



Sebastian Schuon
Stanford University
schuon@cs.stanford.edu

Christian Theobalt
Stanford University
theobalt@cs.stanford.edu

James Davis
UC Santa Cruz
davis@cs.ucsc.edu

Sebastian Thrun
Stanford University
thrun@stanford.edu



LidarBoost: Depth Superresolution for ToF 3D Shape Scanning

Abstract

Depth maps captured with time-of-flight cameras have very low data quality: the image resolution is rather limited and the level of random noise contained in the depth maps is very high. Therefore, such flash lidars cannot be used out of the box for high-quality 3D object scanning.

To solve this problem, we present LidarBoost, a 3D depth superresolution method that combines several low resolution noisy depth images of a static scene from slightly displaced viewpoints, and merges them into a high-resolution depth image. We have developed an optimization framework that uses a data fidelity term and a geometry prior term that is tailored to the specific characteristics of flash lidars. We demonstrate both visually and quantitatively that LidarBoost produces better results than previous methods from the literature.



Validation

We compared our method to a state-of-the-art 2D super-resolution method (IBSR [1]) applied to the depth images and a SR method that combines depth and intensity images (Diebel's MRF[2]). In both cases our 3D method that is tailored to ToF sensors did much better.

Formulation

$$\sum_{k=1}^K \|T_k .* W_k .* (D_k - X)\|_2 + \lambda \sum_{u,v} \|\nabla X_{u,v}\|_2$$

Data Term

- k k-th Depth Image
- u, v Footprint Regularization Term
- D_k Input Samples
- T_k Position Samples are enforced
- X Superresolution output
- W_k Confidence Weight

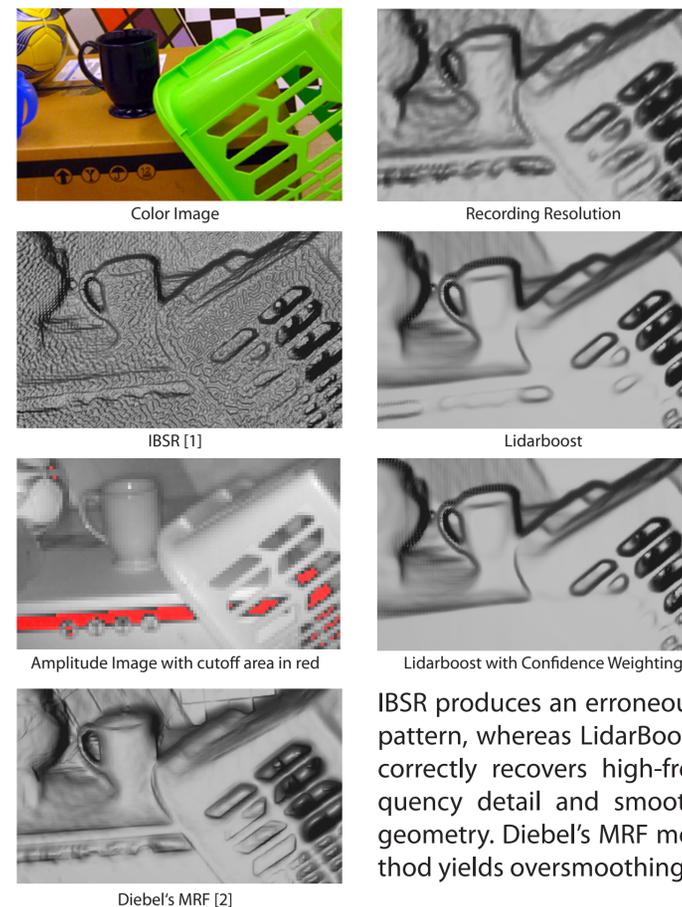
Prior Term

- λ Tradeoff Parameter
- Bilateral, Rotation Invariant Gradient Prior:

$$E_{\text{regular}}(X) = \sum_{u,v} \|\nabla X_{u,v}\|_2 = \sum_{u,v} \left\| \begin{pmatrix} G_{u,v}(0,1) \\ G_{u,v}(1,0) \\ \vdots \\ G_{u,v}(l,m) \end{pmatrix} \right\|_2$$

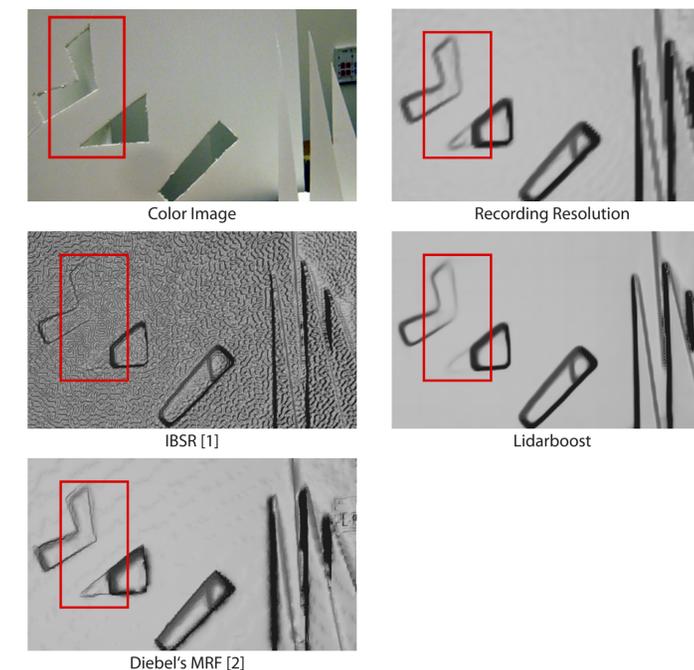
$$G_{u,v}(l,m) = \frac{X(u,v) - X(u+l,v+m)}{\sqrt{l^2 + m^2}}$$

Results I: Fine Details



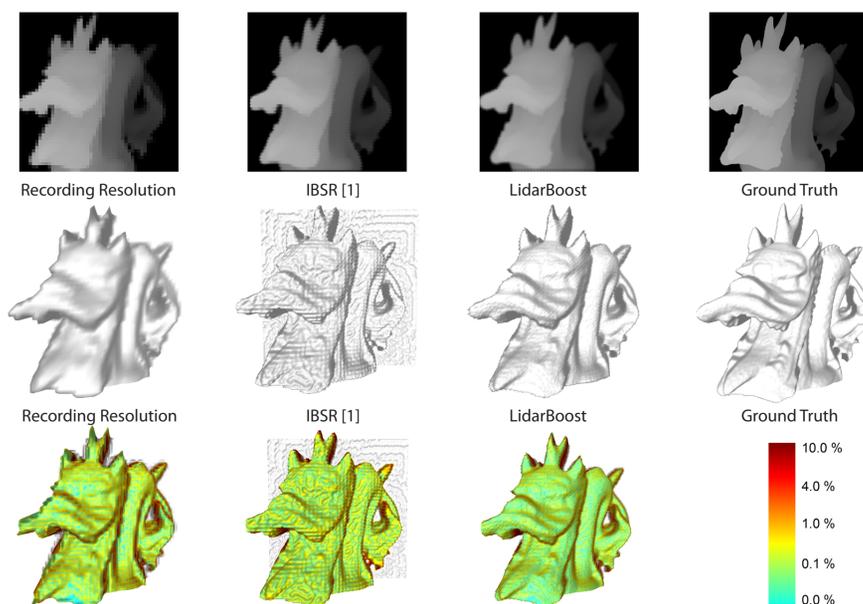
IBSR produces an erroneous pattern, whereas LidarBoost correctly recovers high-frequency detail and smooth geometry. Diebel's MRF method yields oversmoothing.

Results II: Gain In Resolution



IBSR demonstrates increased resolution at the edges, but some aliasing and the strong pattern remains. LidarBoost reconstructs the edges much more clearly and there is hardly a trace of aliasing. MRF upsampling oversmooths the depth edges and in some places allows the low resolution aliasing to persist.

Quantitative Validation On Synthetic Data - Case I: No Noise

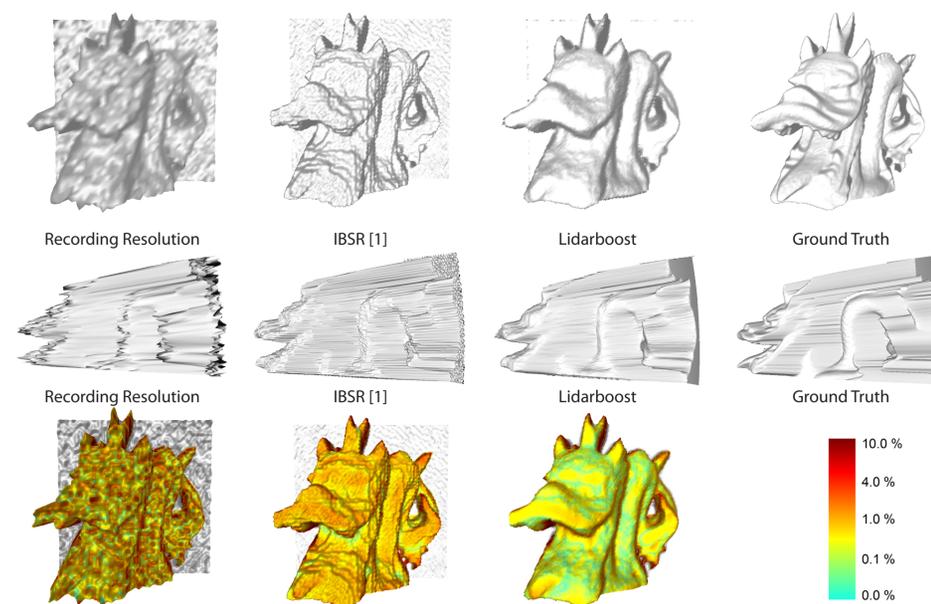


For test purposes, low resolution depth images of the Stanford dragoon have been created synthetically. Here, no noise has been added and the mean square error visualized in the color plots.

As shown in the 3D models, IBSR recovers the overall structure, but exhibits a noise pattern. LidarBoost recovers the structure almost perfectly and yields a smooth surface. Also quantitatively LidarBoost is superior to IBSR, as indicated by the relative MSE indicate both shown in the color coded plots and the following table:

	No Noise $var = 0$	Medium Noise $var = 0.7$	Stark Noise $var = 5$
LR	157.6	161.7	203.9
IBSR	83.8	89.9	127.0
LidarBoost	70.6	72.5	82.9

Quantitative Validation On Synthetic Data - Case II: Stark Noise



This second synthetic set has severe noise added (variance 5.0, 4x upsampling). With rendered 3D geometry in frontal view, LidarBoost shows shows best upsampling result, while IBSR suffers from staircasing.

Also in a lateral view it is apparent that LidarBoost's reconstruction is closest to ground truth. This is verified by the color coded error plots.

[1] S. Schuon, C. Theobalt, J. Davis, and S. Thrun. High-quality scanning using time-of-flight depth superresolution. CVPR TOF Workshop 2008
 [2] J. Diebel and S. Thrun. An application of markov random fields to range sensing. In Advances in Neural Information Processing Systems 18, pages 291–298. 2006.