



# Probabilistic multi-modal depth estimation based on camera–LiDAR sensor fusion

Johan S. Obando-Ceron<sup>3</sup> · Victor Romero-Cano<sup>1</sup> · Sildomar Monteiro<sup>2</sup>

Received: 17 February 2022 / Revised: 6 June 2023 / Accepted: 28 June 2023 / Published online: 29 July 2023  
© The Author(s) 2023

## Abstract

Multi-modal depth estimation is one of the key challenges for endowing autonomous machines with robust robotic perception capabilities. There have been outstanding advances in the development of uni-modal depth estimation techniques based on either monocular cameras, because of their rich resolution, or LiDAR sensors, due to the precise geometric data they provide. However, each of these suffers from some inherent drawbacks, such as high sensitivity to changes in illumination conditions in the case of cameras and limited resolution for the LiDARs. Sensor fusion can be used to combine the merits and compensate for the downsides of these two kinds of sensors. Nevertheless, current fusion methods work at a high level. They process the sensor data streams independently and combine the high-level estimates obtained for each sensor. In this paper, we tackle the problem at a low level, fusing the raw sensor streams, thus obtaining depth estimates which are both dense and precise, and can be used as a unified multi-modal data source for higher-level estimation problems. This work proposes a conditional random field model with multiple geometry and appearance potentials. It seamlessly represents the problem of estimating dense depth maps from camera and LiDAR data. The model can be optimized efficiently using the conjugate gradient squared algorithm. The proposed method was evaluated and compared with the state of the art using the commonly used KITTI benchmark dataset.

**Keywords** Sensor fusion · CRFs · LiDAR · Monocular camera

## 1 Introduction

Autonomous robots are composed of different modules that allow them to perceive, learn, decide and act within their envi-

ronment. The perception module processes cues that inform the robot about the appearance and geometry of the environment. Particularly when working in outdoor or underground scenarios, these cues must be robust to unseen phenomenon. A fully autonomous robot must execute all operations, monitor itself and be able to handle all unprecedented events and conditions, such as unexpected objects and debris on the road, unseen environments and adverse weather. Therefore, reliable and robust perception of the surrounding environment is one of the key tasks of autonomous robotics.

Among the main inputs for a perception system are the distances of the robot from multiple points in its environment. This input can be obtained directly from a sensor or estimated by a depth estimation module. Depth estimation can be performed by processing monocular camera images, stereo vision, radar or LiDAR (Light Detector and Ranging) sensors, among others. Although monocular cameras can only be used to generate depth information up-to-a-scale, they are still an important component of a depth estimation system due to their low price and the rich appearance data they provide. Although a monocular camera is small, low-

---

This document is a result of the research project 17INTER-297, funded by Universidad Autonoma de Occidente.  
Johan S. Obando-Ceron undertook this work while he was part of Universidad Autónoma de Occidente.

---

✉ Victor Romero-Cano  
varomero@uao.edu.co

Johan S. Obando-Ceron  
jobando0730@gmail.com

Sildomar Monteiro  
sildomar@ieee.org

<sup>1</sup> Robotics and Autonomous Systems Laboratory, Faculty of Engineering, Universidad Autónoma de Occidente, Valle del Cauca, Cali, Colombia

<sup>2</sup> Autonomy Research Division, Aurora Flight Sciences, Cambridge, MA, USA


<sup>3</sup> University of Montreal, Montréal, QC, Canada

cost and energy efficient, it is very sensitive to changes in the illumination. Additionally, the accuracy and reliability of depth estimation methods based on monocular images is still far from being practical. For instance, the state-of-the-art RGB-based depth prediction methods [12, 31, 36] produce an average error (measured by the root-mean-squared error) of over 50 cm in indoor scenarios (e.g. on the NYU-Depth-v2 dataset [61]). Such methods perform even worse outdoors, with at least 4 meters of average error on the Make3D and KITTI datasets [18, 58]. 3D LiDAR scanners, on the other hand, can provide accurate geometric information about the environment even when illumination changes occur. Even though, due to their active nature, LiDARs are robust to dark or overexposed scenarios, the generated 3D point cloud is sparse and non-uniformly distributed, which decreases its utility for recognition tasks. In general, each type of sensor has its own weaknesses.

To address the potential fundamental limitations of image-based depth estimation, this paper considers the use of sparse depth measurements, along with RGB data, to reconstruct depth in full resolution. Data fusion techniques have been extensively employed for robust perception systems, where fusing and aggregating data from different sensors is required. Although some approaches to robust perception resort to statistical methods for dealing with data outliers [2], the work presented in this paper belongs to the group that tackles the robust perception problem by leveraging the complementary nature of passive and active sensor modalities. Multi-sensor approaches to robotic perception can be categorized according to the level at which the data from the different sensing modalities is fused in order to obtain the estimate of interest. According to [7], data fusion can be made at the level of symbolic estimates (high-level fusion), at the level of features (medium-level fusion) or at the level of raw data (low-level fusion).

In this paper, a low-level fusion method is developed. It explores the complementary relations between the passive and active sensors at the pixel level [8, 14, 55]. The approaches in [51, 54, 63] follow this intuition but require the fused modalities to have similar coverage densities. Our proposed framework provides a procedure for fusing LiDAR and image data independently of the LiDAR data's density. The main contribution of this paper is a depth regression model that takes both a sparse set of depth samples and RGB images as the inputs and predicts a full-resolution depth map. This is achieved by modelling the problem of fusing low-resolution depth images with high-resolution camera images as a conditional random field (CRF).

The intuition behind our CRF formulation is that depth discontinuities in a scene often co-occur with changes in colour or brightness within the associated camera image. Since the camera image is commonly available at much higher resolution, this insight can be used to enhance the resolution and



**Fig. 1** Input features of our framework. We developed a CRF regression model to predict a dense depth image from a single RGB image, and a set of sparse depth samples: (a), (b) and (c) are the input RGB image, a set of sparse depth samples projected on the image plane and the projected surface normals, respectively

accuracy of the depth image. A depth map will be produced by our approach using three features as illustrated in Fig. 1. The first one is an RGB colour image from the camera sensor, top image in Fig. 1 (a). The second one is 2D sparse depth map captured by a LiDAR sensor, middle image in Fig. 1 (b). The third feature is a surface-normal map generated from the sparse depth samples, bottom image in Fig. 1 (c).

The rest of this paper is organized as follows. Section 2 reviews related work on depth estimation. Section 3 explains how the LiDAR points and the camera images were registered. In Sect. 4, we first introduce our CRF-Fusion framework, and then, we provide a detailed explanation of the proposed model: the energy potentials that compose our CRF model and its inference machine. The experimental validation, performed on the KITTI dataset, is reported in Sect. 5. Finally, conclusions and directions for future work are listed in Sect. 6.

## 2 Related work

Depth estimation from monocular images is a long-standing problem in computer vision. Early works on depth estimation using RGB images usually relied on handcrafted features and inference on probabilistic graphical models. Classical methods include shape-from-shading [72] and shape-from-defocus [62]. Other early methods were based on hand-tuned models or assumptions about the orientations of the surfaces [4]. Newer methods treat depth estimation as a machine learning problem, most recently using deep artificial neural

networks [13, 68]. For instance, Saxena et al. [57] estimated the absolute scales of different image patches and inferred a depth image using a Markov random field model. Eigen et al. used a multi-scale convolutional network to regress from colour images to depths [12, 13]. Laina et al. used a fully convolutional network architecture based on ResNet [32]. Liu et al. proposed a deep convolutional neural field model that combines deep networks with Markov random fields [38]. Roy et al. combined shallow convolutional networks with regression forests to reduce the need for large training sets [56]. In [69], the proposed attention model is seamlessly integrated with a CRF, allowing end-to-end training of the entire architecture. This approach benefits from a structured attention model which automatically regulates the amount of information transferred between corresponding features at different scales.

The approach of Li et al. [5] combines deep learning features on image patches with hierarchical CRFs defined on a superpixel segmentation of the image. They use pretrained AlexNet [29] features of image patches to predict depth at the centre of the superpixels. A hierarchical CRF refines the depth across individual pixels. Liu et al. [37] also propose a deep structured learning approach that avoids handcrafted features. They presented a deep structured learning scheme which learns the unary and pairwise potentials of a continuous CRF in a unified deep CNN framework. Liu et al. [39] proposed a discrete-continuous CRF model to take into consideration the relations between adjacent superpixels, e.g. occlusions.

Recent work has also shown the benefit of adopting multi-task learning strategies, e.g. for jointly predicting depth and performing semantic segmentation, ego-motion estimation or surface-normal computation [50, 76]. Some recent papers have proposed unsupervised or weakly supervised methods for reconstructing depth maps [21, 30]. With the rapid development of deep neural networks, monocular depth estimation based on deep learning and computer vision techniques has been widely studied recently and achieved promising performance in terms of accuracy [75]. However, not considering information from other sensors makes the estimate not so robust. In the mentioned literature, there are used different kinds of network frameworks, loss functions and training strategies with just one sensory modality. The architecture proposed in this paper uses two sensory modalities.

Fusing data coming from multiple sensors has the potential to improve the robustness of the depth estimates. Ma et al. [43] use RGB images together with sparse depth information to train a bottleneck network architecture. Compared to imagery-only methods, their approach generates better depth estimation results. Others have investigated depth estimation from colour images augmented with sparse sets of depth measurements using probabilistic graphical models. The techniques described in [3, 11, 71], and [67] are

able to fuse the information from both sources to significantly improve the resolution of low-quality and sparse range images. Wang et al. proposed a multi-scale feature fusion method for depth completion [66] using sparse LIDAR data. Ma et al. proposed two methods: a supervised method for depth completion using a ResNet-based architecture and a self-supervised method which uses the sparse LiDAR input along with pose estimates to add additional training information based on depth and photometric losses [42].

Although recent methods have achieved impressive progress in terms of evaluation metrics such as the pixel-wise relative error, most of them neglect the geometric constraints in 3D space. This component is considered in our CRF model, which makes this approach different from previous fusion methods.

Providing strong cues from surface information is relevant for improving the accuracy of the depth prediction [33]. Recently, Zhang et al. [73] proposed predicting surface normals and occlusion boundaries using a deep network and further used them to help depth completion in indoor scenes. The works in [53, 70] propose end-to-end deep learning systems to produce dense depth maps from sparse LiDAR data and a colour image taken from outdoor on-road scenes, leveraging surface normals as the intermediate representation. Zhang et al. [74] predicted surface normals by leveraging RGB data, leading to a better prior for depth completion. They ultimately combined these predictions with sparse depth input to generate a complete depth map. Our method takes advantage of the surface normals to improve the performance of the proposed model. [70], in particular, uses a diffusion layer to refine the completions. [40] shows that diffusion models can be interpreted as energy-based models (EBMs). Now, CRFs can also be interpreted as energy-based models parametrized as factor graphs. Therefore, there is a relation between the diffusion and the CRF approach as both can be interpreted as EBMs, but with different formulations and assumptions. For the problem of depth estimation, the CRF approach more naturally captures the prior knowledge of the geometric constraints inherent in the physics underlying the LiDAR data measurements, whereas the diffusion approach relies on a carefully crafted similarity metric, which might make the models unstable and hard to train. In other words, we expect a CRF approach to require less training to achieve equivalent performance as the diffusion-based method.

The problem of fusing LiDAR and image data can be approached as a camera pose estimation problem, where the relation between the 3D LIDAR coordinates and the 2D image coordinates is characterized by camera parameters, such as the position, orientation and focal length. In [47], there was proposed an information-theoretic similarity measure to automatically register 2D-Optical imagery with 3D LiDAR scans by searching for a suitable camera transforma-

tion matrix. The fusion of 3D LiDAR data with stereoscopic images is addressed in [23, 44, 48, 49]. The advantage of stereoscopic depth estimation is its ability to produce dense depth maps of the surroundings by using stereo matching techniques. In [49], for example, there was proposed a two-stage cascade deep architecture that first fuses stereo and LiDAR disparity maps and then refines the estimated disparity maps by introducing colour features. In contrast, our method does not rely on a stereo matching algorithm, something which tends to be computationally costly.

### 3 Image and LIDAR point cloud registration

This section gives a brief introduction to the process of aligning an image and a LIDAR point cloud, which allows the projection of LiDAR points onto the image plane. As presented in [19], in a robotic platform equipped with both a LiDAR and a camera, these two sensors are synchronized so simultaneous scans from both sensor are collected. The camera and LiDAR are cross-calibrated so that the point cloud can be projected onto the image plane [19]. Once projected, the LiDAR points are associated with either pixels or groups of pixel, also called superpixels. This section also briefly describes the Simple Linear Iterative Clustering (SLIC) algorithm, by which these superpixels are obtained.

#### 3.1 Point cloud projection

In order to fuse the image and LiDAR data, it is imperative to find mathematical models that represent the spatial correspondence of pixels and 3D points. These models allow us to project the LiDAR point cloud onto the image plane. The projection  $\mathbf{y}$  of a 3D point  $\mathbf{x} = (x, y, z, 1)^T$  in rectified and rotated camera coordinates to a point  $\mathbf{y} = (u, v, 1)^T$  in the  $i$ 'th camera image is given by

$$\mathbf{y} = \mathbf{P}_{rect}^{(i)} \mathbf{x}$$

with

$$\mathbf{P}_{rect}^{(i)} = \begin{pmatrix} f_u^{(i)} & 0 & c_u^{(i)} & -f_u^{(i)} b_x^{(i)} \\ 0 & f_v^{(i)} & c_v^{(i)} & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

being the  $i$ th projection matrix. Here,  $b_x^{(i)}$  denotes the baseline with respect to a reference camera. Note that in order to project the 3D point  $\mathbf{x}$ , in the camera reference coordinates, to a point  $\mathbf{y}$  on the  $i$ th image plane, the rectifying rotation matrix of the reference camera  $\mathbf{R}_{rect}^{(0)}$  must be considered as well.

$$\mathbf{y} = \mathbf{P}_{rect}^{(i)} \mathbf{R}_{rect}^{(0)} \mathbf{x}$$

Here,  $\mathbf{R}_{rect}^{(0)}$  has been expanded into a  $4 \times 4$  matrix by appending a fourth zero-row and column, and setting  $\mathbf{R}_{rect}^{(0)}(4, 4) = 1$ . We also need to register the laser scanner with respect to the camera's coordinate system. The rigid-body transformation from LiDAR coordinates to camera coordinates is given by

$$\mathbf{T}_{velo}^{cam} = \begin{pmatrix} \mathbf{R}_{velo}^{cam} & \mathbf{t}_{velo}^{cam} \\ 0 & 1 \end{pmatrix}$$

Finally, a 3D point  $\mathbf{x}$  in the LiDAR coordinate system gets projected to a point  $\mathbf{y}$  in the  $i$ th camera image:

$$\mathbf{y} = \mathbf{P}_{rect}^{(i)} \mathbf{R}_{rect}^{(0)} \mathbf{T}_{velo}^{cam} \mathbf{x}$$

Subsequently, as a preprocessing step, points with a negative value of  $z$  are removed. Then, the remaining points can be projected onto the image plane using the projection matrix

$$[x' y' z']^T = \mathbf{y} [x_p \ y_p \ z_p \ 1]^T$$


The projected pixel coordinates of the LIDAR points can be obtained by

$$[x, y] = \begin{bmatrix} x' / z' & y' / z' \end{bmatrix}$$

Figure 2 portrays the projection of a point cloud projection onto the image plane.

#### 3.2 Superpixel segmentation using Simple Linear Iterative Clustering (SLIC)


Considering superpixels rather than individual pixels greatly reduces the complexity of the subsequent image processing tasks. In order to harness the full potential of the use of superpixels, their calculation must be fast, easy to use, and produce high-quality segmentations. Superpixel segmentation algorithms provide an effective way to extract information about local image features. Our framework uses an implementation of the SLIC algorithm to group sparse depth measurements which have been previously projected onto the image plane. SLIC is a simple and parallelizable pixel clustering method, based on the  $k$ -means algorithm, which is used for decomposing an image into a regular grid of visually homogeneous regions or so-called superpixels [1]. As a result, SLIC superpixels provide a regular grouping of image pixels according to their distance both spatially and in the colour space. We use this superpixel segmentation method to assign depth values from a sparse point cloud to all of the pixels within the superpixels. Applying SLIC to the original image provides a segmentation with as many segments as the number of superpixels set as a hyperparameter. Each segment is identified by an ID that allows individual pixels to be assigned to




**Fig. 2** Given an (a) input image and its corresponding depth, where dark blue indicates closer distances, our work is focused on densifying the input sparse depth. We combine the (b) superpixel segmentation with the sparse input depth to obtain our initial superpixel depth. We use a CRF to enhance the resolution and accuracy of the depth image considering information from the 4 neighbours for each superpixel. In (d), we can see how we achieve significant improvements with regard to the sparse input depth. Each pixel inside the superpixels has the same depth value

superpixel segments. Additionally, the coordinates of the segment’s centroid and those of its neighbours are also output by a super-segmentation step. Thanks to the superpixel segmentation, each node in our proposed CRF will be associated with a small number of segments, rather than have the extremely large quantity of individual pixels usually present in high definition images. As provided by the super-segmentation step, the number of neighbours around a segment may vary, depending on the spatial homogeneity of the colour in the image and the expected number of segments. After segmentation, the grid-like structure of the original images is lost. The grid structure of a CRF requires that each node, and therefore each segment, be associated with only four neighbours. Thus, it is necessary to find the four nearest neighbours corresponding to each segment, since the grid-like structure is not well defined.

In order to determine which superpixels are the four closest neighbours, the angles between the centroid of each segment and its neighbours are calculated. We select as neighbours the four segments whose angles have the least difference from 0, 90, 180 and 270 degrees, respectively. For two examples of nodes, depicted as red points, Fig. 3



**Fig. 3** Two examples of superpixels and their neighbours. The red dots in A and B are nodes assigned to two different superpixels, while the green lines represent their corresponding nearest neighbours. Note that the super-segmentation used allows a superpixel node to have more than four neighbours



**Fig. 4** Selected 4 nearest neighbours for superpixel nodes (red dots) in Fig. 3. Dark green lines connect nodes with their selected neighbours

illustrates their closest neighbours, whose centroids are represented by yellow dots.

Figure 4 shows the 4 closest neighbours selected. These neighbours are represented by dark green lines and dark yellow dots. Note that for the superpixels located at the corners, respectively, on the edges, only the 2, respectively, 3, closest neighbours need to be found.

## 4 CRF-based camera-LIDAR fusion for depth estimation

In this paper, depth estimation is formulated as a superpixel-level inference task on a modified conditional random field (CRF). Our proposed model is a multi-sensor extension of the classical pairwise CRF. In this section, we first briefly introduce the CRF model. Then, we show how to fuse the information of an image and a sparse LIDAR point cloud with our novel CRF framework.

### 4.1 Overview

The conditional random field (CRF) is a type of undirected probabilistic graphical model which is widely used for solving labelling problems. Formally, let  $\mathbf{X} = \{X_1, X_2, \dots, X_N\}$  be a set of discrete random variables to be inferred from an observation or input tensor  $\mathbf{Y}$ , which in turn is composed of the observation variables  $c_i$  and  $y_i$ , where  $i$  is an index over superpixels. For each superpixel  $i$ , the variable  $c_i$  corresponds to an observed three-dimensional colour value and

$y_i$  is an observed range measurement. The goal of our framework is to infer the depth of each pixel in a single image depicting general scenes. Following the work of [11, 45], we make the common assumption that an image is composed of small homogeneous regions (superpixels) and consider a graphical model composed of nodes defined on superpixels. Note that our framework is flexible and can estimate depth values on either pixels or superpixels. The remaining question is how to parametrize this undirected graph. Because the interaction between adjacent nodes in the graph is not directed, there is no reason to use a standard conditional probability distribution (CPD), in which one represents the distribution over one node given the others. Rather, we need a more symmetric parametrization. Intuitively, we want our model to capture the affinities between the depth estimates of the superpixels in a given neighbourhood. These affinities can be captured as follows: Let  $\tilde{P}(X, Y)$  be an unnormalized Gibbs joint distribution parametrized as a product of factors  $\Phi$ , where

$$\Phi = \{\phi_1(D_1), \dots, \phi_k(D_k)\},$$

and

$$\tilde{P}(X, Y) = \prod_{i=1}^m \phi_i(D_i).$$

We can then write a conditional probability distribution of the depth estimates  $X$  given the observations  $Y$  using the previously introduced Gibbs distribution, as follows:

$$Pr(X|Y) = \frac{P(X, Y)}{Z(Y)}$$

where

$$Z(Y) = \sum_X \tilde{P}(X, Y).$$

Here,  $Z(Y)$ , also known as ‘the partition function’, works as a normalizing factor which marginalizes  $X$  from  $\tilde{P}(X, Y)$ , allowing the calculation of the probability distribution  $P(X|Y)$ :

$$P(X|Y) = \frac{1}{\sum_X \tilde{P}(X, Y)} \tilde{P}(X, Y).$$

Therefore, similar to conventional CRFs, we model the conditional probability distribution of the data with the following density function:

$$P(\mathbf{X}|\mathbf{y}) = \frac{1}{Z(\mathbf{Y})} \exp(-E(\mathbf{X}, \mathbf{Y}))$$

where  $E$  is the energy function and  $Z$  is the partition function defined by

$$Z(\mathbf{Y}) = \int_{\mathbf{Y}} \exp\{-E(\mathbf{X}, \mathbf{Y})\} d\mathbf{Y}.$$

Since  $Z$  is continuous, this integral equation can be analytically solved. This is different from the discrete case, in which approximation methods need to be applied. To predict the depths of a new image, we solve the following maximum a posteriori (MAP) inference problem:

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x}} P(\mathbf{X}|\mathbf{Y}).$$

To simplify the solution for the energy function, one can take the negative logarithm of the left-hand side and right-hand side of the equation of the probability distribution  $Pr(X|Y)$ : then, the problem of maximizing the conditional probability becomes an energy minimization problem. Therefore, maximizing the probability distribution  $Pr(\mathbf{X}|\mathbf{Y})$  is equivalent to minimizing the corresponding energy function:

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} E(\mathbf{X}, \mathbf{Y}).$$

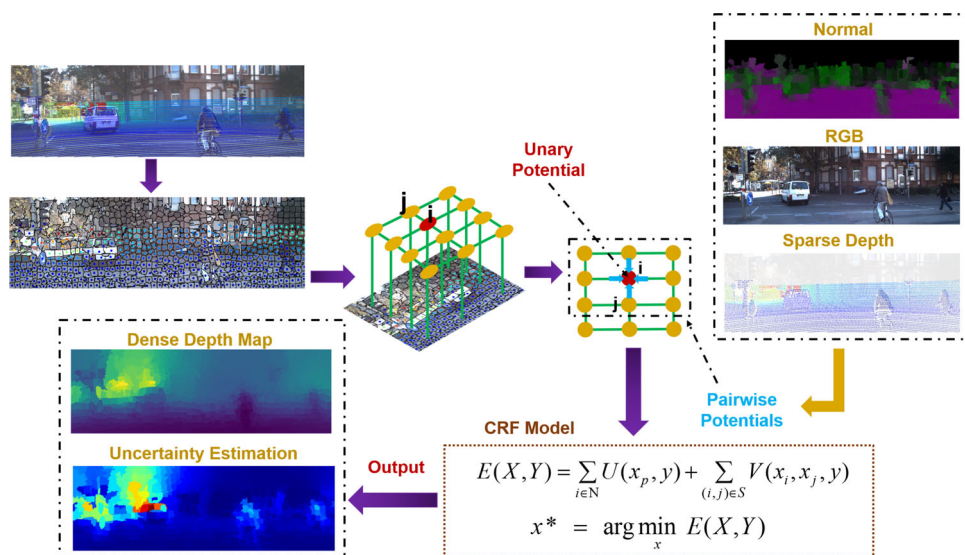
We formulate the energy function as a typical combination of unary potentials  $U$  and pairwise potentials  $V$  over the nodes (superpixels)  $N$  and edges  $S$  of the image  $\mathcal{I}$ :

$$E(\mathbf{X}, \mathbf{Y}) = \sum_{p \in \mathcal{N}} U(x_p, \mathbf{y}) + \sum_{(p,q) \in \mathcal{S}} V(x_p, x_q, \mathbf{y})$$

The unary term  $U$  aims to regress the depth value from a single superpixel. The pairwise term  $V$  encourages neighbouring superpixels with similar appearances to take similar depths [11, 23].

## 4.2 Potential functions

The proposed multi-modal depth estimation model is composed of unary and pairwise potentials. For an input image, which has been over-segmented into  $n$  superpixels, we define a unary potential for each superpixel. The pairwise potentials are defined over the four-neighbour vicinity of each superpixel. The unary potentials are built by aggregating all LiDAR observations inside each superpixel. The pairwise part is composed of similarity vectors, each with  $K$  components, that measure the agreement between different features of neighbouring superpixel pairs. Therefore, we explicitly model the relations between neighbouring superpixels through pairwise potentials. In the following, we describe the details of the potentials involved in our energy function.



**Fig. 5** Illustration of the proposed model. On the top left is a fused view of the image and LiDAR point cloud on superpixels. On the top right are the normal surface map and RGB inputs used in the pairwise potentials. On the top middle is the graph structure of the CRF: The yellow nodes represent the centroids of the image superpixels, and the

green branches the connections between them. The outputs of the unary part and the pairwise part are then fed to the CRF structured loss layer, which minimizes the corresponding energy function. On the bottom left is the probabilistic output, a dense depth map and uncertainty estimation map (see text for details)

### 4.2.1 Unary potential

The unary potential is constructed from the LiDAR sensor measurements by considering the least square loss between the estimated  $x_i$  and observed  $y_i$  depth values:

$$\Phi(\mathbf{x}, \mathbf{y}) = \sum_{i \in \mathcal{L}} \sigma_i (x_i - y_i)^2$$

$$\Phi(\mathbf{x}, \mathbf{y}) = \|\mathbf{W}(\mathbf{x} - \mathbf{y})\|^2$$

where  $\mathcal{L}$  is the set of indices for which a depth measurement is available, and  $\sigma_i$  is a constant weight placed on the depth measurements. This potential measures the quadratic distance between the estimated range  $X$  and the measured range  $Y$ , where available. Finally, in order to write the unary potential in a more efficient matrix form, we define the diagonal matrix  $W$  with entries

$$W_{i,i} = \begin{cases} \sigma_i & \text{if } i \in \mathcal{L} \\ 0 & \text{otherwise} \end{cases}$$

### 4.2.2 Colour pairwise potential

We construct a pairwise potential from  $K$  types of similarity observations, each of which enforces smoothness by exploiting colour consistency features of the neighbouring

superpixels. This pairwise potential can be written as

$$\Psi^c(\mathbf{x}, \mathbf{I}) = \sum_i \sum_{j \in \mathcal{N}(i)} e_{i,j} (x_i - x_j)^2$$

$$\Psi^c(\mathbf{x}, \mathbf{I}) = \|\mathbf{S}\mathbf{x}\|^2$$

where  $I$  is an RGB image,  $\mathcal{N}(i)$  is the set of horizontal and vertical neighbours of  $i$ , and each row of  $S$  represents the weighting factors for pairs of adjacent range nodes. As the edge strength between nodes, we use an exponentiated  $L_2$  norm of the difference in pixel appearance.

$$e_{i,j} = \exp - \frac{\|\mathbf{c}_i - \mathbf{c}_j\|^2}{\sigma_d^2}$$

where  $\mathbf{c}_i$  is the RGB colour vector of pixel  $i$  and  $\sigma_d$  is a tuning parameter. A small value of  $\sigma_d$  increases the sensitivity to changes in the image. Thanks to this potential, the lack of content or features in the RGB image is considered by our model as indicative of a homogeneous depth distribution, in other words, a planar surface.

### 4.2.3 Surface-normal pairwise potential

The mathematical formulation of this potential is similar to the previous colour potential. However, the surface-normal potential considers surface-normal similarities instead of colour. The weighting factors  $nr_{i,j}$  for this case are formulated using the cosine similarity, which is a measure of the

similarity between two nonzero vectors of an inner product space that employs the cosine of the angle between them. The cosine of 0 is 1, and it is less than 1 for any angle in the interval  $[0, \pi]$  radians. It is thus a measurement of orientation instead of magnitude [60]. The cosine of two nonzero vectors can be found by using the Euclidean dot product formula:

$$\mathbf{A} \cdot \mathbf{B} = \|\mathbf{A}\| \|\mathbf{B}\| \cos \theta$$

Therefore, the cosine similarity can be expressed by

$$\cos(\theta) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\| \|\mathbf{B}\|} = \frac{\sum_{i=1}^n A_i B_i}{\sqrt{\sum_{i=1}^n A_i^2} \sqrt{\sum_{i=1}^n B_i^2}}$$

where  $A_i$  and  $B_i$  are the components of vectors  $A$  and  $B$ , respectively. Finally, we define our surface-normal potential by the following equations.

$$\Psi^n(\mathbf{x}, \mathbf{In}) = \sum_i \sum_{j \in \mathcal{N}(i)} nr_{i,j} (x_i - x_j)^2$$

$$\Psi^n(\mathbf{x}, \mathbf{In}) = \|\mathbf{P}\mathbf{x}\|^2$$

$$nr_{i,j} = \frac{\sum_{l=1}^n In_l In_j}{\sqrt{\sum_{l=1}^n In_l^2} \sqrt{\sum_{l=1}^n In_j^2}}$$

### 4.2.4 Depth pairwise potential

This pairwise potential encodes a smoothness prior over depth estimates which encourages neighbouring superpixels in the image to have similar depths. Usually, pairwise potentials are only related to the colour difference between pairs of superpixels. However, depth smoothness is a valid hypothesis which can potentially enhance depth inference. To enforce depth smoothness, a distance-aware Potts model was adopted. Neighbouring points with smaller distances are considered to be more likely to have the same depth. The mathematical formulation of this potential is similar to the colour pairwise potential, as it follows the Potts model:

$$\Psi^d(\mathbf{x}, \mathbf{D}) = \sum_i \sum_{j \in \mathcal{N}(i)} e_{i,j} (x_i - x_j)^2$$

and the weighting factor  $dp_{i,j}$  for this case is formulated as

$$dp_{i,j} = \exp - \frac{\|\mathbf{p}_i - \mathbf{p}_j\|^2}{\sigma_p^2}$$

where  $\mathbf{p}_i$  is the 3D location vector of the LiDAR point  $i$  and  $\sigma_p$  is a parameter controlling the strength of enforcing close points to have similar depth values.

### 4.2.5 Uncertainty potential:

Depth uncertainty estimation is important for refining depth estimation [16, 65], and in safety critical systems [28]. It allows an agent to identify unknowns in an environment in order to reach optimal decisions. Our method provides uncertainties for the estimates of the pixel-wise depths by taking into account the number of LiDAR points present for each superpixel. The uncertainty potential is similar to the unary potential. It is constructed from the number of LiDAR points projected onto a superpixel and employs the following least square loss:

$$U^c(\mathbf{x}, \mathbf{y}) = \sum_{i \in \mathcal{L}} \sigma_i (x_i - unc_i)^2$$

$$U^c(\mathbf{x}, \mathbf{y}) = \|\mathbf{W}(\mathbf{x} - \mathbf{unc})\|^2$$

where  $\mathbf{unc}$  is defined as follows:

$$\mathbf{unc}_{i,i} = \begin{cases} \sigma_i & \text{if P projected on SPx is 0} \\ \psi i & \text{if P projected on SPx is } >0 \text{ and } <2 \\ mean & \text{otherwise} \end{cases}$$

where  $P$  is a 3D point and  $SPx$  is a superpixel. In locations with accurate and sufficiently many LiDAR points, the model will produce depth predictions with a high confidence. This uncertainty estimation provides a measure of how confident the model is about the depth estimation. This results in an overall better performance, since uncertain estimates with high uncertainty can be neglected by higher-level tasks that use the estimated depth maps as an input.

### 4.3 Optimization

With the unary and the pairwise potentials defined, we can now write the energy function as

$$E(\mathbf{X}, \mathbf{Y}) = (\alpha) \Phi(\mathbf{x}, \mathbf{y}) + (\beta) \Psi^c(\mathbf{x}, \mathbf{I}) \dots + \dots (\gamma) \Psi^n(\mathbf{x}, \mathbf{In}) + (\delta) \Psi^d(\mathbf{x}, \mathbf{In}) \tag{1}$$

The scalars  $\alpha, \beta, \gamma, \delta \in [0,1]$  are weightings for the four terms. We may further expand the unary and pairwise potentials to

$$\Phi(\mathbf{x}, \mathbf{y}) = \alpha(\mathbf{x}^T \mathbf{W}^T \mathbf{W} \mathbf{x} - 2\mathbf{z}^T \mathbf{W}^T \mathbf{W} \mathbf{x} + \mathbf{z}^T \mathbf{W}^T \mathbf{W} \mathbf{z}) \tag{2}$$

$$\Psi^c(\mathbf{x}, \mathbf{In}) = \beta(\mathbf{x}^T \mathbf{S}^T \mathbf{S} \mathbf{x}) \tag{3}$$

$$\Psi^n(\mathbf{x}, \mathbf{In}) = \gamma(\mathbf{x}^T \mathbf{P}^T \mathbf{P} \mathbf{x}) \tag{4}$$

$$\Psi^d(\mathbf{x}, \mathbf{In}) = \delta(\mathbf{x}^T \mathbf{D}^T \mathbf{D} \mathbf{x}) \tag{5}$$



We shall pose the problem as one of finding the optimal range vector  $\mathbf{x}^*$  such that:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \{E(\mathbf{X}, \mathbf{Y})\}$$

Substituting equations 2, 3, 4 and 5 into 1 and solving for  $x$  reduces the problem to:  $\mathbf{Ax} = \mathbf{b}$  where

$$\begin{aligned} \mathbf{A} &= \alpha(\mathbf{W}^T\mathbf{W}) + \beta(\mathbf{S}^T\mathbf{S}) + \gamma(\mathbf{P}^T\mathbf{P}) + \delta(\mathbf{D}^T\mathbf{D}) \\ \mathbf{b} &= \alpha(\mathbf{W}^T\mathbf{Wz}) \end{aligned}$$

All we need to do to perform the optimization is to solve a large sparse linear system. The methods for solving sparse systems are divided into two categories: direct and iterative. Direct methods are robust but require large amounts of memory as the size of the problem grows. On the other hand, iterative methods provide better performance but may exhibit numerical problems [10, 20]. In the present paper, the fast algorithm conjugate gradient squared proposed by Hestenes and Stiefel [25, 64] is employed to solve the energy minimization problem.

#### 4.4 Pseudo-code

Algorithm 1 provides the complete pseudo-code for our proposed framework, which is previously illustrated in Fig. 5. In this algorithm, lines 1 to 5 perform the preprocessing, which includes gathering the multi-modal raw data, building a connection graph between pairs of adjacent superpixels and projecting the clustered LiDAR points onto the image space. Lines 6 to 11 constitute the core of the approach. They include constructing the cost function using different potentials (unary and pairwise) to obtain the complete CRF for the depth estimation. The objective of the pairwise potentials is to smooth the depth regressed from the unary part based on the neighbouring superpixels. The pairwise potential functions are based on standard CRF vertex and edge feature functions studied extensively in [52] and other papers. Our model uses both the content information of the superpixels and relation information between them to infer the depth.

### 5 Results and discussion

We evaluate our approach on the raw sequences of the KITTI benchmark, which is a popular dataset for single image depth map prediction. The sequences contain stereo imagery taken from a car driving in an urban scenario. The dataset also provides 3D laser measurements from a Velodyne laser scanner, which we use as ground truth measurements (projected into the stereo images using the given intrinsics and extrinsics in KITTI). This dataset has been used to train and

---

#### Algorithm 1: UAOFusion Network: A multi-modal CRF-based method for Camera–LiDAR depth estimation

---

**Input:** RGB Image and Sparse 3D Point Cloud  
**Output:** Dense Point Cloud

- 1 Compute point cloud normal vectors
- 2 Perform superpixel segmentation
- 3 Build a set of edges and nodes considering the definition of the 4-neighbourhood.
- 4  $z \leftarrow$  Project clustered PCL onto image plane
- 5 Initialize uncertainty depth map for the whole image
- 6 **for** each superpixel (node) **do**
- 7 Calculate unary potential
- 8 Calculate Colour pairwise potential
- 9 Calculate Surface-normal pairwise potential
- 10 Calculate Depth pairwise potential
- 11 Calculate Uncertainty potential
- 12 Infer 2D dense depth map

---

evaluate the state-of-the-art methods and allows quantitative comparisons. First, we evaluate the prediction accuracy of our proposed method with different potentials in Sect. 5.2. Second, in Sect. 5.3 we explore the impact on the depth estimation of the number of sparse depth samples and the number of superpixels. Third, Sect. 5.5 compares our approach to state-of-the-art methods on the KITTI dataset. Lastly, in Sects. 5.6 and 5.7, we demonstrate two use cases of our proposed algorithm, one for creating LiDAR super-resolution from sensor data provided by the KITTI dataset and another one for a dataset collected in the context of this work.

#### 5.1 Evaluation metrics

We evaluate the accuracy of our method in depth prediction using the 3D laser ground truth on the test images. We use the following depth evaluation metrics: root-mean-squared error (RMSE), mean absolute error (MAE) and mean absolute relative error (REL), among which RMSE is the most important indicator and chosen to rank submissions on the leader-board since it measures error directly on depth and penalizes on further distance where depth measurement is more challenging. These metrics were used by [13, 17, 30, 53] to estimate the accuracy of monocular depth prediction.

$$\begin{aligned} RMSE &= \sqrt{\frac{1}{|T|} \sum_{d \in T} \|\hat{d} - d\|^2} \\ MAE &= \frac{1}{T} \sum_{d \in T} \|\hat{d} - d\|^2 \\ REL &= \frac{1}{T} \sum_{d \in T} \left( \frac{|\hat{d} - d|}{\hat{d}} \right) \end{aligned}$$

**Table 1** Depth completion errors [mm] after adding pairwise potentials (lower is better)

Algorithm	Potential functions	RMSE
Ours	I	865.31
Ours	II	854.24
Ours	III	849.39

**Table 2** Depth completion errors [mm] for different number of superpixels (lower is better)

Algorithm	#Superpixels	RMSE
Ours	1200	1370.27
Ours	2400	1050.55
Ours	5500	849.39

Here,  $d$  is the ground truth depth,  $\hat{d}$  is the estimated depth, and  $T$  denotes the set of all points in the test set images. In order to compare our results with those of Eigen et al. [13] and Godard et al. [21], we crop our image to the evaluation crop applied by Eigen et al. We also use the same resolution of the ground truth depth image and cap the predicted depth at 80 m [21].


## 5.2 Architecture evaluation

This section presents an empirical study of the impact on the accuracy of the depth prediction of different choices for the potential functions and hyperparameters. In the first experiment, we compare the impact of sequentially adding our proposed pairwise potentials. We first evaluate a model with only unary and colour pairwise potentials. Then, we added the surface-normal pairwise potential, and finally, the depth pairwise potential is included. As shown in Table 1, the RMSE is improved after adding each pairwise potential.

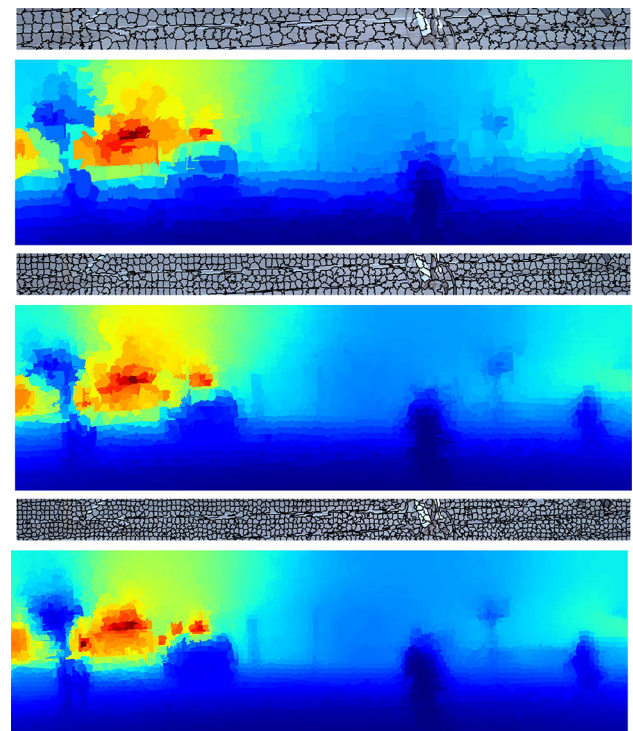
## 5.3 The number of superpixels

In this section, we explore the relation between the prediction accuracy and the number of available depth samples and the number of superpixels.

As displayed in Fig. 5, a greater number of superpixels yields better results in error measurements. Although a larger number of sparse depth observations improves the quality of the depth map, the performance converges when the number of superpixels is more than 5000, which is about 1.5% of the total number of pixels. We ran an exhaustive evaluation of our method for a different amount of superpixels. Figure 8 clearly shows that our method's error decreases with an increased number of superpixels.



**Fig. 6** Qualitative evaluation of the impact of the pairwise potentials defined as CRF terms. In row order: 1st: pairwise potential I, penalizes dissimilar depth estimates of neighbouring pixels which have similar colours in the RGB image, 2nd: pairwise potential II, penalizes the depth differences between neighbouring superpixels whose normal surface vectors have large cosine similarities, and 3rd: pairwise potential III, penalizes neighbouring superpixels with large observed depth differences



**Fig. 7** Visual comparison of dense depth maps produced by the CRF framework when varying the size of the superpixels. From top to bottom, 1200, 2400 and 5500 superpixels




Fig. 8 Convergence for different number of superpixels

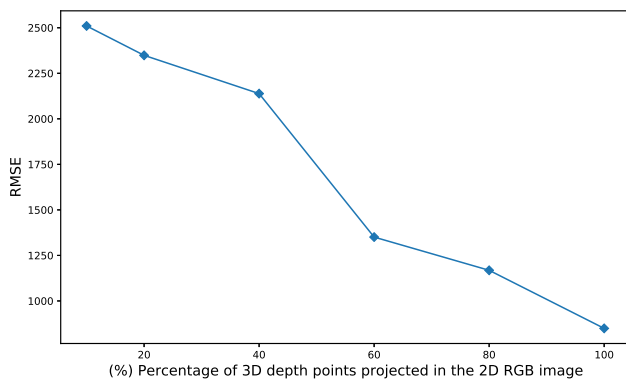


Fig. 9 Convergence for different number of 3D depth points projected into 2D RGB images. The amount of 3D depth points used is represented as a percentage (%). 100% means we are using all the 3D depth points projected into the 2D image

### 5.4 Sub-sampling 3D depth points

We performed a quantitative analysis of the impact of observed 3D point sparsity on the error of our proposed method, by decreasing the amount of 3D points considered during inference. As it is shown in Fig. 9, our method’s error decreases with an increased number of 3D depth points, enabling a better dense depth map estimation. From Fig. 9, we can argue that the amount of 3D depth points is really important for an accurate depth map estimation. Even though our estimation error increases with the sparsity of the depth observations, our method manages to provide state-of-the-art performances even when the depth observations are sampled down to 40%.

### 5.5 Algorithm evaluation for depth completion

This is a more challenging dataset than other datasets for depth estimation: the distances in the KITTI dataset are larger than in other datasets, e.g. NYU-Depth-V2 dataset. Hence, the KITTI odometry dataset is more challenging for the depth estimation task. The performance of our method and those

Table 3 Depth completion errors [mm] by different methods on the test set of KITTI depth completion benchmark (lower is better)


Algorithm	RMSE	MAE
Semantically guided depth upsampling [59]	2312.57	605.47
Convolutional spatial propagation network (CSPN) [9]	1019.64	279.46
Hierarchical multi-scale sparsity-invariant network (HMS-Net) [27]	841.77	253.47
End-to-end sparse-to-dense network (S2DNet) [22]	830.57	247.85
Self-supervised sparse-to-dense network [42]	814.73	249.95
<b>Ours</b>	<b>849.39</b>	<b>263.31</b>

of other existing methods on the KITTI dataset are shown in Table 3. Table 3 shows that the proposed method outperforms other depth map estimation approaches which are well accepted in the robotics community. Our model relies on the number of superpixels and the resolution of input data sources. This means that the model’s performance will increase if we increase the number of superpixels, the image resolution and the density of the LiDAR data.

### 5.6 Algorithm evaluation for LiDAR super-resolution

We present another demonstration of our method in super-resolution of LiDAR measurements. 3D LiDARs have a low vertical angular resolution and thus generate a vertically sparse point cloud. We use all measurements in the sparse depth image and RGB images as input to our framework. An example is shown in Fig. 4. The cars are much more recognizable in the prediction than in the raw scans.

On the other hand, starting from a LiDAR super-resolution map we can generate a 3D reconstruction of the scene. The reconstruction of three-dimensional (3D) scenes has many important applications, such as autonomous navigation [24], environmental monitoring [46] and other computer vision tasks [26]. Therefore, a dense and accurate model of the environment is crucial for autonomous vehicles. In fact, imprecise representations of the vehicle’s surrounding may lead to unexpected situations that could endanger the passengers. In this paper, the 3D modelling is generated using a combination of image and range data as a sensor fusion approach that takes the strengths of each in order to overcome their limitations. Images normally have higher resolution and more visual information than range data, and range data are noisy, sparse, and have less visual information, but already contain 3D information. The qualitative and quantitative results presented here suggest that our system provides 3D recon-




**Fig. 10** Depth completion and uncertainty estimates of our approach on the KITTI raw test set. From top to bottom: RGB and raw depth projected onto the image; high-resolution depth map; raw uncertainty; and estimated uncertainty map

**Table 4** Depth estimation errors [m] by different methods on the test set of KITTI depth estimation benchmark (lower is better)

Algorithm	RMSE[m]	REL	Log10
Multi-modal Auto-Encoders [6]	7.14	0.179	–
Residual of residual network [34]	4.51	0.113	0.049
Residual Up-Projection [15]	3.67	0.072	–
Sparse-to-dense [43]	3.37	0.073	–
CSPN [9]	3.24	0.059	–
S2DNet [22]	3.11	0.069	0.038
<b>Ours</b>	3.59	0.072	0.041

structions of reasonable quality. Following [43], we use a random subset of 2000 images from the test sequences for evaluation. We take the bottom part  $912 \times 228$  due to there being no depth at the top area and only evaluate the pixels with ground truth. The performance of our approach and state-of-the-art depth completion methods is presented in Table 4.

In Table 4, the RMSE value of sparse-to-dense, CSPN and S2DNet methods is slightly better (lower) than our approach. All these three methods solve the depth estimation




**Fig. 11** LiDAR super-resolution. Creating dense point clouds from sparse raw measurements. From top to bottom: RGB image, raw depth map, predicted depth and ground truth depth map. Distant cars are almost invisible in the raw depth map, but are easily recognizable in the predicted depth map


task through deep learning algorithms which require large amounts of data, which can be very costly in practice, and additional techniques like data augmentation to improve their performance. For example, S2DNet uses a dataset composed of 1,070,568 images, while our approach only used 42% of that amount of data plus the LiDAR information. Additionally, these deep learning methods make use of advanced computing resources which limit their use in real-world applications. Our method is very competitive with deep learning models as shown in Table 4, without demanding a large amount of data or additionally strategies. Moreover, we highlight the fact that our method's evaluation does not consider the full image data. As displayed in Fig. 8, our method converges to a better solution when we use 6000 superpixels instead of 5500. We do not evaluate further because the trend is clear, as much superpixels we use, a better dense depth map is predicted.

### 5.7 Depth estimation with very sparse LiDAR data: The UAO LiDAR-RGB Dataset

Thus far, we have sampled the depth from high-quality LiDAR depth maps, but in practice, sparse depth inputs may come from less reliable sources. Therefore, we provide a qualitative evaluation of this model on our own well-calibrated LiDAR and RGB dataset. We use a 16-beam LiDAR along with a Stereo Labs Zed Mini camera with  $1280 \times 720$  resolution. This dataset enables us to demonstrate the stability and robustness of the proposed model in




**Fig. 12** Indoor LiDAR super-resolution. Creating dense point clouds from sparse raw measurements and colour. From top to bottom: RGB image, raw depth map and predicted depth



**Fig. 13** Outdoor LiDAR super-resolution. Creating dense point clouds from sparse raw measurements and colour. From top to bottom: RGB image, raw depth map and predicted depth

particularly challenging scenarios. The scenes were recorded with a low resolution of the camera and the LiDAR sensor in comparison with the KITTI benchmark. Notably, the proposed algorithm is able to estimate a dense depth map of indoor and outdoor environments using colour and sparse depth data. The experimental results are shown in Fig. 12, Fig. 13 and Fig. 14. Dark red indicates farther distances, and dark blue indicates closer distances.

Despite the lower number of LiDAR channels, the proposed method has provided accurate depth information even under challenging outdoor conditions, as shown in Fig. 14. In this scene, there is lot of variability in terms of the light and shadows generated by the environment and the weather itself. After a close look at Fig. 12, Fig. 13 and Fig. 14, it is noticeable that no depth observations from the LiDAR are available at the top and bottom locations of the colour image. After inference, the depth estimates, shown in the bottom images, at the above locations are consistent with the information provided by the image. We can conclude that the framework proposed here works reliably for the depth prediction task. Additionally, it also solves the depth completion problem, as it is able to deal with highly sparse input point clouds projected onto the image space.



**Fig. 14** Outdoors LiDAR super-resolution. Creating dense point clouds from sparse raw measurements and colour. From top to bottom: RGB image, raw depth map and predicted depth

## 6 Conclusions and future work

In this paper, we described an innovative approach to fuse information from different sensor modalities, e.g. cameras and LiDAR, in order to probabilistically estimate a dense point cloud. Our approach achieves good performance in single image depth map prediction on the popular KITTI dataset. It is able to predict detailed depth maps on thin and distant objects. It also reasonably estimates the depth in parts of the image in which there is no ground truth available for supervised learning. The qualitative and quantitative results presented here suggest that our system provides 2D depth maps of reasonable quality, which depends on the density of the laser measurements and the number of superpixels selected. We believe that this method opens up an important avenue for research into multi-sensor fusion and the more general 3D perception problems, which might benefit substantially from sparse depth samples.

There are various variants of CRFs, each with its own pros and cons. A higher-order CRF, for example, considers interactions between more than two variables at a time. This can lead to more accurate predictions. A dynamic CRF is another type of CRF that is designed to work with sequences of data, such as time series or videos. This can be useful for depth estimation using Camera–LiDAR fusion, as it can take into account temporal information in the data. However, both higher-order and dynamic CRFs require more computational resources and are more difficult to train than pairwise, also known as Markov CRFs, which only consider interactions between adjacent variables. Since our work focuses on robotics applications, in which computational performance is a hard constraint, instead of trying to capture long-term dependencies in the data, we leverage the potential of modelling both appearance and geometric constraints from multi-modal sensor observations. A promising future research direction may tackle the challenges of developing efficient and real-time inference solutions for higher-order CRFs in the context of multi-modal depth estimation solutions. Another promising research venue is the parallel implementation of our method, which will bring benefits when deploying it in real robotic platforms. Additionally, the inclusion of other sensor modalities, such as radar [35, 41], could be explored as a way of improving our system's robustness to challenging environmental conditions.

**Acknowledgements** This work was supported by Universidad Autónoma de Occidente (UAO). The authors would like to thank the Research incubator in robotics and autonomous systems (RAS), the Research group on remote and distributed control systems (GITCoD) at UAO, Walter Mayor, Nicolas Llanos Neuta and Juan Carlos Perafán for their feedback and helpful discussions.

**Funding** Open Access funding provided by Colombia Consortium

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Susstrunk, S.: Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Trans. Pattern Anal. Mach. Intell.* **34**, 2274–2282 (2012)
- Agamennoni, G., Furgale, P., Siegwart, R.: Self-tuning m-estimators 4628–4635 (2015)
- Andreasson, H., Triebel, R., Lilienthal, A.: Vision-based interpolation of 3d laser scans (2006)
- Bansal, A., Russell, B., Gupta, A.: Marr revisited: 2d-3d alignment via surface normal prediction (2016). [arXiv:1604.01347](https://arxiv.org/abs/1604.01347)
- Bo Li, Chunhua Shen, Yuchao Dai, van den Hengel, A., Mingyi He: Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs 1119–1127 (2015)
- Cadena, C., Dick, A., Reid, I.: Multi-modal auto-encoders as joint estimators for robotics scene understanding (2016)
- Castanedo, F.: A review of data fusion techniques. *TheScientific-WorldJournal* **2013**, 704504 (2013)
- Ceron, J.S.O., Cano, V.R., Neuta, N.L., Toro, W.M.: Probabilistic perception system for object classification based on camera-lidar sensor fusion. In: *LatinX in AI Research at ICML* (2019)
- Cheng, X., Wang, P., Yang, R.: Depth estimation via affinity learned with convolutional spatial propagation network (2018). [arXiv:1808.00150](https://arxiv.org/abs/1808.00150)
- Davis, T.A., Rajamanickam, S., Sid-Lakhdar, W.M.: A survey of direct methods for sparse linear systems. *Acta Numerica* **25**, 383–566 (2016)
- Diebel, J., Thrun, S.: An application of markov random fields to range sensing 291–298 (2006)
- Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture (2014). [arXiv:1411.4734](https://arxiv.org/abs/1411.4734)
- Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network (2014). [arXiv:1406.2283](https://arxiv.org/abs/1406.2283)
- Eitel, A., Springenberg, J.T., Spinello, L., Riedmiller, M., Burgard, W.: Multimodal deep learning for robust rgb-d object recognition (2015). [arXiv:1507.06821](https://arxiv.org/abs/1507.06821)
- Fu, C., Mertz, C., Dolan, J.: Lidar and monocular camera fusion: On-road depth completion for autonomous driving. In: *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 273–278 (2019)
- Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation (2018). [arXiv:1806.02446](https://arxiv.org/abs/1806.02446)
- Gansbeke, W.V., Neven, D., Brabandere, B.D., Gool, L.V.: Sparse and noisy lidar completion with RGB guidance and uncertainty. *CoRR abs/1902.05356* (2019). [arXiv:1902.05356](https://arxiv.org/abs/1902.05356)

18. Geiger, A.: Are we ready for autonomous driving? the kitti vision benchmark suite 3354–3361 (2012)
19. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* **32**, 1231–1237 (2013)
20. George, A., Gilbert, J., Liu, J.: Graph theory and sparse matrix computation (2012)
21. Godard, C., Aodha, O.M., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency (2016). [arXiv:1609.03677](https://arxiv.org/abs/1609.03677)
22. Hambarde, P., Murala, S.: S2dnet: Depth estimation from single image and sparse samples. *IEEE Transactions on Computational Imaging* **6**, 806–817 (2020)
23. Harrison, A., Newman, P.: Image and sparse laser fusion for dense scene reconstruction 219–228 (2009)
24. Hecht, J.: Cutting-Edge Cinema. *Optics & Photonics News* **29**, 28 (2018)
25. Hestenes, M.R., Stiefel, E.: Methods of conjugate gradients for solving linear systems. *J Res NIST* **49**, 409–436 (1952)
26. Horaud, R., Hansard, M., Evangelidis, G., Clément, M.: An overview of depth cameras and range scanners based on time-of-flight technologies. *Machine Vision and Applications* **27** (2016)
27. Huang, Z., Fan, J., Cheng, S., Yi, S., Wang, X., Li, H.: Hmsnet: Hierarchical multi-scale sparsity-invariant network for sparse depth completion. *IEEE Transactions on Image Processing* **PP**, 1–1 (2019)
28. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? (2017). [arXiv:1703.04977](https://arxiv.org/abs/1703.04977)
29. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks 1097–1105 (2012)
30. Kuznetsov, Y., Stückler, J., Leibe, B.: Semi-supervised deep learning for monocular depth map prediction (2017). [arXiv:1702.02706](https://arxiv.org/abs/1702.02706)
31. Laina, I., Ruppelrecht, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks (2016a). [arXiv:1606.00373](https://arxiv.org/abs/1606.00373)
32. Laina, I., Ruppelrecht, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks (2016b). [arXiv:1606.00373](https://arxiv.org/abs/1606.00373)
33. Lee, B., Jeon, H., Im, S., Kweon, I.S.: Depth completion with deep geometry and context guidance 3281–3287 (2019)
34. Liao, Y., Huang, L., Wang, Y., Kodagoda, S., Yu, Y., Liu, Y.: Parse geometry from a line: Monocular depth estimation with partial laser observation (2016). [arXiv:1611.02174](https://arxiv.org/abs/1611.02174)
35. Lin, J.T., Dai, D., Gool, L.V.: Depth estimation from monocular images and sparse radar data. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 10233–10240 (2020). <https://doi.org/10.1109/IROS45743.2020.9340998>
36. Liu, F., Shen, C., Lin, G.: Deep convolutional neural fields for depth estimation from a single image (2014a). [arXiv:1411.6387](https://arxiv.org/abs/1411.6387)
37. Liu, F., Shen, C., Lin, G.: Deep convolutional neural fields for depth estimation from a single image (2014b). [arXiv:1411.6387](https://arxiv.org/abs/1411.6387)
38. Liu, F., Shen, C., Lin, G., Reid, I.: Learning depth from single monocular images using deep convolutional neural fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**, 2024–2039 (2016)
39. Liu, M., Salzmann, M., He, X.: Discrete-continuous depth estimation from a single image 716–723 (2014)
40. Liu, N., Li, S., Du, Y., Torralba, A., Tenenbaum, J.B.: Compositional visual generation with composable diffusion models (2022). [arXiv:2206.01714](https://arxiv.org/abs/2206.01714)
41. Long, Y., Morris, D., Liu, X., Castro, M., Chakravarty, P., Narayanan, P.: Radar-camera pixel depth association for depth completion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12507–12516 (2021)
42. Ma, F., Cavalheiro, G.V., Karaman, S.: Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera (2018). [arXiv:1807.00275](https://arxiv.org/abs/1807.00275)
43. Ma, F., Karaman, S.: Sparse-to-dense: Depth prediction from sparse depth samples and a single image (2017). [arXiv:1709.07492](https://arxiv.org/abs/1709.07492)
44. Maddern, W., Newman, P.: Real-time probabilistic fusion of sparse 3d lidar and dense stereo 2181–2188 (2016)
45. Mahmoudi, M., Sapiro, G.: Sparse representations for range data restoration. *Trans. Img. Proc.* **21**, 2909–2915 (2012)
46. Mallet, C., Bretar, F.: Full-waveform topographic lidar: State-of-the-art. *ISPRS Journal of Photogrammetry and Remote Sensing* **64**, 1–16 (2009)
47. Mastin, A., Kepner, J., Fisher, J.: Automatic registration of lidar and optical images of urban scenes 2639–2646 (2009)
48. Nickels, K., Castano, A., Cianci, C.M.: Fusion of lidar and stereo range for mobile robots (2003)
49. Park, K., Kim, S., Sohn, K.: High-precision depth estimation with the 3D LiDAR and stereo fusion. In: IEEE International Conference on Robotics and Automation, Brisbane. pp. 2156–2163 (2018)
50. Peng Wang, Xiaohui Shen, Zhe Lin, Cohen, S., Price, B., Yuille, A.: Towards unified depth and semantic prediction from a single image 2800–2809 (2015)
51. Piniés, P., Paz, L.M., Newman, P.: Too much tv is bad: Dense reconstruction from sparse laser with non-convex regularisation 135–142 (2015)
52. Qin, T., Liu, T.Y., Zhang, X.D., Wang, D.S., Li, H.: Global ranking using continuous conditional random fields 1281–1288 (2008)
53. Qiu, J., Cui, Z., Zhang, Y., Zhang, X., Liu, S., Zeng, B., Pollefeys, M.: Deeplidar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image (2018). [arXiv:1812.00488](https://arxiv.org/abs/1812.00488)
54. Romero-Cano, V., Nieto, J.I.: Stereo-based motion detection and tracking from a moving platform 499–504 (2013)
55. Romero-Cano, V., Vignard, N., Laugier, C.: Xdvision: Dense outdoor perception for autonomous vehicles 752–757 (2017)
56. Roy, A., Todorovic, S.: Monocular depth estimation using neural regression forest 5506–5514 (2016)
57. Saxena, A., Chung, S.H., Ng, A.Y.: Learning depth from single monocular images 1161–1168 (2005)
58. Saxena, A., Sun, M., Ng, A.Y.: Make3d: Learning 3d scene structure from a single still image. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**, 824–840 (2009)
59. Schneider, N., Schneider, L., Pinggera, P., Franke, U., Pollefeys, M., Stiller, C.: Semantically guided depth upsampling (2016). [arXiv:1608.00753](https://arxiv.org/abs/1608.00753)
60. Sidorov, G., Gelbukh, A., Gómez-Adorno, H., Pinto, D.: Soft similarity and soft cosine measure: Similarity of features in vector space model. *Computación y Sistemas* **18** (2014)
61. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from rgbd images 746–760 (2012)
62. Suwajanakorn, S., Hernández, C., Seitz, S.M.: Depth from focus with your mobile phone. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3497–3506 (2015)
63. Tanner, M., Piniés, P., Paz, L.M., Newman, P.: What lies behind: Recovering hidden shape in dense mapping 979–986 (2016)
64. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Jarrod Millman, K., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C., Polat, İ., Feng, Y., Moore, E.W., Vand erPlas, J., Laxalde, D., Perktold, J., Cimman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., Contributors, S...: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272. 2020

65. Walz, S., Gruber, T., Ritter, W., Dietmayer, K.: Uncertainty depth estimation with gated images for 3d reconstruction (2020)
66. Wang, B., Feng, Y., Liu, H.: Multi-scale features fusion from sparse lidar data and single image for depth completion. *Electronics Letters* **54**, 1375–1377 (2018)
67. Xiang, R., Zheng, F., Su, H., Zhang, Z.: 3ddepthnet: Point cloud guided depth completion network for sparse depth and single color image (2020). [arXiv:2003.09175](https://arxiv.org/abs/2003.09175)
68. Xie, J., Girshick, R., Farhadi, A.: Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks (2016). [arXiv:1604.03650](https://arxiv.org/abs/1604.03650)
69. Xu, D., Wang, W., Tang, H., Liu, H., Sebe, N., Ricci, E.: Structured attention guided convolutional neural fields for monocular depth estimation (2018). [arXiv:1803.11029](https://arxiv.org/abs/1803.11029)
70. Xu, Y., Zhu, X., Shi, J., Zhang, G., Bao, H., Li, H.: Depth completion from sparse LiDAR data with depth-normal constraints, in: *IEEE International Conference on Computer Vision*, pp. 2811–2820 (2019). <https://doi.org/10.1109/ICCV.2019.00290>, [arXiv:1910.06727](https://arxiv.org/abs/1910.06727)
71. Yang, Q., Yang, R., Davis, J., Nister, D.: Spatial-depth super resolution for range images 1–8 (2007)
72. Zhang, R., Tsai, P.S., Cryer, J.E., Shah, M.: Shape from shading: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **21**, 690–706 (1999)
73. Zhang, Y., Funkhouser, T.: Deep depth completion of a single rgb-d image (2018a). [arXiv:1803.09326](https://arxiv.org/abs/1803.09326)
74. Zhang, Y., Funkhouser, T.: Deep depth completion of a single rgb-d image (2018b). [arXiv:1803.09326](https://arxiv.org/abs/1803.09326)
75. Zhao, C., Sun, Q., Zhang, C., Tang, Y., Qian, F.: Monocular depth estimation based on deep learning: An overview. *Science China Technological Sciences* **63**, 1612–1627 (2020)
76. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video (2017). [arXiv:1704.07813](https://arxiv.org/abs/1704.07813)



**Victor Romero-Cano** received a BS degree in mechatronic engineering from Universidad Autónoma de Occidente, Colombia, in 2007, an MSc degree in electrical engineering from the Universidade de Sao Paulo, Brazil, in 2010 and a PhD from the Australian Centre for Field Robotics, University of Sydney, Australia, in 2015. He joined INRIA as a postdoctoral research associate within the CHROMA team in 2016. Since 2017, he has been a lecturer and researcher at the Universidad



**Sildomar Monteiro** received an MSc degree from the Instituto Tecnológico de Aeronáutica, Brazil, and a PhD from the Tokyo Institute of Technology, Japan, in 2007. He is a research scientist working on AI, machine learning, computer vision and robotics for autonomous systems. He enjoys leading and mentoring teams to develop rigorous approaches for challenging real-world problems.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Johan S. Obando-Ceron** received the BS degree in mechatronics engineering from the Universidad Autónoma de Occidente (Autonomous University of the West), Colombia, in 2017, where he worked for three years as a research assistant. His current research interests include reinforcement learning, as well as statistics and optimization theory and its applications.