

# SyntStereo2Real: Edge-Aware GAN for Remote Sensing Image-to-Image Translation while Maintaining Stereo Constraint

Vasudha Venkatesan<sup>1\*</sup>Daniel Panangian<sup>2</sup>Mario Fuentes Reyes<sup>2</sup>Ksenia Bittner<sup>2†</sup><sup>1</sup>University of Freiburg

venkatev@informatik.uni-freiburg.de

<sup>2</sup>German Aerospace Center (DLR)

{daniel.panangian,mario.fuentesreyes,ksenia.bittner}@dlr.de

## Abstract

In the field of remote sensing, the scarcity of stereo-matched and particularly lack of accurate ground truth data often hinders the training of deep neural networks. The use of synthetically generated images as an alternative, alleviates this problem but suffers from the problem of domain generalization. Unifying the capabilities of image-to-image translation and stereo-matching presents an effective solution to address the issue of domain generalization. Current methods involve combining two networks—an unpaired image-to-image translation network and a stereo-matching network—while jointly optimizing them. We propose an edge-aware GAN-based network that effectively tackles both tasks simultaneously. We obtain edge maps of input images from the Sobel operator and use it as an additional input to the encoder in the generator to enforce geometric consistency during translation. We additionally include a warping loss calculated from the translated images to maintain the stereo consistency. We demonstrate that our model produces qualitatively and quantitatively superior results than existing models, and its applicability extends to diverse domains, including autonomous driving.

## 1. Introduction

The challenges in obtaining ground truth images in the remote sensing domain stem from the difficulty in capturing matching images due to temporal changes, sparse measurements and a significantly large baseline. The correspondence tasks like disparity estimation or stereo reconstruction for these images, can be both cumbersome and expensive. The concept of using synthetic data for training deep

\*Work done during the completion of master's thesis at DLR

†Corresponding author



Figure 1. Examples of aerial scene translated by SyntStereo2Real. Our model can produce semantic-consistent realistic translations.

neural networks arises from the persistent challenges posed by data scarcity, privacy concerns, and the overall difficulty in acquiring authentic data. Synthetic data provides essential ground truth such as accurate labels and stereo disparity maps information for training machine learning models. While the synthetic data is obtained from a simulation of real-world scenario, it may not perfectly represent the complexities and variations in real-world data. This can result

in domain shift, where the model struggles to generalize to real world data. Unpaired image-to-image translation algorithms have been used to address the problem of domain shift [6, 14, 18]. They provide promising results to reduce the domain gap between the domains. However they can alter the structural information of the image as shown in Fig. 2. This can pose as a serious challenge when training for downstream tasks as the translated images do not align with their corresponding labels.

Our approach focuses on the specific task of translating synthetic images to realistic domain while maintaining the stereo constraints, which means that pixels do not move within or to another epipolar line and we particularly address the two-view image case. Some of the existing methods such as StereoGAN [15] have addressed this task for autonomous driving datasets with joint optimization of image translation and disparity estimation networks. Images from remote sensing domain are rich with diverse content. Existing methods suffer from the problem of an increased likelihood of hallucinations and discrepancies in disparity.

We address this problem using a *lightweight edge-aware GAN network*, that performs unpaired image-to-image translation while maintaining the disparity values. At first, the edge maps of input images are obtained from Sobel operator and are provided as an additional input along with image pairs from both domains to the generator. The encoder of the generator computes the content and edge code separately from the input image and its edge map and is added together as content edge code. The content edge code is provided to the decoder along with a random style code to generate images of different domain as shown in Fig. 3. The use of edge maps ensures that the structure of the image is retained and not lost in translation and thus prevents the matching algorithm fail due to blurred boundaries. Additionally, we use a warping loss, where we warp the left translated image with its respective disparity map and compare it to the right translated image to enforce stereo consistency. Extensive experiments across multiple datasets demonstrate our method outperforms the existing methods quantitatively and qualitatively. Moreover, we use a single lightweight network to perform optimization on two tasks without the use of any pre-trained networks. To sum up, our main contributions are:

- Developing a lightweight framework for *image-to-image translation of stereo pairs* considering a consistent translation of left and right images that preserves the matching. By including edge maps and a warping loss, we improve the matching features of the generated pairs.
- Results show that the quality of the translated images leads to a better disparity prediction than other state-of-the-art translation methods.
- By testing on data of remote sensing and autonomous driving tasks, we demonstrate that our approach works

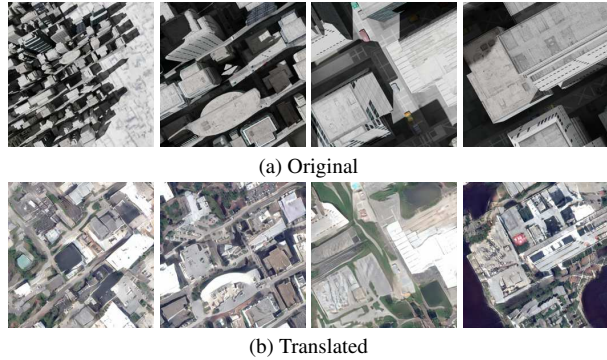


Figure 2. Aerial images translated using CUT [18]. The model tends to hallucinate when translating images with diverse scenes, where the target distribution is more likely to be unbalanced.

with a variety of datasets.

## 2. Related work

### 2.1. Stereo matching

Semi-global matching [3], a classical stereo method uses pixel-wise matching cost for computing the disparities between two images. It produces an approximate global optimal solution and is still one of the best performing techniques for disparity estimation in certain domains. MC-CNN [25] introduced deep learning based techniques to estimate disparity using a deep siamese network where matching cost for cost volume is computed from the network and cost aggregation is carried out through average pooling, followed by additional refinement. GANet [26] has optimized the cost aggregation by introducing semi-global aggregation (SGA - guided) layer which aggregates matching cost in multiple directions and local guided aggregation (LGA) to recover disparities at thin structures and object edges. The use of warping operation as a constraint for stereo matching has proven to be useful. It ensures that the estimated disparities produce geometrically consistent results. Nonetheless, all of these methods suffer from a significant domain gap when applied to data from a different acquisition nature.

Other disparity algorithms focus on tackling the domain gap while learning the matching. CFNet [21] proposes a fused cost volume representation in a cascade design for a more robust learning that can be applied to different domains. DSMNet [27] applies a domain normalization to the input images leading to sharper disparity maps. GrafNet [13] transforms the features with a U-shape network before inputting them to the cost volume, improving the domain adaptation. AdaStereo [22] transforms the color distribution into the target domain while training. In other cases, networks are designed to have a robust domain generalization on unseen data such as RAFT-Stereo [12], which is based on convolutional gated recurrent units (GRUs) or the

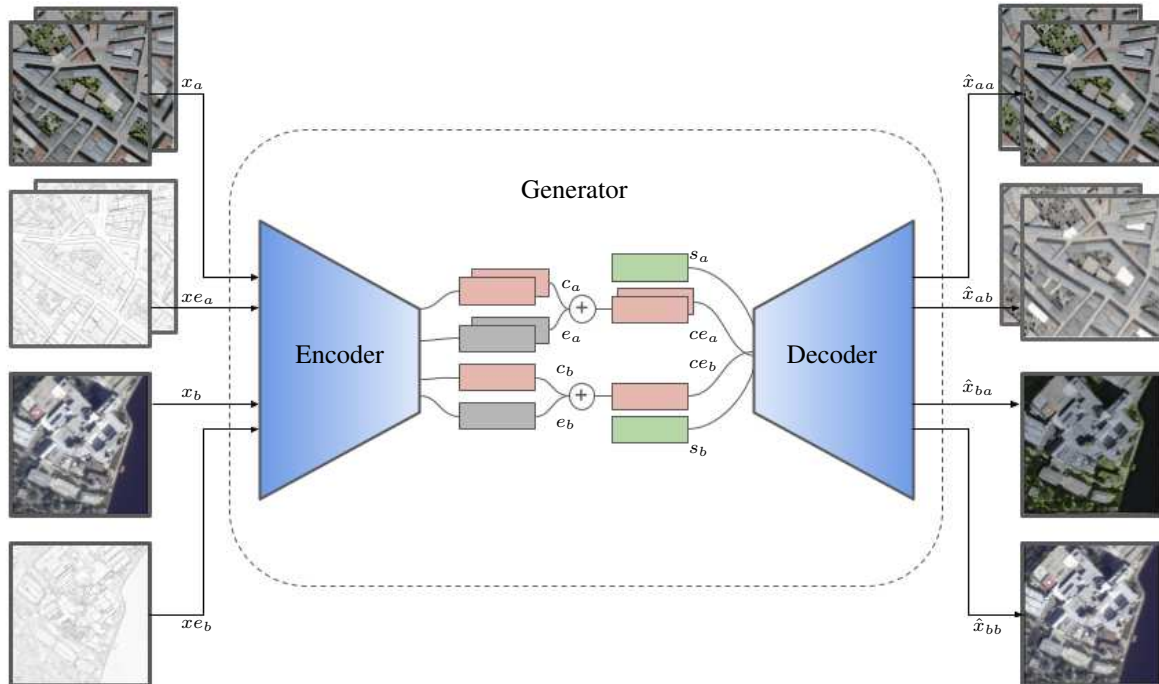


Figure 3. Illustration of the generator architecture in an autoencoder with edge map integration. The image along with its corresponding edge map is encoded and added together as content edge code before applying it as an input to the decoder. The decoder merges the content-edge code with style code from every domain to generate content that is contextually fitting.  $x_{c_a}, x_{c_b}$  represents the input images from both domains (content),  $x_{e_a}, x_{e_b}$  represents the corresponding edge maps.  $c_a, c_b, e_a, e_b$  represents the content and edge code from encoder for both domains.  $s_a, s_b$  are the randomly initialized style code before the training.  $x_{a_a}, x_{a_b}, x_{b_a}, x_{b_b}$  represents the respective output images from the decoder.




Figure 4. Pairs of translated images. For the translated left-view images 4a and 4c, the corresponding right-view images 4b and 4d are also displayed. As can be seen from the images, a semantic-consistent translation is applied to both the left and right-view images.

case of IGEV-Stereo [23], where geometry encoding improves the results in both stereo and multi-view networks. FC-stereo [28] computes two losses (selective whitening and contrastive feature) to preserve the stereo consistency between images and helped existing networks to generalize better while training only on synthetic data. Still, we consider that an offline adaptation of the dataset might lead to a good domain generalization without compromising the capabilities of the matching algorithms themselves.

## 2.2. Unpaired image-to-image translation

Unpaired image-to-image translation aims to learn the mapping from a source image domain to a target image domain without paired training data. CycleGAN [29] has been a pioneer in solving this task by identifying the key mappings in unpaired data from two different domains. The authors introduced cycle consistency loss to constrain the one-to-one mapping space by reconstructing the original image back from the translated image. This loss, in conjunction with adversarial loss and identity loss, plays

a pivotal role in image-to-image translation, leading to remarkable visual results. The CUT [18] model extends this concept for one-sided image translation with a contrastive loss. It is calculated using negative samples obtained from the same input, thus enabling faster training. UNIT [14] carries out unsupervised image-to-image translation under the assumption that images from both domains consist of a shared latent space. The model uses weight sharing between the layers of generators and discriminators to learn the joint distribution of data. MUNIT [6] extends this architecture to handle multiple styles using the disentanglement principle to obtain content and style code separately. The content code from the image is combined with a random style code from cross-domains to obtain diverse styled images. However, these GAN-based methods produce visually appealing image translations but sometimes fail to maintain semantic consistency between source and translated image especially given large number of features. In certain domains such as remote sensing as shown in Fig. 2, CUT [18] still suffers from the problem of hallucination.

### 2.3. Synthetic-to-real translation

The task of translating synthetic images to realistic while preserving semantic consistency has been an active research topic with multiple applications such as semantic segmentation, stereo matching and pseudo label learning. A straightforward approach relies on auxiliary information which is extracted from a task network to track changes in the source and target domains. CyCADA [4] leverages a method that preserves semantic consistency by constraining on a cycle consistent task-loss. It uses an additional loss which tracks the discrepancy between segmentation maps predicted by a pre-trained segmentation network from the generated images and the ground truth maps. Chen et al. [2] extend the method further by incorporating depth maps. Semantic-aware Grad-GAN [10] introduces a soft gradient-sensitive objective and a semantic aware discriminator for domain adaptation of virtual to real urban scenarios. To address alterations affecting object boundaries in generated images, their method involves applying the Sobel filter to both the image and its corresponding semantic map for deviation tracking. StereoGAN [15] is specifically designed for the task of translating synthetic images to realistic domain while maintaining the stereo constraints. It utilizes a CycleGAN for image translation and a DispNet [16] for disparity estimation. SDA [11] utilizes the spatial feature transform to fuse features of edge maps with source images. Different than the previous works, Secogan [8] utilized content disentanglement architecture from MUNIT for translating synthetic images of autonomous driving datasets to realistic domain. Instead of relying on a task network, it performs content disentanglement by employing fixed style codes in the generator, making the model computationally effective

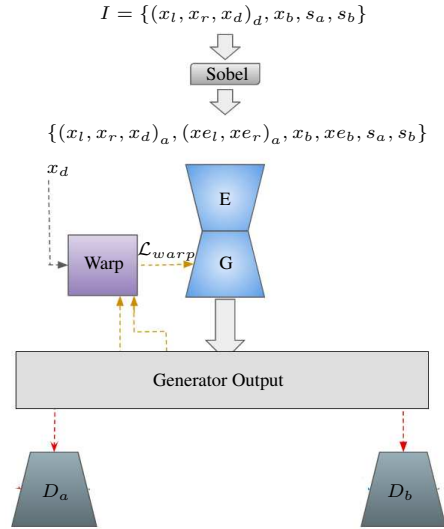


Figure 5. Illustration of the GAN-based model architecture featuring multiple loss functions. The design incorporates a combination of adversarial, reconstruction, cycle and warping losses. Adversarial loss promotes realistic image generation, while reconstruction loss ensures faithful reproduction of input data, cycle loss enforces the correct mapping between domains and warping loss enforces geometrical stereo constraints.


while preserving semantic consistency.

The task of translating images to a realistic style while maintaining the content structure for stereo matching is a dual optimization task. Although the existing networks address this issue, they suffer when applied to remote sensing images due to large baselines (which implies more occlusion), different acquisition times for left and right images, and city growing. The models developed are predominantly applied in the field of autonomous driving and struggle with achieving domain generalisation. Another challenge is the training of existing models tends to become computationally expensive, as it is a combination of two deep learning networks, one for image translation and the latter for stereo matching. The number of parameters required for training is high and can slow the training process. We address both of the above concerns in our work by employing a single edge-aware image translation GAN model trained additionally with warping loss to enforce the stereo constraints.

### 3. Method

We carry out the translation of synthetic to realistic domain images under the assumption that both domains share universal features that describe the elements in the scene (such as buildings, roads, vegetation), as well as distinctive features specific to the particular domain, focusing on visual attributes like appearance or style.

Given a synthetic *left-right-disparity* tuple



(a) Driving to KITTI2015

(b) SyntCities to US3D

Figure 6. Comparison of image translations: The first row showcases original synthetic images, the second row presents images translated using StereoGAN, and the third row exhibits images translated using our SyntStereo2Real.

$(x_l, x_r, x_d)_a \in \mathcal{X}_a$  denoting the stereo pair of left and right image with its corresponding disparity for source domain, a real image  $x_b \in \mathcal{X}_b$  representing the target domain, and two randomly sampled style codes  $s_a, s_b$  for each domain, our model synthesizes a realistic stereo matched pair of the synthetic image.

Our work draws inspiration from MUNIT [6] and Secogan [8] to learn disentangled representations from two domains without supervision. Similar to [6], our translation model consists of an autoencoder (encoder  $E$  and decoder  $G$ ) as a generator for both domains. The encoder factorizes each input into latent content code  $c_i (i = a, b)$ , where  $c_i = E(x_i)$ . Style code is initialized before the training using normal distribution as  $s_i = (\gamma_i, \beta_i)$  [8], for each domain and remains constant during the training. Here  $\gamma_i$  and  $\beta_i$  represents the mean and standard deviation of the normal distribution. Edge maps of the corresponding input images are obtained from the Sobel operator  $x_{e_i} = SO(x_i)$  and are given as additional input to preserve structural information. The encoder generates the latent edge code  $e_i = E(x_{e_i})$  from the edge maps. The edge code is added to the content code as content-edge code  $ce_i = c_i + e_i$  and is provided as an input to the decoder as shown in Fig. 3. The decoder generates the output image by swapping the content and style codes. The discriminator distinguishes the original image to the generated image by adversarial training. Since we have a real and synthetic domain, we have two discriminators  $D_A$  and  $D_B$ .

Multiple losses help in constraining and generating im-

ages in a meaningful manner in GAN based networks. Figure 5 shows an overview of the losses used in the training of the model. A reconstruction loss:

$$\mathcal{L}_{rec}^{aa}(E, G) = E_{x_a \sim \mathcal{X}_a} \|G(E(x_a), s_a) - x_a\|_1, \quad (1)$$

ensures that the model generates accurate reconstruction of images after content disentanglement.

In image-to-image translation, it is essential that the generated images in the target domain are not only realistic but also faithfully represent the original content. Cycle consistency loss [29]:

$$\mathcal{L}_{cycle}^{aba}(E, G) = E_{x_a \sim \mathcal{X}_a} \|G(E(x_{ab}), s_a) - x_a\|_1, \quad (2)$$

enforces this constraint by calculating the loss between original image and the transformation of original image to another domain ( $x_{ab}$ ), and transform it back again to original domain ( $x_{aba}$ ).

Since we use a GAN based approach to train the model, we use an adversarial loss:

$$\mathcal{L}_{adv}^a(E, G, D_a) = E_{x_a \sim \mathcal{X}_a} \log D_a(p(x_a)) + E_{x_b \sim \mathcal{X}_b} \log(1 - D_a(p(x_{ba}))), \quad (3)$$

which matches the data distribution of translated images to the distribution of target domain. The adversarial loss is employed by both the discriminator and generator, whereas the other mentioned loss exclusively guides the training of the generator. Since we use a patch based discriminator, the  $p$  in Eq. (3) refers to random patches of image.

Considering the images from one domain are synthetically generated, we assume to have access to additional information like ground truth labels, disparity maps, and segmentation masks. Warping loss as an additional constraint can be a useful addition, especially in tasks where the images are later used for training disparity estimation models. We compute the warping loss as

$$\begin{aligned} \mathcal{L}_{\text{warp}} = & \lambda_1 \cdot \mathcal{L}_1(G(E(x_{r_a}), s_b)) \\ & - W(G(E(x_{l_a}), s_b), x_d) \\ & + \lambda_2 \cdot (1 - \text{SSIM}(G(E(x_{r_a}), s_b) \\ & - W(G(E(x_{l_a}), s_b), x_d))), \end{aligned} \quad (4)$$

which compares the warped left image  $W(x_{l_a}, x_d)$ , which has undergone translation, and the right image after translation  $x_{r_{ab}}$ . We use a combination of  $\mathcal{L}_1$  loss and SSIM loss for calculating the warping loss.

The corresponding losses from other domain  $\mathcal{L}_{rec}^{bb}$ ,  $\mathcal{L}_{cycle}^{bab}$  and  $\mathcal{L}_{adv}^b$  are calculated in a similar manner. Therefore, the overall loss function for the generator is given by

$$\begin{aligned} \min_{E,G} \max_{D_a, D_b} \mathcal{L}(E, G, D_a, D_b) = & \lambda_3 \cdot (\mathcal{L}_{rec}^{aa} + \mathcal{L}_{rec}^{bb}) + \\ & \lambda_4 \cdot (\mathcal{L}_{cyc}^{aba} + \mathcal{L}_{cyc}^{bab}) + \\ & \lambda_5 \cdot (\mathcal{L}_{adv}^a + \mathcal{L}_{adv}^b) + \\ & \mathcal{L}_{warp}. \end{aligned} \quad (5)$$

## 4. Experiments


### 4.1. Network Architecture and Training

The autoencoder with a pair of encoder and decoder for generator is based on MUNIT architecture [6]. The discriminators are implemented using PatchGAN [7] architecture. The input to the network consists of images from both the domains and their corresponding edge maps and the output consists of translated images of both domains with the style and content interchanged. We recommend the use of Sobel operator to obtain edge maps due to its simplicity and effectiveness, which was compared in [11] with other existing edge detectors. The Sobel operator employs two  $3 \times 3$  convolution masks: one for estimating the gradient in the x-direction and the other for the y-direction. This design aligns effectively with CUDA architecture, allowing individual threads to apply the  $3 \times 3$  convolution masks to their assigned pixel and its neighboring pixels in the image.


The network model is implemented using PyTorch [19] and the training is carried out for 100 epochs with a batch size of 4. The hyperparameter values for  $\lambda_1$ ,  $\lambda_2$ ,  $\lambda_3$ ,  $\lambda_4$  and  $\lambda_5$  in Eq. (4) and Eq. (5) are set to 1, 1, 0.8, 10 and 10 respectively. We use stochastic mini batch gradient descent with Adam optimizer [9]. Beta coefficients of Adam are set to 0.5 and 0.999 respectively.




(a) Reference




(b) Ground Truth



(c) Baseline



(d) StereoGAN



(e) SyntStereo2Real(ours)

Figure 7. Results of disparity estimation from the AANet for the KITTI 2015 dataset. Three models are computed for the image shown in (a) RGB reference image, (b) Ground truth, (c) Model trained on Driving (baseline), (d) Model trained on Driving translated using StereoGAN (e) Model trained on Driving translated using SyntStereo2Real (ours).

### 4.2. Datasets

We use two sets of datasets from different application areas to study the generalization capabilities of *SyntStereo2Real* architecture. For remote sensing data, we use SyntCities dataset [20] for synthetic data and Urban semantic 3D dataset [1] for real domain data. SyntCities is a large dataset set consisting of synthetically generated images of aerial imagery. It is specially developed to train deep learning networks for disparity estimation, providing accurate disparity

Datasets	Metrics	Baseline	StereoGAN	SyntStereo2Real(ours)
<b>SyntCities to US3D</b>	MAD ↓	1.801	1.520	<b>1.319</b>
	3px-acc% ↑	63.097	66.765	<b>69.906</b>
	1px-acc% ↑	30.790	33.619	<b>35.928</b>
<b>Driving to KITTI</b>	MAD ↓	0.721	0.626	<b>0.575</b>
	3px-acc% ↑	88.871	89.646	<b>91.373</b>
	1px-acc% ↑	61.832	64.271	<b>65.892</b>

Table 1. Comparison of metrics for SyntCities to US3D and Driving to KITTI. The table illustrates the performance across datasets, showcasing results for the SyntCities (baseline), StereoGAN, and SyntStereo2Real (ours). Bold values highlight superior performance in MAD reduction and accuracy enhancement.

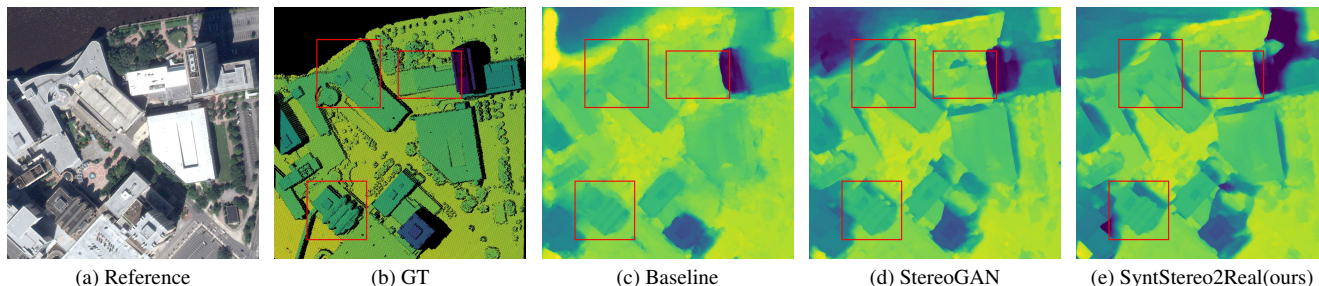


Figure 8. Results of disparity estimation from the AANet for the US3D dataset. Three models are computed for the image shown in (a) RGB reference image, (b) Ground truth, (c) Model trained on SyntCities (Baseline), (d) Model trained on SyntCities translated using StereoGAN (e) Model trained on SyntCities translated using SyntStereo2Real (ours).

Model	$n_{params}$
StereoGAN	54M
SyntStereo2Real(ours)	11M

Table 2. Comparison of the number of learnable parameters to train model between StereoGAN and SyntStereo2real models.

ground truth and different baselines. It consists of 8800 pairs of images resembling architectures of three cities: *New York, Paris and Venice*, with size of each image being  $1024 \times 1024$ . We use 1000 tuples of images taken evenly from all the three cities for training. The Urban3D (US3D) dataset consists of satellite images taken from WorldView3 mission and ground truth disparities are derived from aerial LiDAR data. In this dataset, a significant portion of the images primarily consists of vegetation with limited urban content. To address this, we filtered images based on label data, retaining only those images that contain a minimum of 15% building-related content. We randomly selected 1000 samples each of size  $1024 \times 1024$  for training.

For autonomous driving data, we use the Driving dataset from SceneFlow [16] for synthetic domain and KITTI2015 [17] dataset in real domain. We use the complete dataset from Driving consisting of 4400 images of size  $540 \times 960$  and the 160 training images each of size

$385 \times 1242$  provided by KITTI2015 benchmark. We re-size the images to  $512 \times 512$  for remote sensing dataset and  $256 \times 512$  for autonomous driving dataset during training due to memory and time constraints.

### 4.3. Evaluation metrics

We compare the two models based on performance of stereo matching and number of learnable parameters (compactness of the model architecture) required to train the model.

We acquire translated images and assess their performance on disparity estimation by training them on a disparity network. Specifically, we employ AANet [24] for the training and evaluation of estimation. In case of SyntCities to US3D we trained the model for 400 epochs and Driving to KITTI 2015 for 120 epochs, as this is a larger dataset. In both the cases we used a batch size of 20 and the maximum disparity was set to 192. To evaluate the predicted disparity maps, we removed the areas where the ground truth is not defined. 60 samples from US3D were used for testing and 40 for KITTI 2015 (no overlapping with the training samples). The cases where the original data (before translation) is taken as input is named as baseline.

Given the scarcity of models specializing in synthetic-to-real domain adaptation with stereo constraints, we conduct a comparative analysis of our model against StereoGAN. The results are given in Tab. 1. We use MAD (Median Ab-

Metrics	No Edge and No Disp	With Edge	With Disp	With Edge and Disp
<b>MAD</b> ↓	1.779	1.755	1.646	<b>1.319</b>
<b>3px-acc%</b> ↑	62.670	62.887	63.847	<b>69.906</b>
<b>1px-acc%</b> ↑	31.503	31.853	32.600	<b>35.929</b>

Table 3. Ablation studies. Here the *Edge* refers to the addition of edge information along with input image and *Disp* refers to the additional use of warping loss to enforce disparity constraints.

solute Deviation) [5], 3px accuracy percentage and 1px accuracy percentage for evaluation of stereo matching. MAD is a robust statistic, being resilient to outliers in a dataset compared to standard deviation because it is calculated by obtaining the median of the absolute difference of pixels and not the squared mean as in standard deviation. 3px accuracy represents the percentage of pixels in the disparity map for which the estimated disparity is within a range of  $\pm 3$  pixels from the ground truth disparity and 1px refers to the same metric but for a 1 pixel range.

#### 4.4. Quantitative Results

As indicated in Tab. 1, our approach demonstrates enhancements, showcasing a notable improvement with respect to StereoGAN of +3.14% in 3px accuracy and +2.30% in 1px accuracy for remote sensing images. Besides, the model exhibits improvements of +1.727% in 3px accuracy and +1.621% in 1px accuracy for autonomous driving datasets. Please note that the ground truth in the KITTI dataset is sparse and can not be evaluated for all the pixels. Despite that, we can visually compare the reconstruction capabilities for not labelled pixels. The disparity maps illustrated in Fig. 7 and Fig. 8 highlights a more complete prediction without empty regions. Comparing the number of parameters in Tab. 2, our model has a significantly smaller number of learnable parameters for training, making it ideal for applications with limited storage and processing capabilities.

#### 4.5. Qualitative Results

Figure 6 displays the results of translation of StereoGAN and our network SyntStereo2Real. The main challenge in translating in remote sensing images are maintaining the structural information for all resolution of images. StereoGAN, while proficient in certain aspects of disparity estimation, fails in the translation of shadows by hallucinating green patches instead of building shadows. Our model effectively captures and reproduces the content such as architectural details of building rooftops, bridges and roads, improving shadow handling, and preserving epipolar geometry simultaneously. Our method shows consistent prediction of disparity maps for complete objects without empty gaps or unclear boundaries. We also demonstrate our approach performs well in different application domains beyond au-

tonomous driving.

#### 4.6. Ablation Studies

In the Tab. 3, various configurations of the model are evaluated based on the presence or absence of edge information and warping loss for disparity. Firstly, the inclusion of edge information results in a decrease in the Mean Absolute Deviation (MAD), indicating improved results in predicting deviations from the ground truth. This decrease, coupled with a corresponding increase in both 3px accuracy and 1px accuracy indicates the importance of addition of edge maps. Similarly, addition of warping loss helps in improving the accuracy and MAD of the model significantly. We also experimented incorporating edges using Spatial feature transform (SFT), where the model learns the weight of edges along with image. This network produced images with strong focus on edges and lost other details making the data unsuitable for training. Thus the the ablation study demonstrates that incorporating both edge information and disparity significantly improves the model’s performance across all evaluated metrics for the used datasets.

#### 5. Conclusion

In this paper, we propose a light-weight edge based GAN model designed for unpaired image-to-image translation of synthetic-to-real data, while adhering to stereo constraints. Our approach leverages the importance of edge maps along with input images to retain the structural information while translation. Additionally, we incorporate a warping loss to maintain the accuracy of disparities on translated images. The integration of these two crucial elements yields state-of-the-art results in a single synthetic to real image translation network.

#### 6. Acknowledgement

The authors would like to thank the Johns Hopkins University Applied Physics Laboratory and IARPA for providing the data used in this study, and the IEEE GRSS Image Analysis and Data Fusion Technical Committee for organizing the Data Fusion Contest.



## References

- [1] Marc Bosch, Kevin Foster, Gordon Christie, Sean Wang, Gregory D Hager, and Myron Brown. Semantic stereo for incidental satellite images. In *2019 IEEE Winter Conference on Applications of Computer Vision*, pages 1524–1532, 2019. 6
- [2] Yuhua Chen, Wen Li, Xiaoran Chen, and Luc Van Gool. Learning semantic segmentation from synthetic data: A geometrically guided input-output adaptation approach. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1841–1850, 2019. 4
- [3] Heiko Hirschmuller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):328–341, 2008. 2
- [4] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. Pmlr, 2018. 4
- [5] Joachim Höhle and Michael Höhle. Accuracy assessment of digital elevation models by means of robust statistical methods. *ISPRS Journal of Photogrammetry and Remote Sensing*, 64(4):398–406, 2009. 8
- [6] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018. 2, 4, 5, 6
- [7] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6
- [8] Mert Keser, Artem Savkin, and Federico Tombari. Content disentanglement for semantically consistent synthetic-to-real domain adaptation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3844–3849. IEEE, 2021. 4, 5
- [9] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 6
- [10] Peilun Li, Xiaodan Liang, Daoyuan Jia, and Eric P. Xing. Semantic-aware grad-gan for virtual-to-real urban scene adaptation. In *British Machine Vision Conference*, 2018. 4
- [11] Xing Li, Yangyu Fan, Zhibo Rao, Guoyun Lv, and Shiya Liu. Synthetic-to-real domain adaptation joint spatial feature transform for stereo matching. *IEEE Signal Processing Letters*, 29:60–64, 2022. 4, 6
- [12] Lahav Lipson, Zachary Teed, and Jia Deng. RAFT-Stereo: Multilevel recurrent field transforms for stereo matching. In *2021 International Conference on 3D Vision (3DV)*, pages 218–227. IEEE, 2021. 2
- [13] Biyang Liu, Huimin Yu, and Guodong Qi. Graftnet: Towards domain generalized stereo matching with a broad-spectrum and task-oriented feature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13012–13021, 2022. 2
- [14] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. *Advances in Neural Information Processing Systems*, 30, 2017. 2, 4
- [15] Rui Liu, Chengxi Yang, Wenxiu Sun, Xiaogang Wang, and Hongsheng Li. StereoGAN: Bridging synthetic-to-real domain gap by joint optimization of domain translation and stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12757–12766, 2020. 2, 4
- [16] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 4040–4048, 2016. 4, 7
- [17] Moritz Menze, Christian Heipke, and Andreas Geiger. Joint 3D estimation of vehicles and scene flow. In *ISPRS Workshop on Image Sequence Analysis (ISA)*, 2015. 7
- [18] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *Computer Vision – ECCV 2020*, pages 319–345. Cham, 2020. Springer International Publishing. 2, 4
- [19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 2019. 6
- [20] Mario Fuentes Reyes, Pablo D’Angelo, and Friedrich Fraundorfer. SyntCities: A large synthetic remote sensing dataset for disparity estimation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:10087–10098, 2022. 6
- [21] Zhelun Shen, Yuchao Dai, and Zhibo Rao. Cfnets: Cascade and fused cost volume for robust stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13906–13915, 2021. 2
- [22] Xiao Song, Guorun Yang, Xinge Zhu, Hui Zhou, Yuexin Ma, Zhe Wang, and Jianping Shi. Adastereo: An efficient domain-adaptive stereo matching approach. *International Journal of Computer Vision*, pages 1–20. 2
- [23] Gangwei Xu, Xianqi Wang, Xiaohuan Ding, and Xin Yang. Iterative geometry encoding volume for stereo matching. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [24] Haofei Xu and Juyong Zhang. AANet: Adaptive aggregation network for efficient stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1959–1968, 2020. 7
- [25] Jure Zbontar and Yann LeCun. Stereo matching by training a convolutional neural network to compare image patches. *Journal of Machine Learning Research*, 17:1–32, 2016. 2
- [26] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. GA-Net: Guided aggregation net for end-to-end stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 185–194, 2019. 2

- [27] Feihu Zhang, Xiaojuan Qi, Ruigang Yang, Victor Prisacariu, Benjamin Wah, and Philip Torr. Domain-invariant stereo matching networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 420–439. Springer, 2020. [2](#)
- [28] Jiawei Zhang, Xiang Wang, Xiao Bai, Chen Wang, Lei Huang, Yimin Chen, Lin Gu, Jun Zhou, Tatsuya Harada, and Edwin R Hancock. Revisiting domain generalized stereo matching networks from a feature consistency perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13001–13011, 2022. [3](#)
- [29] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2223–2232, 2017. [3](#), [5](#)