

Joint Graph-based Depth Refinement and Normal Estimation

Mattia Rossi*, Mireille El Gheche*, Andreas Kuhn†, Pascal Frossard*

*École Polytechnique Fédérale de Lausanne, †Sony Europe B.V.

Abstract

Depth estimation is an essential component in understanding the 3D geometry of a scene, with numerous applications in urban and indoor settings. These scenarios are characterized by a prevalence of human made structures, which in most of the cases are either inherently piece-wise planar or can be approximated as such. With these settings in mind, we devise a novel depth refinement framework that aims at recovering the underlying piece-wise planarity of those inverse depth maps associated to piece-wise planar scenes. We formulate this task as an optimization problem involving a data fidelity term, which minimizes the distance to the noisy and possibly incomplete input inverse depth map, as well as a regularization, which enforces a piece-wise planar solution. As for the regularization term, we model the inverse depth map pixels as the nodes of a weighted graph, with the weight of the edge between two pixels capturing the likelihood that they belong to the same plane in the scene. The proposed regularization fits a plane at each pixel automatically, avoiding any a priori estimation of the scene planes, and enforces that strongly connected pixels are assigned to the same plane. The resulting optimization problem is solved efficiently with the ADAM solver. Extensive tests show that our method leads to a significant improvement in depth refinement, both visually and numerically, with respect to state-of-the-art algorithms on the Middlebury, KITTI and ETH3D multi-view datasets.

1. Introduction

The accurate recovery of depth information in a scene represents a fundamental step for many applications, ranging from 3D imaging to the enhancement of machine vision systems and autonomous navigation. Typically, *dense depth estimation* is implemented either using active devices such as *Time-Of-Flight* cameras, or via *dense stereo matching* methods that rely on two [41, 13, 38, 2] or more [9, 15, 10, 32, 14, 40] images of the same scene to compute its geometry. Active methods suffer from noisy measurements, possibly caused by light interference or multiple reflections, therefore they can benefit from a post-processing

step to refine the depth map. Similarly, dense stereo matching methods have a limited performance in untextured areas, where the matching becomes ambiguous, or in the presence of occlusions. Therefore, a stereo matching pipeline typically includes a refinement step to fill the missing depth map areas and remove the noise.

In general, the refinement step is guided by the image associated to the measured or estimated depth map. The depth refinement literature mostly focuses on enforcing some kind of first order smoothness among the depth map pixels, possibly avoiding smoothing across the edges of the guide image, which may correspond to object boundaries [1, 40, 37]. Although depth maps are typically piece-wise smooth, first order smoothness is a very general assumptions, which does not exploit the geometrical simplicity of most 3D scenes. Based on the observation that most human made environments are characterized by planar surfaces, some authors propose to enforce second order smoothness by computing a set of possible planar surfaces a priori and assigning each depth map pixel to one of them [24]. Unfortunately, this refinement strategy imposes to select a finite number of plane candidates a priori, which may not be optimal in practice and lead to reduced performance.

In this article we propose a depth map refinement framework, which promotes a piece-wise planar arrangement of scenes without any a priori knowledge of the planar surfaces in the scenes themselves. We cast the depth refinement problem into the optimization of a cost function involving a data fidelity term and a regularization. The former penalizes those solutions deviating from the input depth map in areas where the depth is considered to be reliable, whereas the latter promotes depth maps corresponding to piece-wise planar surfaces. In particular, our regularization models the depth map pixels as the nodes of a weighted graph, where the weight of the edge between two pixels captures the likelihood that their corresponding points in the 3D scene belong to the same planar surface.

Our contribution is twofold. On the one hand, we propose a graph-based regularization for depth refinement that promotes the reconstruction of piece-wise planar scenes explicitly. Moreover, thanks to its underneath graph, our regularization is flexible enough to handle non fully piece-wise

planar scenes as well. On the other hand, our regularization is defined in order to estimate the normal map of the scene jointly with the refined depth map.

The proposed depth refinement and normal estimation framework is potentially very useful in the context of large scale 3D reconstruction [10, 32, 14, 40, 39, 19, 20], where the large number of images to be processed requires fast dense stereo matching methods, whose noisy and potentially incomplete depth maps can benefit from a subsequent refinement [14, 40, 39, 19, 20]. It is also relevant in the 3D reconstruction fusion step, when multiple depth maps must be merged into a single point cloud and the estimated normals can be used to filter out possible depth outliers [32, 20]. We test our framework extensively and show that it is effective in both refining the input depth map and estimating the corresponding normal map.

The article is organized as follows. Section 2 provides an overview on the depth map refinement literature. Section 3 motivates the novel regularization term and derives the related geometry. Section 4 presents our problem formulation and Section 5 presents our full algorithm. In Section 6 we carry out extensive experiments to test the effectiveness of the proposed depth refinement and normal estimation approach. Section 7 concludes the paper.

2. Related work

Depth refinement methods fall mainly into three classes: local methods, global methods and learning-based methods.

Local methods are characterized by a greedy approach. Tosi et al. [37] adopt a two step strategy. First, the input disparity map is used to compute a binary confidence mask that classifies each pixel as reliable or not. Then, the disparity at the pixels classified as reliable is kept unchanged and used to infer the disparity at the non reliable ones, using a wise interpolation heuristic. In particular, for each non reliable pixel, a set of anchor pixels with a reliable disparity is selected and the pixel disparity is estimated as a weighted average of the anchor disparities. Besides its low computational requirements, the method in [37] suffers two major drawbacks. On the one hand, pixels classified as reliable are left unchanged: this does not permit to correct possible pixels misclassified as reliable, which may bias the refinement of the other pixels. On the other hand, the method in [37], and local methods in general, cannot take advantage of the reliable parts of the disparity map fully, due to their local perspective.

Global methods rely on an optimization procedure to refine each pixel of the input disparity map jointly. Barron and Poole [1] propose the *Fast Bilateral Solver*, a framework that permits to cast arbitrary image related ill posed problems into a global optimization formulation, whose prior resembles the popular bilateral filter [36]. In [37] the Fast Bilateral Solver has been shown to be effective in the disparity

refinement task, but its general purposefulness prevents it from competing with specialized methods, even local ones like [37]. Global is also the disparity refinement framework proposed by Park et al. [24], which can be broken down into four steps. First, the input reference image is partitioned into super-pixels and a *local* plane is estimated for each one of them using RANSAC. Second, super-pixels are progressively merged into macro super-pixels to cover larger areas of the scene and a new *global* plane is estimated for each of them. Then, a *Markov Random Field (MRF)* is defined over the set of super-pixels and each one is assigned to one of four different classes: the class associated the local plane of the super-pixel, the class associated to the global plane of the macro super-pixel to which the super-pixel belongs, the class of pixels not belonging to any planar surface, or the class of outliers. The MRF employs a prior that enforces connected super-pixel to belong to the same class, thus promoting a global consistency of the disparity map. Finally, the parameters of the plane associated to each super-pixel are slightly perturbed, again within a MRF model, to allow for a finer disparity refinement. This method is the closest to ours in flavour. However, the a priori detection of a finite number of candidate planes for the whole scene biases the refinements toward a set of plane hypotheses that may either not be correct, as estimated on the input noisy and possibly incomplete disparity map, or not be rich enough to cover the full geometry of the scene.

Finally, recent learning based methods typically rely on a deep neural network which, fed with the noisy or incomplete disparity map, outputs a refined version of it [35, 25]. In [35] the task is split into three sub-tasks, each one addressed by a different network and finally trained end to end as a single one: detection of the non reliable pixels, gross refinement of the disparity map and fine refinement. Instead, Knöbelreiter and Pock [25] revisit the work of Cherabier et al. [4] in the context of disparity refinement. First, the disparity refinement task is cast into the minimization of a cost function, hence a global optimization, whose minimizer is the desired refined disparity map. However, the cost function is partially parametrized, rather than fully handcrafted. Then, the cost function solver can be unrolled for a fixed number of iterations, thus obtaining a network structure, and the parametrized cost function can be learned. Once the network parameters are learned, the disparity refinement requires just a network forward pass. Both the methods in [35] and [25] permit a fast refinement of the input disparity. However, due to their learning-based nature, they can fall short easily in those scenarios which differ from the ones employed at training time, as shown for the method in [25], which performs remarkably well in the Middlebury benchmark [31] training set, while quite poorly in the test set of the same dataset.

Our graph-based depth refinement framework, instead,

does not rely on any training procedure. It adopts a global approach which permits to compute a pair of consistent depth and normal maps jointly. Moreover, it does not need any a priori knowledge of the possible planar surfaces in the scene, as it automatically assigns a plane to each pixel based on its neighbors in the underneath graph. Finally, the proposed framework does not call for a separate handling of pixels belonging to planar surfaces and not, again thanks to the graph underneath.

3. Depth map model

In this section we investigate the relationship between a plane in the 3D space and its 2D depth map. In particular, we show that a plane imaged by a camera has an *inverse* depth map described by a plane as well, thus motivating a piece-wise planar model for the inverse depth map of those scenes where planar structures are prevalent. Let us consider a plane \mathcal{P} in the 3D scene in front of a pinhole camera. For the sake of simplicity and w.l.o.g., hereafter we assume the scene coordinate system to coincide with the camera one. In particular, the camera coordinate system is assumed to be left handed, with the Z axis aligned with the camera optical axis and pointing outside the camera. The plane \mathcal{P} can be described uniquely by a pair (P_0, \mathbf{n}_0) , where $P_0 = (X_0, Y_0, Z_0) \in \mathcal{P} \subset \mathbb{R}^3$ is a point of the plane and $\mathbf{n}_0 = (a_0, b_0, c_0) \in \mathbb{R}^3$, with $\|\mathbf{n}_0\|_2 = 1$, is the *plane normal* vector defining the orientation of the plane itself. Therefore, for all and only the points $P = (X, Y, Z) \in \mathcal{P}$, the following equation holds true:

$$\langle \mathbf{n}_0, (X, Y, Z) - (X_0, Y_0, Z_0) \rangle = 0, \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the dot product. Equivalently, Eq. (1) can be rewritten as follows:

$$a_0 X + b_0 Y + c_0 Z - \rho_0 = 0, \quad (2)$$

where $\rho_0 = \langle \mathbf{n}_0, (X_0, Y_0, Z_0) \rangle \in \mathbb{R}$. In the pinhole model, the 3D point $P = (X, Y, Z)$ is projected into the camera image plane at the pixel image coordinates $(x, y) \in \mathbb{R}^2$:

$$x = \frac{X}{Z} f^x + c^x, \quad y = \frac{Y}{Z} f^y + c^y, \quad (3)$$

where $(c^x, c^y) \in \mathbb{R}^2$ are the coordinates of the camera center of projection and f^x, f^y are the horizontal and vertical focal lengths, respectively. Solving for X and Y in Eq. (3), the plane equation in Eq. (2) can be expressed as a function of the image coordinates (x, y) and the corresponding depth $Z = Z(x, y)$:

$$\left(a_0 \frac{(x - c^x)}{f^x} + b_0 \frac{(y - c^y)}{f^y} + c_0 \right) Z - \rho_0 = 0. \quad (4)$$

Similarly, in image coordinates, ρ_0 reads as follows:

$$\rho_0 = \left(a_0 \frac{(x_0 - c^x)}{f^x} + b_0 \frac{(y_0 - c^y)}{f^y} + c_0 \right) Z_0,$$

where (x_0, y_0) is the projection of the point P_0 into the camera image plane.

Let us introduce the vector $u(x_0, y_0) = (u_0^x, u_0^y) \in \mathbb{R}^2$:

$$u_0^x = \frac{a_0}{\rho_0 f^x}, \quad u_0^y = \frac{b_0}{\rho_0 f^y}. \quad (5)$$

Using the vector $u(x_0, y_0)$ and introducing the *inverse depth* $d(x, y) = 1/Z(x, y)$ permits to rewrite Eq. (4) as follows:

$$d(x, y) = d(x_0, y_0) + \langle u(x_0, y_0), (x - x_0, y - y_0) \rangle. \quad (6)$$

A proof is provided in the supplementary material [30]. Eq. (6) can be interpreted as a first order Taylor expansion of the inverse depth map at the image coordinate (x_0, y_0) , such that $u(x_0, y_0)$ should be close to $\nabla d(x_0, y_0)$. In particular, Eq. (6) shows that the inverse depth $d(x, y)$ of every point $P \in \mathcal{P}$ is described by a plane, which passes through the point $(x_0, y_0, d(x_0, y_0))$ and has a plane normal vector $(u(x_0, y_0), -1) \in \mathbb{R}^3$.

We showed that the plane \mathcal{P} is represented equivalently by the pair (P_0, \mathbf{n}_0) in the scene domain or by the pair $((x_0, y_0, d(x_0, y_0)), u(x_0, y_0))$ in the camera domain. Eq. (5) permits to move from the scene domain to the camera one by computing $u(x_0, y_0)$ when \mathbf{n}_0 is given. To recover \mathbf{n}_0 when $u(x_0, y_0)$ is given instead, we can solve the following non linear system in the variables a_0, b_0, c_0 :

$$\begin{cases} u_0^x = (\rho_0 f^x)^{-1} a_0 & (7a) \\ u_0^y = (\rho_0 f^y)^{-1} b_0 & (7b) \\ a_0^2 + b_0^2 + c_0^2 = 1, & (7c) \end{cases}$$

where the constraints in Eqs. (7a)–(7b) refer to Eq. (5) and Eq. (7c) is the normal unitary constraint $\|\mathbf{n}_0\|_2 = 1$. The closed form solution of the system is provided in the supplementary material [30]. In what follows, we present our optimization problem to jointly estimate the normal map u and the refined inverse depth map d .

4. Depth map refinement problem

Given an image $I \in \mathbb{R}^{H \times W}$ we are interested in recovering the corresponding depth map Z when only a noisy and possibly incomplete estimate \bar{Z} is available. We assume that \bar{Z} is provided together with a confidence mask m with entries in $[0, 1]$. In particular, the confidence map is such that, $\forall i = (x, y) \in \{0, \dots, W-1\} \times \{0, \dots, H-1\}$, $m_i = 0$ when the entry \bar{Z}_i is considered completely inaccurate, while $m_i = 1$ when \bar{Z}_i is considered highly accurate.¹ In the following, we focus on estimating the refined inverse depth map $d = 1/Z$ and the corresponding normal map,

¹A wide variety of algorithms addressing pixel-wise confidence prediction exist in the literature, either based on hand-crafted features or learning-based [27]. In practice, also the simple stereo reprojection error could be adopted [25].

given the initial estimate $\bar{d} = 1/\bar{Z}$ and the mask m . We consider the following optimization problem:

$$d^*, u^* = \underset{d, u}{\operatorname{argmin}} f(d) + \lambda g(d, u) \quad (8)$$

where $f(\cdot)$ is a data term, $g(\cdot)$ is a regularization term for piece-wise planar functions and $\lambda \in \mathbb{R}_{\geq 0}$ is a scalar weight factor. The refined depth map is eventually computed as $Z^* = 1/d^*$, while the 3D normal map n^* is obtained from u^* via the close form solution of the system in Eq. (7).

In more details, the data fidelity term $f(\cdot)$ enforces consistency between the inverse depth map estimate d and the input inverse depth map \bar{d} . We adopt a data term of the following form:

$$f(d) = \sum_i |d_i - \bar{d}_i| m_i,$$

which enforces that the estimated inverse depth map d is close to \bar{d} at those pixels i where the latter is considered accurate, i.e., where m_i tends to one.²

Then, the regularization term $g(\cdot)$ enforces the inverse depth map d to be piece-wise planar, according to the model developed in Section 3. In particular, we choose to model the inverse depth map as a weighted graph, where each pixel represents a node and where the weight of the edge between two pixels can be interpreted as the likelihood that the corresponding two points in the 3D scene belong to the same plane. Namely, if the image looks locally similar at two different pixels, the probability is large for these pixels to belong to the same physical object, hence the same plane. The regularization term parametrizes the inverse depth at each pixel with a different plane, but it enforces strongly connected pixels in the graph, i.e., those pixels connected by an edge with high weight, to share the same plane parametrization. Specifically, our regularization term $g(\cdot)$ encompasses two terms balanced by a scalar weight $\alpha \in \mathbb{R}_{>0}$ and reads as follows:

$$g(d, u) = \sum_i \sqrt{\sum_{j \sim i} w_{ij}^2 (d_j - d_i - \langle u_i, j - i \rangle)^2} \quad (9a)$$

$$+ \alpha \sum_i \sum_{j \sim i} w_{ij} \|u_j - u_i\|_2, \quad (9b)$$

where $\{j \sim i\}$ describes the direct neighbours of i in the graph and $w_{ij} \in \mathbb{R}_{>0}$ is the weight associated to the edge between the pixel i and its neighbour pixel j .

The first term of the regularization in Eq. (9a) enforces the following constraint between the pixel i and its neighboring pixel j , for every $j \sim i$:

$$d_j = d_i + \langle u_i, j - i \rangle,$$

²The quality of the confidence map m can affect the quality of the refined depth map. However, in the case of missing confidence, i.e., m constant, our formulation in Eq. (8) still promotes piece-wise planar scenes.

which requires the inverse depth map in the neighbourhood of the pixel i to be approximated by the plane \mathcal{P}_i whose orientation is given by the vector $u_i \in \mathbb{R}^2$. This constraint recalls Eq. (6) and it is weighted by the likelihood w_{ij} that i and j are the projections of two points P_i and $P_j \in \mathbb{R}^3$ belonging to the same plane \mathcal{P}_i in the 3D scene. However, using only Eq. (9a) does not guarantee that the plane \mathcal{P}_j , fitted at the pixel j , and the plane \mathcal{P}_i are the same (e.g., the plane normal vectors u_j and u_i coincide). Therefore, the second term of the regularization in Eq. (9b) enforces that the two planes fitted at i and j , with orientations u_i and u_j , respectively, are consistent with each other when P_i and P_j are considered likely to belong to the same plane \mathcal{P} .

We conclude by observing that Eq. (9b) can be interpreted as a generalization of the well-known anisotropic *Total Variation (TV)* regularization [6], typically referred to as *Non Local Total Variation (NLTV)* [12] in general graph settings. In fact, the quantity $\|u_j - u_i\|_2$ can be interpreted as the magnitude of the derivative of u at the node i in the direction of the node j [34], so that Eq. (9b) enforces a piece-wise constant signal u on the graph, which enforces the signal d to be piece-wise planar. This corresponds to the depth map model of Section 3.

5. Depth refinement algorithm

In this section we present the structure of the graph underneath the regularization in Eqs. (9a)–(9b). Then, we detail the optimization algorithm adopted to find the solution of the joint depth refinement and normal estimation problem presented in Eq. (8).

5.1. Graph construction

We assume that areas of the image I sharing the same texture correspond to the same object and likewise to the same planar surface in the 3D scene. Based on this assumption, we associate a weight to the graph edge (i, j) , which quantifies our confidence about the two pixels i and j to belong to the same object. Formally, first we define a $B \times B$ pixels search window centered at the pixel i . Then, for each pixel j in the window we compute the following weight:

$$w_{ij} = \exp\left(-\frac{\|Q_i - Q_j\|_F^2}{2\sigma_{int}^2}\right) \exp\left(-\frac{\|i - j\|_2^2}{2\sigma_{spa}^2}\right), \quad (10)$$

where $Q_i \in \mathbb{R}^{Q \times Q}$ is a patch centered at the pixel i of the image I , $\|\cdot\|_F$ denotes the Frobenius norm and $\sigma_{int}, \sigma_{spa} \in \mathbb{R}_{>0}$ are tunable parameters. The first exponential in Eq. (10) has a high weight, hence high likelihood, when the values of the image pixels in two patches centered at i and j are similar; it is low otherwise [3, 17, 8, 29]. The second exponential then makes the weight decay as the Euclidean distance between i and j increases.

After the weights associated to all the pixels in the considered $B \times B$ search window have been computed, we design a K *Nearest Neighbours* graph by keeping only the $K \in \mathbb{N}$ edges with the largest weights. Limiting the number of connections at each pixel to K reduces the computation during the minimization of the problem in Eq. (8), on the one hand, and it avoids weak edges that may connect pixels belonging to different objects, on the other one.

5.2. Solver

The problem in Eq. (8) is convex, but non smooth. Multiple solvers specifically tailored for this class of problems exist, such as the *Forward Backward Primal Dual (FBPD)* solver [5]. However, the convergence of these methods calls for the estimation of multiple parameters before the actual minimization takes place, such as the operator norm of the implicit linear operator associated to the regularization term in Eqs. (9a)–(9b), which can be very time demanding. Therefore, we decide to solve the problem in Eq. (8) using *Gradient Descent with momentum*, in particular *ADAM* [16], as we empirically found it to be considerably faster (time-wise) than FBPD in our scenario.

Overall, our algorithm consists of two tasks: the graph construction and the solution of the problem in Eq. (8) with ADAM. Resorting to [7], the graph construction has a complexity $O(HWB \log_2 B)$ regardless of the patch size Q . The complexity of a single ADAM iteration is $O(HWK)$, with $K \ll B^2$, and it is due to the gradient computation.

Finally, to further speed up ADAM convergence, we adopt a multi-scale approach. The noisy and possibly incomplete inverse depth map \bar{d} is progressively down-sampled by a factor $r \in \mathbb{N}$ to get $\bar{d}^\ell \in \mathbb{R}^{\lfloor H/r^\ell \rfloor \times \lfloor W/r^\ell \rfloor}$ with $\ell = 0, \dots, L-1$ and $L \in \mathbb{N}$ the number of scales. An instance of the problem in Eq. (8) is solved for each \bar{d}^ℓ and the solution at the scale ℓ is up-sampled and scaled by a factor r to initialize the solver at the scale $\ell-1$.³ The scaling is a consequence of the relation $u^{\ell-1} = r^{-1}u^\ell$. In fact, the up/down-sampling operations emulate a change of the pixel area, while the camera sensor area remains constant. We refer to the supplementary material [30] for a formal proof.

6. Experimental results

We test the effectiveness of our joint depth refinement and normal estimation framework on the training splits of the *Middlebury v3* stereo dataset [31] at quarter resolution, of the *KITTI 2015* stereo dataset [23] and of the *ETH3D Multi-View Stereo (MVS)* dataset [33] at half resolution. Since these datasets come with ground truth depth maps but lack ground truth normals, we provide numerical results for

the depth refinement part of the framework, while we provide only visual results for the normal estimation part.

Regarding the ground truth normal map n_{gt} , we approximate it by solving the system in Eq. (7) with $u = \nabla d_{gt}$, where the gradient is computed using a 5×5 pixels Gaussian derivative kernel with standard deviation $\sigma = 0.2$ pixels. The small standard deviation permits to recover fine details, as the ground truth inverse depth map d_{gt} is not affected by noise. Although this does not permit a numerical evaluation, it permits to appreciate the normals estimated by our framework.

6.1. Middlebury and KITTI datasets

Similarly to the recent disparity refinement method in [37], we refine the disparity maps computed via *Semi-Global Matching (SGM)* [13] and census-based *Block Matching (BM)* [41]. We compare our framework to the disparity refinement method recently proposed in [37], as it also relies on a confidence map and, most importantly, it showed to outperform many other widely used disparity refinement methods, e.g., [26, 22, 21, 42, 1], on both the Middlebury and the KITTI datasets. Moreover, since our new regularization in Eqs. (9a)–(9b) resembles NLTGV [28], we compare to NLTGV as well. In particular, we replace $g(\cdot)$ with NLTGV in our problem formulation in Eq. (8).

It is crucial to observe that, originally, NLTGV was introduced in the context of optical flow [28] as a general purpose regularization, without any ambition to connect the geometry of the optical flow and the geometry of the underneath scene. Here instead, we aim at modeling the joint piece-wise planarity of the inverse depth map and of the underneath scene explicitly. In fact, the mixed $\ell_{1,2}$ -norms employed in both the terms of our regularization, as opposed to the simple ℓ_1 -norm of NLTGV, are chosen to make our regularization more robust in its global plane fitting.⁴

The SGM and BM disparity maps to refine are provided by the authors in [37], who provided also their refined disparity maps and binary confidence maps. In order to carry out a fair comparison, these confidence maps are used by all the methods considered in the experiments. As described in [37], the considered binary confidence maps are the result of a learning-based framework trained on a split of the *KITTI 2012* stereo dataset [11], therefore there is no bias toward the Middlebury and KITTI 2015 datasets.

Since our framework assumes a depth map at its input, we convert the disparity map to be refined into a depth map and we then convert the refined depth map back to the disparity domain, in order to carry out the numerical evaluation. The evaluation involves the *bad pixel* metric, which is the percentage of pixels with an error larger than a predefined disparity threshold, together with the *average absolute*

³All up-sampling and down-sampling operations are performed using nearest neighbor interpolation.

⁴A thorough analysis of the differences between the proposed regularization and NLTGV is provided in the supplementary material [30].

error (*avgerr*) and the *root mean square error* (*rms*). We carry out the evaluation on all the pixels with an available ground truth, regardless of the occlusions.

For a fair comparison, in the graph construction we adopt the same parameters for both NLTGV and our framework: weight parameters $\sigma_{int} = 0.07$ and $\sigma_{spa} = 3$ pixels, search window size $B = 9$, patch size $P = 3$ and maximum number of per pixel connections $K = 20$. For both, we also set the number of scales $L = 2$ and $r = 2$. Instead, the multipliers λ and α in front and inside the regularization $g(\cdot)$, respectively, are the result of a grid search and listed below.

Middlebury dataset The Middlebury training dataset [31] consists of 15 indoor scenes carefully crafted to challenge modern stereo algorithms. Some scenes contain multiple untextured planar surfaces, which represent a hard challenge for stereo methods but are compliant with the model underneath our framework; other scenes are inherently non piece-wise planar instead. Due to its variety, the Middlebury dataset permits to evaluate the flexibility of our framework to different settings.

For NLTGV we set $\lambda = 7.5$ and $\alpha = 50$, regardless of the scale. For our framework and SGM disparity maps at the input, we set $\lambda = 15$ and 25 at the low and high scales, respectively; for BM disparity maps at the input instead, we set $\lambda = 10$ and 20 at the low and high scales, respectively; we set $\alpha = 3.5$ regardless of the input disparity map.

The results of our experiments on the Middlebury dataset are presented in Table 1. When BM is considered, our framework outperforms the method in [37] and NLTGV in all the considered metrics. Similarly, when SGM is considered, our framework outperforms the method in [37] and NLTGV in four of the five metrics; in the *bad 1px* metric, where the best error is achieved by the method in [37], our result is comparable. Moreover, in the most common *bad 2px* metric, our framework always provides the best error regardless of the input disparity map. Clearly, some scenes in the dataset are far from fulfilling our piece-wise planar assumption, e.g., *Jadeplant* and *Pipes*: these affect the average results in Table 1 and mask the large improvement exhibited by our framework in those scenes which fulfill the assumption even partially.

In Figure 1 we provide the results of our experiments on the scene *Piano*, when the stereo method BM is considered. The normal map associated to the input BM disparity map and to the one refined by the method in [37] are computed with the same approach adopted for the ground truth normal map, while employing $\sigma = 5$ pixels in order to handle the noise. In fact, the input BM disparity is significantly noisy, especially in the walls surrounding the piano. The method in [37] manages to decrease the error in some areas of the surrounding walls: however, since no global consistency is considered, the result is a speckled error. Instead,

Table 1. Disparity refinement on the Middlebury dataset [31]. The first column specifies the stereo method whose disparity map is refined. The second column provides the error metric used in the evaluation: *bad px* thresholds, the *average absolute error* (*avgerr*) and the *root mean square error* (*rms*). All the pixels with a ground truth disparity are considered. The columns from four to six report the error of the disparity maps refined by the method in [37], NLTGV [28], our method. The best result for each row is in bold.

	Err. metric	Input	[37]	[28]	Ours
SGM [13]	bad 0.5px	41.33	39.14	36.57	35.70
	bad 1px	28.90	25.58	26.02	25.71
	bad 2px	23.48	19.55	19.88	19.25
	avgerr	4.06	3.32	3.31	2.87
	rms	9.75	8.27	7.99	6.86
BM [41]	bad 0.5px	47.48	39.01	38.49	35.01
	bad 1px	37.56	25.83	28.28	25.40
	bad 2px	33.98	20.61	22.03	19.41
	avgerr	8.41	3.48	3.35	2.79
	rms	17.32	8.58	7.91	6.97

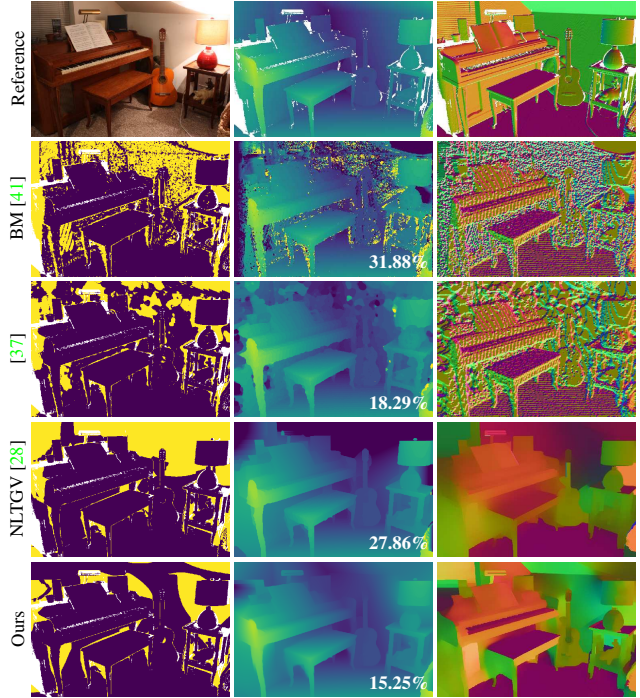


Figure 1. Middlebury [31] scene *Piano*. The first row hosts, from left to right, the reference image and the ground truth disparity and normal maps. Each other row hosts, from left to right, the *bad 2px* disparity error mask and the disparity and normal map. The second row refers to the stereo method BM [41], whose disparity is refined by the method in [37], NLTGV [28] and ours, in the rows three to five, respectively. The pixels in the error maps are color coded: error within 2px in dark blue, error larger than 2px in yellow, missing ground truth in white. The *bad 2px* error percentage is reported on the bottom right corner of each disparity map.

our method manages to approximate the surrounding walls better, using multiple planes. Finally, NLTGV fails to capture the geometry of the surrounding wall, as its relying on

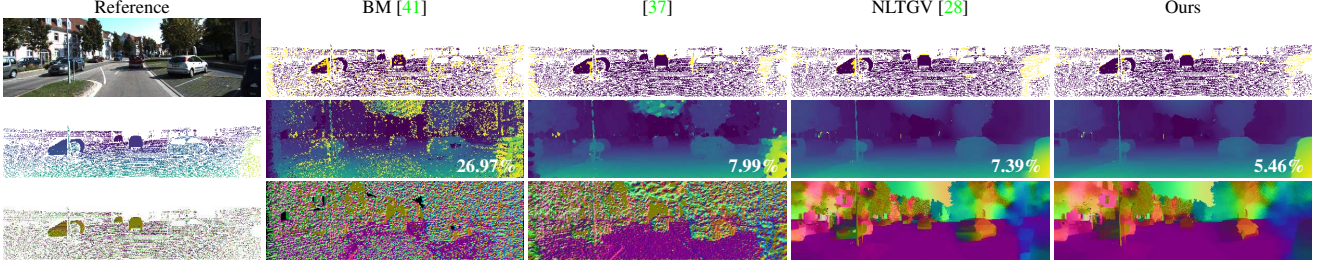


Figure 2. KITTI [23] scene 126. The first column hosts, from top to bottom, the reference image and the ground truth disparity and normal maps. Each other column hosts, from top to bottom, the *bad 3px* disparity error mask and the disparity and normal map. The second column refers to the stereo method BM [41], whose disparity is refined by the method in [37], NLTGV [28] and ours, in the columns three, four and five, respectively. The pixels in the error maps are color coded: error within 3px in blue, error larger than 3px in yellow, missing ground truth in white. The *bad 3px* error percentage is reported on the bottom right corner of each disparity map.

a simple ℓ_1 -norm makes it more sensible to outliers than our mixed $\ell_{1,2}$ -norm when trying to fit planes.

KITTI dataset The KITTI 2015 training dataset [23] consists of 200 scenes captured from the top of a moving car. As a consequence, the prevalent content of each scene are the road, possible vehicles and possible buildings at the two sides of the road. At a first glance, this content may seem to match our piece-wise planar assumption. In practice, however, the buildings at the sides of the road are mostly occluded by the vegetation, which is far from piece-wise planar. We select 20 scenes randomly and test our framework on them, in order to analyze its flexibility.

For NLTGV we set $\lambda = 7.5$ and $\alpha = 15$ regardless of the scale. For our framework we set $\lambda = 10$ and 20 at the lowest and highest scales, respectively, while we set $\alpha = 15$ regardless of the scale.

The results of our experiments on the KITTI dataset are presented in Table 2. Regardless of the considered metric and input stereo method, NLTGV outperforms the method in [37], while our framework outperforms all the others. Moreover, when the most common *bad 3px* error is considered, our framework improves the SGM and BM disparity maps by more than 4.57% and 31.75%, respectively.

In Figure 2 we provide the results of our experiments on the scene 126, when the stereo method BM is considered. The method in [37], NLTGV and our framework manage all to reduce sensibly the high amount of noise that affects the input disparity map, represented by the yellow speckles. However, only NLTGV and our framework manage to preserve fine details like the pole on the left side of the image, which appears broken in the disparity map associated to [37]. Finally, our framework provides the sharpest disparity map, as NLTGV exhibits some disparity bleeding at object boundaries. This is visible on the car at the bottom right corner of the image, both by observing the disparity maps and the error masks. This is also confirmed by the numerical results, as our *bad 3px* error is significantly lower.

Table 2. Disparity refinement on the KITTI dataset [23]. The first column specifies the stereo method whose disparity map is refined. The second column specifies the considered error metric: *bad px* thresholds, the *average absolute error* (*avgerr*) and the *root mean square error* (*rms*). All the pixels with a ground truth disparity are considered. The columns from four to six report the error of the disparity maps refined by the method in [37], NLTGV [28] and our method, respectively. The best result for each row is in bold.

	Err. metric	Input	[37]	[28]	Ours
SGM [13]	bad 2px	14.25	11.58	10.49	9.82
	bad 3px	10.11	7.65	6.07	5.54
	avgerr	21.12	2.97	1.62	1.51
	rms	46.50	8.49	7.91	7.88
BM [41]	bad 2px	40.96	16.75	11.09	10.54
	bad 3px	38.15	12.80	6.76	6.40
	avgerr	1.94	1.63	1.30	1.22
	rms	5.46	4.52	3.43	3.15

6.2. ETH3D dataset

Large scale 3D reconstruction methods [10, 32, 18, 14, 40, 39, 20] estimate the depth map of a reference image from a large number of input images of the same scene, by leveraging geometric and photometric constraints, and subsequently fuse them to produce a model of the scene itself. Large scale 3D reconstruction methods can largely benefit from a refinement of the estimated depth maps and can exploit the corresponding normal maps during the fusion step. In order to demonstrate the suitability of our joint depth refinement and normal estimation framework on high resolution images from challenging MVS configurations, we test it on the training split of the ETH3D dataset [33], a popular benchmark for large scale 3D reconstruction algorithms, involving both indoor and outdoor sequences.

The ground truth ETH3D depth maps are very sparse, but characterized by twice the resolution adopted in our tests. Therefore, similarly to [14], we back project the sparse ground truth depth maps to half resolution in order to get denser ones, to be used in our evaluation. In [20], the authors propose a novel deep-network-based confidence prediction framework for depth maps computed by MVS algo-

Table 3. Refinement of MVS-derived [20] depth maps from the ETH3D training dataset [33]. The table is divided into a top and a bottom sub-table, with their first columns specifying the test scenes, whose number of images is specified in brackets. The top sub-table reports the percentage of pixels with an error exceeding a predefined threshold: 2cm and 5cm. The bottom sub-table reports the *average absolute error* (*avgerr*) and the *root mean square error* (*rms*) in the second and third columns, respectively. For each scene and error metric, the best result is in bold.

	2cm			5cm		
	Input	[28]	Ours	Input	[28]	Ours
Pipes (14)	18.16	11.17	10.71	14.18	7.64	7.10
Delivery (44)	24.15	19.20	18.33	12.05	6.41	5.80
Office (26)	47.54	39.34	38.59	39.32	30.23	29.13
avg.	29.95	23.24	22.54	21.85	14.76	14.08

	avgerr			rms	
	Input	[28]		Input	[28]
Pipes (14)	0.347	0.119	0.082	2.090	0.685
Delivery (44)	0.233	0.025	0.023	26.06	0.227
Office (26)	0.330	0.183	0.167	1.218	0.409
avg.	0.303	0.109	0.091	9.303	0.440

gorithms, hence in the context of large baselines and severe occlusions. For our experiments, in order to estimate the confidence map m , we re-train the network proposed in [20] jointly on the synthetic dataset proposed in the same work and on the dense ground truth depth maps of the ETH3D training split. For an unbiased evaluation, we extract three sequences of the ETH3D training split (Pipes, Office, Delivery Area) from the training procedure and use them exclusively for our evaluation. We compare our refined depth and normal maps against the depth and normal maps derived by the *PatchMatch*-based [2] MVS method presented in [20].

For both NLTGV and our framework we set $\lambda = 7.5$ and $\alpha = 7.5$ regardless of the scale, adopt the graph parameters selected for Middlebury and KITTI, set the number of scales $L = 4$ and $r = 2$. The continuous confidence map provided by the trained network is binarized with a 0.5 threshold.

Table 3 compares the input MVS depth map with those refined by NLTGV and our method. The top part of the table reports the percentage of pixel, computed over all the pixels of all the images in the sequence, with an error within a given threshold: 2cm and 5cm. On average, our method outperforms NLTGV and manages to improve the input depth maps by more than 7% when the 2cm threshold, the most common in the ETH3D benchmark, is considered. In the bottom part of the same table, we provide also the *average absolute error* (*avgerr*) and the *root mean square error* (*rms*). The *rms* metric is very sensitive to outliers and, especially in the *Delivery Area* sequence, it highlights our improvement over the input depth map. Finally, a visual example is provided in Figure 3 for the sequence *Pipes*, which is characterized by multiple untextured planar surfaces, representing a hard challenge for MVS meth-

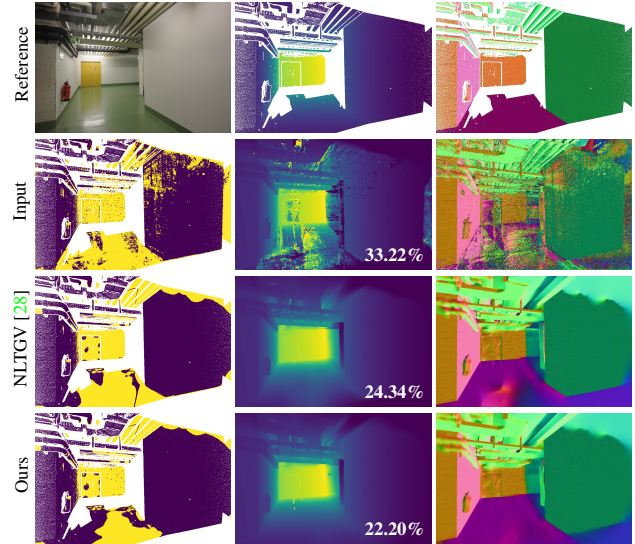


Figure 3. ETH3D [33] *Pipes*. The first row hosts, from left to right, the reference image and the ground truth depth and normal maps. Each other row hosts, from left to right, the 2cm error map, the depth and normal maps. The second row refers to the MVS method [20], whose depth is refined by NLTGV [28] and our method in the rows three and four, respectively. The pixels in the error maps are color coded: error within 2cm in blue, error larger than 2cm in yellow, missing ground truth in white. The error percentage associated to the 2cm threshold is reported on the bottom right corner of each depth map.

ods. Our method targets exactly these scenarios instead: in fact, it manages to refine the input depth map by capturing the main planes, as exemplified by the estimated normal map. Moreover, it manages to fit better planes than NLTGV, which fails to capture the correct floor orientation.

7. Conclusions

In this article, we presented a variational framework to address the problem of depth map refinement. In particular, we cast the problem into the minimization of a cost function involving a data fidelity and a graph-based regularization term. The latter enforces piece-wise planar solutions explicitly, by estimating the depth map and the corresponding normal map jointly, as most human made environments exhibit a planar bias. Moreover, the graph-based nature of the regularization makes our framework flexible enough to handle non fully piece-wise planar scenes as well. We showed that the proposed framework outperforms state-of-the-art depth refinement methods when the considered scene meets our piece-wise planar assumption, and it leads to competitive results otherwise. Interesting perspectives include the a priori segmentation of the reference image into *planar* and *non planar* areas, so that the strength of the regularization could be adapted accordingly.

References

- [1] Jonathan T. Barron and Ben Poole. The fast bilateral solver. In *European Conference on Computer Vision (ECCV)*, Amsterdam, The Netherlands, 2016. 1, 2, 5
- [2] Michael Bleyer, Christoph Rhemann, and Carsten Rother. PatchMatch stereo - Stereo matching with slanted support windows. In *British Machine Vision Conference (BMVC)*, Dundee, UK, 2011. 1, 8
- [3] Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. A review of image denoising algorithms, with a new one. *SIMUL*, 4:490–530, 2005. 4
- [4] Ian Cherabier, Johannes L. Schönberger, Martin R. Oswald, Marc Pollefeys, and Andreas Geiger. Learning priors for semantic 3D reconstruction. In *European Conference on Computer Vision (ECCV)*, Munich, Germany, 2018. 2
- [5] Laurent Condat. A primal-dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms. *Journal of Optimization Theory and Applications*, 158(2):460–479, Aug. 2013. 5
- [6] Laurent Condat. Discrete total variation: New definition and minimization. *SIAM Journal on Imaging Sciences*, 10(3):1258–1290, Aug. 2017. 4
- [7] Jérôme Darbon, Alexandre Cunha, Tony F. Chan, Stanley Osher, and Grant J. Jensen. Fast nonlocal filtering applied to electron cryomicroscopy. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI)*, Paris, France, 2008. 5
- [8] Alessandro Foi and Giacomo Boracchi. Foveated self-similarity in nonlocal image filtering. In Bernice E. Rogowitz, Thrasyvoulos N. Pappas, and Huib de Ridder, editors, *Human Vision and Electronic Imaging XVII*, volume 8291, pages 296–307. International Society for Optics and Photonics, SPIE, 2012. 4
- [9] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multi-view stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 32(8):1362–1376, 2010. 1
- [10] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *IEEE International Conference on Computer Vision (ICCV)*, Santiago, Chile, 2015. 1, 2, 7
- [11] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? The kitti vision benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Providence, RI, USA, 2012. 5
- [12] Guy Gilboa and Stanley Osher. Nonlocal operators with applications to image processing. *Multiscale Modeling & Simulation*, 7(3):1005–1028, Nov. 2009. 4
- [13] Heiko Hirschmüller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 30(2):328–341, Feb 2008. 1, 5, 6, 7
- [14] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. DeepMVS: Learning multi-view stereopsis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA, 2018. 1, 2, 7
- [15] Michal Jancosek and Tomas Pajdla. Multi-view reconstruction preserving weakly-supported surfaces. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Colorado Springs, CO, USA, 2011. 1
- [16] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 2015. 5
- [17] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *International Conference on Neural Information Processing Systems (NIPS)*, Granada, Spain, 2011. 4
- [18] Andreas Kuhn, Heiko Hirschmüller, Daniel Scharstein, and Helmut Mayer. A TV prior for high-quality scalable multi-view stereo reconstruction. *International Journal of Computer Vision (IJCV)*, 124(1):2–17, 2017. 7
- [19] Andreas Kuhn, Shan Lin, and Oliver Erdler. Plane completion and filtering for multi-view stereo reconstruction. In *German Conference on Pattern Recognition (GCPR)*, Dortmund, Germany, 2019. 2
- [20] Andreas Kuhn, Christian Sorzmann, Mattia Rossi, Oliver Erdler, and Friedrich Fraundorfer. DeepC-MVS: Deep confidence prediction for multi-view stereo reconstruction. *CoRR*, abs/1912.00439, 2020. 2, 7, 8
- [21] Ziyang Ma, Kaiming He, Yichen Wei, Jian Sun, and Enhua Wu. Constant time weighted median filtering for stereo matching and beyond. In *IEEE International Conference on Computer Vision (ICCV)*, Washington, DC, USA, 2013. 5
- [22] Stefano Matocchia. A locally global approach to stereo correspondence. In *IEEE International Conference on Computer Vision Workshops (ICCVW)*, Kyoto, Japan, 2009. 5

- [23] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, USA, 2015. 5, 7
- [24] Min-Gyu Park and Kuk-Jin Yoon. As-planar-as-possible depth map estimation. *Computer Vision and Image Understanding (CVIU)*, 181:50–59, Apr. 2019. 1, 2
- [25] Knöbelreiter Patrick and Pock Thomas. Learned collaborative stereo refinement. In *German Conference on Pattern Recognition (GCPR)*, Dortmund, Germany, 2019. 2, 3
- [26] Simon Perreault and Patrick Hebert. Median filtering in constant time. *IEEE Transactions on Image Processing (TIP)*, 16(9):2389–2394, Sep. 2007. 5
- [27] Matteo Poggi, Fabio Tosi, and Stefano Mattoccia. Quantitative evaluation of confidence measures in a machine learning world. In *IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, 2017. 3
- [28] René Ranftl, Kristian Bredies, and Thomas Pock. Non-local total generalized variation for optical flow estimation. In *European Computer Vision Conference (ECCV)*, Zurich, Switzerland, 2014. 5, 6, 7, 8
- [29] Mattia Rossi, Mireille El Gheche, and Pascal Frossard. A nonsmooth graph-based approach to light field super-resolution. In *IEEE International Conference on Image Processing (ICIP)*, Athens, Greece, 2018. 4
- [30] Mattia Rossi, Mireille El Gheche, Andreas Kuhn, and Pascal Frossard. Joint graph-based depth refinement and normal estimation. *CoRR*, abs/1912.01306, 2020. 3, 5
- [31] Daniel Scharstein, Heiko Hirschmüller, York Kitajima, Greg Krathwohl, Nera Nešić, Xi Wang, and Porter Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *German Conference on Pattern Recognition (GCPR)*, Münster, Germany, 2014. 2, 5, 6
- [32] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, Amsterdam, The Netherlands, 2016. 1, 2, 7
- [33] Thomas Schöps, Johannes L. Schönberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, Hawaii, USA, 2017. 5, 7, 8
- [34] David I. Shuman, Sunil K. Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):16. 83–98, Apr. 2013. 4
- [35] Gidaris Spyros and Komodakis Nikos. Detect, replace, refine: Deep structured prediction for pixel wise labeling. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, 2017. 2
- [36] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *IEEE International Conference on Computer Vision (ICCV)*, Bombay, India, 1998. 2
- [37] Fabio Tosi, Matteo Poggi, and Stefano Mattoccia. Leveraging confident points for accurate depth refinement on embedded systems. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Long Beach, CA, USA, 2019. 1, 2, 5, 6, 7
- [38] Oliver Woodford, Philip Torr, Ian Reid, and Andrew Fitzgibbon. Global stereo reconstruction under second-order smoothness priors. *IEEE Transactions on Pattern Analysis Machine Intelligence (TPAMI)*, 31(12):2115–2128, Dec. 2009. 1
- [39] Qingshan Xu and Wenbing Tao. Multi-scale geometric consistency guided multi-view stereo. *CoRR*, abs/1904.08103, 2019. 2, 7
- [40] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. MVSNet: Depth inference for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, Munich, Germany, 2018. 1, 2, 7
- [41] Ramin Zabih and John Woodfill. Non-parametric local transforms for computing visual correspondence. In *European Conference on Computer Vision (ECCV)*, Stockholm, Sweden, 1994. 1, 5, 6, 7
- [42] Qi Zhang, Li Xu, and Jiaya Jia. 100+ times faster weighted median filter. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Columbus, OH, USA, 2014. 5