONCLUSIONS

The use of CS reconstruction algorithms for Lissajous nodding Lidar was presented. This method allows for very small small amounts of data, and thus very small data collection times, to produce high fidelity depth data. Experimental and simulation results show the feasibility of this approach in both indoor and outdoor scenarios.

ACKNOWLEDGMENT

The work of Sathishchandra and Andersson was supported in part by the NSF through CMMI-1562031. The work of Benson and Clayton was supported in part by the NSF through OISE-1658696.

REFERENCES

[1] M. Maimone, Y. Cheng, and L. Matthies. Two years of visual odometry on the mars exploration rovers. *Journal of Field Robotics*, 24(3):169–186, 2007.

[2] D. Scaramuzza. Performance evaluation of 1-point-ransac visual odometry. *Journal of Field Robotics*, 95(1):792–811, July 2011.

[3] F. Fraundorfer and D. Scaramuzza. Visual odometry: Part ii: Matching, robustness, optimization, and applications. *Robotics & Automation Magazine, IEEE*, 19(2):78–90, 2012.

[4] Stefan Kohlbrecher, Oskar Von Stryk, Johannes Meyer, and Uwe Klingauf. A flexible and scalable slam system with

(a) Outdoor panorama

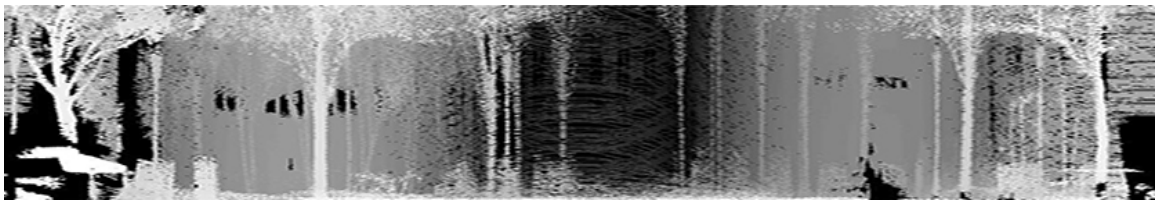

(b) Raster scan

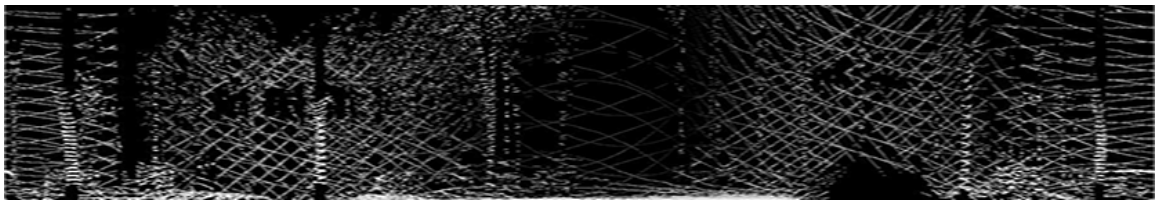

(c) 4Hz Lissajous scan

(d) Nearest neighbor interpolation

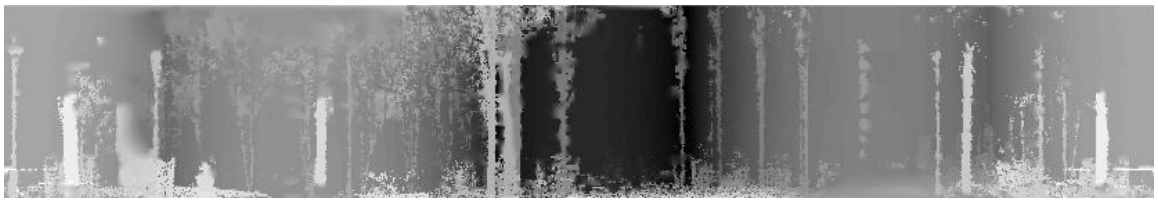

(e) CS reconstruction

Figure 7: Results for the outdoor courtyard scan showing a) An image of the courtyard space for reference, b) A high resolution raster scan, c) The 2.5% Lissajous scan, d) Reconstruction results for nearest neighbor interpolation and e) Reconstruction results for the compressive sensing reconstruction algorithm.

full 3d motion estimation. In *2011 IEEE International Symposium on Safety, Security, and Rescue Robotics*, pages 155–160. IEEE, 2011.

[5] Jorge Fuentes-Pacheco, José Ruiz-Ascencio, and Juan Manuel Rendón-Mancha. Visual simultaneous localization and mapping: a survey. *Artificial Intelligence Review*, 43(1):55–81, 2015.

[6] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-time loop closure in 2d lidar slam. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1271–1278. IEEE, 2016.

[7] David Geller. Orbital rendezvous: When is autonomy required? *Journal of guidance, control, and dynamics*, 30(4):974–981, 2007.

[8] Farzin Amzajerdian, Michael Vanek, Larry Petway, Diego Pierrottet, George Busch, and Alexander Bulyshev. Utiliza-

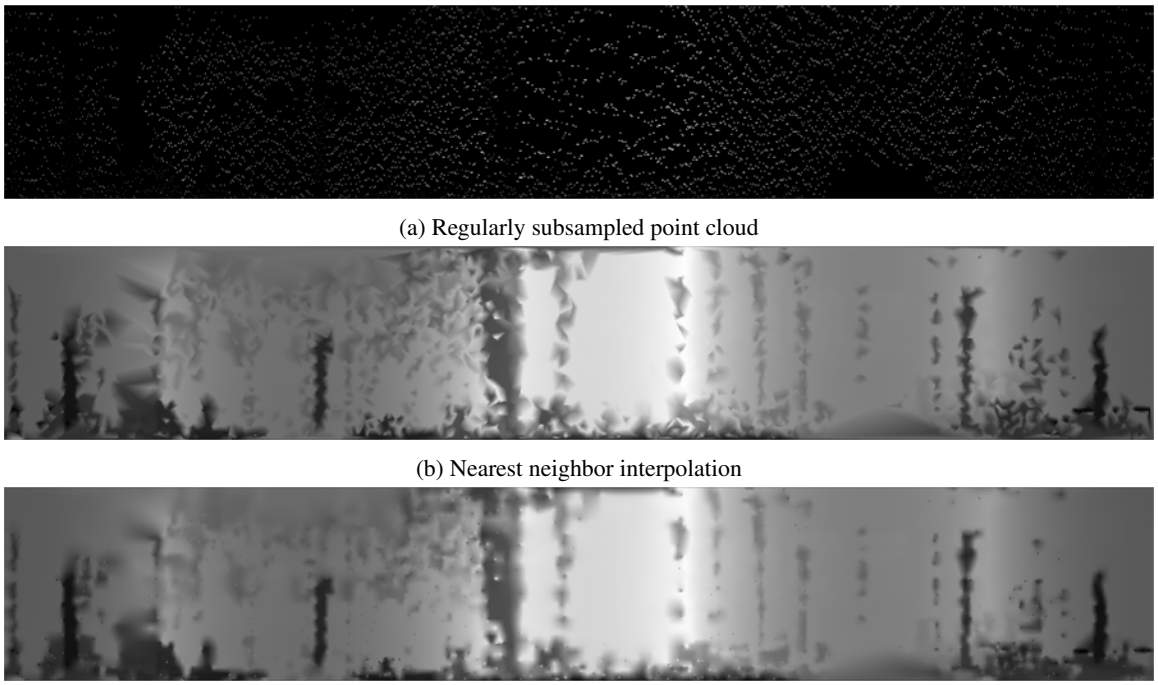

(a) Regularly subsampled point cloud

(b) Nearest neighbor interpolation

(c) CS reconstruction

Figure 8: Reconstruction results for the outdoor courtyard scene using 10% regularly spaced downsampling of the 2.5% Lissajous scan shown in Fig.3d (for a net downsampling of 0.25%). (a) Sub-sampled data used in the reconstructions. (b) Nearest neighbor interpolation. (c) CS reconstruction. These should be compared to the ground truth in 7b.

tion of 3d imaging flash lidar technology for autonomous safe landing on planetary bodies. In *Quantum Sensing and Nanophotonic Devices VII*, volume 7608, page 760828. International Society for Optics and Photonics, 2010.

[9] Kirk MacTavish and Timothy D Barfoot. Towards hierarchical place recognition for long-term autonomy. In *ICRA Workshop on Visual Place Recognition in Changing Environments*, pages 1–6, 2014.

[10] Velodyne LiDAR. Velodyne lidar. `https://velodynelidar.com/`.

[11] Ouster, Inc. Ouster lidar. `https://www.ouster.io/`.

[12] Sungsik Huh, David Hyunchul Shim, and Jonghyuk Kim. Integrated navigation system using camera and gimbaled laser scanner for indoor and outdoor autonomous flight of uavs. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 3158–3163. IEEE, 2013.

[13] Jesús Morales, Jorge L Martinez, Anthony Mandow, Alejandro Pequeño-Boter, and Alfonso Garcia-Cerezo. Design and development of a fast and precise low-cost 3d laser rangefinder. In *Mechatronics (ICM), 2011 IEEE International Conference on*, pages 621–626. IEEE, 2011.

[14] Oliver Wulf and Bernardo Wagner. Fast 3d scanning methods for laser measurement systems. In *International conference on control systems and computer science (CSCS14)*, pages 2–5, 2003.

[15] Ankit Desai and Daniel Huber. Objective evaluation of scanning ladar configurations for mobile robots. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 2182–2189. IEEE, 2009.

[16] Tomoaki Yoshida, Kiyoshi Irie, Eiji Koyanagi, and Masahiro Tomono. 3d laser scanner with gazing ability. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 3098–3103. IEEE, 2011.

[17] Hartmut Surmann Simone Frintrop, Andres Nuchter. Visual attraction for object recognition in spatial 3d data. 2004.

[18] Michael Benson, Jonathan Nikolaidis, and Garrett M Clayton. Lissajous-like scan pattern for a nodding multi-beam lidar. In *ASME 2018 Dynamic Systems and Control Conference*, pages V002T24A007–V002T24A007. American Society of Mechanical Engineers, 2018.

[19] Yufan Luo and Sean B Andersson. A continuous sampling pattern design algorithm for atomic force microscopy images. *Ultramicroscopy*, 196:167–179, January 2019.

[20] Yufan Luo and Sean B Andersson. A Compressive

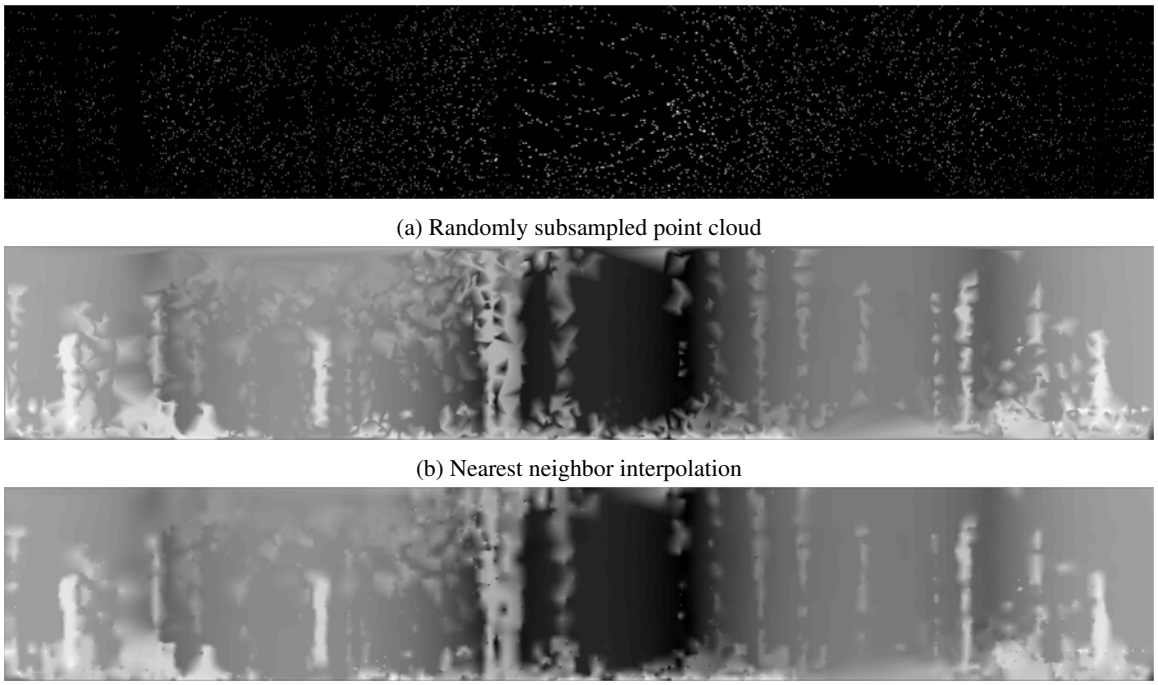

(a) Randomly subsampled point cloud

(b) Nearest neighbor interpolation

(c) CS reconstruction

Figure 9: Reconstruction results for the outdoor courtyard scene using 10% regularly spaced downsampling of the 2.5% Lissajous scan shown in Fig.3d (for a net downsampling of 0.25%). (a) Sub-sampled data used in the reconstructions. (b) Nearest neighbor interpolation. (c) CS reconstruction. These should be compared to the ground truth in 7b.

Sensing-Based Pixel Sharing Algorithm for High-Speed Atomic Force Microscopy. In *IEEE Conference on Decision and Control*, pages 2834–2839, December 2016.

[21] Yufan Luo and Sean B Andersson. A comparison of reconstruction methods for undersampled atomic force microscopy images. *Nanotechnology*, 26(50):505703, December 2015.

[22] Fangchang Ma, Luca Carlone, Ulas Ayaz, and Sertac Karaman. Sparse sensing for resource-constrained depth reconstruction. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 96–103. IEEE, 2016.

[23] Gianfranco Forlani, Carla Nardinocchi, Marco Scaioni, and Primo Zingaretti. Complete classification of raw lidar data and 3d reconstruction of buildings. *Pattern analysis and applications*, 8(4):357–374, 2006.

[24] Zheng Wang and Tony Schenk. Building extraction and reconstruction from lidar data. *International Archives of Photogrammetry and Remote Sensing*, 33(B3/2; PART 3):958–964, 2000.

[25] Aparajithan Sampath and Jie Shan. Segmentation and reconstruction of polyhedral building roofs from aerial lidar point clouds. *IEEE Transactions on geoscience and remote sensing*, 48(3):1554–1567, 2009.

[26] Emmanuel J Candès. Compressive sampling. In *International Congress of Mathematicians*, 2006.

[27] Emmanuel J Candès and Justin Romberg. Sparsity and incoherence in compressive sampling. *Inverse Problems*, 23(3):969–985, April 2007.