Fast Minimax Path-Based Joint Depth Interpolation

Longquan Dai, Feihu Zhang, Xing Mei, and Xiaopeng Zhang

Abstract—**We propose a fast minimax path-based depth interpolation method. The algorithm computes for each target pixel varying contributions from reliable depth seeds, and weighted averaging is used to interpolate missing depths. Compared with state-of-the-art joint geodesic upsampling method which selects the K nearest seeds to interpolate missing depths with $O(Kn)$ complexity, our method does not need to limit the number of seeds to K and reduces the computational complexity to $O(n)$. In addition, the minimax path chooses a path with the smallest maximum immediate pairwise pixel difference on it, so it tends to preserve sharp depth discontinuities better. In contrast to the results of previous depth upsampling algorithms, our approach can provide accurate depths with fewer artifacts.**

Index Terms—**Depth map, minimax path, upsampling.**

I. INTRODUCTION

RECENTLY, low-cost depth sensors such as ToF camera and Micosoft Kinect have been widely used in various applications and gradually shape a new man-machine interactive way. However, limited by the current sensing techniques, the depth map suffers from all sorts of problems. Specifically, the resolution of a depth image is very low and some depths in the depth map are lost due to occlusion or other degradation. Thus the quality is worse than the traditional optical photo. To obtain a high-quality depth map as the optical photo, we propose a novel joint depth interpolation algorithm.

Our algorithm interpolates the depths of target pixels from the depths of seeds under the guidance of a registered color image for a corrupted depth map, where seeds stand for the the pixels with observed depths and target pixels denote the pixels whose depths are lost. The guidance information of color image can help our method produce high quality result because color transition usually suggests depth transition and thus the guidance image's structure can indicate the structure of the depth map D for edge-preserving results. Specifically, in our algorithm, we extract all-pairs minimax paths from the guidance image I and map them onto D to present the structure of D. After that, the

missing depths of target pixels are figured out from the reliable depths of seeds along these minimax paths.

On the graph representation of an image, the minimax path [1] between two nodes is the path linking the two nodes with the smallest intermediate length of the longest edge. It does not cross the color edges of I, which are typically also the depth boundaries of D. Otherwise, the crossing boundary edge will increase the maximal length of edges in the path. Since the minimax path is sensitive to the underlying structure of the guidance image, we employ the length of the minimax path, which links a seed and a target pixel, to calculate the contribution of the target pixel received from the seed in a geometric aware manner.

II. RELATED WORK

The most related works are the depth upsampling methods which focus on upscaling depth maps. Generally, these upsampling methods can be roughly categorized as global methods and local methods. Global methods [2], [3], [4], [5] utilize the optimization framework that punishes a large cost for coupled pixels that have similar colors but different depths. Global methods usually produce high-quality upsampling results, but the computational cost is too heavy for real time processing.

Local methods [6], [7], [8], [9], [10], [11] are based on the weighted average scheme introduced by Kopf *et al.* [8] and often have fast implementations. Our approach, which is inspired by a recently proposed aggregation method [12], also falls into this category. Specifically, Yang proposed an $O(n)$ non-local cost aggregation method. However, the algorithm is not suitable for interpolation. We make an advance and present an efficient interpolation implementation with $O(n)$ complexity by designing a novel reliable seed choosing strategy and introducing a geometric-aware similarity metric, where n denotes the total number of image's pixels.

Euclidean distance [8] is the simplest similarity measure, but it is not aware of the structure of the data points that lie on a curved manifold. On the graph representation of these points, people propose a set of computation intensive link-based distances [6], [13], [14], [15], [16] to conquer the weakness of Euclidean distance. Here, we prefer the length of minimax path (or the minimax distance) which is an efficient and geometric-aware affinity metric.

III. JOINT INTERPOLATION ALGORITHM

A. Essential Ingredients of an Ideal Joint Interpolation

Here we identify what the essential ingredients of the ideal joint interpolation method are by disclosing the shortcomings
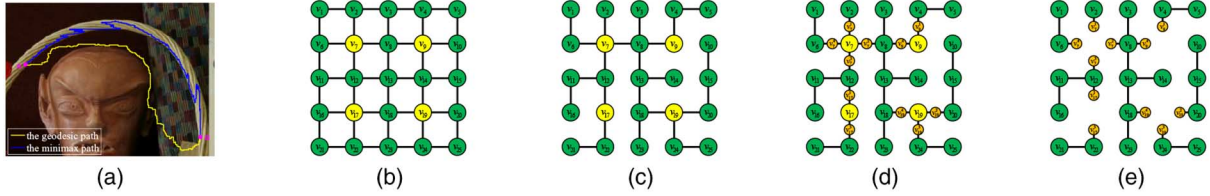
Fig. 1. The comparison between minimax path and geodesic path, and the procedure demonstration of obtaining an interpolation tree. (a) shows that the geodesic chooses the crossing edges path which has a large jump at the color boundaries as the path is almost on the flatten area and the length on flatten area is nearly zero. In contrast, the minimax path is more sensitive to color edges and is able to follow the structure of guidance image. (b) the graph representation G of guidance image, where the green nodes are target pixels and yellow nodes denote seeds. (c) a MST of G. (d) the obtained interpolation tree by inserting orange auxiliary nodes into the MST. (e) the auxiliary nodes divide the interpolation tree into several parts by removing seeds.

of local upsampling methods. Kopf *et al.* [8] introduced the first weighted average based upsampling formula (1).

$$D(p) = \frac{\sum_{q \in \Omega_p \bigcap S} w_{p,q}^e D(q)}{\sum_{q \in \Omega_p \bigcap S} w_{p,q}^e} \quad p \in \Omega \backslash S \qquad (1)$$

where Ω denotes the entire image domain, S stands for the seeds and $w_{p,q}^e = \exp(-\frac{d_e^2(p,I(p)),(q,I(q))}{\sigma_e^2})$. However, the user-specified window Ω_p and the Euclidean distance d_e [8] defined on the five-dimensional space-color volume [17] are suboptimal to take advantage of the geometric structure of guidance image I because both Ω_p and d_e are not adaptive to the underlying geometric structure of the guidance image I. Liu *et al.* [6] substituted the Euclidean distance d_e [8] with the geodesic distance d_g [6] and replaced the user-specified window selecting method with the K nearest seeds choosing strategy. The interpolation formula (2) is still problematic.

$$D(p) = \frac{\sum_{q \in \mathcal{K}_p} w_{p,q}^g D(q)}{\sum_{q \in \mathcal{K}_p} w_{p,q}^g} \quad p \in \Omega \backslash S \qquad (2)$$

where $w_{p,q}^g = \exp(-\frac{d_g^2(\pi_{p,q}^g)}{\sigma_g^2})$ and $\mathcal{K}_p$ denotes the K nearest seeds of the target pixel p. Specifically, Liu *et al.* employ an iterative algorithm with $O(Mn)$ complexity to calculate the geodesic distance, as the computational cost for extracting all-pairs geodesic paths is $\log(n^2 \log n)$ [8]. But accurate approximation demands M iterations, which is still inefficient. Additionally, it is reasonable to expect that the quality of results will increase with the number K of reliable seeds. Unfortunately, Liu's $O(Kn)$ algorithm depends on K too.

Above all, we argue that an ideal joint interpolation should have following properties:

1) An ideal interpolation method should utilize a reliable seed choosing strategy which can automatically catch all the reliable seeds for any distribution.

2) An ideal interpolation method should employ a geometric-aware similarity metric, which is sensitive to the underlying geometric structure of guidance image, to weight the contributions from reliable seeds.

3) The implementation of an ideal interpolation method should be high-efficiency, especially for real-time applications, thus an algorithm with $O(n)$ complexity should be given the primary consideration.

B. The Geometric Aware Minimax Distance

Similar to the well-known geodesic distance, the minimax distance (i.e the length of the minimax path) is another kind of geometric-aware metric which can be used to measure the similarities between seeds and target pixels. The minimax path is defined on the graph representation $G(V, E)$ (Fig. 1(b)) of the guidance image I, where each pixel of I corresponds to a node $p \in V$ and the first order neighborhood (v_1, v_2) of pixels v_1 and v_2 denotes an edge $e_{v_1, v_2} \in E$ with length $d(e_{v_1, v_2}) = \|I(v_1) - I(v_2)\|_1$. Thus the minimax path $\pi_{p,q}^m$ that links pixels p and q is $\pi_{p,q}^m \in \arg \min_{\pi_{p,q} \in \Pi_{p,q}} (d_\infty(\pi_{p,q}^m))$, where Π_{pq} denotes all the possible paths which connect the initial node $v_{a_0} = p$ and the destination vertex $v_{a_m} = q$ such that $\{e_{v_{a_0}, v_{a_1}}, e_{v_{a_1}, v_{a_2}}, \ldots, e_{v_{a_{m-1}}, v_{a_m}}\} \subseteq E$, for any path $\pi_{v_{a_0}, v_{a_m}} = (v_{a_0}, v_{a_1}, \ldots, v_{a_m})$, $v_{a_i} \in V$. Moreover, $d_\infty(\pi_{v_{a_0}, v_{a_m}}) = \max_{i \in \{0, \cdots, m-1\}} \{d(e_{v_{a_i}, v_{a_{i+1}}})\}$. Using the same symbols, the minimax distance $d(\pi_{a_0, a_m}^m)$ of π_{a_0, a_m}^m can be written as $d(\pi_{a_0, a_m}^m) = \sum_{i=0}^{m-1} d(e_{v_{a_i}, v_{a_{i+1}}})$.

The minimax path is more sensitive to color transition than geodesic. Specifically, the minimax path, according to the definition, minimizes the maximal length of edges of a path whereas geodesic prefers the shortest path which tolerates large color transition on the path, as shown in Fig. 1(a). More importantly, unlike the fastest extracting all-pairs geodesic paths algorithm with $O(n^2 \log n)$ complexity [6], Hu [18] proves that all-pairs minimax paths (shown in Fig. 1(c)) can be represented by a minimum spanning tree (MST) which has an efficient extracting algorithm with $O(\log n)$ complexity [19].

C. A Reliable Seed Choosing Strategy

A MST of the graph representation $G(V, E)$ of a depth map connects all points on the graph G and the path linking two nodes along the MST is the minimax path between them [18]. We find that the seeds on the graph $G(V, E)$ separate the graph into L different subgraphs $\Omega_i, 1 \leq i \leq L$ along a MST of the graph and thus form the boundary B_i of these subgraphs. In Fig. 1, the yellow points demonstrate the boundary seeds. We can observe that the path that links any two interior nodes of two different subgraphs along the MST must pass through at least one boundary node (or seed). For example, the path linking v_6 and v_8 passes through v_7. Then, for each target pixel in the subgraph Ω_i, using the depths of the boundary points B_i to interpolate the depths of the target pixels in the subgraph Ω_i is more reliable than using other seeds because boundary points are the closest seeds to the target pixels in the sense of minimax distance, in other words, the minimax path from a non-boundary seed to any target pixel in Ω_i must pass through a boundary seed of B_i. We call these boundary seeds B_i as the reliable seeds of

the target pixels in Ω_i and only use their depths to interpolate the target pixels in Ω_i.

The reliable seed choosing strategy has three major advantages. First, it automatically catches all reliable seeds for any target pixel and thus our interpolation algorithm can employ all reliable depth information to interpolate the depth of each target pixel. Second, it is not sensitive to the seed distribution and can be used to various seed distributions other than the uniform distribution taken by upsampling. Last but not least, the strategy can be integrated into a linear time joint interpolation algorithm without extra cost.

D. A Linear Complexity Joint Interpolation Algorithm

Our interpolation formula is expressed as (3), which both takes the minimax distance as the similarity metric and employs our reliable seed choosing strategy to choose the interpolating seeds for each target pixel.

$$D(p) = \frac{\sum_{q \in B_i} w_{p,q}^m D(q)}{\sum_{q \in B_i} w_{p,q}^m} \quad p \in \Omega_i \backslash S \tag{3}$$

where $w_{p,q}^m = \exp(-\frac{d(\pi_{p,q}^m)}{\sigma_m})$ and $1 \le i \le L$.

The formula can be computed in a linear complexity by transforming the MST into an interpolation tree and updating depths along the interpolation tree T. Let $N(s)$ denote the neighborhood nodes which are connected to the seed s by the MST. For each edge $e_{s,p}, p \in N(s)$, we insert an intermediate point s_p into the middle of s and p, as shown in Fig. 1(d), and assign $d(e_{s,s_p}) = \infty, d(e_{s_p,p}) = d(e_{s,p})$. Here, all auxiliary points are denoted as A. We can observe from Fig. 1(e) that the interpolation tree T is divided into L subtrees T_i by removing all seeds, and the depths of auxiliary nodes $\widetilde{B}_i = A \bigcap_{T_i}$ determine the depths of the target pixels in the T_i. Then (3) can be rewritten as

$$D(p) = \frac{\sum_{q \in \widetilde{B}_i} w_{p,q}^t D(q)}{\sum_{q \in \widetilde{B}_i} w_{p,q}^t} \quad p \in T_i \backslash A, \ 1 \le i \le L \tag{4}$$

where $w_{p,q} = \exp(-\frac{d(\pi_{p,q}^t)}{\sigma_m})$ and $\pi_{p,q}^t$ is the path that connects p to q along the interpolation tree T.

In order to obtain the updating formulae with linear complexity, we put $\widetilde{C}(q) = D(q), \widetilde{W}(q) = 1$, for $q \in S$; $\widetilde{C}(q) = D(s), \widetilde{W}(q) = 1$, for $q \in A$; $\widetilde{C}(q) = 0, \widetilde{W}(q) = 0$, for a target pixel q, then (4) can be further transformed into following forms.

$$C(p) = \sum_{q \in T} w_{p,q}^t \widetilde{C}(q) \quad W(p) = \sum_{q \in T} w_{p,q}^t \widetilde{W}(q) \tag{5}$$

$$D(p) = C(p)/W(p) \quad p \in T_i \backslash A, 1 \le i \le L \tag{6}$$

In (5), we replace $\widetilde{B}_i$ with T. Although the values of $C(p)$ and $W(p)$ for $p \in T_i \backslash A$ receive supports from all other nodes on the interpolation tree T according to (5), the depths of target pixels in T_i only depend on the depths of auxiliary nodes $\widetilde{B}_i$ as the length of edge between the seed and its auxiliary points is set as infinity.

Yang [12] shows that the non-local cost aggregation form defined on the tree as $C(p)$ and $W(p)$ can be efficiently figured out

by traversing the tree structure of the interpolation tree T in two sequential passes. Initially, we visit the nodes of T_p using the breadth first traversal algorithm to assign each node v_i an order i based on the visit sequence with the property that if v_i is the father of v_j, then $i \le j$, where $1 \le i, j \le \#T_p$, $\#T_p$ denotes the nodes number of T_p. In the first pass, while the tree is traced from $v_{\#T_p}$ to v_1, $C^\uparrow(v_i)$ and $W^\uparrow(v_i)$ are not updated until all its children have been updated by the updating formulae (7) (8).

$$C^\uparrow(P(v_i)) = C^\uparrow(P(v_i)) + w_{v_i,P(v_i)}^t C^\uparrow(v_i) \tag{7}$$

$$W^\uparrow(P(v_i)) = W^\uparrow(P(v_i)) + w_{v_i,P(v_i)}^t W^\uparrow(v_i) \tag{8}$$

where $C^\uparrow(v)$ and $W^\uparrow(v)$ are initialized as $\widetilde{C}(v)$ and $\widetilde{W}(v)$, $P(v)$ records the father of node v, both $C^\uparrow(v)$ and $W^\uparrow(v)$ are the intermediate aggregated results. In the second pass, we employ the updating formulae (9) (10) to renew $C(v)$ and $W(v)$ when the tree is traversed from v_1 to $v_{\#T_p}$. At last, we use (6) to estimate final interpolation depths.

$$C(v_i) = w_{P(v_i),v_i}^t \widetilde{C}(P(v_i)) + (1 - w_{v_i,P(v_i)}^{t^2}) C^\uparrow(v_i) \tag{9}$$

$$W(v_i) = w_{P(v_i),v_i}^t \widetilde{W}(P(v_i)) + (1 - w_{v_i,P(v_i)}^{t^2}) W^\uparrow(v_i) \tag{10}$$

The computational complexity of our method consists of two parts: the cost of extracting a MST from the graph representation G of the guidance image I and the cost of updating (7) (8) (9) (10). In (7) (8) (9) (10), $w_{P(v_i),v_i}^t$ only depends on the edge weights of the interpolation tree T and can be pre-computed in linear complexity. Thus, only a total of 6 addition/subtraction operations and 7 multiplication/division operations are required for each node. So we can complete the interpolation task with $O(n)$ complexity disregarding the number K of reliable seeds. The interpolation tree T can be easily transformed from a MST with linear time by traversing the MST and inserting auxiliary nodes. To extract a MST, Karger *et al.* [20] present a linear time randomized algorithm. Moreover, with a linear number of processors, the parallel algorithm [19] for the minimum spanning tree problem can solve the problem in $O(\log n)$ time. Above all, we conclude that the overall complexity of our algorithm is $O(n)$.

IV. EXPERIMENTS

We implement our algorithm using Matlab on a Desktop computer with 4 GB memory. Parameter sensitivity and computation analysis are offered. We also exhibit the quantitative and visual evaluation of interpolation results.

A. Quantitative and Visual Evaluation

We compare our method with five methods which are JBL [8], JG [6], ATGV [22], NLA [12] and AR [4] respectively. For fair comparison, we subscribe the 2X, 4X, 8X, 16X upsampling results of Art and Book of the Middlebury Stereo Datasets from original authors. In the experiments, the low resolution depth images are obtained by downsampling and we use the percentage of bad matching pixels (PBP) [21] to evaluate the performance at discontinuity regions and continuous areas,
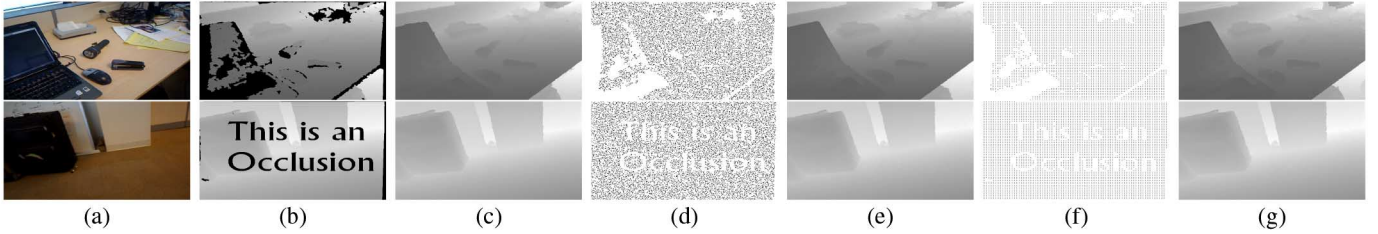
Fig. 2. Visual illustration of the interpolation results for real word data. (a) the registered guidance images. (b) the depth maps which are used as seed for structure missing. (c) the restoration results for structure missing. (d) the seed distribution for simultaneous 5% random missing and structure missing. (e) the inpainted results of (d). (f) the seed distribution for simultaneous 8X upsampling and structure missing. (g) the upsampled results of (f).

TABLE I
WE COMPARE THE PROPOSED ALGORITHM WITH STATE-OF-THE-ART
METHODS ON THE MIDDLEBURY DATASET USING PERCENTAGE OF BAD
MATCHING PIXELS [21], WHERE THE DISPARITY ERROR TOLERANCE IS 1. DISC
DENOTES THE PERCENTAGE OF BAD MATCHING PIXELS AT DISCONTINUOUS
REGIONS AND CONT STANDS FOR THE PERCENTAGE OF BAD MATCHING
PIXELS AT CONTINUOUS AREAS

		Art				Book			
		2X	4X	8X	16X	2X	4X	8X	16X
AR	DISC	0.15	0.24	0.32	0.41	0.12	0.23	0.34	0.39
	CONT	0.03	0.06	0.10	0.17	0.03	0.05	0.10	0.19
ATGV	DISC	0.13	0.20	0.28	0.36	0.10	0.22	0.30	0.33
	CONT	0.01	0.04	0.10	0.15	0.01	0.03	0.09	0.17
NLA	DISC	0.17	0.27	0.39	0.48	0.22	0.33	0.41	0.49
	CONT	0.01	0.04	0.07	0.12	0.02	0.04	0.08	0.18
JBL	DISC	0.16	0.24	0.39	0.49	0.19	0.27	0.41	0.49
	CONT	0.01	0.03	0.06	0.12	0.01	0.02	0.05	0.16
JG	DISC	0.11	0.19	0.30	0.36	0.06	0.18	0.18	0.43
	CONT	0.03	0.05	0.08	0.12	0.03	0.05	0.09	0.17
Ours	DISC	0.04	0.09	0.16	0.22	0.02	0.06	0.09	0.15
	CONT	0.01	0.03	0.07	0.12	0.01	0.04	0.08	0.13

where the discontinuity regions are obtained by dilating 1-pixel wide edge of ground truth and the rest of image domain Ω are considered as continuous areas. In the experiments, we assign $\sigma_m = 0.05$ consistently.

The quantitative DISC (i.e the PBP at discontinuous regions) and CONT (i.e the PBP at continuous areas) evaluations are listed in Table I. We find that the CONT of different methods are almost same, and the DISC indices, which reflect the edge preserving ability, distinguish different methods. Benefitting from the geometric-aware minimax path, our method achieves the lowest DISC ranks among all the methods as we expected.

We also use our method to process the depth map of Microsoft Kinect. We input the registered guidance images in Fig. 2(a) and the deteriorated depth maps in Fig. 2(b). The restoration results are illustrated in Fig. 2(c). We also exhibit the results of 5% random missing inpainting and 8X upsampling in Fig. 2(e), 2(g) respectively. The seed distributions of both kinds of deterioration, which suffer from structure missing as Fig. 2(b) does, are shown in Fig. 2(d), 2(f) respectively. From Fig. 2, we conclude that our method can produce satisfactory results for real world data and can be used in the working environment.

B. Parameter Sensitivity & Computation Analysis

The parameter setting of our method is simpler than other depth upsampling methods because the only parameter is σ_m that adjusts the similarity of coupled nodes. With varying σ_m, we record the PBP of our method to draw the performance curve of PBP and report the statistical data in Fig. 3. The figure exhibits that the performance of our algorithm at the discontinuous regions and continuous areas are rather robust in a wide range
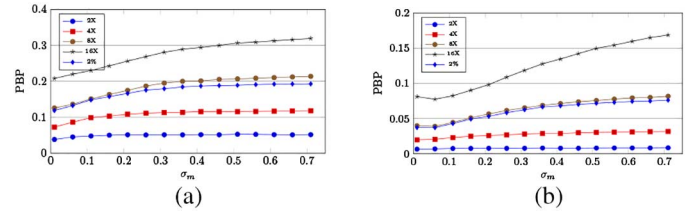

Fig. 3. PBP performance curves of discontinuous regions and continuous areas are shown as a function of σ_m. The proposed algorithm renders good performance for a wide range of σ_m under 2X, 4X, 8X, 16X upsampling rates and 2% random missing inpainting (a) Discontinuous regions (b) Continuous areas.
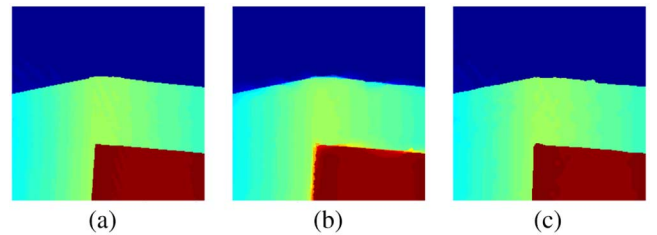

Fig. 4. Close-ups of the 8X upsampled result of book. For clarity, we visualize the image intensities using a color map. We can observes that our method can keep the sharp depth edges satisfactorily (a) Ground truth (b) Yang's (c) Ours.

of σ_m, therefore we do not need to fine tweak the scale factor σ_m for pursuing stable results.

We use a 690×560 image to test the efficiency of our method. Our method's computational time is stable at 8 fps for 2% random missing inpainting and 2x, 4x, 8x upsampling. In contrast to the fastest upsampling method [6] which achieves 0.03 fps for exact method and 3 fps for approximation approach, our algorithm outperforms it significantly.

C. Comparison with Yangs Method

This section is devoted to explain why our method is better than Yang's, since our method is derived from and similar to the method of Yang [12]. Fig. 4 shows the comparison results of Book. We can observe that our method keeps the sharp depth edges much better. We own this achievement to the reliable seed choosing strategy. Specifically, Yang's aggregation method does not distinguish the different pixels and thus each pixel receives the contribution of all other pixels indiscriminately. On the contrary, our method only takes advantage of reliable seeds (or the closest points in the sense of minimax distance). In this way, our method successfully rules out the negative effects of the unreliable depths that are far from the target pixels.

References

[1] M. Pollack, "The maximum capacity through a network," *Oper. Res.*, vol. 8, pp. 733–736, 1960.

[2] J. Diebel and S. Thrun, "An application of markov random fields to range sensing," in *NIPS 2005*, 2005.

[3] J. Park, H. Kim, Y.-W. Tai, M. Brown, and I. Kweon, "High quality depth map upsampling for 3d-tof cameras," in *ICCV 2011*, 2011.

[4] J. Yang, X. Ye, and Li, "Depth recovery using an adaptive color-guided auto-regressive model," in *ECCV 2012*, 2012.

[5] D. Ferstl, C. Reinbacher, R. Ranftl, M. Ruether, and H. Bischof, "Image guided depth upsampling using anisotropic total generalized variation," in *ICCV*, 2013.

[6] M.-Y. Liu, O. Tuzel, and Y. Taguchi, "Joint geodesic upsampling of depth images," in *CVPR*, 2013.

[7] J. Dolson, J. Baek, C. Plagemann, and S. Thrun, "Upsampling range data in dynamic environments," in *CVPR*, 2010.

[8] J. Kopf, M. F. Cohen, D. Lischinski, and M. Uyttendaele, "Joint bilateral upsampling," in *SIGGRAPH*, 2007.

[9] Q. Yang, R. Yang, J. Davis, and D. Nister, "Spatial-depth super resolution for range images," in *CVPR*, 2007.

[10] D. Chan, H. Buisman, C. Theobalt, and S. Thrun, A. Cavallaro and H. Aghajan, Eds., "A noise-aware filter for real-time depth upsampling," in *ECCV Workshop on Multi-camera and Multi-modal Sensor Fusion Algorithms and Applications*, Marseille, France, 2008, pp. 1–12.

[11] B. Huhle, T. Schairer, P. Jenke, and W. Straer, "Fusion of range and color images for denoising and resolution enhancement with a non-local filter," *Comput. Vis. Image Understand.*, vol. 114, no. 12, pp. 1336–1345, Dec. 2010.

[12] Q. Yang, "A non-local cost aggregation method for stereo matching," *Comput. Vis. Patt. Recognit.*, 2012.

[13] F. Fouss, A. Pirotte, J.-M. Renders, and M. Saerens, "Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation," *IEEE Trans. Knowl. Data Eng.*, vol. 19, no. 3, pp. 355–369, 2007.

[14] L. Yen, M. Saerens, A. Mantrach, and M. Shimbo, "A family of dissimilarity measures between nodes generalizing both the shortest-path and the commute-time distances," in *Proc. KDD*, 2008, pp. 785–793.

[15] M. Saerens, Y. Achbany, F. Fouss, and L. Yen, "Randomized shortest-path problems: Two related models," *Neural Comput.*, vol. 21, no. 8, pp. 2363–2404, Aug. 2009.

[16] A. Mantrach and Yen, "The sum-over-paths covariance kernel: A novel covariance measure between nodes of a directed graph," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 32, no. 6, pp. 1112–1126, Jun. 2010.

[17] A. Adams, N. Gelfand, J. Dolson, and M. Levoy, "Gaussian kd-trees for fast high-dimensional filtering," *ACM Trans. Graph.*, vol. 28, no. 3, pp. 21:1–21:12, Jul. 2009.

[18] T. C. Hu, "The maximum capacity route problem," *Oper. Res.*, vol. 9, pp. 898–900, 1961.

[19] K. W. Chong, Y. Han, and T. W. Lam, "Concurrent threads and optimal parallel minimum spanning trees algorithm," *J. ACM*, vol. 48, no. 2, pp. 297–323, Mar. 2001.

[20] R. E. T. David, R. Karger, and P. N. Klein, "A randomized linear-time algorithm to find minimum spanning trees," *J. ACM*, vol. 42, pp. 321–328, 1995.

[21] D. Scharstein and R. Szeliski, "A taxonomy and evaluation of dense two-frame stereo correspondence algorithms," *Int. J. Comput. Vision*, vol. 47, no. 1–3, pp. 7–42, Apr. 2002.

[22] D. Ferstl, C. Reinbacher, R. Ranftl, M. Ruether, and H. Bischof, "Image guided depth upsampling using anisotropic total generalized variation," in *IEEE Int. Conf. Computer Vision*, Dec. 2013, pp. 993–1000.