

Dense Disparity Estimation in Ego-motion Reduced Search Space[★]

Luka Fućek, Ivan Marković, Igor Cvišić, Ivan Petrović^{*}

** University of Zagreb, Faculty of Electrical Engineering and Computing, Croatia (e-mail: luka.fucek@fer.hr, ivan.markovic@fer.hr, igor.cvisic@fer.hr, ivan.petrovic@fer.hr).*

Abstract: Depth estimation from stereo images remains a challenge even though studied for decades. The KITTI benchmark shows that the state-of-the-art solutions offer accurate depth estimation, but are still computationally complex and often require a GPU or FPGA implementation. In this paper we aim at increasing the accuracy of depth map estimation and reducing the computational complexity by using information from previous frames. We propose to transform the disparity map of the previous frame into the current frame, relying on the estimated ego-motion, and use this map as the prediction for the Kalman filter in the disparity space. Then, we update the predicted disparity map using the newly matched one. This way we reduce disparity search space and flickering between consecutive frames, thus increasing the computational efficiency of the algorithm. In the end, we validate the proposed approach on real-world data from the KITTI benchmark suite and show that the proposed algorithm yields more accurate results, while at the same time reducing the disparity search space.

Keywords: stereo vision, visual odometry, disparity estimation.

1. INTRODUCTION

Depth estimation plays an important role in many autonomous systems in automotive industry, augmented reality, and robotics in general. Besides time-of-flight cameras, which can be limited by external lighting and range, stereo cameras are often used as a primary vision sensor for depth estimation. Even though using cameras requires significant computing efforts to obtain disparity maps (DM), stereo cameras are a very popular choice due to their commercial availability and efficiency.

In order to obtain valuable depth information, images from calibrated and synchronized stereo cameras are used. Early research was focused on local methods that considered only neighboring pixels to find a stereo match. Lack of rich texture in the stereo images resulted with semi-dense DM with poor accuracy. To tackle these problems, pixel-wise global methods penalizing discontinuities were introduced, where a depth map is sought minimizing a global energy function, e.g., works of Kolmogorov and Zabih (2001); Klaus et al. (2006); Yang et al. (2009). However, global methods incur high computational and memory costs and in Hirschmüller (2008) semi-global matching (SGM) algorithm was proposed. Therein the computational complexity of global optimization methods was reduced by reducing the domain of considered pixels to several linear paths in the images. Even though the

DM obtained by SGM is dense and more accurate than the ones obtained with local matching techniques, they still lack temporal information and can introduce depth flickering among neighboring frames.

As opposed to standard approaches that extract depth information from a single stereo image, Dobias and Sara (2011) predict the DM using ego-motion and DM of the previous frame. They take predicted disparities as they are and fill the rest of the DM using a traditional stereo algorithm, while checking the validity of predicted disparities without any method to refine them. Additionally, Jiang et al. (2014) detect moving objects in the scene and avoid predictions based on ego-motion in these regions of the image. However, their algorithm accumulates propagated disparity error with time, since they also take predicted disparities without any refinement. To address this problem, the authors discard the predicted DM every 100 frames to start from a fresh SGM-based DM. From the viewpoint of accuracy, theoretically this result does not outperform SGM, since unreasonable disparities are propagated into future frames until the reset is performed. From estimation perspective, Vaudrey et al. (2008) integrate previous frames with the current one using an iconic (pixel-wise) Kalman filter as introduced by Matthies et al. (1989). Their work extends the idea of integrating stereo iconically to provide more information with a higher certainty. They extend the Kalman filter model by introducing *disparity rate* in the depth direction. However, they constrain their model to motion in longitudinal direction and, thereby, neglect all movements in lateral and vertical directions. Their algorithm performs well only in scenes where most of the movement is in the longitudinal direction, such as highway traffic scenes. Morales and Klette (2013) aim

[★] This work has been supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 688117 and by the Croatian Science Foundation under contract No. I-2406-2016. This research has also been carried out within the activities of the Centre of Research Excellence for Data Science and Cooperative Systems supported by the Ministry of Science, Education and Sports of the Republic of Croatia.

to improve the disparity estimation of objects that are static with respect to the ground or moving longitudinally away from the ego-vehicle. However, they also constrain their motion model (cf. Franke et al. (2005)). Jakob Engel, Jörg Stücker (2015) consider ego-motion and form keyframes that contain depth information for simultaneous localization and mapping (SLAM). They integrate the disparity information from current frame into the assigned keyframe. Since the focus of their work was SLAM, DM generated by their approach are semidense. Žbontar and LeCun (2015) use convolutional neural networks (CNN) to learn a similarity measure on small image patches and compute matching costs. SGM is then used to optimize the results. Mayer et al. (2016) formulated the problem as a supervised learning task that can be solved with CNN. They proposed three synthetic stereo video datasets for training of large networks and presented a CNN for real-time disparity estimation that provides state-of-the-art results using high power graphics processing units.

In this paper we propose to generalize the disparity prediction model and increase the accuracy of depth estimation in comparison to standard methods like the SGM. We focus our work on stable, precise and fast spatio-temporal reconstruction, thus constraining the usecase of the proposed method to static scenes. Although this can be seen as a limitation, in fact, this approach will form the base for dense stereo detection of dynamic objects and can also be used in applications where static scenes are predominant, e.g., space exploration robotics. We track pixels from the previous stereo frame to the current one directly in the disparity domain using ego-motion estimation. In order to avoid the need to introduce any other sensors, we obtain ego-motion using the visual odometry algorithm proposed by Cvišić and Petrović (2015). For each pixel we deterministically compute the displacement based on ego-motion and stochastically track the value of its disparity while updating its uncertainty through time with Kalman filtering. Disparity of each pixel is estimated by combining the newly matched (measured) DM and predicted DM. We perform stereo matching using a custom SGM on a reduced disparity search space, based on the predicted DM and its uncertainty to reduce the computing complexity and number of matched outliers, while producing denser and more accurate DM. In the end, we validate the proposed approach on real-world data from the KITTI dataset (Menze and Geiger (2015))

2. ALGORITHM DESCRIPTION

As an input for stereo disparity estimation, we use a sequence of stereo image pairs, captured using a pair of calibrated, rectified, and synchronized cameras. The proposed algorithm relies on using the previous and current stereo image, I^{k-1} and I^k , with the accompanying DM of the previous frame, D^{k-1} , in order to estimate the DM of the current frame D^k . Each stereo frame at time instant k consists of a left and right image, I_L^k and I_R^k , and the accompanying left and right DM, D_L^k and D_R^k . We use two consecutive stereo images, I^{k-1} and I^k , for *ego-motion estimation* of the stereo rig.

Given the estimated displacement calculated from I^{k-1} and I^k , and the DM D^{k-1} from the previous frame, we

can predict the DM of current frame $D^{k|k-1}$. This *disparity prediction* serves as the base for a *stereo matching* technique, e.g., in the present paper we use SGM, to reduce the disparity search space, thus consequently reducing the computing efforts and the number of outliers in the measured (matched) DM D_m^k . To increase the accuracy and produce a denser DM, the predicted DM, $D^{k|k-1}$, is updated with the measured DM, D_m^k , within the framework of the Kalman filter. This step produces the updated disparity $D^{k|k}$ and we denote it as the *disparity update*. The algorithm can be summed up as follows:

- (1) Ego-motion estimation
- (2) Disparity prediction
- (3) Stereo matching
- (4) Disparity update.

2.1 Ego-motion estimation

In order to predict current DM $D^{k|k-1}$ based on previous frame DM D^{k-1} , an ego-motion estimation is needed. Transformation $T_{k-1,L}^k$ represents the homogeneous transformation from the current coordinate frame of the left camera, \mathcal{F}_L^k , to the previous frame \mathcal{F}_L^{k-1} . This transformation can be obtained using several different methods. In the present paper, we choose the visual odometry approach.

Disparity space. We perform prediction directly in disparity space and, therefore, we transform $T_{k-1,L}^k$ from Euclidean space to disparity space. Let w , M and Γ be defined as

$$\omega = \begin{bmatrix} x \\ y \\ d \\ 1 \end{bmatrix}, M = \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}, \Gamma = \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 0 & fb \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad (1)$$

where x, y and d are coordinates in the disparity space, X, Y and Z are coordinates in the Euclidean space camera frame \mathcal{F}_L^k , b is the baseline of stereo rig and f is the focal length of the cameras. Γ represents a projective transformation between homogeneous coordinates M in Euclidean space and homogeneous coordinates ω in disparity space (Demirdjian and Darrell (2001)). Analogous to the standard homogeneous transformation in Euclidean space

$$M^k = T_{k-1}^k M^{k-1} \quad (2)$$

the transformation of coordinates in disparity space is defined as

$$\omega^{k*} \simeq H_{k-1}^k \omega^{k-1} \quad (3)$$

where \simeq denotes equality up to a scale factor and H is defined as

$$H_{k-1}^k = \Gamma T_{k-1}^k \Gamma^{-1}. \quad (4)$$

2.2 Disparity prediction

Using the ego-motion H_{k-1}^k we transform each pixel from D^{k-1} to estimate $D^{k|k-1}$. The transformation is performed directly in the disparity space using (3). Note that in (3) the result of each pixel transformation is only equal up to a scale factor. Since ω is expected to be homogeneous as in (1), we scale ω^{k*} by dividing it with the value of its fourth member to get a homogeneous coordinate form ω^k .

Uncertainty prediction. While new x , y and d coordinates of each pixel are deterministically calculated by (3), each disparity d^{k-1} in D^{k-1} is associated with an accompanying variance p_d^{k-1} , forming a variance map P^{k-1} that includes left and right DM variances P_L^{k-1} and P_R^{k-1} . Variance of each disparity is predicted using the motion model applying the same displacement

$$p_d^{k|k-1} = (\Phi^{k-1})^2 p_d^{k-1} + q^{k-1}, \quad (5)$$

where q^{k-1} denotes the variance of estimated ego-motion in the disparity space, and Φ^{k-1} denotes the system model of current transformation for a pertaining pixel. The variance of estimated ego-motion depends on the precision of the used odometry algorithm. How to compute odometry variance and transform it to disparity space is out of the scope of the present paper, and in lieu of a time varying q^{k-1} , we use an empirically determined constant q . Since we already know $d^{k|k-1}$, for variance estimation we compute the system model as the following ratio

$$\Phi^{k-1} = \frac{d^{k|k-1}}{d^{k-1}} \quad (6)$$

thus avoid the need for computing Φ^{k-1} analytically.

Disparity refinement. When transforming each pixel from D^{k-1} , it is likely that multiple pixels from frame $k-1$ will result in the same x and y coordinates in $D^{k|k-1}$. If there are no outliers in the prediction process, pixels with the highest disparity value $d^{k|k-1}$ are the closest ones to the camera and are most likely the pixels not being occluded by other pixels. Therefore, we select these pixels and discard the others.

Since we take a deterministic approach to x and y coordinates propagation described with (3), incorrect disparity predictions are expected near object edges. Using ego-motion information T_{k-1}^k and previous disparity d^{k-1} we predict the observation of $d^{k|k-1}$. In case of bad ego-motion estimation, prediction $d^{k|k-1}$ could result with wrong disparity value. If the disparity is placed away from depth discontinuities, disparity prediction error will be small enough and *stereo matching* in a close interval around this disparity will easily correct the wrong prediction (see Section 2.3). On the other hand, if we consider disparity $d^{k|k-1}$ that is placed on the depth discontinuity (edge), prediction based on bad ego-motion could result with disparity error that can not be corrected in *stereo matching* phase. To address this problem, we reject all the predicted pixels near disparity discontinuities.

Additionally, if the stereo rig is moving forward, gaps in the predicted DM appear. This phenomenon can be compared to the *zooming* effect when lack of information results in holes in the predicted disparity (confer Fig. 3b for an illustration of this effect). We address this problem with task-specific interpolation that fills the *invalid* disparities caused by this phenomenon. We use horizontally neighboring pixels and use their values to determine the value of an *invalid* disparity. If d_i is an *invalid* predicted disparity in an image row, its value is determined as

$$\begin{cases} d_i = \frac{d_{i-1} + d_{i+1}}{2}, & \text{if } f_s(d_{i-1}, d_{i+1}) = 1 \\ \text{invalid,} & \text{else} \end{cases} \quad (7)$$

where f_s represents a similarity function described as

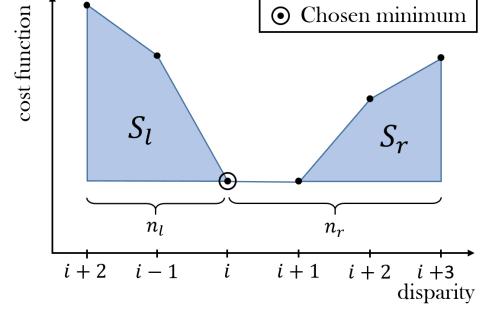


Fig. 1: Variance estimation by counting neighboring disparities from left and right side of the cost function around its minimum.

$$f_s(a, b) = \begin{cases} 1, & \text{if } |a - b| < \gamma_f \\ 0, & \text{else} \end{cases} \quad (8)$$

where γ_f denotes a threshold value defining if a and b (in our case, d_{i-1} and d_{i+1}) are similar enough. In other words, we compare neighboring disparities and, if they are similar enough, we use their values to fill the *invalid* pixels. We take the same approach for filling invalid values based on vertical neighboring pixels.

2.3 Stereo matching

Once the predicted DM, $D^{k|k-1}$, and its variance map, $P^{k|k-1}$, is available, we use this information to additionally reduce the computational complexity of the process of stereo matching. For the stereo matching we use SGM with eight-path configuration. For each path an energy function that penalizes disparity changes among neighboring pixels is minimized to find the optimal disparity value for each pixel. Unlike its usual implementation, we reduce the disparity search space based on the predicted DM $D^{k|k-1}$ and $P^{k|k-1}$. Instead of searching the whole disparity space for every pixel at every iteration of stereo matching, we only consider small environment centered around the predicted disparity $d^{k|k-1}$. The reduced search interval is defined as $[d^{k|k-1} \pm 3\sqrt{p^{k|k-1}}]$.

Matching uncertainty. To estimate the variance of each matched pixel, we refer to the approach described in Wedel et al. (2011) where authors observe the slope from left and right side of the cost function around its minimum. The slope serves as a quality measure of the estimated disparity. If the slope is low, the disparity is not estimated precisely and the variance is higher. If the slope is high, the disparity is estimated precisely and the variance is lower. Wedel et al. (2011) have shown that uncertainty determined using this technique is correlated to the true variance of the errors based on comparison with ground truth. The disadvantage of this approach lies in the case when there are two or more neighboring pixels with minimal cost. This way the slope is horizontal and the uncertainty is infinite which often is not the case.

We propose to expand the variance estimation approach by counting neighboring disparities n_l and n_r while sums S_l and S_r of their costs are lower than S_{max} (Fig. 1). In our experiments, we determined that for matching window [3x3], $S_{max} = 10$ performs well over a wide range of scenes. The matching variance r^k is then defined as

$$r^k = (n_l + n_r). \quad (9)$$

By considering multiple nearby matching cost, we avoid infinite uncertainties. Additionally, the variance value does not depend on the immediately neighboring costs alone, but also on the costs that fit into sum S_{max} .

Alternatively, innovation variance $s^k = p^{k|k-1} + r^k$ could be used for disparity search space reduction. However, in order to determine s^k , we would need to first calculate the matching variance r^k . Since r^k depends on the disparity value and nearby optimization function costs, which are available only after the matching process, we would need to rerun the matching process (on a wider interval), which would induce additional computational load.

Temporal stability. In Jiang et al. (2014), disparities are propagated from previous moment $k-1$ to current moment k as they are. If an unreasonable value (outlier) is present in the previous frame, it will also be propagated to the current frame. Since there is no mechanism to forget these unreasonable values, they will be propagated to every future frame and the number of outliers in the final DM will increase with time. Jiang et al. (2014) address this by forgetting the predicted disparity every 100 frames. In this work, we aim at continuous and constant improvement in accuracy and computing time, hence no predicted pixel is taken as it is. Predicted value is only used to reduce the disparity search space and its variance is tracked for further filtering as described in Section 2.4.

2.4 Disparity update

Combining the predicted DM $D^{k|k-1}$ and the matched DM D_z^k , i.e. the measured disparity, both with their respective variances, we update the predicted disparity using Kalman filter on the pixel level. As described in Section 2.2 we predict the new x and y coordinates of each disparity deterministically, while disparity d is estimated by Kalman filtering. First we determine the Kalman gain:

$$K_k = \frac{p^{k|k-1}}{p^{k|k-1} + r^k} \quad (10)$$

where r_k denotes the variance of the matched disparity (measurement) as described in Subsection 2.3.1. The final disparity is estimated using the standard Kalman filter update equation:

$$d^k = d^{k|k-1} + K_k(d_z^k - d^{k|k-1}) \quad (11)$$

and the variance is updated as follows

$$p^k = (1 - K_k)p^{k|k-1}. \quad (12)$$

Disparity consistency check. Disparity consistency check is performed for two reasons. First, outliers present in the current DM can be propagated to the next frame. Second, even though propagated disparity is only used to reduce search space, the real minimum of the cost function could be outside of the reduced search space. We address this problem by using several filtering techniques for rejecting outliers like *left-right consistency check* and *sum of absolute difference (SAD) check* on left and right images. Using this method we reject occluded pixels that are often poorly matched and which introduce outliers being propagated to future frames. In the next frame, for the rejected pixels, the search is then performed in the

whole search space. Please confer Hu and Mordohai (2010) for an overview of alternative methods.

3. EXPERIMENTAL VALIDATION

To test and validate the proposed algorithm we used the KITTI dataset (Menze and Geiger (2015)), which includes 200 training and testing scenes. Every scene consist of 20 stereo sequences of real world images, thus making them appropriate for validation of algorithms that rely on image sequences rather than a single image frame. Since KITTI dataset is captured on the streets of Karlsruhe, moving objects are present in the majority scenes. In order to evaluate the proposed algorithm, we selected scenes where no moving objects were present. Current implementation treats the scene as static and any moving objects will introduce prediction errors since their movement is neglected, thereby resulting with bad final DM estimation. As discussed in Section 1, this limitation is in fact a base for future dense scene flow estimation and moving object detection. Furthermore, it can also be used in applications where moving objects are not dominant or even non-existent, like space exploration robotics.

As described in Section 2.3 we use SGM with disparity estimation in the reduced search space based on ego-motion. Demonstration of an arbitrary optimized path with matching costs is shown in Fig. 2. Figure 2b shows optimized matching costs of the path shown in Fig. 2a as a product of our basic SGM implementation. Next, Fig. 2c highlights the chosen minima, i.e., chosen disparity values on the whole path. The idea of the proposed algorithm is to reduce the disparity search space based on valid disparity predictions as shown in Fig. 2d. Solid gray color represents the ignored disparity search space. After the search space reduction, minima can be located faster while reducing the chance of false disparity detection. Figure 2e highlights the reduced search space chosen minima, from which reduction of required computation effort can be seen.

In the first iteration, whole disparity search space is considered since there is no disparity prediction available. In the following frames the disparity prediction is available and SGM can be performed on the reduced disparity search space. Our experiments show that less than 50% of whole disparity search space is considered in most scenes. For the scene shown in Fig. 2, the considered disparity search space is 100%, 49.05%, 47.24% and 46.05% for frames 0, 1, 2 and 3, respectively. As the prediction becomes denser the whole computing process becomes faster due to the reduced search space.

Figure 3 demonstrates the whole disparity estimation process. First we take the DM of the previous frame (Fig. 3a) and form a prediction based on ego-motion (Fig. 3b). Here, the “zooming” effect is easily noticeable. We successfully negate that effect by interpolation method described in Section 2.2 as shown in Fig. 3c. Next, left-right consistency check is performed on left and right DM to compare their results and discard all disparity predictions that are not below the similarity threshold. This method is used to reduce a number of outliers caused by incorrect ego-motion propagation or poorly matched disparities from previous frames. The most noticeable rejected area is visible on the left side of Fig. 3d where all



Fig. 2: Disparity search space reduction on an optimized matching cost function. (a) shows the left image of stereo frame with highlighted path example, (b) shows optimized matching costs on the path, (c) highlights the chosen minima (disparities) (d) shows optimized matching costs with the reduced disparity search space (e) highlights the chosen minima (disparities) after search space reduction.

the disparities that cannot be observed with both left and right cameras are marked as invalid. Fig. 3d also represents the final DM prediction used for further processing. Fig. 3e shows the results of our implementation of an eight-path SGM algorithm with reduced disparity search space. It is noticeable that predicted and matched DM (Figs. 3d and 3e) are very similar. The main differences are in the areas of occluded pixels due to different perspectives of previous and current stereo frames. Other noticeable difference is near object edges where edge rejection filtering described in Section 2.2 is applied after the prediction step. These rejections are necessary to reduce the number of outliers that could propagate to future frames and result with wrong disparity search interval in the matching process. Finally, we update the DM as described in Section 2.4. The results are shown in Fig. 3f.

We additionally demonstrate the benefits of this approach by referring to Fig. 4, where it is visible that the updated DM combines best disparities from both the matched and predicted DM. As elaborated in Section 2.2, rejected disparities near disparity discontinuities can be seen in the predicted DM in Fig. 4. While useful for preserving the edges and shape of foreground objects, this rejection can result with holes in the predicted background disparities. On the other hand, newly matched DM in Fig. 4 shows denser background reconstruction, but is unable to reconstruct foreground objects properly (the traffic light pole). By using the predicted and matched DM, we can combine best attributes of both reconstructions. This way, a denser map with less outliers is generated.

In order to evaluate the proposed approach, we used KITTI benchmarking scripts, where a bad disparity refers to disparities with error greater than 3 pixels or 5% relative to the ground truth disparities acquired by the 3D laser range sensor. Direct output of the evaluation for four different scenes is shown in Fig. 5. As Table 1 shows, our approach results with less bad disparities than the classical eight-path SGM, while considering less than 50% of the whole disparity search space. This way we increased both computation speed and accuracy of DM estimation.

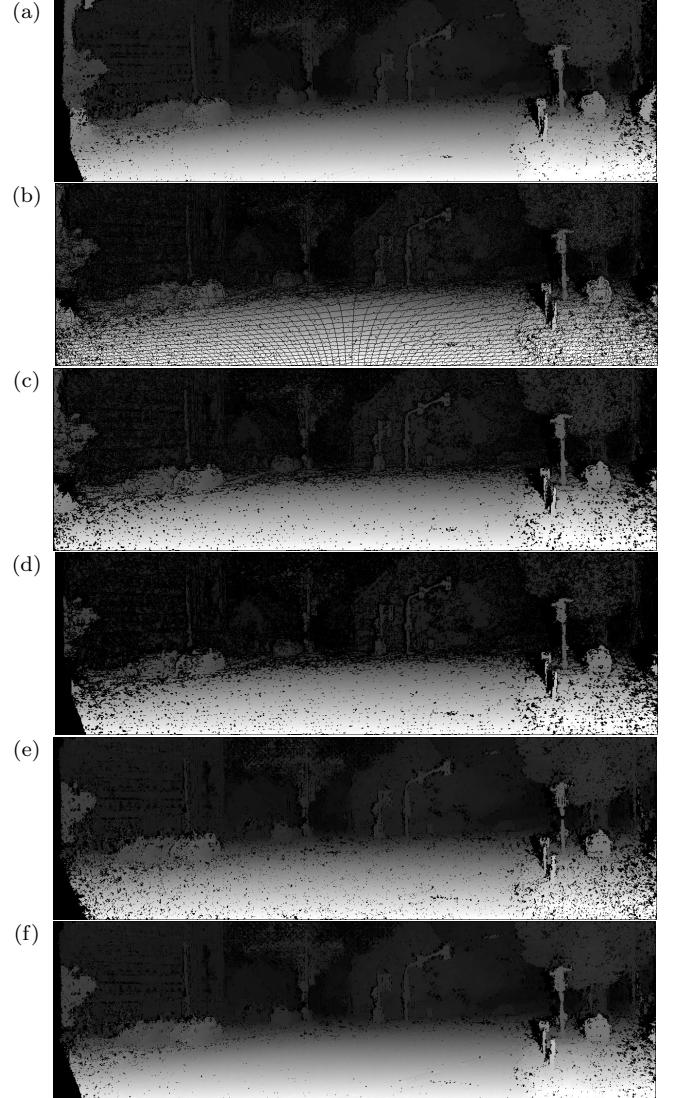


Fig. 3: Demonstration of disparity processing. Images show left DM of (a) previous frame, (b) prediction of current frame based on ego-motion and previous DM, (c) prediction of current frame interpolated due to “zooming” effect, (d) prediction of current frame filtered with left-right consistency check, (e) SGM matching result on reduced search disparity space, (f) updated (d) with (e).

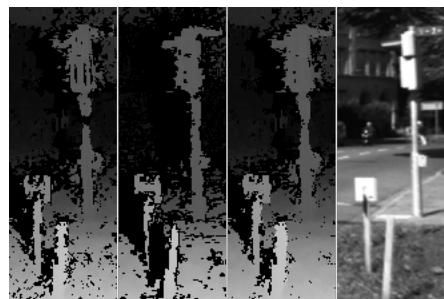


Fig. 4: The comparison of (1) matched DM, (2) predicted DM, (3) updated DM and (4) intensity image.

Our experiments indicate that this also implies about 50% reduction of execution time when compared to base SGM.

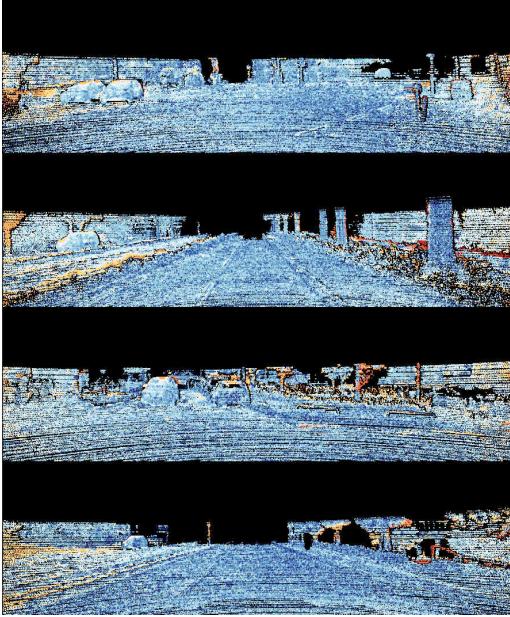


Fig. 5: Direct output of KITTI benchmarking script for scene flow training scenes 130, 84, 87 and 146 respectively (red - bad disparity, blue - good disparity).

Scene	SGM	Proposed	Proposed, interpolated
130.	14.87%	10.19%	4.33%
84.	15.55%	11.79%	6.34 %
87.	18.13%	12.24%	6.51%
148.	12.25%	9.75 %	5.17 %

Table 1: Comparison of the proposed and basic SGM stereo matching accuracy on the KITTI benchmark scenes

CONCLUSION AND FURTHER WORK

In this paper we have presented a stable, accurate, and efficient spatio-temporal disparity estimation algorithm. The proposed approach is based on using ego-motion between consecutive frames to transform the DM from the previous frame to the current one. The transformed disparity, i.e., the predicted disparity, is used as a reference in the disparity search space, which is reduced based on the predicted disparity uncertainty. The newly matched disparity from the reduced search space is then used as a measurement within the Kalman filter in order to update the predicted disparity. This results with reduced computational effort and an increase in the accuracy of estimated DM compared to the basic SGM. We constrained our use-case to scenes without moving objects, since our primary aim was to achieve accurate reconstruction of the static parts of the scene, which will then serve as a base for dense scene flow estimation and moving object detection. Moreover, the proposed approach can be used in applications with no moving objects, e.g., in space exploration robotics. We tested the algorithm on the KITTI benchmark and shown that it can achieve better accuracy than the basic SGM implementation, while reducing the disparity search space.

REFERENCES

Cvišić, I. and Petrović, I. (2015). Stereo odometry based on careful feature selection and tracking. *European Conference on Mobile Robots (ECMR)*.

- Demirdjian, D. and Darrell, T. (2001). Motion estimation from disparity images. In *IEEE International Conference on Computer Vision (ICCV)*.
- Dobias, M. and Sara, R. (2011). Real-time global prediction for temporally stable stereo. In *IEEE International Conference on Computer Vision (ICCV)*, 704–707.
- Franke, U., Rabe, C., Badino, H., and Gehrig, S. (2005). 6D-Vision : Fusion of Stereo and Motion for Robust Environment Perception. *Lecture Notes in Computer Science*, 3663, 216–223.
- Hirschmüller, H. (2008). Stereo Processing by Semi-global Matching and Mutual Information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2), 328–341.
- Hu, X. and Mordohai, P. (2010). Evaluation of stereo confidence indoors and outdoors. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1466–1473.
- Jakob Engel, Jörg Stückler, D.C. (2015). Large-Scale Direct SLAM with Stereo Cameras. In *International Conference on Intelligent Robots and Systems (IROS)*.
- Jiang, J., Cheng, J., Chen, B., and Wu, X. (2014). Disparity prediction between adjacent frames for dynamic scenes. *Neurocomputing*, 142, 335–342.
- Klaus, A., Sormann, M., and Karner, K. (2006). Segment-Based Stereo Matching Using Belief Propagation and a Self-Adapting. In *International Conference on Pattern Recognition (ICPR)*.
- Kolmogorov, V. and Zabih, R. (2001). Computing visual correspondences with occlusions using graph cuts. In *International Conference for Computer Vision*, 508–515.
- Matthies, L., Kanade, T., and Szeliski, R. (1989). Kalman Filter-based algorithms for estimating depth from image sequences. *International Journal of Computer Vision*, 3, 209–236.
- Mayer, N., Ilg, E., Häusser, P., Fischer, P., Cremers, D., Dosovitskiy, A., and Brox, T. (2016). A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Menze, M. and Geiger, A. (2015). Object scene flow for autonomous vehicles. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Morales, S. and Klette, R. (2013). Kalman-filter based spatio-temporal disparity integration. *Pattern Recognition Letters*, 34(8), 873–883.
- Vaudrey, T., Badino, H., and Gehrig, S. (2008). Integrating disparity images by incorporating disparity rate. *Lecture Notes in Computer Science*, 4931, 29–42.
- Wedel, A., Brox, T., Vaudrey, T., Rabe, C., Franke, U., and Cremers, D. (2011). Stereoscopic Scene Flow Computation for 3D Motion Understanding. *International Journal of Computer Vision*, 95(1), 29–51.
- Yang, Q., Wang, L., Yang, R., Stewénius, H., and Nistér, D. (2009). Stereo matching with color-weighted correlation, hierarchical belief propagation, and occlusion handling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(3), 492–504.
- Žbontar, J. and LeCun, Y. (2015). Stereo Matching by Training a Convolutional Neural Network to Compare Image Patches. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2002.