

MIT Open Access Articles

Sparse sensing for resource-constrained depth reconstruction

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Ma, Fangchang, Luca Carlone, Ulas Ayaz, and Sertac Karaman. "Sparse Sensing for Resource-Constrained Depth Reconstruction." 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (October 2016).

As Published: <http://dx.doi.org/10.1109/IROS.2016.7759040>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <http://hdl.handle.net/1721.1/116275>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Sparse Sensing for Resource-Constrained Depth Reconstruction

Fangchang Ma, Luca Carlone, Ulas Ayaz, Sertac Karaman

Abstract—We address the following question: *is it possible to reconstruct the geometry of an unknown environment using sparse and incomplete depth measurements?* This problem is relevant for a resource-constrained robot that has to navigate and map an environment, but does not have enough on-board power or payload to carry a traditional depth sensor (e.g., a 3D lidar) and can only acquire few (point-wise) depth measurements. In general, reconstruction from incomplete data is not possible, but when the robot operates in man-made environments, the depth exhibits some regularity (e.g., many planar surfaces with few edges); we leverage this regularity to infer depth from incomplete measurements. Our formulation bridges robotic perception with the compressive sensing literature in signal processing. We exploit this connection to provide formal results on exact depth recovery in 2D and 3D problems. Taking advantage of our specific sensing modality, we also prove novel and more powerful results to completely characterize the geometry of the signals that we can reconstruct. Our results directly translate to practical algorithms for depth reconstruction; these algorithms are simple (they reduce to solving a linear program), and robust to noise. We test our algorithms on real and simulated data, and show that they enable accurate depth reconstruction from a handful of measurements, and perform well even when the assumption of structured environment is violated.

I. INTRODUCTION

In the last two decades, robot perception witnessed dramatic advances, leading to many working solutions that are steadily transitioning to industrial practice. A large body of research focused on the development of techniques to perform inference from data produced by “information-rich” sensors (e.g., high-resolution cameras, 2D and 3D laser scanners). A variety of approaches has been proposed to perform geometry reconstruction using these sensors, see [1], [2], [3] and the references therein. On the other extreme of the sensor spectrum, applications and theory have been developed to cope with the case of minimalistic sensing [4], [5], [6], [7]. In this latter case the sensor data is not metric (e.g., the sensor cannot measure distances or angles), but is more often binary in nature (e.g., binary detection of landmarks), and the goal is to infer the topology of the (usually planar) environment rather than its geometry.

This work investigates a relatively unexplored region between these two extremes of the sensor spectrum. Our goal is to estimate a depth profile (i.e., a laser scan in 2D, or a depth image in 3D, see Fig. 1) from sparse and incomplete measurements. Contrarily to the literature on minimalistic sensing, we provide tools to recover complete geometric information, while requiring much fewer data compared to standard high-resolution depth sensors.

Our interest towards depth estimation from incomplete measurements is motivated by navigation of resource-

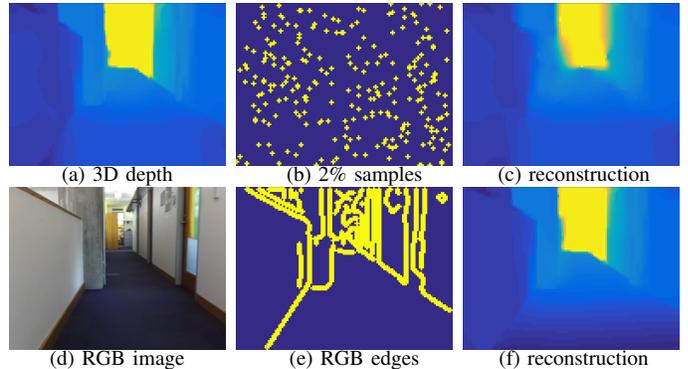


Fig. 1. We show how to reconstruct an unknown depth profile (a) from a handful of samples (b). Our reconstruction is shown in (c). Our results also apply to traditional stereo vision and enable accurate reconstruction (f) from few depth measurements (e) corresponding to the edges in the RGB image (d). Figures (a) and (d) are obtained from a ZED stereo camera.

constrained robots that do not have enough on-board power or payload to carry traditional sensors. Our overarching goal is two-fold. First, we want to establish theoretical conditions under which depth reconstruction is possible. Second, we want to develop practical inference algorithms for depth estimation. The combination of these two results would enable, for instance, advanced navigation of miniaturized robots (e.g., the *robot bee* [8]), among other applications.

The first question to answer is: *is it really possible to reconstruct a depth profile from incomplete information?* In general, the answer is negative, since the depth can be very adversarial (e.g., 2D laser scan in which each beam is drawn from a uniform distribution), and we would not be able to recover the depth from a small set of measurements. Fortunately, when the robot operates in structured environments (e.g., indoor, urban scenarios) the depth data exhibits some regularity. Intuitively, man-made scenarios include many planar surfaces with few edges and corners. This work shows how to leverage this regularity to recover a depth profile from a handful of measurements.

Related work. Our study of depth reconstruction from sparse sensor data is related to the literature on minimalistic sensing. Suri *et al.* [4], Derenick *et al.* [5], and Tovar *et al.* [7] use binary measurements of the presence of landmarks, to infer the topology of the environment. Marinakis *et al.* [9] reconstruct the topology of a sensor network from unlabeled observations. O’Kane *et al.* [10] and Erickson *et al.* [11] investigate localization from contact sensors. Milford [12] addresses minimality in vision-based place recognition.

Our approach is also motivated by the recent interest in fast perception and dense 3D reconstruction. The idea of leveraging priors on the structure of the environment has been investigated in early work in computer vision for single-view 3D reconstruction and feature matching [13], [14]. More recently, Pillai *et al.* [15] propose an approach to

F. Ma, L. Carlone, U. Ayaz, and S. Karaman are with the Laboratory for Information & Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, USA, {fcmal, lcarlone, uayaz, sertac}@mit.edu

speed-up stereo reconstruction by computing the disparity at a small set of pixels. Piniés *et al.* [16] propose an approach to compute a dense depth map from a sparse point cloud. The latter work is related to our proposal with three main differences. First, the work [16] uses an energy minimization approach which requires parameter tuning (the authors use Bayesian optimization to learn such parameter); our approach is parameter free and only assumes bounded noise. Second, we use a 2nd-order difference operator to promote depth regularity, while [16] considers alternative costs, including nonconvex regularizers. Finally, by recognizing connections with the *cosparsity* model in compressive sensing, we provide theoretical foundations for the reconstruction problem.

Finally, our work is related to the literature on *compressive sensing* (CS) [17]. CS revolutionized signal processing by showing that a signal can be reconstructed from a small set of measurements if it is *sparse* in some domain. CS results impacted many research areas, including image processing [18]), data compression, and 3D reconstruction [19]. Most CS literature assumes that the signal to recover (e.g., z) is sparse: this setup is called the *synthesis model*. Very recent work considers the case in which the signal becomes sparse after a transformation is applied (i.e., given a matrix D , the vector Dz is sparse). The latter setup is called *the analysis (or cosparsity) model* [20], [21].

Contribution. Section III formulates our problem and presents our first contribution. Here we recognize that the “regularity” of a depth profile is captured by a specific function (the ℓ_0 -norm of the 2nd-order differences of the depth profile). We also show that by relaxing the ℓ_0 -norm to the (convex) ℓ_1 -norm, our problem falls within the cosparsity model in CS. Section IV contains the core contribution of the paper. We provide formal conditions for exact depth recovery. We also show that when exact recovery is not possible, we can fully describe the geometry of the depth profiles produced by our approach. Section V provides practical algorithms for depth reconstruction. Section VI reports experimental results on simulated and real data. The experiments confirm our theoretical findings and show that the proposed approach is extremely resilient to noise and works well even when the regularity assumptions are violated. While our main motivation is sparse sensing, we show that our results have potential to reduce the computational load associated with processing of data from standard sensors (e.g., a stereo camera). Section VII discusses future research. All proofs, together with extra visualizations, are given in the supplemental material [22]. We use the notation “SM-2.1” to recall a specific section (Section 2.1 in the example) in [22].

II. PRELIMINARIES AND NOTATION

We use uppercase letters for matrices, e.g., D , and lowercase letters for vectors and scalars, e.g., $z \in \mathbb{R}^n$. Sets are denoted with calligraphic fonts, e.g., \mathcal{M} . The cardinality of a set \mathcal{M} is denoted with $|\mathcal{M}|$. For a set \mathcal{M} , the symbol $\bar{\mathcal{M}}$ denotes its complement. For a vector $z \in \mathbb{R}^n$ and a set of indices \mathcal{M} , $z_{\mathcal{M}}$ is the subvector of z corresponding to the entries of z with indices in \mathcal{M} . In particular, z_i is the i -th entry. The symbols $\mathbf{1}$ (resp. $\mathbf{0}$) denote a vector of all ones

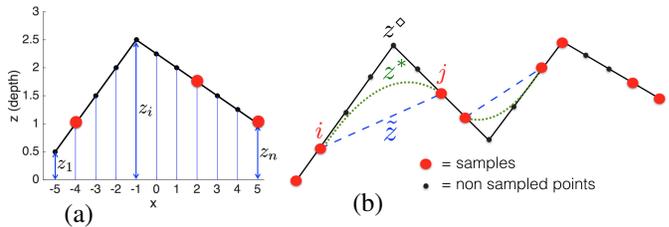


Fig. 2. (a) Example of 2D depth profile. Our goal is to reconstruct the full signal (black solid line) from sparse samples (red dots); (b) When we do not sample the corners and the neighboring points, (L1_D) admits multiple minimizers, which obey the conditions of Proposition 14.

(resp. zeros). The support set of a vector is denoted with

$$\text{supp}(z) = \{i \in \{1, \dots, n\} : z_i \neq 0\}.$$

We denote with $\|z\|_2$ the Euclidean norm and we also use the following norms: $\|z\|_\infty \doteq \max_{i=1, \dots, n} |z_i|$ (ℓ_∞ -norm); $\|z\|_1 \doteq \sum_{i=1, \dots, n} |z_i|$ (ℓ_1 -norm); $\|z\|_0 \doteq |\text{supp}(z)|$ (ℓ_0 -norm, i.e., the number of nonzero elements in z). The sign vector $\text{sign}(z)$ of $z \in \mathbb{R}^n$ is a vector with entries: $\text{sign}(z)_i = 1$ if $z_i > 0$, $\text{sign}(z)_i = -1$ if $z_i < 0$, and zero otherwise.

For a matrix D and an index set \mathcal{M} , we denote with $D_{\mathcal{M}}$ the submatrix of D containing only the rows of D with indices in \mathcal{M} . Sometimes, given two sets \mathcal{I} and \mathcal{J} , we also use the notation $D_{\mathcal{I}, \mathcal{J}}$ which is a sub-matrix of D including only rows \mathcal{I} and columns \mathcal{J} . The identity matrix is denoted with \mathbf{I} . Given a matrix $A \in \mathbb{R}^{p \times n}$, we use the following matrix operator norm $\|A\|_{\infty \rightarrow \infty} \doteq \max_{i, \dots, p} \|A_i\|_1$.

In the rest of the paper we use the *cosparsity model* where we assume that the application of an *analysis operator* D produces a vector with small number of nonzero entries. The following definitions formalize this concept.

Definition 1 (Cosparsity). *A vector $z \in \mathbb{R}^n$ is said to be cosparse with respect to a matrix $D \in \mathbb{R}^{p \times n}$ if $\|Dz\|_0 \ll p$.*

Definition 2 (D -support and D -cosupport). *Given a vector $z \in \mathbb{R}^n$ and a matrix $D \in \mathbb{R}^{p \times n}$, the D -support of z is the set of indices corresponding to the nonzero entries of Dz , i.e., $\mathcal{I} = \text{supp}(Dz)$. The D -cosupport is the complement of \mathcal{I} , i.e., the indices of the zero entries of Dz .*

III. PROBLEM FORMULATION: RESOURCE-CONSTRAINED DEPTH ESTIMATION

We want to reconstruct 2D depth profiles (i.e., a scan from a 2D laser range finder) and 3D depth profiles (i.e., a depth image produced by a kinect or a stereo camera) from partial and incomplete measurements. We formalize this problem as follows, by first considering the 2D and the 3D cases separately and then reconciling them under a unified framework.

2D Depth Reconstruction. Here we want to recover a depth profile $z^\diamond \in \mathbb{R}^n$. One can imagine that the vector z^\diamond includes (unknown) depth measurements at discrete angles; this is what a standard planar range finder would measure.

In our problem, due to sensing constraints, we do not have direct access to z^\diamond , but we measure

$$y = Az^\diamond + \eta \quad (1)$$

where the matrix $A \in \mathbb{R}^{m \times n}$ with $m \ll n$ is the *measurement matrix*, and η represents measurement noise. The structure of A is formalized in the following definition.

Definition 3 (Sample set and sparse sampling matrix). A sample set $\mathcal{M} \subseteq \{1, \dots, n\}$ is the set of entries of the profile that are measured. A matrix $A \in \mathbb{R}^{m \times n}$ is called a (sparse) sampling matrix (with sample set \mathcal{M}), if $A = \mathbf{I}_{\mathcal{M}}$.

Recall that $\mathbf{I}_{\mathcal{M}}$ is the subset of rows of the identity matrix at indices \mathcal{M} . It follows that $Az = z_{\mathcal{M}}$, i.e., the matrix A selects a subset of entries of z . Since $m \ll n$, we have much fewer measurements than unknowns, hence z^\diamond cannot be recovered from y , without further assumptions.

We assume that the profile is z^\diamond sufficiently regular, which means that it contains few ‘‘corners’’, e.g., Fig. 2(a). Corners are produced by changes of slope: considering 3 consecutive points at coordinates $[x_{i-1}, z_{i-1}]$, $[x_i, z_i]$, and $[x_{i+1}, z_{i+1}]$, there is a corner at i if $\frac{z_{i+1}-z_i}{x_{i+1}-x_i} - \frac{z_i-z_{i-1}}{x_i-x_{i-1}} \neq 0$. In the following we assume that $x_i - x_{i-1} = 1$ for all i : this comes without loss of generality since the full profile is unknown and we can reconstruct it at arbitrary resolution (i.e., at arbitrary x). The definition of corner hence becomes:

Definition 4 (Corner set). Given a 2D profile $z \in \mathbb{R}^n$, the corner set $\mathcal{C} \subseteq \{2, \dots, n-1\}$ is the set of indices i such that $z_{i-1} - 2z_i + z_{i+1} \neq 0$.

Intuitively, $z_{i-1} - 2z_i + z_{i+1}$ is the discrete equivalent of the 2nd-order derivative at z_i . We call $z_{i-1} - 2z_i + z_{i+1}$ the *curvature* at sample i : if this quantity is zero, the neighborhood of i is flat (the three points are collinear); if it is negative, the curve is locally concave; if it is positive, it is locally convex. To make notation more compact, we introduce the 2nd-order difference operator:

$$D \doteq \begin{bmatrix} 1 & -2 & 1 & 0 & \dots & 0 \\ 0 & 1 & -2 & 1 & \dots & 0 \\ \vdots & 0 & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 & -2 & 1 \end{bmatrix} \in \mathbb{R}^{(n-2) \times n} \quad (2)$$

Then a profile with few corners is one for which Dz^\diamond is sparse. In fact, the ℓ_0 -norm of Dz^\diamond counts exactly the number of corners of a profile: $\|Dz^\diamond\|_0 = |\mathcal{C}|$.

When operating in indoor environments, it is reasonable to assume that z^\diamond has few corners. Therefore, we want to exploit this regularity assumption and the partial measurements y in (1) to reconstruct z^\diamond . Let us start from the case in which $\eta = \mathbf{0}$ in (1). In this case, a reasonable way to reconstruct the profile z^\diamond is to solve the following optimization problem:

$$\min_z \|Dz\|_0 \quad \text{subject to } Az = y \quad (\text{L0})$$

which looks for the profile z that is consistent with the measurements (1) and contains the smallest number of corners. Unfortunately, problem (L0) is NP-hard. In this work we study the following relaxation of problem (L0):

$$\min_z \|Dz\|_1 \quad \text{subject to } Az = y \quad (\text{L1}_D)$$

which can be rephrased as a linear program, and can be solved efficiently. Sections IV provides conditions under which (L1_D) recovers the solution of (L0). Problem (L1_D) falls in the class of the cosparsity models in CS [21].

In presence of measurement noise (1), and assuming bounded error $\|\eta\|_\infty \leq \varepsilon$, the ℓ_1 -minimization becomes:

$$\min_z \|Dz\|_1 \quad \text{subject to } \|Az - y\|_\infty \leq \varepsilon \quad (\text{L1}_D^\varepsilon)$$

We assume that the infinity norm of the noise η is bounded, since this naturally reflects the sensor model in our robotic application (i.e., bounded error in each laser beam).

3D Depth Reconstruction. Here we want to recover a 3D depth profile $Z^\diamond \in \mathbb{R}^{r \times c}$ (a depth map, as the one in Fig. 1(a)), using incomplete measurements. As in the 2D setup, we do not have direct access to Z^\diamond , but we can only take (point-wise) $m \ll r \times c$ measurements in the form:

$$y_{i,j} = Z_{i,j}^\diamond + \eta_{i,j} \quad (3)$$

where $\eta_{i,j} \in \mathbb{R}$ represents measurement noise. Each measurement is a noisy sample of the depth of Z^\diamond at pixel (i, j) .

We assume that Z^\diamond is sufficiently regular, which intuitively means that the depth profile contains many planar regions and few ‘‘edges’’. We define the edges as follows.

Definition 5 (Edge set). Given a 3D profile $Z \in \mathbb{R}^{r \times c}$, the vertical edge set $\mathcal{E}_V \subseteq \{2, \dots, r-1\} \times \{1, \dots, c\}$ is the set of indices (i, j) such that $Z_{i-1,j} - 2Z_{i,j} + Z_{i+1,j} \neq 0$. The horizontal edge set $\mathcal{E}_H \subseteq \{1, \dots, r\} \times \{2, \dots, c-1\}$ is the set of indices (i, j) such that $Z_{i,j-1} - 2Z_{i,j} + Z_{i,j+1} \neq 0$. The edge set \mathcal{E} is the union of the two sets: $\mathcal{E} \doteq \mathcal{E}_V \cup \mathcal{E}_H$.

Intuitively, (i, j) is not in the edge set if the 3×3 patch centered at (i, j) is planar, while $(i, j) \in \mathcal{E}$ otherwise. As in the 3D case we introduce 2nd-order difference operators to compute the vertical differences $Z_{i,j-1} - 2Z_{i,j} + Z_{i,j+1}$ and the horizontal differences $Z_{i-1,j} - 2Z_{i,j} + Z_{i+1,j}$:

$$D_V Z^\diamond \in \mathbb{R}^{(r-2) \times c}, \quad Z^\diamond D_H^T \in \mathbb{R}^{r \times (c-2)} \quad (4)$$

where the matrices D_V and D_H are the same as the one defined (2), but with suitable dimensions; each entry of the matrix $D_V Z^\diamond$ contains the vertical (2nd-order) differences at a pixel, while $Z^\diamond D_H^T$ collects the horizontal differences.

Following the same reasoning of the 2D case, we obtain the following ℓ_1 -norm minimization

$$\min_z \left(\| \text{vec}(D_V Z) \|_1 + \| \text{vec}(Z D_H^T) \|_1 \right) \quad (5)$$

subject to $Z_{i,j} = y_{i,j}$ for each measured pixel (i, j)

where $\text{vec}(M)$ denotes the (column-wise) vectorization of a matrix M , and we assumed noiseless measurements.

To establish a counterpart of the 2D case in 3D, we define $n \doteq r \times c$, and the vectorized version of the 3D depth image: $z^\diamond \doteq \text{vec}(Z^\diamond) \in \mathbb{R}^n$. Moreover, we stack all measurements in a vector $y \in \mathbb{R}^m$. Using standard properties of the vectorization operator, we can rewrite problem (5) as:

$$\min_z \|\Delta z\|_1 \quad \text{subject to } Az = y \quad (\text{L1}_\Delta)$$

where the matrix A is the same of Definition 3 (with sample set $\mathcal{M} \subseteq \{1, \dots, n\}$) and the ‘‘regularization’’ matrix Δ is a suitable sparse matrix (full derivation is given in SM-3.1).

In presence of noise, we define an error vector $\eta \in \mathbb{R}^m$ which stacks the noise terms in (3) for each (i, j) , and assume $\|\eta\|_\infty \leq \varepsilon$. The noisy ℓ_1 -minimization becomes:

$$\min_z \|\Delta z\|_1 \quad \text{subject to } \|Az - y\|_\infty \leq \varepsilon \quad (\text{L1}_\Delta^\varepsilon)$$

Comparing $(L1_D)$, $(L1_\Delta)$, and $(L1_D^\varepsilon)$, $(L1_\Delta^\varepsilon)$, it is clear that in 2D and 3D we solve the same optimization problems, and the only difference lies in the matrices D and Δ .

IV. DEPTH RECOVERY FROM SPARSE SENSING

This section contains the key technical results of this paper. Section IV-A provides sufficient analytic and geometric conditions for exact recovery of a depth profile. Section IV-B completely characterizes the solution set of problems $(L1_D)$ and $(L1_\Delta)$, which will guide algorithm design (Section V).

A. Sufficient Conditions for Exact Recovery

Recent results on cosparsity in compressive sensing provide sufficient conditions for exact recovery of a cosparsity profile z^\diamond , from measurements $y = Az^\diamond$ (Proposition 6 below holds for a generic A). After presenting the result we discuss why this is not the condition a roboticist would like to use.

Proposition 6 (Exact Recovery [20]). *Consider a vector $z^\diamond \in \mathbb{R}^n$ with D -support \mathcal{I} and D -cosupport \mathcal{J} . Define $\bar{m} \doteq n - m$. Let $N \in \mathbb{R}^{\bar{m} \times n}$ be a matrix whose rows span the null space of the matrix A . Let $(\cdot)^\dagger$ denote the Moore-Penrose pseudoinverse of a matrix. If the following holds:*

$$C_{er} \doteq \|(N(D_{\mathcal{J}})^\dagger)^\dagger N(D_{\mathcal{I}})^\dagger\|_{\infty \rightarrow \infty} < 1 \quad (6)$$

then problem $(L1_D)$ recovers z^\diamond exactly.

Despite being very general, Proposition 6 provides an algebraic condition. In our depth estimation problem, we would rather like geometric conditions that suggest an optimal sampling strategy to recover the profile. Our contribution of this section is to rephrase Proposition 6 in geometric terms.

We provide a first result for the 2D case.

Proposition 7 (Exact Recovery of 2D depth profiles). *Let $z^\diamond \in \mathbb{R}^n$ be a 2D depth signal with corner set \mathcal{C} . Assuming noiseless measurements (1), the following hold:*

- (i) *if the sampling set \mathcal{M} is the union of the corner set and the first and last entries of z^\diamond , then $C_{er} = 1$;*
- (ii) *if the sampling set \mathcal{M} includes the corners and their neighbors (adjacent entries), then $C_{er} = 0$ and problem $(L1_D)$ recovers z^\diamond exactly.*

Proposition 7 shows that we can recover the original profile if we measure the neighborhood of each corner. When we sample only the corners, intuitively, one might still hope to recover the profile z^\diamond , since the condition $C_{er} < 1$ is only sufficient for exact recovery. However, it turns out that in our problem, one can find counterexamples with $C_{er} = 1$ in which ℓ_1 -minimization fails to recover z^\diamond (SM-2.2).

We derive a similar condition in 3D problems.

Proposition 8 (Exact Recovery of 3D depth profiles). *Let Z^\diamond be a 3D depth signal with edge set \mathcal{E} . Assume noiseless measurements. If the sampling set \mathcal{M} includes the edges and their (vertical and horizontal) neighbors (adjacent pixels), then $C_{er} = 0$, and $(L1_\Delta)$ recovers $\text{vec}(Z^\diamond)$ exactly.*

In the experiments, we show that these initial results already unleash interesting applications. For instance, in stereo vision problems, we could locate the position of edges from the RGB images and recover the depth in a neighborhood of the edge pixels. Then, the complete depth profile can be recovered (at arbitrary resolution) via $(L1_\Delta)$.

B. Algebraic Optimality Conditions

The exact recovery conditions of Proposition 7 and 8 are quite restrictive if we do not have prior knowledge of the position of the corners or edges. Now we provide more powerful results that do not require sampling corners/edges. Empirically, we observe that without sampling all edges, the optimization problems $(L1_D)$ and $(L1_\Delta)$ admit multiple solutions, i.e., multiple signals z attain the optimal cost. We address the following questions in this section: *what is the optimal solution set \mathcal{S}^* of problems $(L1_D)$ and $(L1_\Delta)$? Is the ground truth signal z^\diamond within this optimal solution set?*

In this section, we derive a general algebraic condition for a 2D signal (resp. 3D) to be in the solution set of $(L1_D)$ (resp. $(L1_\Delta)$). Section IV-C and Section IV-D translate this algebraic condition into a geometric constraint on the curvature of the signals in the solution set.

Proposition 9 (2D Optimality). *Let A be the sampling matrix and \mathcal{M} be the sample set. Given a profile $z \in \mathbb{R}^n$ which is feasible for $(L1_D)$, z is a minimizer of $(L1_D)$ if and only if there exists a vector $u \in \mathbb{R}^{n-2}$ such that*

$$(D^\top)_{\bar{\mathcal{M}}} u = \mathbf{0} \text{ and } u_{\mathcal{I}} = \text{sign}(Dz)_{\mathcal{I}} \text{ and } \|u\|_\infty \leq 1 \quad (7)$$

where $\bar{\mathcal{M}}$ is the set of entries of z that we do not sample (i.e., the complement of \mathcal{M}).

The proof is based on the subdifferential of the ℓ_1 -minimization problem. An analogous result holds in 3D.

Corollary 10 (3D optimality). *A given profile Z is in the set of minimizers of $(L1_\Delta)$ if and only if the conditions of Proposition 9 hold, using Δ instead of D in (7).*

C. Geometric Conditions for Optimality in 2D

In this section we derive necessary and sufficient geometric conditions for z^\diamond to be in the solution set of $(L1_D)$. Using these findings we obtain two practical results: (i) we can bound how far any solution z^* of $(L1_D)$ is from the ground truth signal z^\diamond ; (ii) we can design a general algorithm that recovers z^\diamond also when the conditions of Proposition 7 fail.

To introduce our results, we need the following definition.

Definition 11 (2D Sign Consistency). *Let $s_k = \text{sign}(z_{k-1} - 2z_k + z_{k+1})$ (sign of the curvature at k). A 2D depth signal z is sign consistent if, for any two consecutive samples $i < j \in \mathcal{M}$, one of the two conditions holds:*

- (i) *no sign change: for any two integers k, h , with $i \leq k, h \leq j$, if $s_k \neq 0$ and $s_h \neq 0$, then $s_k = s_h$;*
- (ii) *sign change only at the boundary: for any integer k , with $i < k < j$, $s_k = 0$;*

While the definition is quite technical, it has a clear geometric interpretation. A signal z is sign consistent, if its curvature does not change sign (i.e., it is either convex or concave) within each interval between consecutive samples (see SM-2.3 for an example).

It turns out that picking pairs of consecutive samples makes it easier to recover the depth profile and to analyze the recovery performance (Theorem 13 below). Therefore we define the notion of ‘‘twin samples’’.

Definition 12 (Twin samples). *A twin sample is a pair of consecutive samples, e.g. $(i, i + 1)$ with $i \in \{1, \dots, n - 1\}$.*

Theorem 13 (2D Sign Consistency \equiv Optimality). *Let z be a 2D signal which is feasible for problem (L1_D). Assume that the sample set includes only twin samples and we sample the “boundary” of the signal, i.e., z_1 , and z_n . Then, z is optimal for (L1_D) if and only if it is sign consistent.*

Theorem 13 provides a tight geometric condition for a signal to be in the optimal solution set, i.e., when our ground truth signal is among the minimizers of (L1_D).

Proposition 14 (z^\diamond and 2D Optimal Solution Set). *Let z^\diamond be the ground truth generating noiseless measurements (1). Assume that we sample the boundary of z^\diamond and the sample set includes a twin sample in each linear segment in z^\diamond . Then, z^\diamond is in the set of minimizers of (L1_D). Moreover, denote with \tilde{z} the naive solution obtained by connecting consecutive samples with a straight line (linear interpolation). Then, any optimal solution z^* lies between z^\diamond and \tilde{z} , i.e., for any index $i \in \{1, \dots, n\}$, it holds that if $z_i^\diamond \leq \tilde{z}_i$ then $z_i^\diamond \leq z_i^* \leq \tilde{z}_i$, and if $\tilde{z}_i \leq z_i^\diamond$ then $\tilde{z}_i \leq z_i^* \leq z_i^\diamond$.*

With regard to Fig. 2(b), Proposition 14 states that any optimal solution z^* (e.g., the dotted green line in figure) should lie between the true depth z^\diamond (solid black line) and the naive solution \tilde{z} (dashed blue line). Moreover, the sign of the curvature of z^* cannot change between consecutive samples. Now that we have a complete understanding of the solution set of (L1_D), we can get two desirable results. First, we obtain error bounds that tell us how far is any z^* from z^\diamond . We provide this bound in SM-6.5. Second, and more interestingly, we devise an algorithm based on (L1_D) which exactly recovers a 2D profile also when we are not able to sample the corners. This algorithm is given in Section V.

D. Geometric Conditions for Optimality in 3D

Here we provide a sufficient condition for a 3D signal to be in the solution set of (L1_Δ). We start by introducing a specific sampling strategy (the analogous of the twin samples in the 2D case) which allows us to discuss optimality.

Definition 15 (Grid samples and Patches). *Given a 3D signal $Z \in \mathbb{R}^{r \times c}$, a grid sample set includes pairs of consecutive rows and consecutive columns of Z , and the boundary (first and last two rows, first and last two columns) of Z . This sampling strategy divides the image in rectangular patches, i.e., the set of non-sampled pixels surrounded by a pair of sampled rows and columns.*

If we have K patches, and we denote with $\bar{\mathcal{M}}_i$ the pixels in patch i , then the union $\mathcal{M} \cup \{\bar{\mathcal{M}}_i\}_{i=1}^K$ includes all the pixels in the image. We can now extend the notion of sign consistency to the 3D case.

Definition 16 (3D Sign Consistency). *Let $Z \in \mathbb{R}^{r \times c}$ be a 3D signal. Let \mathcal{M} be a grid sampling set and $\{\bar{\mathcal{M}}_i\}_{i=1}^K$ be the corresponding patches. Let $Z_{[\bar{\mathcal{M}}_i]}$ be restriction of Z to its entries in $\bar{\mathcal{M}}_i$. Then, Z is called 3D sign consistent if for all $i = \{1, \dots, K\}$, the nonzero entries of $\text{sign}(\text{vec}(DZ_{[\bar{\mathcal{M}}_i]}))$ are all +1 or -1, and the nonzero entries of $\text{sign}(\text{vec}(Z_{[\bar{\mathcal{M}}_i]}D^T))$ are all +1 or -1, where D is 2nd-order difference operator (2) of suitable dimension.*

In words, 3D sign consistency requires the profile not to change the sign of its curvature (along the horizontal and

Algorithm 1: Exact recovery of 2D depth profiles.

```

input : Measurements  $y$ , and sample set  $\mathcal{M}$ , including boundary and
         twin samples
output: Original signal  $z^\diamond$ 

/* solve  $\ell_1$ -minimization */
1 create matrices  $A$  (Definition 3) and  $D$  (eq. (2));
2 solve  $(f^*, z^*) = \min_z \|Dz\|_1$  subject to  $Az = y$ ;
/* populate a vector of signs  $s \in \{-1, 0, +1\}^n$  */
3 for consecutive twin samples  $(i-1, i), (j, j+1)$  do
4   foreach  $k \in \{i+1, \dots, j-1\}$  do
5      $s_k = \text{sign}((z_{j+1}^* - z_j^*) - (z_i^* - z_{i-1}^*))$ 

/* recover  $z^\diamond$  within the solution set */
6  $z^\diamond = \text{argmin}_z s^T z$  subject to  $Az = y, \|Dz\|_1 \leq f^*$ ;
7 return  $z^\diamond$ .

```

vertical directions) within each patch. We can now present a sufficient condition for Z^\diamond to be in the solution set of (L1_Δ).

Proposition 17 (3D Sign Consistency \Rightarrow Optimality). *Let $Z \in \mathbb{R}^{r \times c}$ be a 3D signal, feasible for problem (L1_Δ). Assume the sample set \mathcal{M} is a grid sample set. Then Z is in the set of minimizers of (L1_Δ) if it is 3D sign consistent.*

Roughly speaking, if our grid sampling is “fine” enough to capture all changes in the sign in the curvature, then Z^\diamond is among the solutions of (L1_Δ).

V. ALGORITHMS

The formulations discussed so far, namely (L1_D), (L1_D^ε), (L1_Δ), (L1_Δ^ε), directly translate into algorithms: each optimization problem can be solved using standard linear programming routines and returns a depth estimate.

In this section we describe other two algorithms. The first algorithm, given in Section V-A, solves 2D problems and is inspired by Proposition 14. The second one, given in Section V-B, solves 3D problems and is a variant of (L1_Δ).

A. Enhanced Recovery in 2D problems

Proposition 14 tells us that any optimal solution of (L1_D) lies between the naive solution and the ground truth profile z^\diamond (recall Fig. 2(b)). Algorithm 1 is based on a simple idea. If the true profile is concave between two samples (*cf.* with the first corner in Fig. 2(b)), then we should look for a profile within the optimal set of (L1_D) that has depth as large as possible within that interval (i.e., as close as possible to z^\diamond). If the shape is convex (second corner in Fig. 2(b)) we look for a solution having small depth, as this is close to z^\diamond .

Algorithm 1 first solves problem (L1_D) and computes an optimal solution z^* and the corresponding optimal cost f^* (lines 1-2). Let us skip lines 3-5 for the moment and take a look at line 6: the constraints in this optimization problem include the same constraint of line 2 ($Az = y$) plus an additional constraint that restricts z to be optimal for the problem in line 2 ($\|Dz\|_1 \leq f^*$). Therefore, we only remain to design an objective function that “searches” a solution close to z^\diamond within this optimal set. We use a simple linear objective $s^T z$, where $s \in \{0, \pm 1\}^n$ is a vector of signs that rewards the entries of the signal z to be as small as possible when $s_k = +1$, or as large as possible when $s_k = -1$. The last outstanding question is how to set these signs. This is done in lines 3-5. For any consecutive twin samples $(i-1, i), (j, j+1)$ the algorithm looks at the slope difference between

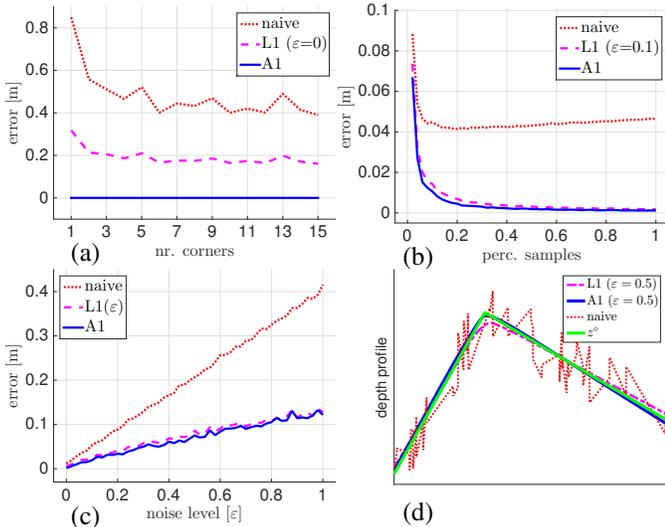


Fig. 3. (a) Recovery errors for A1, L1, and naive. A twin sample is acquired in each linear segment and measurements are noiseless. (b) Recovery errors from uniformly sampled noisy measurements, with increasing number of samples. (c) Recovery errors from uniformly sampled noisy measurements (5% of the depth data points) and increasing noise level. (d) Example of ground truth profile (green) and recovered profiles.

the last two samples (i.e., $z_{j+1}^* - z_j^*$) and the first two samples ($z_i^* - z_{i-1}^*$). If this difference is negative, then the function z^* is expected to be concave between the samples. In this case the sign s_k for any point k between the samples is set to -1 ; otherwise the signs are set to $+1$. We prove the following.

Corollary 18 (Exact 2D Reconstruction in Algorithm 1). *Under the assumptions of Proposition 14, Algorithm 1 recovers the 2D depth profile z^\diamond exactly.*

While Algorithm 1 is designed for noiseless samples, in the experiment we also test a noisy variant by substituting the constraints in lines 2 and 6 with $\|Az - y\|_\infty \leq \varepsilon$.

B. Enhanced Recovery in 3D problems

In the formulations (L1 $_{\Delta}$) and (L1 $_{\Delta}^{\varepsilon}$) we used the matrix Δ to encourage “flatness” or regularity of the depth profiles. In this section we propose a slightly different regularization matrix that performs even better in practice. The idea is the following: while the matrix Δ only penalizes changes in curvature along the horizontal direction ($Z_{i,j-1} - 2Z_{i,j} + Z_{i,j+1}$) and vertical direction ($Z_{i-1,j} - 2Z_{i,j} + Z_{i+1,j}$), we can also penalize changes along the diagonal. More formally, for each pixel (i, j) , we also penalize the following term:

$$\frac{1}{4} (Z_{i+1,j+1} - Z_{i+1,j-1} - Z_{i-1,j+1} + Z_{i-1,j-1}) \quad (8)$$

which is the commonly used discrete version of the cross partial derivative [23, eq. 2.47]. Adding this penalty term for each pixel (i, j) only adds extra rows to the matrix Δ , without altering problems (L1 $_{\Delta}$) and (L1 $_{\Delta}^{\varepsilon}$) otherwise.

VI. EXPERIMENTS

This section supports our theoretical derivation using real and synthetic data. Empirical evidence shows that our recovery techniques perform very well in practice, in both noisy and noiseless scenarios. For our tests, we use CVX/MOSEK as parse/solver for optimization. The average recovery error for an estimate z^* is defined as $\frac{1}{n} \|z^* - z^\diamond\|_1$ (average depth

error, in meters). Results are averaged over 50 runs unless specified otherwise. When possible we compare our results with a standard interpolation scheme (Matlab’s command `interp1`). The results obtained with linear interpolation are denoted with “naive”, according to the terminology of Proposition 14. The results produced by (L1 $_D$) and (L1 $_{\Delta}$) are denoted with L1 (the distinction between 2D and 3D will be clear from the context). Finally, the estimate produced by the noisy problems (L1 $_D^{\varepsilon}$) and (L1 $_{\Delta}^{\varepsilon}$) are denoted with L1($\varepsilon = \cdot$), where in parenthesis we specify the noise level.

A. Simulated depth profiles

Exact and stable recovery of 2D profiles. For this simulated experiment we create random piecewise linear depth profiles of size $n=2000$, with given number of corners.

Fig. 3(a) compares naive, L1, and the estimate of Algorithm 1 (label: A1). These results consider noiseless measurements and sample set including a twin sample in each linear region (these are the assumptions of Proposition 14); reconstruction errors are shown for profiles with increasing number of corners. As predicted by Corollary 18, A1 recovers the original signal exactly (zero error). naive has large errors, while the L1 estimate falls between the two.

Fig. 3(b) considers a more realistic situation: since in practice we do not know where the corners are, in this case we uniformly sample depth measurements and we consider noisy measurements with $\varepsilon = 0.1\text{m}$. As this percentage goes to 1 (100%), we sample all entries of the depth profile. We consider profiles with 3 corners in this test. The figure shows that for increasing number of samples, our approach largely outperforms the naive approach. A1 improves over L1 even in presence of noise, while the improvement is not as substantial as in the noiseless case of Fig. 3(a). Fig. 3(b) also shows that the error committed by naive does not improve when adding more samples. This can be understood from Fig. 3(d), which shows an example of 2D profile (in green) and the reconstructions produced by naive, L1, and A1. naive linearly interpolates the samples, hence even when measuring all depth data, it still produces a jagged line. It is easy to show that when measurement noise is uniformly distributed in $[-\varepsilon, +\varepsilon]$ (as in our tests), the error converges to $\varepsilon/2$ for increasing number of samples. On the other hand, L1 and A1 correctly smooths the noise out.

Fig. 3(c) considers a fix amount of samples (5%) and tests the three approaches for increasing measurement noise ε . Our techniques (L1, A1), are very resilient to noise and degrade gracefully even in presence of large noise (e.g., $\varepsilon = 1\text{m}$).

The CPU times required by L1 and A1 are around 0.2 and 0.4 second (see SM-2.4 for extra timing results).

Stable recovery of 3D profiles. For the tests in this section we create random piecewise planar 3D depth profiles of size $n = 100 \times 100$. Besides the naive and L1 approaches, we also test the variant described in Section V-B, which we denote with L1diag. Fig. 4(a) shows the reconstruction error for increasing percentage of measurements. We also include the neighbors of each samples (as suggested in Proposition 8), since we found that this improves our reconstruction. The interested reader can find the results

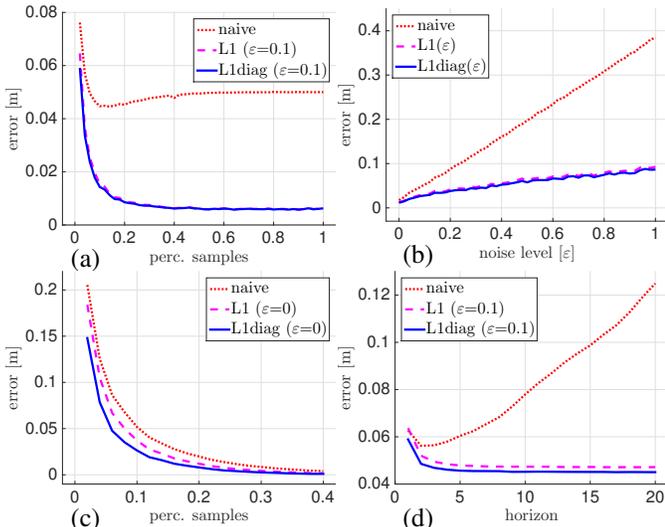


Fig. 4. (a) Recovery errors from uniformly sampled noisy measurements, with increasing number of samples. (b) Errors from uniformly sampled noisy measurements (5% of the depth data points) and increasing noise. (c) Gazebo: errors for increasing number of samples. (d) Error with samples collected over a receding horizon, for increasing horizon length.

without including neighboring pixels in SM-3.3. As in the 2D case, the proposed approaches are able to filter out noise. Similar comments apply to Fig. 4(b), which tests reconstruction with fixed number of samples (5%) and increasing noise.

For a 100×100 image, the CPU times required by L1 and L1diag are around 3.5 and 8 seconds, respectively (see SM-3.5 for more timing results). While the CPU time is relatively large in our Matlab implementation, we leave computational aspects (e.g., use a dedicated linear solver) to future work.

B. Real Data and Applications

2D mapping from sparse measurements. We use the *Stage* simulator to simulate a robot equipped with a laser scanner with only 10 beams, moving in a 2D scenario. The robot is in charge of mapping the scenario; we assume the trajectory to be given. Our approach works as follows: we feed the 10 samples measured by our “sparse laser” to algorithm A1; A1 returns a full scan (covering 180 degrees with 180 scans in our tests), which we feed to a standard mapping routine (we use `gmapping` [24] in our tests). Fig. 5 compares the occupancy grid map produced by a standard mapping algorithm based on a conventional laser scan, against the occupancy grid map reconstructed from our 10-beam laser. Fig. 5(c) shows that we are able to do a fairly accurate reconstruction from very partial information. As a technical remark, we note that, before feeding the 10 laser measurements to A1, we express the corresponding points in Cartesian coordinates: the original data given by the laser is in polar coordinates and piecewise linear signals do not remain piecewise linear in polar coordinates (see SM-2.1).

Sparse 3D depth reconstruction. We consider two sets of profiles (these are the ground truth profiles that we want to reconstruct). These profiles are produced by rendering 20 full depth image in Gazebo and by collecting 1000 full disparity images using a ZED stereo camera. Representative images are given in Fig. 1, Fig. 6, and SM-4-SM-5.

A remark is now in order. Each pixel in a depth image identifies a direction and the corresponding entry is the

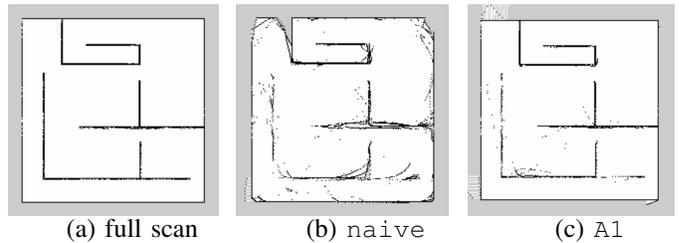


Fig. 5. (a) Mapping using a conventional laser scanner, (b-c) Mapping using scans reconstructed from 10 measurements by naive and A1.

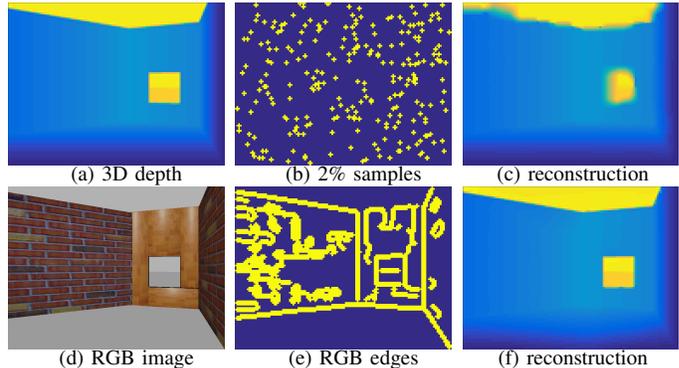


Fig. 6. Gazebo: (a) full depth profile, (b) uniformly drawn sparse samples, and (c) reconstruction using L1diag. (d) RGB image, (e) edges in the RGB image, (f) L1diag reconstruction from the depth at the RGB edges.

depth of a 3D point observed along that direction. This can be understood as a spherical parametrization of the observed 3D points. Ideally, as in the 2D case, we should express each point in Cartesian coordinates before applying our algorithms, since this spherical parametrization distorts planes into curved patches (see SM-3.2 for an example). In practice we observed that our algorithm works well directly over the original depth images, hence in this paper we skip the conversion to Cartesian coordinates altogether. The interested reader can find the results obtained by working in Cartesian coordinates in the SM-3.2-SM-3.4 (with some discussion on orthogonal and perspective projection).

Fig. 4(c) reports the error for increasing number of samples, averaged over all ZED images. Contrarily to the noisy test of Fig. 4(a), here we do not add extra noise (the original depth image that we sample is already noisy as it is the result of stereo reconstruction). For this reason naive, L1, and L1diag all converge to zero error, while in the regime with small samples L1diag still provides the best reconstruction. Example of profiles reconstructed from %2 of the samples are shown in the first row of Fig. 1 and Fig. 6.

Sparse Stereo. While so far we discussed reconstruction of depth profiles from few sparse samples, our results also have useful applications for conventional sensors, as stereo cameras. Standard stereo vision algorithms compute the disparity at each pixel and this is computationally expensive. On the other hand, Proposition 8 proves that we only need to measure the depth in the neighborhood of the edges to recover the full depth. This suggests the following *sparse stereo* approach: we compute the pixel location of the edges via edge detection on the RGB image; under the assumption that changes of depth cause changes in appearance, the RGB edges are a good proxy for the edges in the depth profile. Then, we compute the disparity only at those edges

(usually $< 10\%$ of the image), and we reconstruct the entire depth using ℓ_1 minimization. A similar insight is proposed in [15], where the authors use an interpolation approach to reconstruct an approximate depth profile. Table I shows the reconstruction errors (mean and standard deviation) when the profile is reconstructed from the RGB edges. Again, Lldiag largely outperforms naive interpolation. Average reconstruction time is 1.2s for L1, and 1.8s for Lldiag.

		Gazebo	ZED
error (mean & std) [m]	naive	0.102 (0.034)	0.165 (0.180)
	L1	0.072 (0.027)	0.144 (0.150)
	Lldiag	0.056 (0.022)	0.111 (0.111)

TABLE I
SPARSE STEREO RESULTS

VII. CONCLUSION AND FUTURE WORK

We discussed theoretical conditions and practical algorithms to recover 2D and 3D depth profiles from sparse and incomplete samples. Our formulation leads to constrained ℓ_1 -minimization problems: we provide sufficient conditions under which this minimization recovers the original signal and we provide a complete algebraic and geometric characterization of the set of minimizers. Our algorithms are very effective in practice: the capability of properly modeling measurement noise enables a performance boost with respect to interpolation-based naive approaches; they are also easy to implement since they require only solving a linear program. We provide many application examples, including 2D and 3D mapping, and sparse stereo reconstruction.

This work opens many new research avenues. On the algorithmic side, we are currently investigating a temporal smoothing formulation: rather than using only the samples acquired at time t to perform depth reconstruction, we also use past samples acquired at times $t-h, t-h+1, \dots, t-1$ (h is a given time horizon) after roto-translating them to the current frame. Since the roto-translation between past and current frames is estimated from noisy odometry, past samples are less accurate than the current ones, and we can correctly model this behavior by assigning a different ε to each measurement in $(L1_{\Delta}^{\varepsilon})$. Preliminary results from this receding horizon strategy are shown in Fig.4(d) for increasing length of the horizon. We obtain the figure by simulating a robot observing a depth profile and moving towards it at 0.5m per frame. The intra-frame odometry noise is assumed bounded by 0.1m. naive is not able to model measurement noise and the reconstruction gets worse when including past noisy samples. On the other side, temporal smoothing improves the performance of L1 and Lldiag (for large horizons the error stops improving as past measurements become too noisy to be useful).

On the implementation side, we plan to test the use of more efficient linear solvers, possibly exploiting the structure of our minimization problems. We also plan to extend the range of applications of our algorithms. For instance, from Fig. 4(a) it is clear that our approach is very effective for point cloud denoising. Moreover, it can be useful to increase the resolution of existing sensors and possibly compensate for partial sensor failures. One last outstanding questions is:

can we use the insight of this paper to design motion policies that can actively improve depth reconstruction?

REFERENCES

- [1] R. Newcombe, S. Lovegrove, and A. Davison, "DTAM: Dense tracking and mapping in real-time," in *Intl. Conf. on Computer Vision (ICCV)*, Barcelona, Spain, Nov 2011, pp. 2320–2327.
- [2] R. Mur-Artal, J. Montiel, and J. Tardós, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [3] T. Whelan, S. Leutenegger, R. Salas-Moreno, B. Glocker, and A. Davison, "ElasticFusion: Dense SLAM without a pose graph," in *Robotics: Science and Systems (RSS)*, 2015.
- [4] S. Suri, E. Vicari, and P. Widmayer, "Simple robots with minimal sensing: From local visibility to global geometry," *Intl. J. of Robotics Research*, vol. 27, 2008.
- [5] J. Derenick, A. Speranzon, and R. Ghrist, "Homological sensing for mobile robot localization," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 572–579.
- [6] B. Tovar, L. Freda, and S. M. LaValle, "Learning combinatorial map information from permutations of landmarks," *Intl. J. of Robotics Research*, vol. 30, no. 9, pp. 1143–1156, 2011.
- [7] B. Tovar, F. Cohen, L. Bobadilla, J. Czarnowski, and S. M. LaValle, "Combinatorial filters: Sensor beams, obstacles, and possible paths," *ACM Transactions on Sensor Networks*, vol. 10, no. 3, 2014.
- [8] R. Wood, "RoboBees project," Harvard University, 2015. [Online]. Available: <http://robobeeseas.harvard.edu/>
- [9] D. Marinakis and G. Dudek, "Occam's razor applied to network topology inference," *IEEE Trans. Robotics*, vol. 24, no. 2, pp. 293–306, April 2008.
- [10] J. M. O'Kane and S. M. LaValle, "Localization with limited sensing," *IEEE Trans. Robotics*, vol. 23, pp. 704–716, 2007.
- [11] L. Erickson, J. Knuth, J. O'Kane, and S. LaValle, "Probabilistic localization with a blind robot," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2008, pp. 1821–1827.
- [12] M. Milford, "Vision-based place recognition: how low can you go?" *Intl. J. of Robotics Research*, vol. 32, pp. 766–789, 2013.
- [13] T. Kanade, "Recovery of the 3-dimensional shape of an object from a single view," *Artificial Intelligence*, vol. 17, 1981.
- [14] W. Hong, A. Yang, K. Huang, and Y. Ma, "On symmetry and multiple view geometry: Structure, pose and calibration from a single image," *Intl. J. of Computer Vision*, no. 3, 2004.
- [15] S. Pillai, S. Ramalingam, and J. Leonard, "High-performance and tunable stereo reconstruction," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2016.
- [16] P. Piniés, L. Paz, and P. Newman, "Too much TV is bad: dense reconstruction from sparse laser with non-convex regularisation," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Seattle, WA, USA, May 2015.
- [17] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Boston: Birkhäuser, 2013.
- [18] D. Needell and R. Ward, "Stable image reconstruction using total variation minimization." *SIAM J. Imaging Sci.*, vol. 6, no. 2, pp. 1035–1058, 2013.
- [19] D. Reddy, A. C. Sankaranarayanan, V. Cevher, and R. Chellappa, "Compressed sensing for multi-view tracking and 3-D voxel reconstruction," in *Intl. Conf. on Image Processing (ICIP)*, 2008, pp. 221–224.
- [20] S. Nam, M. Davies, M. Elad, and R. Gribonval, "The cosparsity analysis model and algorithms," *Applied and Computational Harmonic Analysis*, vol. 34, no. 1, pp. 30–56, 2013.
- [21] M. Kabanava and H. Rauhut, "Cosparsity in compressed sensing," in *Compressed Sensing and Its Applications*, H. Boche, R. Calderbank, G. Kutyniok, and J. Vybiral, Eds. Springer, 2015, pp. 315–339.
- [22] F. Ma, L. Carlone, U. Ayaz, and S. Karaman, "Sparse sensing for resource-constrained depth reconstruction," 2016, intl. Conf. on Intelligent Robots and Systems, supplemental material, <https://dl.dropboxusercontent.com/u/4309261/Ma16iros-SM.pdf>.
- [23] *Numerical Partial Differential Equations for Environmental Scientists and Engineers: A First Practical Course*. Boston, MA: Springer US, 2005, ch. Finite Difference Calculus, pp. 11–19. [Online]. Available: http://dx.doi.org/10.1007/0-387-23620-1_2
- [24] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with Rao-Blackwellized particle filters," *IEEE Trans. Robotics*, vol. 23, no. 1, pp. 34–46, 2007.