

CNN-Based Dense Image Matching for Aerial Remote Sensing Images

Shunping Ji, Jin Liu, and Meng Lu

Abstract

Dense stereo matching plays a key role in 3D reconstruction. The capability of using deep learning in the stereo matching of remote sensing data is currently uncertain. This article investigated the application of deep learning-based stereo methods in aerial image series and proposed a deep learning-based multi-view dense matching framework. First, we applied three typical convolutional neural network models, MC-CNN, GC-Net, and DispNet, to aerial stereo pairs and compared the results with those of the SGM and a commercial software, SURE. Second, on different data sets, the generalization ability of each network is evaluated by using direct transfer learning with models pretrained on other data sets and by fine-tuning with a small number of target training data. Third, we present a deep learning-based multi-view dense matching framework where the multi-view geometry is introduced to further refine matching results. Three sets of aerial images as the main data sets and two open-source sets of street images as auxiliary data sets are used for testing. Experiments show that, first, the performance of deep learning-based stereo methods is slightly better than traditional methods. Second, both the GC-Net and the MC-CNN have demonstrated good generalization ability and can obtain satisfactory results on aerial images using a pretrained model on several available stereo benchmarks. Third, multi-view geometry constraints can further improve the performance of deep learning-based methods, which is better than that of the multi-view-based SGM and SURE.

Introduction

Reconstructing terrestrial 3D scenes from stereo or multi-view aerial/satellite images has been a core problem in photogrammetry and remote sensing. The key technique is to obtain the correspondent points for each pixel in the stereo image pairs, which is commonly called dense stereo matching. A conventional process of dense stereo matching includes four steps (Scharstein and Szeliski 2002). The first is to calculate matching cost. Typical matching costs include luminance difference, correlation coefficient, and mutual information related to a pixel's values or distributions. Given a search area along an epipolar line, one of these costs is calculated pixel by pixel, and the minimum cost corresponds to the optimal matching candidate. These empirical designed costs could be heavily affected by the nontexture area, mirror reflection, and repeated pattern (Kendall *et al.* 2017). The second step is to aggregate the matching costs. The cost aggregation is commonly a weighted sum of all of the matching costs within a given neighborhood. However, traditional methods, such as Graph-Cut (Boykov *et al.* 2001) and SGM (Furukawa and Ponce 2010), could oversimplify the cost aggregation, for

example, using an empirical fixed window size and treating the pixels in a neighborhood independently. The third step is disparity calculation. At each pixel, the disparity value corresponds to the minimum cost. Interpolation can be used to achieve subpixel accuracy. The last step is parallax refinement including a series of postprocessing techniques, such as left and right consistency checking, median filter smoothing, and subpixel enhancement. Finally, the dense disparity map could be converted to a depth map to reconstruct a 3D scene.

Dense matching has been extensively studied. We classify them into conventional methods and deep learning-based methods. Graph Cut (Boykov *et al.* 2001) is a widely accepted global stereo matching algorithm introduced early in the 21st century. It uses graph theory, especially graph cut, to solve the problem of 2D energy minimization. Global matching algorithms such as Graph Cut are computationally expensive and are not suitable for large-volume remote sensing images. In 2008, a semiglobal matching method (SGM) with higher matching efficiency has been proposed (Hirschmüller 2008). The SGM considers the 2D cost aggregation as 16 1D cost aggregations and performs dynamic programming to solve the minimum cost. The Patch-Match (Bleyer *et al.* 2011) algorithm uses the local correlation of the image and assumes that the areas around the matching points also match each other. In addition to stereo, multi-view geometry is often used in dense matching for stronger constraints. The PMVS (Patch-Based Multi-View Stereo) algorithm (Furukawa and Ponce 2010) extracts feature points and retrieves the surrounding patches centered at the feature points and performs patch matching to obtain quasi-dense matching points. The SURE algorithm proposed by Rothmel *et al.* (2012) extends the stereo SGM to multi-view image matching. Multi-view geometry fusion is added to merge the redundant depth estimation values to achieve mutual restraint.

Deep learning, especially *convolutional neural networks* (CNNs), have been widely applied in image processing. It has been shown that CNN-based methods (LeCun *et al.* 2015; Schmidhuber 2015) can not only improve the accuracy of image recognition and classification but also increase the efficiency of online operations. More important, empirical and manual feature engineering can be replaced by the powerful representation learning ability of deep learning. From 2015, some studies have started applying deep learning to dense stereo matching to replace the empirically designed matching costs, cost aggregation, or the whole matching procedure. The matching results obtained on the computer vision benchmarks have gradually exceeded the traditional methods in speed and accuracy.

Two strategies are commonly used for CNN-based dense matching methods: (1) the end-to-end prediction from image to disparity image and (2) applying CNN to learn parts of the four standard steps of stereo matching. For example, the MC-CNN network (Zbontar and LeCun 2015) automatically learns

Shunping Ji and Jin Liu are with the School of Remote Sensing and Information Engineering, Wuhan University, Wuhan, China, 430079 (jishunping@whu.edu.cn).

Meng Lu is with the Department of Physical Geography, Faculty of Geoscience, Utrecht University, Utrecht, The Netherlands.

Photogrammetric Engineering & Remote Sensing
Vol. 85, No. 6, June 2019, pp. 415–424.
0099-1112/19/415-424

© 2019 American Society for Photogrammetry
and Remote Sensing
doi: 10.14358/PERS.85.6.415

the matching cost (i.e., step 1) via a Siamese CNN structure. SGM-Net (Seki and Pollefeys 2017) introduces a CNN learning penalty terms in the standard process of SGM.

The end-to-end learning strategy predicts disparity images directly from stereo pairs. For example, DispNet (Mayer *et al.* 2016) is a typical fully convolutional network (FCN). In the encoding stage, the network extracts high-level features of stereo images layer by layer; in the decoding stage, the network restores the feature map from coarse to the original image resolution to produce the disparity map. GC-Net (Kendall *et al.* 2017) makes full use of the geometric information and semantic information between pixels. 3D volume with context information, consisting of 2D feature maps cross disparities extracted by 2D CNN, is convoluted by a series of 3D kernels and finally flattened into a 2D disparity image. The pyramid stereo matching network (PSM-Net) (Chang and Chen 2018) is a pyramid stereo matching network consisting of spatial pyramid pools and 3D convolutional layers. It combines the global background information into stereo matching to achieve reliable estimation of occlusion areas, textureless areas, or pattern repeated areas. The cascade residual learning (Pang *et al.* 2017) method concatenates two improved DispNet networks. The first network obtains the initial disparity value between stereo pairs; the second network uses the residuals of the previous stage to train a finer disparity map. In Shaked and Wolf (2017), a new highway network structure is proposed; multi-level residual skip connections and composite loss function is applied. All of these methods operate in a supervised manner, requiring high-precision disparity maps as labels for training. An exception was presented by Zhong *et al.* (2017), who designed a CNN to learn disparity maps directly from the stereo pair without training samples based on the assumption that the left disparity map (based on the left image) and the right disparity map (based on the right image) learned by the network are inverted.

Although the deep learning-based stereo methods have been applied to match close-range images and have achieved improved results compared to conventional methods, it has not been applied to remote sensing images. In addition, the capability of it has not been compared with the mainstream photogrammetric algorithms. The main objective of this work is to comprehensively investigate the application of deep learning-based stereo methods on aerial remote sensing images and to compare them with conventional methods, including commercial software. The second contribution is to introduce multi-view geometry for deep learning-based dense matching for the first time. Multi-view geometry gives more constraints between co-visible images and could be more robust than a stereo vision. In addition, we evaluate the generalization ability of deep learning in aerial stereo matching; that is, can the model trained on the available benchmark data sets be directly applicable to aerial imagery? It is highly relevant, as in the existing and upcoming aerial image data sets, accurately labeled samples are commonly lacking or insufficient.

Method

CNN for Learning Only Matching Cost

The similarity score calculated by the normalized correlation coefficient, intensity difference, or cross entropy is shown to be incompetence in some complex situations. Matching cost learned by CNNs has shown advantages over those empirical designs on close-range data (Han *et al.* 2015; Zagoruyko and Komodakis 2015; Aguilera *et al.* 2016). In this study, we use the MC-CNN (especially the fast structure) (Zbontar and LeCun 2015) for evaluating the performance of a CNN-based matching cost.

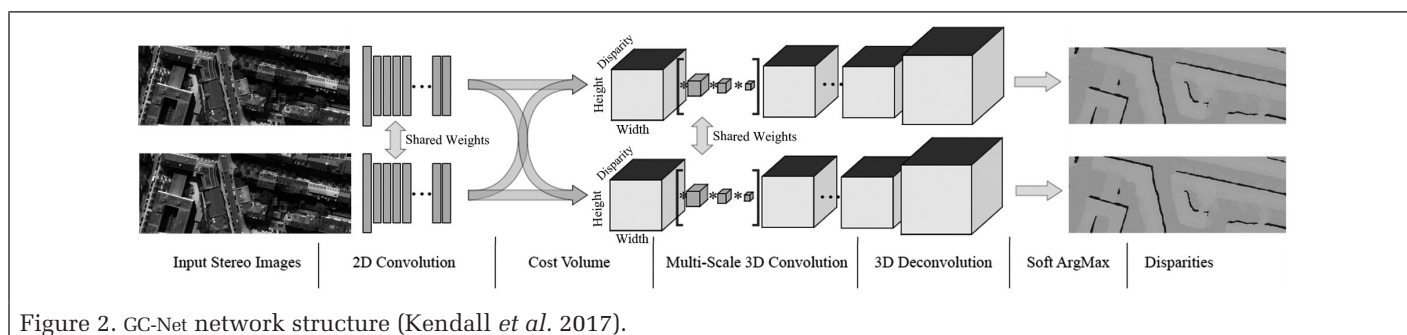
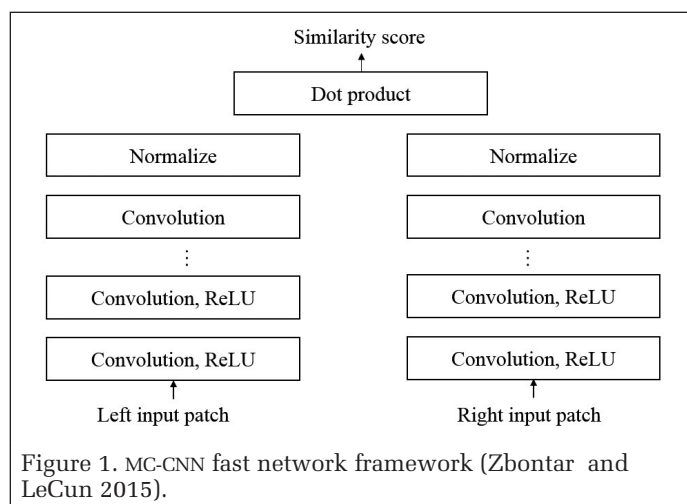
The fast MC-CNN structure uses Siamese convolutional networks with shared weights to extract feature vectors from the input stereo tiles. The dot product operator measures the similarity between the two extracted and normalized feature vectors. The network structure is shown in Figure 1.

In this article, the number of convolution layers is set to 4, and the convolution kernel size is set to 3×3 . The MC-CNN calculates the loss values of a pair of positive and negative samples and trains the network by minimizing a hinge loss function. The hinge loss of the positive and negative samples is defined as $\max(0, m + s-, s+)$ where $s+$ is the output of the positive sample, $s-$ is the output of the negative sample, and the tolerance m is set to 0.2.

Based on the initial disparity map that is computed with the learned matching score, a series of postprocessing steps including cost aggregation (Zhang *et al.* 2009), semiglobal matching, left and right consistency check, subpixel enhancement, median filtering, and bilateral filtering are applied to ensure the best matching results. The process bears a resemblance to SGM.

Special CNN for End-to-End Stereo Matching

We use the GC-Net (Kendall 2017) as an end-to-end stereo benchmark algorithm and evaluate its performance on aerial data sets. Its core concept of geometry and context combination is to treat the parallax as the third dimension orthogonal



to the image plane. The feature maps, which are learned from a series of 2D convolutional layers, consist of a 3D tensor across each disparity. The 3D convolution then learns geometric and semantic features to obtain an optimal disparity map (i.e., a curved surface cutting through the 3D tensors). The GC-Net network structure is shown in Figure 2. Several 2D convolutions are applied to the stereo pairs with shared weights to extract feature maps. The m feature maps of the last layer with size of $w \times h$ are concatenated across $0 \sim n$ disparity to constructed m feature volumes with size of $w \times h \times (n + 1)$. 3D convolution and deconvolution are executed to learn a series of 3D feature maps. The size of the final 3D maps is $W \times H \times n$, where H and W are the length and width of the original image, respectively. The last step is to flatten the 3D feature map to a 2D disparity map with a Soft Argmin operation. The predicted disparity map and its corresponding label are utilized to train the optimal parameters by iterating the forward- and backward-propagating process. In our experiment, we used 18 2D convolutional layers with $32 \ 3 \times 3$ kernels; 14 3D convolutional layers, and five 3D deconvolution layers with a kernel size of $3 \times 3 \times 3$.

General FCN for End-to-End Stereo Matching

Although GC-Net can combine context information with geometry and is preferable theoretically, the 3D convolution requires much more memory compared to the 2D convolution. An end-to-end learning from a 2D image pair to a 2D disparity map could be achieved using a general FCN. Among several 2D FCNs that have reported satisfactory performance in stereo matching (Mayer *et al.* 2016), we used the DispNet network for evaluation, which is based on the FlowNet (Flow Estimation Network) (Dosovitskiy *et al.* 2015) and is modified to calculate disparity maps. DispNet is a typical FCN structure consisting of an encoder and a decoder. The encoder has six convolutional layers where the kernel sizes of the first two layers are 7×7 and 5×5 , respectively, and all the other layers are 3×3 . The decoder consists of five up-convolutional layers

with kernel size 4×4 ; each layer is first concatenated with the feature map of the corresponding layer in the encoding step and then merged by a series of convolution operations. The structure is similar to mainstream semantic segmentation CNN, such as U-Net (Ronneberger *et al.* 2015). A schematic diagram of the DispNet network is shown in Figure 3.

Transfer Learning

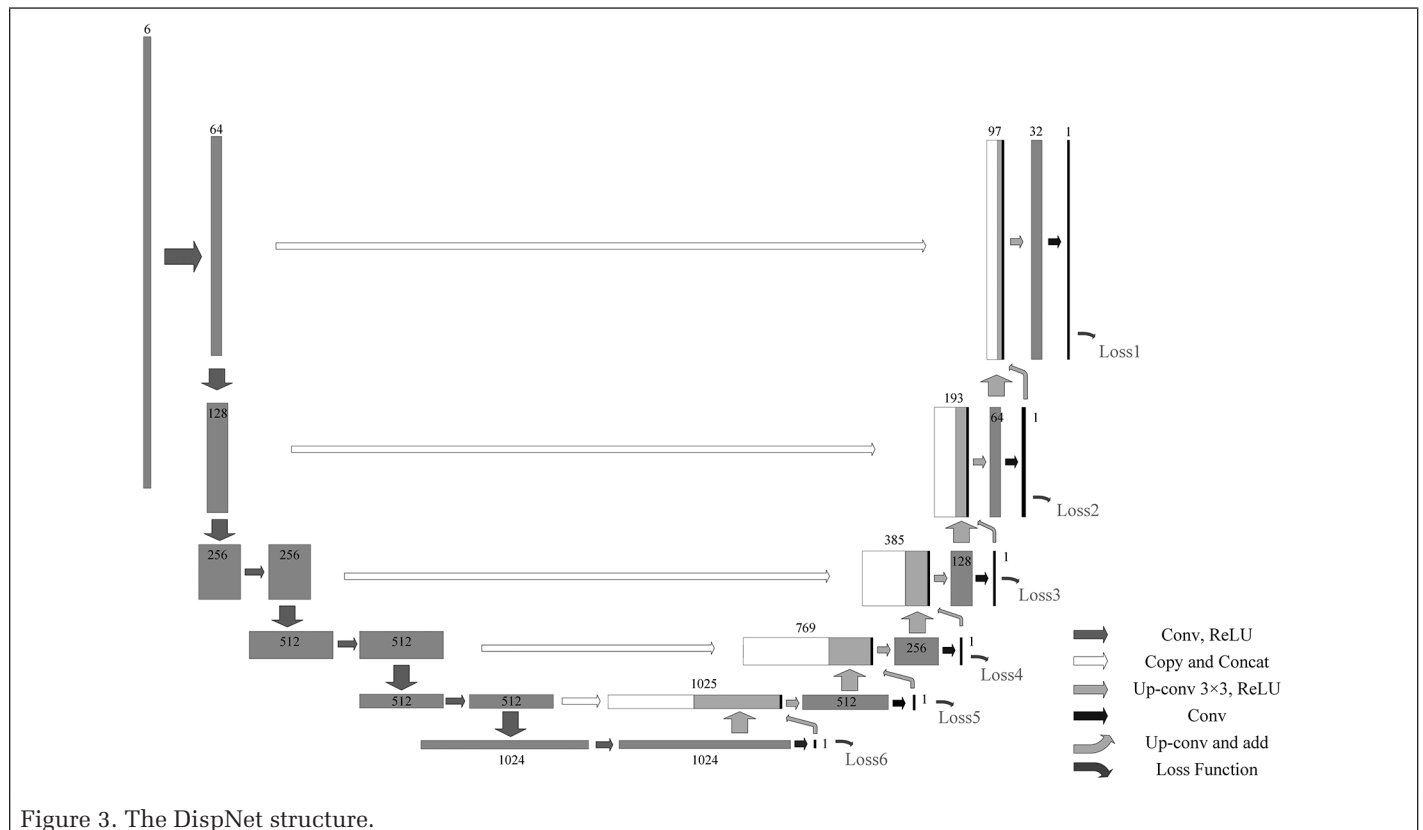
The performance of applying a pre-trained dense matching model on a different data set is a key issue in practice, as we seldom have enough highly accurate disparity maps in target data to train a CNN with multiple parameters. We adopt and evaluate two transfer learning strategies (Pan and Yang 2010) direct model transfer and model fine-tuning—with small samples available in target data set.

A direct model transfer utilizes the model pretrained on the source data set, predicting the target data without any parameter tuning. This method requires high generalization ability of the model. Especially, we evaluate whether the learned features through a CNN for finding pixel correspondence is universal for all sorts of data sources, including close-range images, aerial images, and simulated scene images.

Assuming that the target data set have insufficient samples to train a robust network model, fine-tuning with these samples conditioned on a pretrained model is a common choice. It can reduce the number of iterations required for training a new model and alleviate the problem of insufficient samples. There are two different strategies for fine-tuning: one is to train the parameters of all layers, and the other is to train only the top layers and freeze the bottom layers. As the number of network layers involved in this article is relatively few, we retrain all the parameters in the CNN.

Deep Learning-Based Multi-View Dense Matching

Up to now, most of the CNN-based stereo methods evaluate their performances only at a pixel level (i.e., disparity rather than in depth), which is insufficient, as their accuracy is not



equal. Further, the important multi-view geometry has not been applied to CNN-based stereo matching. We translate the disparity to depth and merge the multi-view matching results on the georeferenced coordinate system for comprehensively understanding the depth accuracy that a CNN-based method could reach.

To generate rectified epipolar stereos of each image pair given a base image, we adopt the method used in Fusiello *et al.* (2000), where only some simple 3×3 2D perspective transformation matrices are required for transferring between the epipolar stereos and between the original image and the rectified one. The relationship between the disparity and the depth of a pixel is

$$Z = \frac{Bf}{d} \quad (1)$$

where B denotes the baseline between the rectified stereos and d , f , and Z are the disparity, focal length, and depth value, respectively. Similarly, the depth of a pixel in a rectified image along the ray is

$$D_b^r = \frac{B\sqrt{(x_b^r)^2 + (y_b^r)^2 + f^2}}{d} \quad (2)$$

where (x^r, y^r, f) is the camera coordinate of a point in the rectified base image. The rectified depth map can be converted into the original image coordinate according to the inverse perspective transformation matrixes. In Figure 4, the depths of the ray in stereo (l, b) and stereo (r, b) are D_1 and D_2 , respectively. Ideally, D_1 and D_2 are equal. According to the average accuracy σ (e.g., 0.3 m) of the depth map obtained by stereo methods, a threshold could be set to discriminate whether the depth values of the two stereos are consistent.

Figure 5 shows the process of the deep learning-based multi-view dense matching. Triple-view images are used for demonstration where I_b represents the base image and I_l and I_r are the left image and the right image, respectively. The two stereos (consisting of I_l, I_b and I_r, I_b , respectively) are first rectified via the corresponding perspective transformation matrixes H for obtaining two epipolar stereos. Then a CNN model is separately applied to the two rectified stereos to output two corresponding disparity maps, $d_{1,b}^r$ and $d_{2,b}^r$. The two disparity maps are converted to depth maps $D_{1,b}^r$ and $D_{2,b}^r$ according to Equation 2. Then the depth maps are transferred to the original image coordinates (Figure 4). Finally, the difference between the values of every overlapped pixel below a given threshold indicating multi-view compatibility and the mean value is used as the final depth; otherwise, the pixel is treated

as an abnormal disparity point, which is filtered out and filled in with its neighborhood values.

In this work, we utilize the multi-view constraint as a postprocessing, as in the conventional methods (Rothermel *et al.* 2012), instead of learning a depth map directly from multi-view stereos. This multi-view constraint could also be embedded into the learning loop (forward and backward propagation) in the end-to-end deep learning methods. However, as is shown in Figure 5, many pixel-wise coordinate transformations and massive computations are involved in the process, which may not be feasible with a single graphics processing unit (GPU).

Data Sets

To evaluate the performance of the deep learning-based stereo methods in aerial images, five data sets (KITTI, Driving, Hangzhou, München, and Vaihingen) are used in this experiment, among which the KITTI and Driving data sets are open source and consist of close-range data. The Hangzhou data set is collected from unmanned aerial vehicles (UAVs) and the München and Vaihingen data sets from traditional aerial photography platforms.

KITTI Data Set

The KITTI data set (Menze and Geiger 2015) consists of street-scene data in the German city of Karlsruhe captured from stereo cameras with a 54-cm baseline mounted on the roof of a car. The ground-truth depth is recorded by a rotating a laser scanner installed behind the left camera, ensuring that 30% of

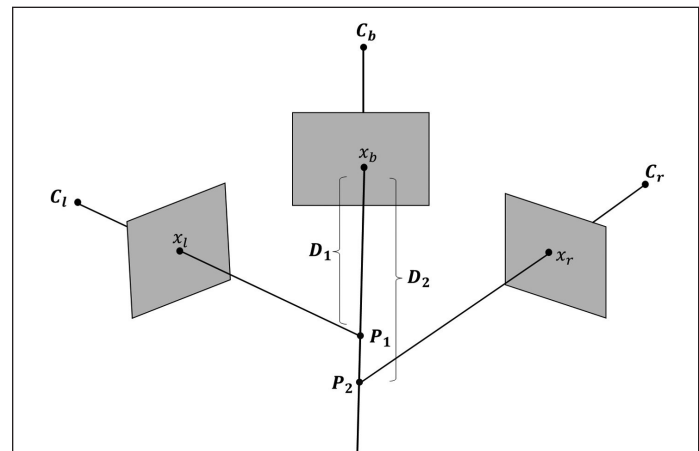


Figure 4. Depths of the ray in stereo (l, b) and stereo (r, b) are D_1 and D_2 in the uniform base image coordinate. The difference of D_1 and D_2 reflects the multi-view consistency.

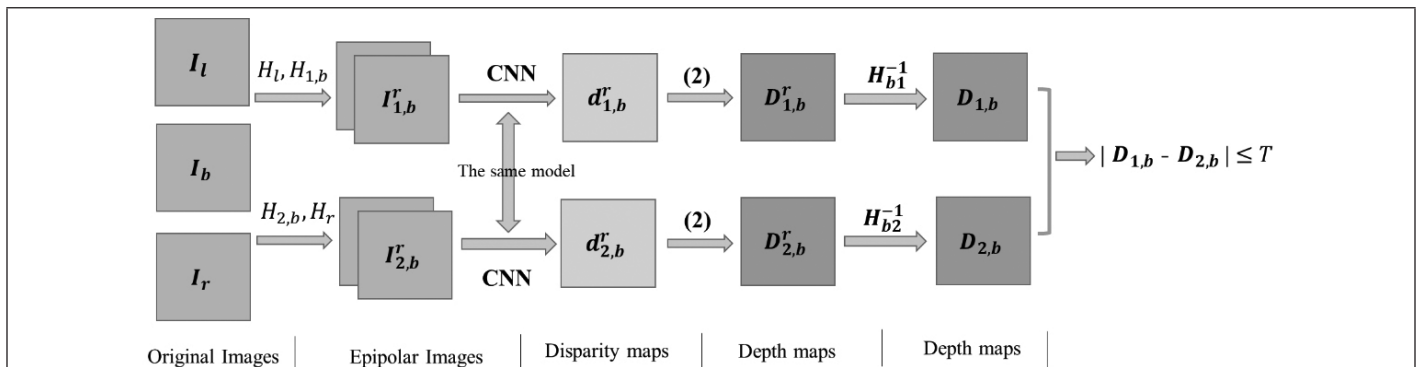


Figure 5. The process of the deep learning-based multi-view dense matching.

pixels have depth value. Both KITTI2012 and KITTI2015 data sets consist of about 200 rectified stereo pairs for training and 200 pairs with unpublic labels for testing with an average image size of 1240×375 pixels. In our experiment, eighty percent of the original training set was used for training and the remaining 20% for testing.

Driving Data Set

The Driving data set (Mayer *et al.* 2016) is a set of virtual street-view images with configured accurate stereo pairs, 3D scenes, and disparity maps. This data set consists of much more data than other existing real data sets and could facilitate the training of large CNNs. The rectified stereo image size is fixed at 960×540 . In the experiment, 300 rectified 960×540 stereo pairs were selected from the whole data set, from which 80% were used as training sets and the remaining 20% as test sets.

Hangzhou Data Set

The Hangzhou data set is a series of UAV images that were taken by a drone at a low altitude of about 640 m above-ground, recording the scene of a mountain village near the city of Hangzhou, China, in August 2017. Twenty aerial images consist of four strips with 80% heading overlap and 60% side overlap, with 0.07-m ground sampling distance (GSD). Lidar point cloud measurements are used as the ground truth with approximate 0.15-m height accuracy. The original 9000×6732 images are rectified into epipolar images. The disparity values of each pixel are calculated from the point cloud. Due to the limitation of the GPU capacity, the rectified aerial images are cropped into subimages of size 1325×354 . After manually removing some image pairs with large visual bias between LiDAR point and images, the remaining 328 pairs of images were used as training sets and 40 pairs as test sets.

München and Vaihingen Data Sets

The München aerial data set consists of 15 $14,114 \times 15,552$ aerial images that were captured in three strips with 80% heading overlap and 80% side overlap, covering urban buildings, roads, and greenbelts. The Vaihingen data set consists of 36 $9,420 \times 14,430$ aerial images of three strips with 60% heading overlap and 60% side overlap, covering flat rural scenes. The München and Vaihingen data sets have a GSD of 0.14 m and 0.23 m, respectively. The ground truth of depth is provided in forms of semidense DSM with 0.1 m and 0.2 m GSD, respectively, which was generated and filtered using the median of the results calculated by seven commercial software and shows high visual accuracy. It is empirically inferred that the DSM height accuracy is between 0.2 m and 0.4 m. Similar to the preprocessing of the Hangzhou data set, the rectified epipolar images are cropped into subimages of 1150×435 and 955×360 , respectively. Finally, the München data set consists of 540 stereo pairs, and the Vaihingen data set consists of 740 pairs. The ratio of the training and test sets is 4:1.

Experiment and Result Analysis

Three experiments were designed to comprehensively evaluate the performance and generalization ability of the deep learning method in aerial remote sensing images. The first tests the performance of deep learning methods using the three aerial data sets Hangzhou, München, and Vaihingen. The results are compared with the SGM and SURE. The second is to test the generalization of stereo dense matching. The models pretrained on computer vision open data sets were applied to aerial imagery. The third one is to extend deep learning-based stereo matching to multi-view matching and evaluate its performance.

Three-pixel-error (3PE, the percentage of pixels with a disparity error less than three pixels) and one-pixel-error (1PE, the percentage of pixels with a disparity error less than one pixel) are used for accuracy assessment. In multi-view matching, the depth map is compared to the ground truth in meters. All deep learning methods are implemented with an NVIDIA Titan Xp 12G GPU.

Comparison of the Deep Learning-Based Methods with Traditional Methods

We evaluated the performance of the MC-CNN, GC-Net, and DispNet on dense matching and compared this with the SGM and SURE. The basic settings of these methods/software are as follows.

For the MC-CNN, at the training stage, the input of model is 128 pairs of positive and negative samples composed of 9×9 image blocks. Small batch gradient descent is adopted to minimize the loss, and the momentum is set to 0.9. The learning rate is 0.002 and was adjusted to 0.0002 after several iterations.

For the GC-Net, which is less effective on the sparse disparity map, the network is trained only on three dense data sets (the data sets that were not processed are represented as “—” in Table 1). The batch size is set to 1, all data are iterated 50 times, and the learning rate is 0.001.

For the DispNet, the batch size is set to 32, and the learning rate is 0.0001 and was gradually decreased during the training process. All training data were iterated 1500 times.

For the SGM, we used the implemented function in Opencv3.0 library, with postprocessing, such as Gaussian smoothing and median filtering.

For the SURE, the inputs are the original aerial images and orientation information, and the output is a 3D model in the format of OSGB. Therefore, experiments were conducted on three sets of aerial images exclusively. We used the 3D workflow mode with the default parameters and settings. To evaluate the accuracy, the corresponding disparities of each point on epipolar images are calculated from the 3D model and compared with the real disparity values.

The performance of the traditional methods and deep learning methods on five data sets is shown in Table 1.

In general, the GC-Net performs the best; the MC-CNN performs comparable to the commercial software SURE in 3PE and slightly lower in the 1PE indicator and far superior to SGM; DispNet performs the worst in 1PE.

Using the GC-Net model, the accuracy on the flat Vaihingen data set is 99.7% (98.0%) and is slightly better than all the other well-performed methods. The München data set has obvious variations in elevation and has more discriminatory power for the comparison of the methods. The 3PE of the GC-Net is about 2% higher than the MC-CNN model (second best), and 1PE is 7.4% higher than the SURE (second best). On the Driving data set, 92.6% (85.7%) of the test accuracy is also much higher than other methods.

The performance of MC-CNN model on all the data sets is stable. The accuracy of each data set is much higher than the SGM: 6.7% (4.6%) higher in the KITTI2015 and 5.7% (7.7%)

Table 1. Accuracy comparison between the traditional and deep learning stereo methods on the five data sets.

Method	Accuracy (3PE/1PE)				
	KITTI2015	Driving	Hangzhou	München	Vaihingen
MC-CNN	0.960/0.778	—	0.953/0.816	0.965/0.867	0.992/0.932
GC-Net	—	0.926/0.857	—	0.984/0.953	0.997/0.980
DispNet	0.937/0.737	0.835/0.547	0.923/0.591	0.883/0.532	0.950/0.710
SGM	0.893/0.732	0.713/0.505	0.896/0.739	0.921/0.859	0.987/0.925
SURE	—	—	0.968/0.831	0.932/0.879	0.990/0.969

higher in the Hangzhou data sets. On the three aerial image data sets Hangzhou, München, and Vaihingen, the MC-CNN is equivalent to the multi-view-based SURE in 3PE and slightly lower (1.5%, 1.2%, and 3.7%, respectively) in 1PE.

The DispNet with a generic FCN structure obtained the worst accuracy on the remote sensing image data sets. The poor results on 1PE especially reflect the limitation of a generic model on dense matching tasks. DispNet is suitable only

for stereos with very small disparity values. For example, DispNet preforms better than SGM on the KITTI2015 data set, which has a maximum disparity of merely 70 pixels, while on remote sensing data sets with large terrain fluctuations, the results become inaccurate and unstable.

Figure 6 shows the predicted disparity maps of all of the methods on the three aerial image data sets. From top to bottom are the stereo image pairs, ground truth, and the

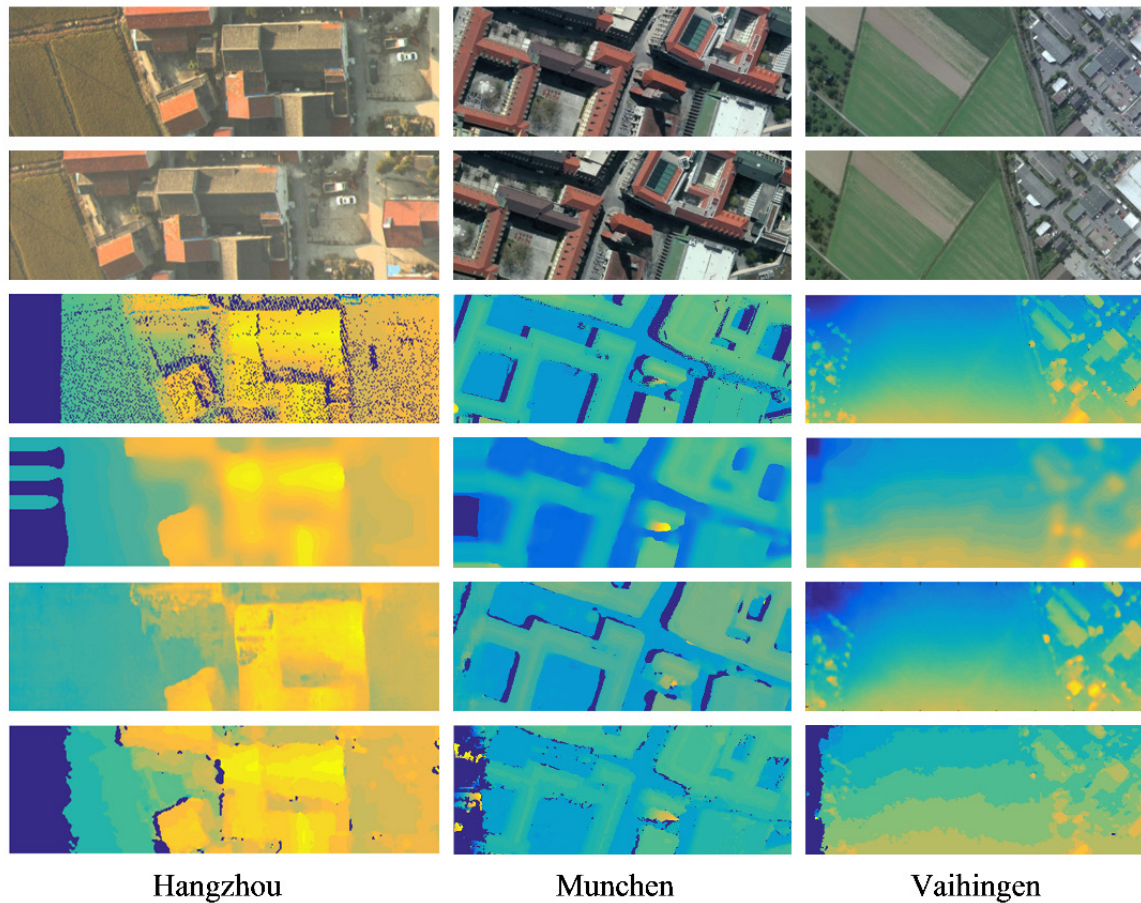


Figure 6. From top to bottom: stereo image pairs, ground truth, and the prediction results of the MC-CNN, GC-Net and SGM methods. For predicting the Hangzhou image with only sparse depth points, the GC-Net model, which requires dense depth points as training samples, is trained on the München data set.

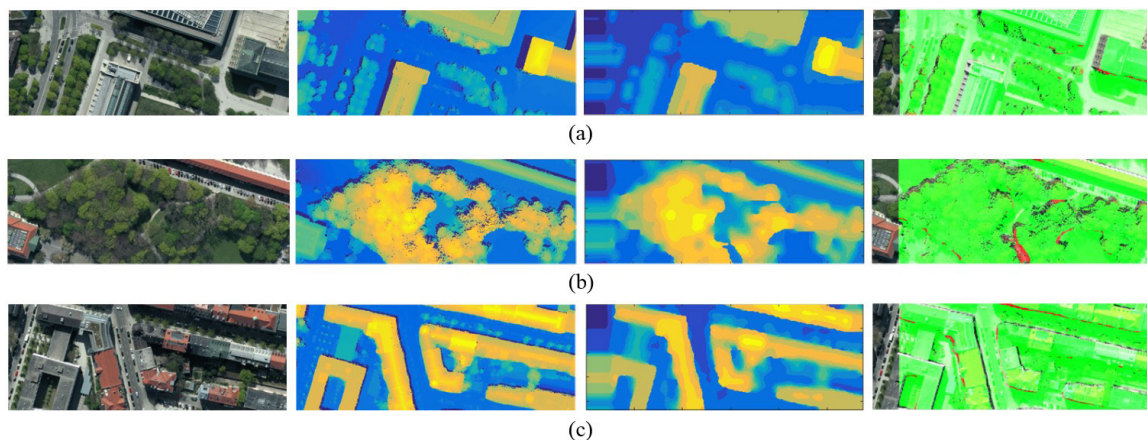


Figure 7. The MC-CNN results on different land cover categories: flat ground (a), trees (b), and buildings (c). From left to right: the left image, ground truth, and the predicted disparity map and residual map. In the residual maps, green indicates disparity error below three pixels, and red indicates disparity error above three pixels.

prediction results of the MC-CNN, GC-Net, and SGM methods. It can be observed that the result of the GC-Net is the closest to the reference map.

Results of applying the MC-CNN method to the München data set are used to analyze the influence of land cover categories. Three image patches, including ground, buildings, and trees, respectively, were selected. The differences between the predicted results and ground truth are shown in Figure 7. Generally, the MC-CNN performs well at all categories and shows no obvious difference. Mismatching appears at the boundary of occluded areas. An obvious mismatching occurred in Figure 7c. It can be observed that the path is partially obscured by tree branches in both or either of the stereo images.

Figure 8 shows the 3D scenes recovered from the dense disparity maps. The result of SURE shows some distortions in the Hangzhou data set. Compared to the other methods, it processes the whole aerial images and is more vulnerable to the accuracy of interior and exterior orientation elements. On the München data set, the results of all the methods are close to the referenced 3D scene and accurate. However, the texture of the buildings' side is much clearer using the SURE, as it utilizes multi-view images, especially those from a side strip. All methods perform well on the Vaihingen data set, which consists of flat landscapes.

Transfer Learning of Deep Learning-Based Stereo Methods

We evaluate direct transfer learning of the MC-CNN and GC-Net (the poorly performed DispNet was omitted). Table 2 shows the test results of the MC-CNN on 3PE and 1PE indicators. The source data set is used for model training and the target data

set for prediction. For example, when using the Hangzhou data itself for training, the 3PE accuracy is 95.3% (bold diagonal elements) on the Hangzhou data set; when using the KITTI2012 for training, the accuracy is 94.4%. The difference between them reflects the degradation of each model using transfer learning (see the last row).

In general, the MC-CNN method has good generalization ability whether the model is trained on close-range images, simulated scenes, or aerial images. The degradation degree varies between 0.2% and 2.2% on 3PE with an average of 0.6% on the three aerial data sets, 0.8% and 5.6% on 1PE, and an average of 2.1% on the aerial data sets. When comparing the 3PE, using a pretrained MC-CNN model is still superior to SGM and equivalent to the SURE software.

Better generalization ability is achieved when the source and target sets are similar. For example, the KITTI2012 model obtained a test accuracy of 95.8% with the KITTI2015, decreasing by only 0.2%. The test accuracy of using the Hangzhou data decreased by 0.5% using the pretrained model on either the München or the Vaihingen data but decreased by 2.2% using the models pretrained on the street-view data. The Vaihingen model shows the worst generalization ability among the three aerial data sets, as the terrain of Vaihingen is relatively flat, and very little information about the dramatic parallax changes is learned.

Table 3 shows the results of direct transfer learning based on the GC-Net model. As only the Driving, München, and Vaihingen data sets contain the dense depth maps, the three data sets are used to train the models, which are then applied to the test data set for prediction.

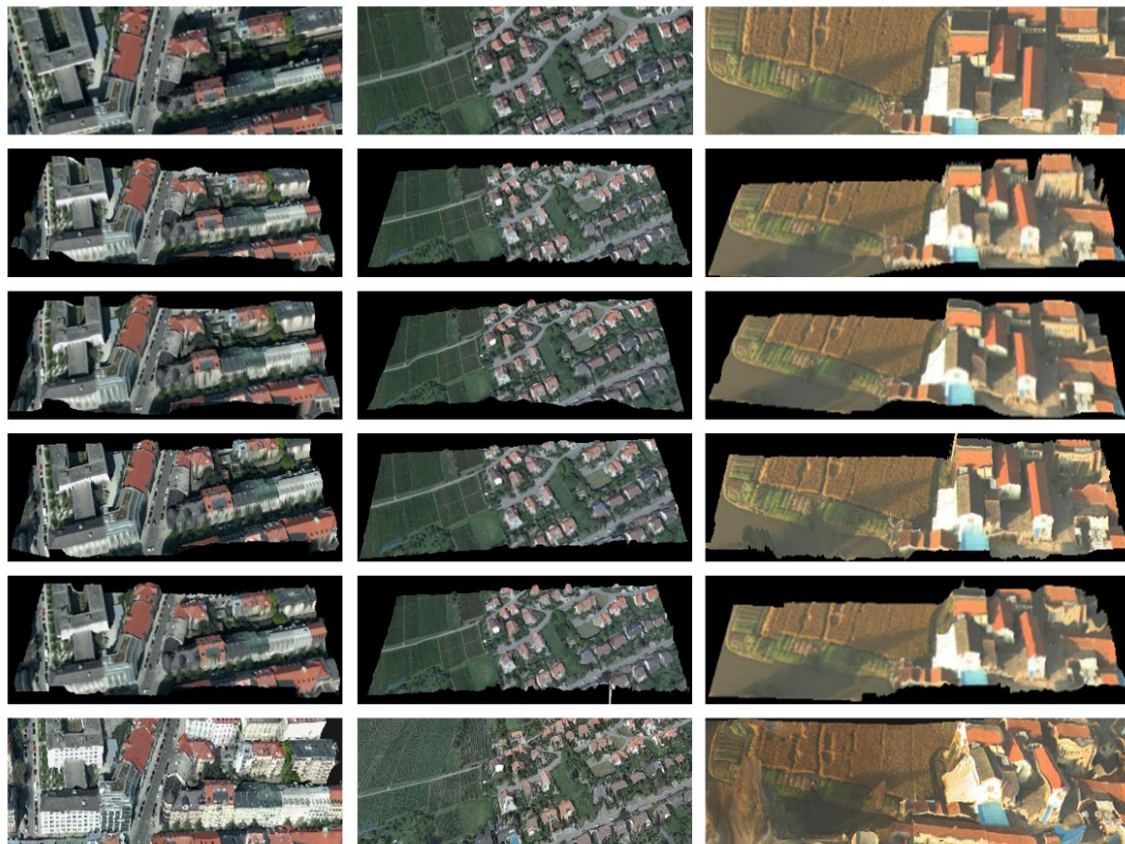


Figure 8. 3D scenes recovered from predicted disparity maps. From top to bottom: the left images, referenced 3D scenes, and the prediction results of the MC-CNN, GC-Net, SGM, and SURE.

Table 2. Pretraining the MC-CNN models on different source data sets and applying them on target data sets.

Target Set	Source Set, Accuracy (3PE/1PE)					
	KITTI2012	KITTI 2015	Hangzhou	München	Vaihingen	Average Degradation
KITTI2012	0.963/0.866	0.957/0.848	0.941/0.856	0.945/0.797	0.946/0.813	−0.016/−0.038
KITTI2015	0.958/0.768	0.960/0.778	0.951/0.761	0.955/0.751	0.953/0.750	−0.006/−0.021
Hangzhou	0.944/0.808	0.942/0.805	0.953/0.816	0.948/0.770	0.940/0.760	−0.010/−0.030
München	0.960/0.854	0.960/0.851	0.960/0.844	0.965/0.867	0.959/0.850	−0.005/−0.017
Vaihingen	0.988/0.919	0.987/0.912	0.987/0.916	0.989/0.922	0.992/0.932	−0.004/−0.015

Table 3. Pretraining the GC-Net models on different source data set and applying them on target data set.

Target Set	Source Set, Accuracy (3PE/1PE)			
	Driving	München	Vaihingen	Average Degradation
Driving	0.926/0.857	0.895/0.808	0.895/0.793	−0.031/−0.057
München	0.969/0.893	0.984/0.953	0.964/0.922	−0.018/−0.046
Vaihingen	0.980/0.881	0.979/0.943	0.997/0.980	−0.018/−0.068
KITTI2015	0.934/0.739	0.881/0.705	0.942/0.743	—/—
Hangzhou	0.911/0.779	0.940/0.799	0.949/0.841	—/—

Table 4. Comparison of stereo and triple-view image matching results on the München data set.

Method	Stereo/Triple-View				
	MC-CNN	GC-Net	DispNet	SURE	SGM
Percentage (< 1 m)	0.922/0.923	0.938/0.938	0.695/0.690	0.902	0.892/0.888
Avg-err (m)	0.590/0.539	0.397/0.392	1.043/1.039	0.635	0.759/0.710

The GC-Net has good generalization ability but slightly worse than the MC-CNN. Compared with the model trained on the same data set, the degradation degree of the model pretrained on other data sets varies between 1.5% and 3% on the 3PE and with an average of 1.8% on the München and Vaihingen data sets, 3.1% and 9.9% on the 1PE, and 5.7% on the two aerial data sets. When transferring between the aerial data sets, the test accuracy decreased by about 2%, while with the MC-CNN, the decrease was only 0.6%.

To sum up, both the pretrained MC-CNN and the GC-Net models can be directly applied to a different data set if the focus is the outlines of 3D scenes with 3PE as an indicator. To identify very fine structures and using the 1PE as an indicator, fine-tuning on small target samples may be required.

Triple-View Geometry for Deep Learning-Based Dense Matching

We evaluate whether using a multi-view geometry could further improve the accuracy of the deep learning-based stereo methods. Table 4 shows the comparison results between stereo and triple-view dense matching on the München data set with two indicators. One indicator is the percentage of the pixels with errors less than 1 m among all pixels (shown as P indicator); the other is the average absolute error of all pixels compared to the georeferenced depth map (shown as A indicator).

Both the stereo MC-CNN and the GC-Net methods have high performance and better than the other methods on the P indicator (Table 4). When triple-view geometry was introduced, the P indicator remains almost the same. It seems that errors larger than 1 m have been mostly eliminated by the postprocessing of the MC-CNN and the integration of the context and geometry information in the GC-Net, respectively. The remaining large errors may be due to the occlusions which cannot be compensated for by the triple-view geometry. Anyway, the parallel is between occlusions and label errors.

In addition, the stereo-based GC-Net performs much better than other methods on the A indicator and reaches 0.397 m. When the multi-geometry is introduced into the model, the GC-Net and MC-CNN improved slightly. The GC-Net outperforms the MC-CNN by 0.137 m, SURE by 0.243 m, and SGM by 0.318 m in depth accuracy, respectively. As the MC-CNN and the SGM

share almost the same structure except the matching cost, they are comparable, and the difference of accuracy (0.171 m) indicates that the matching cost of a simple CNN structure is better than an empirical one. Both show some improvement (0.041 m and 0.049 m, respectively) when the triple-view geometry is introduced. In either case, the DispNet functions the worst.

It could be concluded that when measured by depth instead of parallax, the GC-Net is the best method and that the MC-CNN is slightly better than the SURE. By introducing triple-view geometry, the MC-CNN achieved more significant improvement comparing to the GC-Net.

Figure 9 shows the depth maps of a triple-overlapping area of a whole aerial image generated by the three multi-view methods. It visually demonstrates that the result of the triple-view GC-Net (Figure 9d) is the most similar to the ground truth (Figure 9b). The result of the MC-CNN (Figure 9e) is slightly worse than that of the GC-Net; the result of the SURE shows some holes. From Figure 9c, the holes can be clearly observed. They are caused by the matching failure of the algorithm (or the matching may be not performed pixel-by-pixel) because both the MC-CNN and the GC-Net could find dense and smooth matches in the same area.

Discussion

In this study, we have thoroughly evaluated the performance of deep learning-based dense matching on stereo and multi-view aerial images. To illustrate efficiency, we take the München data set as an example. The overlapped parts of the 14 7072 × 7776 aerial images were almost seamlessly cropped into 300 768 × 384 stereo tiles. It took 5 hours to train the MC-CNN with 80% stereo tile samples; the prediction of a disparity map took 0.6 seconds (about 180 seconds for predicting the whole data set). It took 6.7 hours to train the GC-Net with 80% samples and 0.16 seconds to predict a single disparity map. It took the SURE 4.5 hours to generate the DSM of the aerial images. Thus, the efficiency of the traditional methods and the deep learning-based methods is of the same level. However, if the CNN model has been well pretrained before, the efficiency is much higher than that of a traditional method.

For transfer learning, besides the transfer learning strategy described in the section “Transfer Learning of Deep Learning-Based Stereo Methods”, the other transfer learning strategy conditioned on the target set containing a small number of samples could be considered. Fine-tuning a pretrained model (as initial parameters) with the available samples could improve the performance. Tables 5 and 6 show the results of fine-tuning based on the MC-CNN and GC-Net. DT represents training directly on a target set with random initialized parameters; TL represents model transfer and fine-tuning the parameters with the given samples. The performance of the MC-CNN is evaluated on the Hangzhou data set with a model pretrained on the KITTI2015 (Table 5); the performance of the GC-Net is evaluated on the München data set, with a pretrained model on the Vaihingen (Table 6).

Table 5 shows that an accuracy of 94.3% (79.9%) is obtained when the model is directly trained (DT) with 25 samples. When the sample number is doubled, the test accuracy is improved by only 0.09% (0.012%). This indicates that a large sample size is not necessary to train a satisfactory MC-CNN model. Thus, fine-tuning with a pretrained model (TL) seems to contribute trivially. The test accuracy is 94.9%

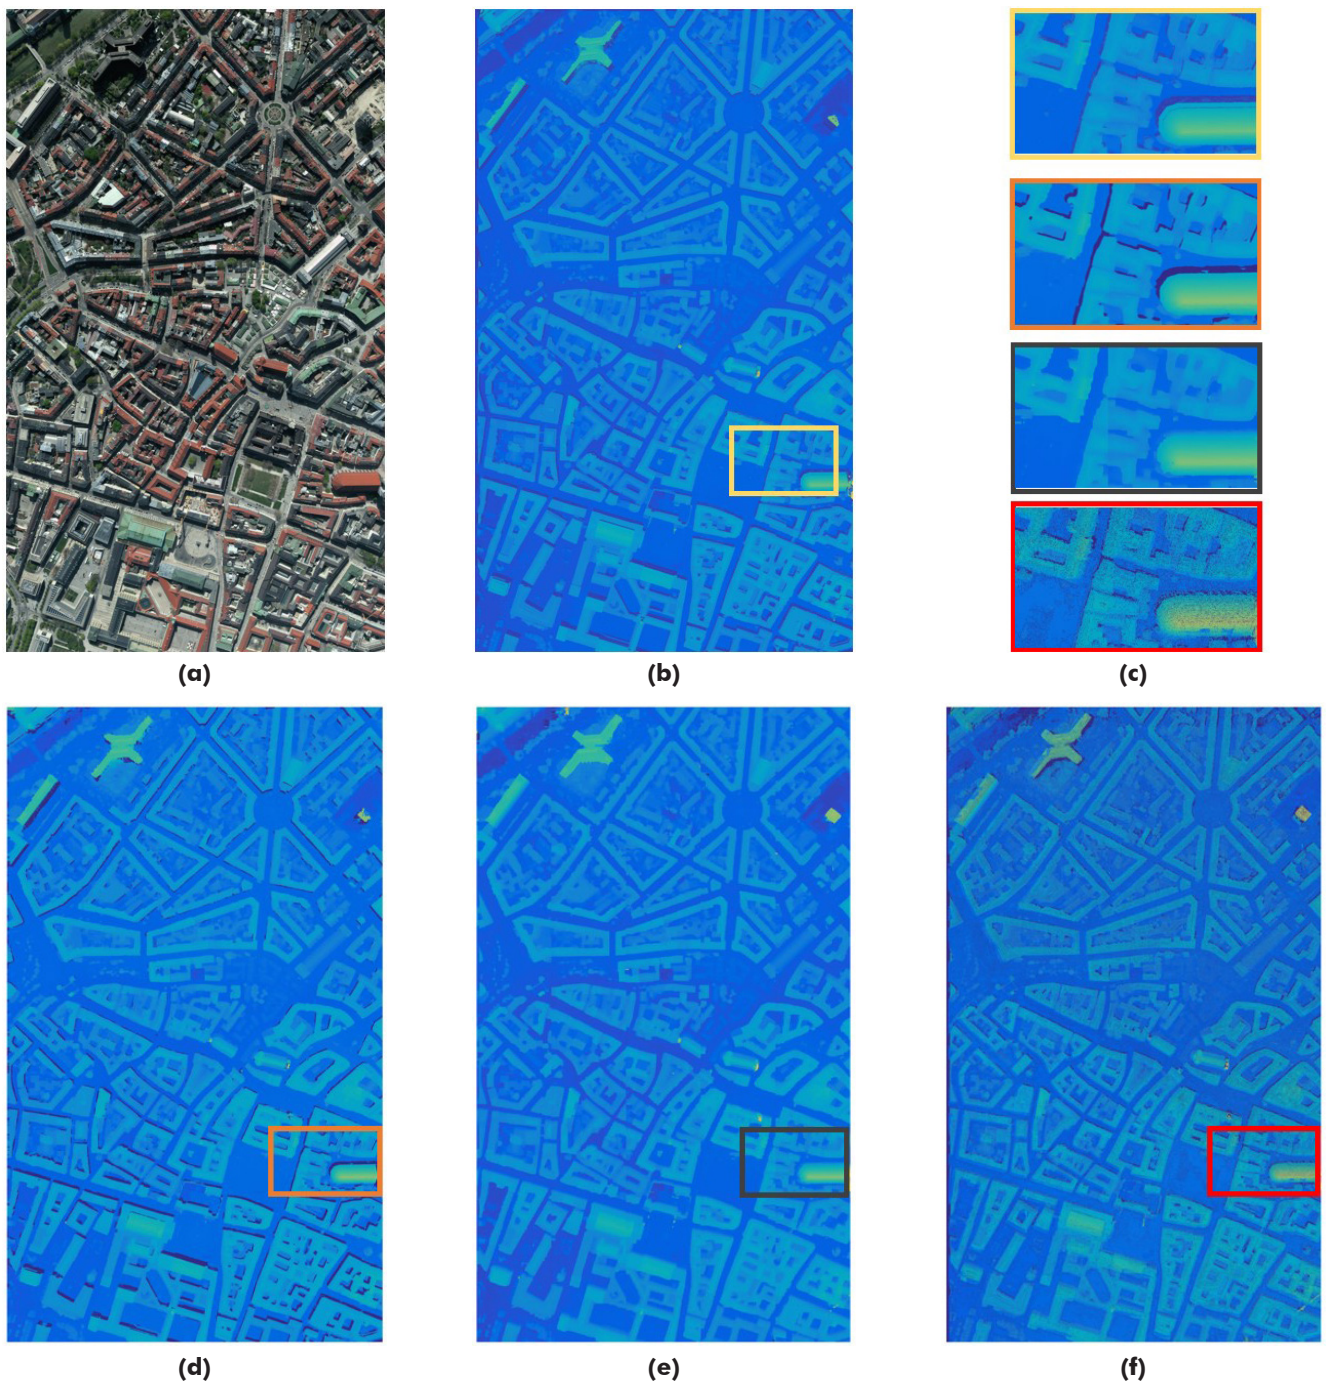


Figure 9. The depth maps of a München image using the multi-view-based methods. (a) a whole aerial image (the triple-overlapping area), (b) ground truth, (d) GC-Net results, (e) MC-CNN results, (f) SURE results and (c) the enlarged parts of the results of the three methods (from top to bottom: ground truth, GC-Net, MC-CNN, and SURE).

in 3PE using 25 training samples for fine-tuning, with a 0.5% improvement compared to the direct training.

In Table 6, when using direct training, the test accuracy is 78.3% (31.6%) with 25 samples; the accuracy reached 90.2% (60.3%) with 50 samples. This indicates that the end-to-end GC-Net requires more training samples than the MC-CNN. When using the pretrained model and fine-tuning strategy, 96.5% (91.6%) accuracy can be achieved with 25 training samples, 18.1% and 60.0% higher than direct training in 3PE and 1PE, respectively.

From the statistical results (Tables 5 and 6), fine-tuning helps improve test accuracy, especially for the end-to-end methods. The fewer the sample number, the greater the

fine-tuning effects. It is also found that fine-tuning can not only improve the accuracy but also reduce the iterations of training.

Conclusion

This study discusses the use of deep learning in the dense matching of aerial images and compares their performance with traditional methods on various data sets, analyzes the generalization ability of deep learning methods, and presents a deep learning-based multi-view dense matching framework. First, in both stereo and triple-view conditions, the end-to-end GC-Net outperforms all the other methods by a large margin. The

MC-CNN, which only learns matching cost, performs slightly better than the SURE. The SGM and the DispNet perform the worst. Second, both the MC-CNN and the GC-Net have shown satisfactory generalization ability, which ensures that a pre-trained model on open training data sets can be directly applied to target aerial images. However, if a high precision is required, a small set of training samples in the target data set could further improve the accuracy through fine-tuning. Finally, for deep learning-based stereo methods, multi-view geometry could further improve the accuracy of the predicted depth map. When available, the use of multi-view information and a pretrained model as initial parameters could significantly improve the performance of remote sensing dense matching.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (41471288).

References

- Aguilera, C.A., Aguilera, F.J., Sappa, A.D., Aguilera, C. and Toledo, R. Learning cross-spectral similarity measures with deep convolutional neural networks, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2016, pp. 1-9.
- Bleyer, M., Rhemann, C. and Rother, C. PatchMatch Stereo-Stereo Matching with Slanted Support Windows, *Proceedings of the British Machine Vision Conference*, 2011, pp. 1-11.
- Boykov, Y.Y. and Jolly, M.P. Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images, *Proceedings of the IEEE International Conference on Computer Vision*, 2001, pp. 105.
- Chang, J.-R. and Chen, Y.-S. Pyramid Stereo Matching Network, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5410-5418.
- Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D. and Brox, T. FlowNet: Learning optical flow with convolutional networks, *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 2758-2766.
- Furukawa, Y. and Ponce, J., 2010. Accurate, dense, and robust multiview stereopsis, *IEEE transactions on pattern analysis machine intelligence*, 32(8):1362-1376.
- Fusiello, A., Trucco, E. and Verri, A., 2000. A compact algorithm for rectification of stereo pairs, *Machine Vision and Applications*, 12(1):16-22.
- Han, X., Leung, T., Jia, Y., Sukthankar, R. and Berg, A.C. Matchnet: Unifying feature and metric learning for patch-based matching, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3279-3286.
- Hirschmuller, H., 2008. Stereo processing by semiglobal matching and mutual information, *IEEE Transactions on pattern analysis and machine intelligence*, 30(2):328-341.
- Kendall, A., Martirosyan, H. and Dasgupta, S., 2017. End-to-end learning of geometry and context for deep stereo regression, *CoRR*, 04309.
- LeCun, Y., Bengio, Y. and Hinton, G., 2015. Deep learning, *Nature*, 521(7553):436.
- Mayer, N., Ilg, E., Hausser, P., Fischer, P., Cremers, D., Dosovitskiy, A. and 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.
- Menze, M. and Geiger, A. Object scene flow for autonomous vehicles, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3061-3070.

Table 5. Fine-tuning results of the MC-CNN on the Hangzhou data set with pretrained model on the KITTI2015 (TL) compared to direct training with the available target samples (DT).

Methods	Samples									
	25 Pairs		50 Pairs		100 Pairs		200 Pairs		300 Pairs	
	DT	TL	DT	TL	DT	TL	DT	TL	DT	TL
Accuracy										
(3PE)	0.943	0.949	0.944	0.948	0.946	0.948	0.951	0.952	0.952	0.953
(1PE)	0.799	0.812	0.801	0.811	0.807	0.812	0.81	0.813	0.811	0.815
Improvement										
(3PE)	0.50%		0.37%		0.14%		0.12%		0.11%	
(1PE)	1.27%		0.97%		0.52%		0.31%		0.35%	

Table 6. Fine-tuning results of the GC-Net on the München data with pretrained model on the Vaihingen (TL) compared to direct training with the available target samples (DT).

Methods	Samples									
	25 Pairs		50 Pairs		100 Pairs		200 Pairs		250 Pairs	
	DT	TL	DT	TL	DT	TL	DT	TL	DT	TL
Accuracy										
(3PE)	0.783	0.965	0.902	0.947	0.928	0.961	0.959	0.977	0.972	0.978
(1PE)	0.316	0.916	0.603	0.899	0.881	0.931	0.904	0.942	0.925	0.944
Improvement										
(3PE)	18.10%		4.50%		3.20%		1.80%		0.60%	
(1PE)	60.00%		29.60%		5.00%		2.80%		1.90%	

- Pan, S.J. and Yang, Q., 2010. A survey on transfer learning, *IEEE Transactions on knowledge data engineering*, 22(10):1345-1359.
- Pang, J., Sun, W., Ren, J.S., Yang, C. and Yan, Q. Cascade Residual Learning: A Two-Stage Convolutional Neural Network for Stereo Matching, *Proceedings of the ICCV Workshops*, 2017.
- Ronneberger, O., Fischer, P. and Brox, T. U-net: Convolutional networks for biomedical image segmentation, *Proceedings of the International Conference on Medical image computing and computer-assisted intervention*, 2015, (Springer), pp. 234-241.
- Rothermel, M., Wenzel, K., Fritsch, D. and Haala, N. SURE: Photogrammetric surface reconstruction from imagery, *Proceedings of the LC3D Workshop*, Berlin, 2012, pp. 2.
- Scharstein, D. and Szeliski, R., 2002. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms, *International Journal of computer vision*, 47(1-3):7-42.
- Schmidhuber, J., 2015. Deep learning in neural networks: An overview, *Neural networks*, 61:85-117.
- Seki, A. and Pollefeys, M. SGM-Nets: Semi-global matching with neural networks, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Honolulu, HI, USA, 2017, pp. 21-26.
- Shaked, A. and 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.
- Zagoruyko, S. and Komodakis, N. Learning to compare image patches via convolutional neural networks, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 4353-4361.
- Zbontar, J. and 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.
- Zhang, K., Lu, J. and Lafruit, G., 2009. Cross-based local stereo matching using orthogonal integral images, *IEEE transactions on circuits and systems for video technology*, 19(7):1073-1079.
- Zhong, Y., Dai, Y. and Li, H., 2017. Self-supervised learning for stereo matching with selfimproving ability, *arXiv preprint arXiv:1709.00930*.arXiv:1709.00930.