
Article

Not peer-reviewed version

Simultaneous Stereo Matching and Confidence Estimation Network

[Tobias Schmähling](#) ^{*} , [Tobias Müller](#) , Jörg Eberhardt , [Stefan Elser](#)

Posted Date: 3 July 2024

doi: 10.20944/preprints202407.0313.v1

Keywords: Stereo Vision; Confidence; Multi-task Learning; Uncertainty



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Article

Simultaneous Stereo Matching and Confidence Estimation Network

Tobias Schmähling¹ , Tobias Müller¹, Jörg Eberhardt¹  and Stefan Elser²

¹ Institute for Photonic Systems Hochschule Ravensburg-Weingarten, University of Applied Sciences, Doggenriedstraße, 88250 Weingarten, Germany; joerg.eberhardt@rwu.de

² Institute for Artificial Intelligence Hochschule Ravensburg-Weingarten, University of Applied Sciences, Doggenriedstraße, 88250 Weingarten, Germany; stefan.elser@rwu.de

* Correspondence: tobias.schmaehling@rwu.de

Abstract: In this paper, we present a multi-task model that predicts disparities and confidence levels in deep stereo matching simultaneously. We do this by combining its successful model for each separate task and obtaining a multitask model that can be trained with a proposed loss function. We show the advantages of this model compared to training and predicting disparity and confidence sequentially. This method enables an improvement of 15% to 30% in the area under the curve (AUC) metric when trained parallel rather than sequential. In addition, the effect of weighting the components in the loss function on the stereo and confidence performance is investigated. By improving the confidence estimate, the practicality of stereo estimators for creating distance images is increased.

Keywords: stereo vision; confidence; multi-task learning; uncertainty

1. Introduction

Stereo vision is a technology to determine the distance with two RGB cameras. A very high accuracy can be achieved by a precise correspondence search between the two images. Reconstructing the geometric configuration of a 3D environment is one of the fundamental and essential problems in computer vision fields [1]. Over the years, researchers have developed a multitude of approaches to tackle the task of stereo matching. These methods have encompassed both handcrafted techniques [2–5] and machine learning-based approaches [6,7]. However, the complexity of the problem arises from several challenging factors, including reflective surfaces, textureless regions, regions with repeated patterns, occlusions [8–10], as well as photometric deformations resulting from variations in illumination and camera specifications [11,12]. However, an incorrect correspondence estimation can result in a large error in distance determination, making it difficult to use stereo vision in practice.

In response to these challenges, there are two different points of view for recognizing errors in stereo vision. The first is to consider a confidence interval, which indicates the probability that the result is correct. In the second, the uncertainty can be considered as the expected error magnitude. Both views are equally valid, we use the confidence view however, which is widely used in stereo vision. [13] Methods for estimating confidence have been developed in parallel with stereo algorithms. There are classical approaches [14–16] as well as those that employ neural networks for estimating the confidence values. However, due to the focus on confidence estimation, stereo vision algorithms based on neural networks are not used to evaluate the confidence methods. Today, deep end-to-end networks outperform classical stereo methods for deriving dense disparity maps, so confidence estimation should also be tested on stereo vision methods using neural networks rather than classical methods.

In this paper, a neural network is presented that combines the state of the art in stereo vision and the confidence estimation of stereo vision. This network is capable of identifying stereo disparity and the corresponding confidence maps. The combined training results is a significant improvement in the confidence estimation compared to training the two parts separately. This is common practice in the use of confidence networks such as the LAF-Net[17]. This method offers an advantage, as the component of the network employed to estimate the disparity values is trained in a way to allow for a subsequent, more accurate estimation of confidence. The network consists of already successfully tested network

components from the AANet [18] as stereo part and from the LAF-Net for the confidence estimation. The proposed network has high potential to advance the field of stereo vision and paving the way for more robust and reliable applications.

2. Related Works

This section provides an overview of the most important related work in the field of stereo vision and stereo confidence estimation.

2.1. Stereo Vision

In computer vision, stereo vision has made significant progress with the adoption of neural networks. Initially, traditional algorithms were partly replaced with learning methods to enhance the effectiveness of the widely used SGM (Semi Global Matching)[5] algorithm. One approach to improve matching costs involves using neural networks to evaluate the matching score between a pair of image matches. Notable examples of such networks are SGM-Net[19] and SGM-Forest[20]. DNet [21] was one of the first approaches that used the 2D architecture of U-Net [22] for stereo vision inspired by the well-known segmentation networks. There are also later approaches [23,24] that use a 2D architecture. Other stereo networks take their inspiration from the classic stereo pipeline. Numerous works [25–28] used a 3D convolutional layer architecture. GC-Net [25] was one of the first approaches to use continuous stereo adaptation with 3D convolutional layers. These frameworks first map the images through a 2D convolution network to obtain dense representations of the features. A 3D cost volume is then constructed over the 2D feature maps, either by concatenation [25] or correlation [29]. After that, the cost volume is filtered through a series of 3D convolutional layers before being mapped to a point-wise depth estimate by a differentiable argument of the minimum (arg-min) operator. There are many variations of this approach, such as using a 3D stacked hourglass to process the cost volume [29] or developing new aggregation layers [30] to improve accuracy. Unlike classical filtering algorithms such as SGM, 3D convolution is a differentiable approximation. The neural networks outperformed traditional methods on datasets such as KITTI [31] and Scene Flow Datasets [32]. The primary drawback of 3D convolutions is their intensive memory and computing requirements. However, to address this issue, certain methods [18,30] have opted to entirely replace the commonly used 3D convolutions with new components, achieving faster inference speeds while maintaining comparable accuracy. One of the works [18] proposed using an intra-scale cost aggregation method based on sparse points with AANet. In addition, the traditional cross-scale cost aggregation algorithm is approximated with neural network layers to handle large texture-free regions. The recent methods try to reduce the impact of erroneous rectification in stereo vision. RAFT-Stereo [33] takes advantage of the iterative refinement of the RAFT optical flow network [34] to develop a network suitable for stereo matching. Li et al. [35] propose CREStereo network with an adaptive group correlation layer (AGCL) because it can happen that the matching points are not perfect in a line. This can be caused by the fact that the images in the real world are not perfectly rectified.

2.2. Confidence Estimation

In parallel with the further development of stereo algorithms, methods for estimating the confidence in the estimated disparity maps are also advancing. Confidence methods attempt to improve classical stereo methods by identifying erroneous pixels. There are purely classical approaches [14–16,36], which were later combined by a random forest [37,38] to estimate the reliability of each pixel.

In later works, the Convolutional Neural Network (CNN) was also used to estimate confidence, but it must be distinguished which data the network receives as input. Some work only on the basis of the disparity map and the RGB image [39,40]; others [17,41], additionally use the cost volume to estimate a confidence level. Because of this additional data, not all stereo methods can be used for evaluation, so classical stereo algorithms such as SGM or AD-Census are often used. Alternatively, neural networks are used, which are similar to classical methods. These neural networks [42] construct

a cost volume by deriving a probability value for each disparity hypothesis. Therefore, the method is very similar to classical stereo methods and can easily be used for confidence estimation. The focus of the confidence methods is on a good confidence estimate, which is why the stereo performance is not so important. An optimal scenario includes the integration of an excellent stereo assessment together with a simultaneous confidence assessment. The work [13] by Poggi et al. is one of the first to apply confidence methods to end-to-end stereo networks. They determine the confidence with different methods from the results of the Guided Aggregation Net (GANet) [30]. GANet is a stereo network and uses a 3D architecture that generates a feature volume that resembles the usual cost volume. The similarity of this feature volume to the cost volume allows for the application of ready-to-exist confidence methods. They showed the effectiveness of confidence estimation through feature volume analysis using an end-to-end stereo network. However, for this, it is necessary that the stereo network has a feature volume that is as large as the resolution of the input. However, this is not the case for all modern stereo networks. For example, AANet uses only 1/3 of the input resolution mapped as feature volume. In this case, the stereo confidence method must be adjusted. Mehlretter [43] used a Bayesian Neural Network to predict both uncertainty and stereo vision based on GCNet[25] stereo network. Chen et al.[44] use the Group-wise Correlation Stereo Network (GwcNet) [29] as base stereo network and create an uncertainty estimation subnetwork that extracts information from the intermediate multi-resolution. The networks are trained with a new loss function based on KL divergence applied on obtained histograms.

3. Method

The objective of this work is to demonstrate how confidence estimation can be improved using a new multitask network. The network is provided with a rectified image pair, and simultaneously estimates the disparity and the confidence value for each pixel. The network consists of two main components, each with a specific task. The first component handles the correspondence search and predicts the disparity, while the second component is responsible for the confidence estimation of the stereo results. These components are interconnected and allow simultaneous training. The information exchange between the components is unidirectional. This means that the output of certain layers of the stereo image component provides the input for the confidence estimate, but not the other way around.

3.1. Stereo Vision Component

The stereo component is based on the structure of AANet [18]. AANet aims to achieve high performance and high speed simultaneously. To accomplish this, the use of complex 3D CNNs is avoided. Instead, the authors employ a scale-internal cost aggregation based on sparse points, which resolves the well-known issue of edge-fattening at disparity discontinuities. Moreover, the traditional cross-scale cost aggregation algorithm is approximated by using a neural network architecture to handle large textural regions. These modules are lightweight and can be integrated into existing architectures, significantly increasing their speed while maintaining accuracy. In this network, different cost function volumes are generated and subsequently merged. The fast response time, while maintaining high accuracy, makes the networks interesting for practical applications.




Figure 1. Proposed network structure, which consists of a combination of the AANet stereo network and the LAF-Net confidence network. A stereo pair is provided in the AANet to predict the disparity map. The output disparity in full resolution, the output tensor from the CSA with 1/3 resolution and the left input images serve as inputs for the LAF-Net component to predict the confidence. Due to the different resolutions, an interpolation layer must be included in the LAF-Net.

3.2. Confidence Estimation Component

The confidence component is inspired by LAF-Net [17]. LAF-Net takes the disparity map from the stereo component, the cost volume, and the left color image as input, and produces a confidence map as output. However, instead of using all the total cost rates, only the k lowest values of the cost rates are used. These k cost rates represent the most likely correct hypotheses with the highest probability. The value of k is set to 7, as it was found to perform effectively by the authors of the network. The LAF-Net is specifically designed for classic stereo methods, where cost volume is the most important factor. Many stereo networks, including AANet, also provide some form of cost volume. In AANet, however, the cost volume is available in three different resolutions: 1/3, 1/6, and 1/12. For confidence estimation, the full resolution cost volume is required and adjustments need to be made. In addition, we do not use the cost volumes directly after correlation, but after the Cross-Scale Aggregation (CSA) layer. This approach gives the network more flexibility to optimize the cost variables for confidence estimation. The selection of a 1/3 resolution was motivated by its proximity to the desired resolution. However, to suit the specified resolution, the LAF-Net architecture requires adjustments. To achieve this, an interpolation layer is introduced following the second convolution block while extracting features from the cost volume. The network with the two components and their connections to each other is shown in Figure 1.

3.3. Training

Our network is trained on the Scene Flow dataset, a large synthetic dataset that provides dense ground truth disparity maps. This dataset comprises more than 39,000 image pairs. Due to the high number of available data, real data is not used for training. Evaluation is also based on real data, without the need to fine-tune the network on real data. For the training we create a loss function, which is a combination of the loss function of the stereo component (ℓ_{Stereo}) and the LAF-Net component (ℓ_{Conf}), defined as follows:

$$L = \ell_{\text{Stereo}}(D, D_{GT}) + \lambda * \ell_{\text{Conf}}(C_{GT}, C) \quad (1)$$

Where D and C are the predicted disparity and confidence maps. D_{GT} and C_{GT} are the respective ground truth values, which serve as labels. The new loss function is a weighted combination of the individual loss functions of the two network components. The weight is determined by the factor λ . The higher its value, the more emphasis is placed on optimizing the confidence estimate. The loss function of stereo is taken from AANet. It is calculated on the different disparity maps with different resolutions. A smooth L1 loss is generated for each disparity. A weighted sum is formed over the individual losses. The original implementation of the LAF-Net network used a classification loss. This means that the network should predict a value of 1 for a pixel if the error is below a threshold, otherwise it should predict 0. To achieve better coordination between the two loss functions. We use a regression loss function instead of a classification. This will minimize the MSE between the C_{GT} and C . Note that the label C_{GT} for this loss function can only be created after the prediction of the disparity map. The label C_{GT} is determined as follows:

$$C_{GT} = e^{-|D_{GT} - D|} \quad (2)$$

This ensures that the value is always between 0 and 1, where 1 corresponds to high confidence and a value close to 0 corresponds to low confidence.

4. Experiments

To demonstrate the benefits of training stereo and confidence in parallel, we initially conducted separate training for each network component, focusing on the Scene Flow datasets. This approach aligns with the traditional methodology employed by most confidence methods. By treating stereo and confidence networks as separate networks, we aimed to establish a comparative baseline.

To assess the benefits of simultaneous training under real conditions without the need for retraining, we evaluate our network using a combination of synthetic and real datasets. Specifically, we utilize Scene Flow [21], a synthetic dataset comprising over 39,000 stereo frames with a resolution of 960x540 pixels, for both training and testing our network. Additionally, we employ the widely used KITTI2015 [45] stereo dataset, to further test our network's performance. Notably, we do not train our network but use the official training set of KITTI2015 for our evaluation, allowing for an unbiased evaluation. Furthermore, we incorporate the RWU3D [46] dataset, a novel dataset designed for the fusion of stereo and Time-of-Flight (ToF) sensors, capturing indoor industrial environments. The dataset includes images with a different resolution, in this work the resolution 960 x 560 pixels is used. For testing purposes, we have selected a subset of 32 images, excluding the calibration scene and scenes intended for characterizing the ToF camera.

4.1. Metrics

To evaluate our network, we employ a variety of evaluation metrics. In this regard, we have adopted the evaluation methodology of LAF-Net, which exhibits subtle differences from the evaluation approach utilized in SEDNet. We prioritize the perspective of trust over uncertainty and have therefore opted for evaluation approach. First, we use the training metrics. These include the Mean Squared Error (MSE) between the ground truth disparity and the predicted disparity, as well as the difference between the predicted confidence and the label for confidence. Second, we use additional evaluation metrics. We employ the Bad3 metric to evaluate the stereo part and use the Area Under the Curve (AUC) to evaluate the confidence part.

4.1.1. Bad3 (Bad Pixels Rate):

The Bad3 metric provides an assessment of the accuracy of our stereo model in estimating disparities. It measures the percentage of pixels for which the disparity error exceeds a threshold of 3. A lower Bad3 value indicates better performance, as it signifies a smaller proportion of pixels with large disparity errors.

4.1.2. Area under the Curve (AUC):

An optimal confidence algorithm produces a value for each pixel that is inverse to its error. The lower the error level, the higher the confidence should be. To evaluate the performance of each confidence measure in identifying correct matches, as outlined in [16], we sort the pixels in a disparity map in descending order of confidence and calculate the error rate (Bad3) on sparse maps generated through iterative sampling (e.g., 5% of pixels at a time) from the density map. This means that first the error of 5% of pixels with the highest confidence is calculated. Next, the error of the 10% of pixels with the highest confidence is calculated and so on. The plot of error rates produces a sparsification curve, which allows us to quantitatively measure the effectiveness of the confidence measure through its AUC. The lower the AUC, the better the performance of the confidence measure. The optimal AUC is achieved when the confidence measure is able to identify all correct matches first and is equal to:

$$AUC_{Opt} = \int_{1-\epsilon}^1 \frac{x - (1 - \epsilon)}{x} dx = \epsilon + (1 - \epsilon) \ln(1 - \epsilon) \quad (3)$$

with ϵ being the *Bad3* computed over the disparity map. Looking at the ROC curve, one can tell how high the error is, for example, if 80% of the pixels with the highest confidence are observed.

4.2. Implementation Details

We implemented the network in PyTorch and trained it with an Nvidia RTX 4090 GPU. For training, we used 35454 stereo pairs from the training set of the Scene Flow dataset. Before input, the raw images are randomly cropped to a size of 288×576 . Adam [47] ($\beta_1 = 0.9$, $\beta_2 = 0.999$) is used as the optimizer. We started with a learning rate of 0.001, which is gradually halved every 10 epochs after 20 epochs. Each configuration is trained for 70 epochs, with a batch size of 4. Firstly, the stereo and confidence parts were trained sequentially, which corresponds to the results of the AANet and LAF-Net, whereby only the 1/3 resolution of the cost function could be used, as the AANet does not have a total cost function.

Secondly, the network is trained in parallel using the proposed loss function. The loss function is used with λ values of 5, 10, and 20. For each training configuration, the network is randomly initialized. It should be noted that the values for λ should be significantly greater than 1, as the values for stereo loss are much greater than those for the condensation loss function. This is due to the fact that the predicted value for the disparity is between 0 and 192, while the predicted value for confidence is only between 0 and 1.

5. Results

Table 1. Evaluation results with different training configurations for Scene Flow, KITTI2015 and RWU3D datasets. Parallel(λ) is our network with the different λ values (5, 10, and 20) for the loss function.

Network	MSE Disparity	MSE Confidence	Bad3 [%]	AUC	AUC_{opt}
Scene Flow					
SEDNet[44]	0.835	-	4.31	0.36	0.19
Sequential	0.680	0.0280	3.21	0.37	0.19
Parallel(5)	0.674	0.0236	3.34	0.32	0.20
Parallel(10)	0.695	0.0220	3.47	0.32	0.21
Parallel(20)	0.720	0.0203	3.71	0.33	0.21
KITTI2015					
SEDNet[44]	10.828	-	35.33	9.55	4.61
Sequential	1.858	0.2705	10.70	3.71	1.23
Parallel(5)	1.878	0.2579	11.20	3.32	1.31
Parallel(10)	1.738	0.2921	10.36	2.76	1.20
Parallel(15)	1.852	0.2679	10.82	2.89	1.27
RWU3D					
SEDNet[44]	11.835	-	39.07	13.47	7.10
Sequential	2.445	0.1547	17.30	6.77	2.32
Parallel(5)	2.681	0.0860	18.65	6.17	2.85
Parallel(10)	2.571	0.0824	17.89	5.22	2.46
Parallel(15)	2.613	0.0657	18.85	5.42	2.71

Table 1 presents the comparison metrics of different settings against SEDNEet. Overall, our networks outperform the SED network, primarily due to the use of AANet as the backbone for stereo matching, which already demonstrates superior stereo performance. However, it's noteworthy how much better our approach generalizes to real data compared to the SED network, without requiring re-training. The results in Table 1 also show that the combination of training confidence and stereo leads to an improved performance in terms of confidence. The results depend on the values used for λ in the loss function. With a small value of 5, the stereo results do not suffer, and the confidence performance is still better. Increasing the value of λ increases the MSE and Bad3 of the stereo, but decreases the MSE of confidence. This clear trend can be seen particularly well with Scene Flow. In principle, this observation can also be made with the other data sets, although this trend sometimes shows outliers. Several factors can explain this phenomenon. Overweighting the confidence components can result in overfitting, which may worsen performance on a new data set. However, adjusting the confidence level slightly can prevent overfitting the stereo and improve, or at the very least maintain, the stereo score. Alternatively, the smaller real test dataset may have introduced some normal noise. To summarize, parallel training leads to a better overall result, which can be easily determined using the AUC value. When the stereo and confidence components are trained together, the AUC values are closer to the AUC_{Opt} value than when they are trained separately. The AUC_{Opt} value is the lower bound and is the value that can be theoretically achieved with optimal confidence estimation. The AUC value is influenced by both the stereo performance and the confidence performance. Therefore, this will be a good indication of which λ value is most appropriate. With the Scene Flow dataset, it is $\lambda = 5$, with KITTI2015 and RWU3D it is $\lambda = 10$ where the lowest AUC is found. Therefore, we can deduce that an optimal value lies between 5 and 10.




Figure 2. The sparsification curve show the rate of bad pixels at threshold 3 when a certain percentage of pixels with the lowest prediction confidence are excluded.

As mentioned above, the AUC is the most interesting value. So let us look at the underlying sparsification curve shown in the figure 2. The curves show the advantage of parallel training with $\lambda = 10$ even more clearly than the figures in the table. The curve shows bad pixels rate when looking at the pixels with the highest confidence. The sparsification indicates the percentage of pixels with low confidence that are ignored. Especially for the two real test data sets RWU3D and KITTI2015, the improvement is even more visible than for the test data from Scene Flow. Therefore, we can assume that our training approach has even led to an increase in generalization ability or, at least, that no overfitting has taken place. In particular, the generalization to a new real dataset is evident when compared to SENet. This improvement is mainly due to the stereo component of the network, which has a strong influence on the AUC curve. Upon detailed examination of the sparsification curve 2, it is possible to identify the bad pixels rate that are reached when only considering the pixels with the highest confidence values. By ignoring 20% of the low confidence pixels identified through sequential training, the error rate of the remaining pixels in the KITTI2015 dataset is reduced to 0.05. With parallel training and improved confidence estimation, this error rate could be further reduced to 0.035. This shows that by masking pixels with a low confidence, the total error of the overlaid pixels is smaller. However, for an application, it is important that the masked pixels are as evenly distributed as possible, which is the case with our confidence estimate, as can be seen in the example image in Figure 3. It is evident that the application of confidence estimation allows the removal of erroneous pixels. By establishing a threshold value of 0.3, any pixel with a confidence estimate lower than this threshold is excluded. This observation aligns with the AUC values, indicating that the filtered pixels are distributed throughout the image and do not signify a substantial loss of critical information.

The strong correlation between error and lower confidence prediction is evident in Figure 4, particularly within the highlighted example region. Regions with height errors often also have a lower confidence, this correlation can be used to obtain a higher confidence for the stereo image in the application.

5.1. Computational Performance and Efficiency

The computational complexity of the proposed network is evaluated by measuring the time required for a single stereo and confidence prediction on a 960x540 image using an Nvidia RTX 4090

GPU. With a processing time of 98 ms, the network achieves a frame rate of more than 10 frames per second, which makes our network interesting for real-time applications. The stereo computer is the most complex of the networks and by using the AANet's computing power instead of other popular stereo networks such as GwcNet [29], a very high response time is still achieved. Additionally, the number of parameters shown in Table 2 is significantly lower than the number of parameters in the SEDNet.

Table 2. Number of parameters.

Architecture	Parameters (M)
SEDNet	6.91
Our Network	
Stereo Component	3.93
Confidence Component	0.54
Total	4.47




Figure 3. The example shows the left RGB images with the predicted disparity. In the error map, the black pixels have an error of more than 3 disparity values. Using the predicted confidence, 20% of the pixels with the lowest confidence were removed, resulting in the disparity map shown in the middle of the second row. The corresponding error map for this disparity map is shown next. Note that the number of black pixels is now lower, indicating that the predicted confidence has removed pixels appropriately.




Figure 4. Example scene from the RWU3D dataset. In the left column the disparity map is visualized, in the center the corresponding error map, and on the right the predicted confidence with whites showing high confidence. The green rectangles show some areas where the confidence is correctly low, as there is also a larger error there.

6. Conclusions

We have demonstrated the feasibility of combining deep stereo networks with a confidence estimation network. In contrast to the stereo techniques employed in evaluating LAF-Net, the stereo output produced by AANet exhibits minimal inconsistencies. However, our network successfully attains a robust confidence estimate. Training the stereo and confidence components simultaneously improves the overall results compared to training them sequentially. This can improve the AUC values by about 15% for the synthetic data sets. When using real data sets, an improvement of about 30% can be observed. It is therefore clear that parallel training offers benefits and is a potential form of multi-task learning. Additionally, our findings reveal that synergizing the training of stereo and confidence networks yields even more promising results. Through the judicious exclusion of pixels with lower confidence scores, we can generate stereo depth images that are notably more reliable. However, there is still scope for improvement to achieve the optimal confidence estimate.

Author Contributions: Conceptualization, T.S., T.M., J.E and S.E.; methodology, T.S.; software, T.S.; validation, T.S., T.M., J.E and S.E.; investigation, T.S.; writing—original draft preparation, T.S.; writing—review and editing, T.M., J.E and S.E.; visualization, T.S.; supervision, J.E and S.E.; All authors have read and agreed to the published version of the manuscript.

Funding: This work is part of the 3D Robust project, which is funded by the Federal Ministry of Education and Research (BMBF)

Data Availability Statement: The data presented in this study are available on request from the corresponding author

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

ToF	Time of Flight
AUC	Area Under the Curve
Bad3	Bad Pixels Rate with threshold of 3
MSE	Mean Squared Error
CSA	Cross-Scale Aggregation

References

1. Scharstein, D.; Szeliski, R. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International journal of computer vision* **2002**, *47*, 7–42.
2. Hannah, M.J. *Computer matching of areas in stereo images.*; Stanford University: Stanford, CA, 1974.
3. Pollard, S.B.; Mayhew, J.E.; Frisby, J.P. PMF: A stereo correspondence algorithm using a disparity gradient limit. *Perception* **1985**, *14*, 449–470.
4. Baker, H.H.; Bolles, R.C.; Woodfill, J.I. Real-time stereo and motion integration for navigation. ISPRS Commission III Symposium: Spatial Information from Digital Photogrammetry and Computer Vision. SPIE, 1994, Vol. 2357, pp. 17–24.
5. Hirschmuller, H. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on pattern analysis and machine intelligence* **2007**, *30*, 328–341.
6. Zbontar, J.; LeCun, Y. Computing the stereo matching cost with a convolutional neural network. Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 1592–1599.
7. Shaked, A.; Wolf, L. Improved stereo matching with constant highway networks and reflective confidence learning. Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 4641–4650.
8. Mei, X.; Sun, X.; Zhou, M.; Jiao, S.; Wang, H.; Zhang, X. On building an accurate stereo matching system on graphics hardware. 2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops). IEEE, 2011, pp. 467–474.
9. Egnal, G.; Wildes, R.P. Detecting binocular half-occlusions: Empirical comparisons of five approaches. *IEEE Transactions on pattern analysis and machine intelligence* **2002**, *24*, 1127–1133.
10. Humenberger, M.; Zinner, C.; Weber, M.; Kubinger, W.; Vincze, M. A fast stereo matching algorithm suitable for embedded real-time systems. *Computer Vision and Image Understanding* **2010**, *114*, 1180–1202.
11. Zabih, R.; Woodfill, J. Non-parametric local transforms for computing visual correspondence. Computer Vision—ECCV'94: Third European Conference on Computer Vision Stockholm, Sweden, May 2–6 1994 Proceedings, Volume II 3. Springer, 1994, pp. 151–158.
12. Heo, Y.S.; Lee, K.M.; Lee, S.U. Robust stereo matching using adaptive normalized cross-correlation. *IEEE Transactions on pattern analysis and machine intelligence* **2010**, *33*, 807–822.
13. Poggi, M.; Kim, S.; Tosi, F.; Kim, S.; Aleotti, F.; Min, D.; Sohn, K.; Mattoccia, S. On the confidence of stereo matching in a deep-learning era: a quantitative evaluation. *IEEE transactions on pattern analysis and machine intelligence* **2021**, *44*, 5293–5313.
14. Scharstein, D.; Szeliski, R. Stereo matching with nonlinear diffusion. *International journal of computer vision* **1998**, *28*, 155–174.
15. Egnal, G.; Mintz, M.; Wildes, R.P. A stereo confidence metric using single view imagery with comparison to five alternative approaches. *Image and vision computing* **2004**, *22*, 943–957.
16. Hu, X.; Mordohai, P. A quantitative evaluation of confidence measures for stereo vision. *IEEE transactions on pattern analysis and machine intelligence* **2012**, *34*, 2121–2133.
17. Kim, S.; Kim, S.; Min, D.; Sohn, K. Laf-net: Locally adaptive fusion networks for stereo confidence estimation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 205–214.
18. Xu, H.; Zhang, J. Aanet: Adaptive aggregation network for efficient stereo matching. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 1959–1968.
19. Seki, A.; Pollefey, M. Sgm-nets: Semi-global matching with neural networks. Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 231–240.

20. Schonberger, J.L.; Sinha, S.N.; Pollefeys, M. Learning to fuse proposals from multiple scanline optimizations in semi-global matching. Proceedings of the European Conference on Computer Vision (ECCV), 2018, pp. 739–755.

21. Mayer, N.; Ilg, E.; Hausser, P.; Fischer, P.; Cremers, D.; Dosovitskiy, A.; Brox, T. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 4040–4048.

22. Gu, Z.; Cheng, J.; Fu, H.; Zhou, K.; Hao, H.; Zhao, Y.; Zhang, T.; Gao, S.; Liu, J. Ce-net: Context encoder network for 2d medical image segmentation. *IEEE transactions on medical imaging* **2019**, *38*, 2281–2292.

23. Liang, Z.; Feng, Y.; Guo, Y.; Liu, H.; Chen, W.; Qiao, L.; Zhou, L.; Zhang, J. Learning for disparity estimation through feature constancy. Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 2811–2820.

24. Tonioni, A.; Tosi, F.; Poggi, M.; Mattoccia, S.; Stefano, L.D. Real-time self-adaptive deep stereo. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 195–204.

25. Kendall, A.; Martirosyan, H.; Dasgupta, S.; Henry, P.; Kennedy, R.; Bachrach, A.; Bry, A. End-to-end learning of geometry and context for deep stereo regression. Proceedings of the IEEE international conference on computer vision, 2017, pp. 66–75.

26. Chang, J.R.; Chen, Y.S. Pyramid stereo matching network. Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 5410–5418.

27. Nie, G.Y.; Cheng, M.M.; Liu, Y.; Liang, Z.; Fan, D.P.; Liu, Y.; Wang, Y. Multi-level context ultra-aggregation for stereo matching. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 3283–3291.

28. Chabra, R.; Straub, J.; Sweeney, C.; Newcombe, R.; Fuchs, H. Stereodrnet: Dilated residual stereonet. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 11786–11795.

29. Guo, X.; Yang, K.; Yang, W.; Wang, X.; Li, H. Group-wise correlation stereo network. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 3273–3282.

30. Zhang, F.; Prisacariu, V.; Yang, R.; Torr, P.H. Ga-net: Guided aggregation net for end-to-end stereo matching. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 185–194.

31. Menze, M.; Geiger, A. Object Scene Flow for Autonomous Vehicles. Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

32. Mayer, N.; Ilg, E.; Häusser, P.; Fischer, P.; Cremers, D.; Dosovitskiy, A.; Brox, T. A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation. IEEE International Conference on Computer Vision and Pattern Recognition (CVPR), 2016. arXiv:1512.02134.

33. Lipson, L.; Teed, Z.; Deng, J. Raft-stereo: Multilevel recurrent field transforms for stereo matching. 2021 International Conference on 3D Vision (3DV). IEEE, 2021, pp. 218–227.

34. Teed, Z.; Deng, J. Raft: Recurrent all-pairs field transforms for optical flow. Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16. Springer, 2020, pp. 402–419.

35. Li, J.; Wang, P.; Xiong, P.; Cai, T.; Yan, Z.; Yang, L.; Liu, J.; Fan, H.; Liu, S. Practical stereo matching via cascaded recurrent network with adaptive correlation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 16263–16272.

36. Kim, S.; Yoo, D.g.; Kim, Y.H. Stereo confidence metrics using the costs of surrounding pixels. 2014 19th International Conference on Digital Signal Processing. IEEE, 2014, pp. 98–103.

37. Haeusler, R.; Nair, R.; Kondermann, D. Ensemble learning for confidence measures in stereo vision. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2013, pp. 305–312.

38. Spyropoulos, A.; Komodakis, N.; Mordohai, P. Learning to detect ground control points for improving the accuracy of stereo matching. Proceedings of the IEEE conference on computer vision and pattern recognition, 2014, pp. 1621–1628.

39. Tosi, F.; Poggi, M.; Benincasa, A.; Mattoccia, S. Beyond local reasoning for stereo confidence estimation with deep learning. Proceedings of the European conference on computer vision (ECCV), 2018, pp. 319–334.

40. Fu, Z.; Ardabilian, M.; Stern, G. Stereo matching confidence learning based on multi-modal convolution neural networks. Representations, Analysis and Recognition of Shape and Motion from Imaging Data:

7th International Workshop, RFMI 2017, Savoie, France, December 17–20, 2017, Revised Selected Papers 7. Springer, 2019, pp. 69–81.

- 41. Mehlretter, M.; Heipke, C. Cnn-based cost volume analysis as confidence measure for dense matching. Proceedings of the IEEE/CVF international conference on computer vision workshops, 2019, pp. 0–0.
- 42. Zbontar, J.; LeCun, Y.; others. Stereo matching by training a convolutional neural network to compare image patches. *J. Mach. Learn. Res.* **2016**, *17*, 2287–2318.
- 43. Mehlretter, M. Joint estimation of depth and its uncertainty from stereo images using bayesian deep learning. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **2022**, *2*, 69–78.
- 44. Chen, L.; Wang, W.; Mordohai, P. Learning the Distribution of Errors in Stereo Matching for Joint Disparity and Uncertainty Estimation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 17235–17244.
- 45. Menze, M.; Heipke, C.; Geiger, A. Joint 3D Estimation of Vehicles and Scene Flow. ISPRS Workshop on Image Sequence Analysis (ISA), 2015.
- 46. Agrawal, A.; Müller, T.; Schmähling, T.; Elser, S.; Eberhardt, J. RWU3D: Real World ToF and Stereo Dataset with High Quality Ground Truth. 2023 Twelfth International Conference on Image Processing Theory, Tools and Applications (IPTA). IEEE, 2023, pp. 1–6.
- 47. Barnes, C.; Shechtman, E.; Finkelstein, A.; Goldman, D.B. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.* **2009**, *28*, 24.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.