

Unpaired Thermal to Visible Spectrum Transfer using Adversarial Training

Adam Nyberg¹[0000-0001-8764-8499], Abdelrahman Eldesokey¹[0000-0003-3292-7153], David Bergström²[0000-0003-2414-4482], and David Gustafsson²[0000-0002-4370-2286]

¹ Computer Vision Laboratory Linköping University, 581 83 Linköping, Sweden

² Swedish Defence Research Agency (FOI), 583 30 Linköping, Sweden

Abstract. Thermal Infrared (TIR) cameras are gaining popularity in many computer vision applications due to their ability to operate under low-light conditions. Images produced by TIR cameras are usually difficult for humans to perceive visually, which limits their usability. Several methods in the literature were proposed to address this problem by transforming TIR images into realistic visible spectrum (VIS) images. However, existing TIR-VIS datasets suffer from imperfect alignment between TIR-VIS image pairs which degrades the performance of supervised methods. We tackle this problem by learning this transformation using an unsupervised Generative Adversarial Network (GAN) which trains on unpaired TIR and VIS images. When trained and evaluated on KAIST-MS dataset, our proposed methods was shown to produce significantly more realistic and sharp VIS images than the existing state-of-the-art supervised methods. In addition, our proposed method was shown to generalize very well when evaluated on a new dataset of new environments.

Keywords: Thermal Imaging · Generative Adversarial Networks · Unsupervised Learning · Colorization

1 Introduction

Recently, thermal infrared (TIR) cameras have become increasingly popular due to their long wavelength which allows them to work under low-light conditions. TIR cameras require no active illumination as they sense emitted heat from objects and map it to a visual heat map. This opens up for many applications such as object detection for driving in complete darkness and event detection in surveillance. In addition, the cost of TIR cameras have gone significantly down while their resolution have improved significantly, resulting in a boost of interest. However, one limitation of TIR cameras is their limited visual interpretability for humans which hinders some applications such as visual-aided driving.

To address this problem, TIR images can be transformed to visible spectrum (VIS) images which are easily interpreted by humans. Figure 1 shows an example of a TIR image, the corresponding VIS image and the VIS image generated




Fig. 1: An example of a TIR image (a), its corresponding VIS image (b) from the KAIST-MS dataset [4] and the VIS image (c) generated by our proposed method using only the TIR image (a) as an input.

directly from the TIR image. This is similar to colorization problems, where grayscale VIS images are mapped to color VIS images. However, transforming TIR images to VIS images is inherently challenging as they are not correlated in the electromagnetic spectrum. For instance, two objects of the same material and temperature, but with different colors in the VIS image, could correspond to the same value in the TIR image. Consequently, utilizing all the available information, i.e. spectrum, shape and context, is very crucial when solving this task. This also requires the availability of enormous amount of data to learn the latent relations between the two spectrums.

In colorization problems, only the chrominance needs to be estimated as the luminance is already available from the input grayscale images. Contrarily, TIR to VIS transformation requires the estimation of both the luminance and the chrominance based on the semantics of the input TIR images. Besides, generating data for learning colorization models is easy as color images could be computationally transformed to grayscale images to create image pairs with perfect pixel-to-pixel correspondences. In contrast, datasets containing registered TIR/VIS image pairs are very few and requires a sophisticated acquisition systems for good pixel-to-pixel correspondence.

The KAIST Multispectral Pedestrian Detection Benchmark (KAIST-MS) [4] introduced the first large-scale dataset with TIR-VIS image pairs. However, it was shown by [1] that the TIR-VIS image pairs in KAIST-MS does not have a perfect pixel-to-pixel correspondence, with a pixel error of up to 16 pixels (5%) in the horizontal direction. This would degrade the performance of supervised learning methods which tries to learn the pixel-to-pixel correspondences between image pairs and leads to corrupted output. To our knowledge there exist no large-scale public dataset of TIR-VIS image pairs with perfect pixel to pixel correspondence. Therefore, the method used for TIR to VIS image transformation need to control for this imperfection.

In this paper, we propose an unsupervised method for transforming TIR images, specifically long-wavelength infrared (LWIR), to visible spectrum (VIS) images. Our method is trained on unpaired images from the KAIST-MS dataset [4] which allows it to handle the imperfect registration between the TIR-VIS image pairs. Qualitative analysis shows that our proposed unsupervised method produces sharp and perceptually realistic VIS images compared to the existing state-of-the-art supervised methods. In addition, our proposed method achieves comparable results to state-of-the-art supervised method in terms of L1 error despite being trained on unpaired images. Finally, our proposed method generalizes very well when evaluated on our new FOI dataset, which demonstrates the generalization capabilities of our method contrarily to the existing state-of-the-art methods.

2 Related Work

Colorizing grayscale images has been extensively investigated in the literature. Scribbles [9] requires the user to manually apply strokes of color to different regions of a grayscale image and neighboring pixels with the same luminance should get the same color. Transfer techniques [16] use the color palette from a reference image and apply it to a grayscale image by matching luminance and texture. Both scribbles and transfer techniques require manual input from the user. Recently, colorization methods based on automatic transformation, i.e., the only input to the method is the grayscale image, have become popular. Promising results have been demonstrated in the area of automatic transformation using Convolutional Neural Networks (CNNs) [3, 5, 8, 17] and Generative Adversarial Networks (GANs) [2, 6, 13] due to their abilities to model semantic representation in images.

In the infrared spectrum, less research has been done on transforming thermal images to VIS images. In [10], a CNN-based method was proposed to transform near-infrared (NIR) images to VIS images. Their method was shown to perform well as the NIR and VIS images are highly correlated in the electromagnetic spectrum. Kniaz *et al.* [7] proposed VIS to TIR transformation using a CNN model as a way to generate synthetic TIR images. The KAIST-MS [11] dataset introduced the first realistic large-scale dataset of TIR-VIS image pairs which opened up for developing TIR-VIS transformation models. Berg *et al.* [1] proposed a CNN-based model to transform TIR images to VIS images trained on the KAIST-MS dataset. However, the imperfect registration of the dataset caused the output from their method to be blurry and corrupted in some cases.

Generative Adversarial Networks (GANs) have shown promising results in unsupervised domain transfer [6, 12, 18, 19]. An unsupervised method does not require a paired dataset, hence, eliminating the need for pixel to pixel correspondence. In [14, 15], GANs have demonstrated a very good performance on transferring NIR images to VIS images. Isola *et al.* [6] has shown some qualitative results from the KAIST-MS dataset as an example for domain transfer. Inspired by [18], we employ an unsupervised GAN to transform TIR images to

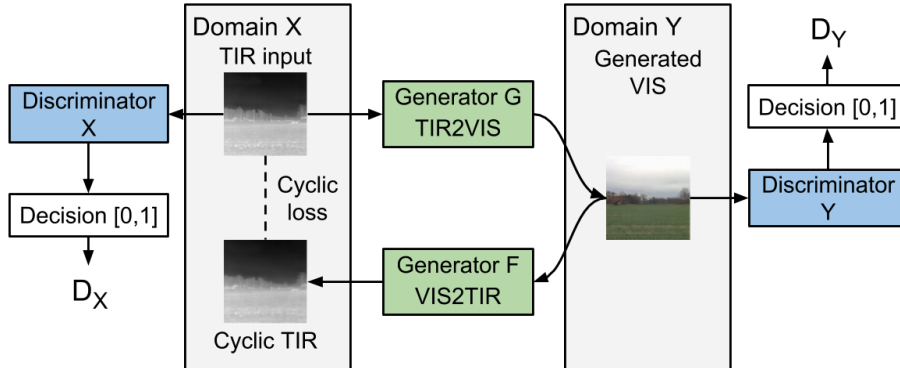


Fig. 2: The unpaired model is mapping between two different domains, X (TIR) and Y (VIS), using $G : X \rightarrow Y$ and $F : Y \rightarrow X$. This model produces three different losses; two adversarial losses based on D_Y and D_X , and one cycle consistency loss.

VIS images which eliminates the deficiencies caused by the imperfect registration in the KAIST-MS dataset as the training set does not need to be paired. Different from [1], our proposed method produces a very realistic and sharp VIS images. In addition, our proposed method is able to generalize very well on unseen data from different environments.

3 Method

Here we describe our proposed approach for transforming TIR images to VIS image while handling data miss-alignment in the KAIST-MS dataset.

3.1 Unpaired TIR-VIS Transfer

Inspired by [18], we perform unsupervised domain transfer between TIR and VIS images. Given TIR domain X with images $\{x_i : x_i \in X\}_{i=1}^N$ and VIS domain Y with images $\{y_j : y_j \in Y\}_{j=1}^M$, we aim to learn two transformations G and F between the two domains as shown in Figure 2. TIR input images are transformed from the thermal domain X to the visible spectrum domain Y using the generator G , while the generator F performs in the opposite direction. An adversarial discriminator D_X aims to discriminate between images x and the transformed images $F(y)$, while another discriminator D_Y discriminates between images y and $G(x)$.

3.2 Adversarial Training

The main objective of a GAN with a cyclic loss is to learn the two transformations $G : X \rightarrow Y$, $F : Y \rightarrow X$ and their corresponding discriminators D_X, D_Y [18].

During training, the generator G_G transforms an image $x \in X$ into a synthetic image \hat{y} . The synthetic image is then evaluated by the discriminator D_Y . The adversarial loss for the G_G is defined as:

$$\mathcal{L}_{G_G}(G_G, D_Y, x) = \mathbb{E}_{x \sim p_{data(x)}} [(D_Y(G_G(x)) - 1)^2] \quad (1)$$

where $p_{data(x)}$ is data distribution in X . The loss value becomes large if the synthetic image \hat{y} was able to fool the discriminator into outputting a value close or equal to one. On the other hand, the discriminator tries to maximize the probability for real images while also minimizing the output on synthetic images, achieved by minimizing the following formula:

$$\mathcal{L}_{D_Y}(D_Y, x, y) = \mathbb{E}_{x \sim p_{data(x)}} [(D_Y(G_G(x)))^2] + \mathbb{E}_{y \sim p_{data(y)}} [(D_Y(y) - 1)^2]. \quad (2)$$

The total adversarial loss for the G transformation is then defined as:

$$\mathcal{L}_{GAN}(G_G, D_Y, x, y) = \mathcal{L}_{D_Y}(D_Y, y) + \min_{G_G} \max_{D_Y} \mathcal{L}_{G_G}(G_G, D_Y, x) \quad (3)$$

A similar loss is utilized to learn the transformation F . To reduce the space of possible transformations, a cycle-consistency loss [18] is employed which ensures that the learned transformation can map only a single input to the desired output. The cycle-consistency loss is defined as:

$$\mathcal{L}_{cyc}(G_G, G_F) = \mathbb{E}_{x \sim p_{data(x)}} [||G_F(G_G(x)) - x||_1] + \mathbb{E}_{y \sim p_{data(y)}} [||G_G(G_F(y)) - y||_1] \quad (4)$$

Combining the above losses gives our total loss which is defined as:

$$\begin{aligned} \mathcal{L}(G_G, G_F, D_X, D_Y) = & \mathcal{L}_{GAN}(G_G, D_Y, x, y) + \\ & \mathcal{L}_{GAN}(G_F, D_X, y, x) + \\ & \lambda \mathcal{L}_{cyc}(G_G, G_F) \end{aligned} \quad (5)$$

where λ is a factor used to control the impact of the cyclic loss.

4 Experiments

For evaluation, we compare our proposed method with the existing state-of-the-art method on TIR to VIS transfer TIR2Lab [1]. The evaluation is performed on the KAIST-MS, the FOI dataset and the generalization capabilities of the evaluated methods are tested by training on the former and evaluating on the latter.

4.1 Datasets

The *KAIST-MS dataset* [4] contains paired TIR and VIS images which were collected by mounting a TIR and a VIS cameras on a vehicle in city and suburban

environments. The KAIST-MS dataset consists of 33,399 training image pairs and 29,179 test image pairs captured during daylight with a spatial resolution of (640×512) . Because the TIR camera used, FLIR A35, is only capable of capturing images with a resolution of (320×256) we resized all images to (320×256) . The images were collected continuously during driving, resulting in multiple image pairs being very similar. To remove redundant images, only every fourth image pair was included in the training set, resulting in 8,349 image pairs. For the evaluation, all image pairs from the test set were used.

The FOI dataset was captured using a co-axial imaging system capable of, theoretically, capturing TIR and VIS images with the same optical axis. Two cameras were used for all data collection, a TIR camera FLIR A65³ and VIS camera XIMEA MC023CG-SY⁴. All images were cropped and re-sized to 320×256 pixels. This system was used to capture TIR-VIS image pairs in natural environments with fields and forests with a training set of 5,736 image pairs and 1,913 image pairs for the test set. The average registration error for all image pairs were between 0.8 and 2.2 pixels.

4.2 Experimental Setup

Since we use the architecture from [18], we crop all training and test images from the center to the size (256×256) . TIR2Lab [1] was trained and evaluated on the full resolution of the images. All experiments were performed on GeForce GTX 1080 Ti graphics card⁵.

KAIST-MS dataset experiments A pretrained model for TIR2Lab [1] trained on the KAIST-MS dataset was provided by the authors. Our proposed model (TIRcGAN) was trained from scratch for 44 epochs using the same hyperparameters as [18]. Those parameters were batch size = 1, $\lambda = 10$ and for the ADAM optimizer we used learning rate = $2e - 4$, $\beta_1 = 0.5$, $\beta_2 = 0.999$ and $\epsilon = 1e - 8$. Some examples from pix2pix model [6] on the KAIST-MS dataset were provided by the authors and are discussed in the qualitative analysis.

FOI dataset experiments When training the TIR2Lab model, we used the same hyperparameters as mentioned in [1], except that we had to train for 750 epochs before it converged. For our TIRcGAN model, we trained for 38 epochs, using the same parameters as in the *KAIST-MS dataset experiments*.

Evaluation Metrics For the quantitative evaluation we use L_1 , root-mean-square error (RMSE), peak signal-to-noise ratio (PSNR) and Structural Similarity (SSIM) calculated between the transformed TIR image and the target VIS image. All metrics were calculated in the RGB color space normalized between 0 and 1 with standard deviation denoted as \pm .

³ <http://www.flir.co.uk/automation/display/?id=56345>

⁴ <https://www.ximea.com/en/products/usb-31-gen-1-with-sony-cmos-xic/mc023cg-sy>

⁵ <https://www.nvidia.com/en-us/geforce/products/10series/geforce-gtx-1080-ti>

| Model | Trained on | Evaluated on | L_1 | RMSE | PSNR | SSIM |
|----------------|-------------|--------------|---------------------------|---------------------------|----------------------------|--------------------|
| TIR2Lab | KAIST-MS | KAIST-MS | 0.13 ± 0.04 | 0.46 ± 0.09 | 14.7 ± 2.20 | 0.64 ± 0.08 |
| TIRcGAN (Ours) | KAIST-MS | KAIST-MS | 0.15 ± 0.05 | 0.21 ± 0.06 | 13.92 ± 2.48 | 0.55 ± 0.10 |
| TIR2Lab | FOI dataset | FOI dataset | 0.12 ± 0.04 | 0.15 ± 0.05 | 17.05 ± 2.52 | 0.81 ± 0.07 |
| TIRcGAN (Ours) | FOI dataset | FOI dataset | 0.11 ± 0.05 | 0.14 ± 0.05 | 17.63 ± 2.83 | 0.77 ± 0.10 |
| TIR2Lab | KAIST-MS | FOI dataset | 0.32 ± 0.11 | 0.39 ± 0.11 | 8.48 ± 2.68 | 0.54 ± 0.06 |
| TIRcGAN (Ours) | KAIST-MS | FOI dataset | 0.20 ± 0.07 | 0.23 ± 0.08 | 13.10 ± 2.82 | 0.60 ± 0.10 |

Table 1: This table shows the qualitative results for the experiments calculated on RGB values normalized between 0 and 1. The standard deviation is denoted as \pm .

4.3 Quantitative Results

Table 1 shows the quantitative results for the state-of-the-art method TIR2Lab [1] on the task of TIR to VIS domain transfer and our proposed method. Our method achieves comparable results to TIR2Lab in terms of L_1 despite the fact that our proposed method is unsupervised. On the other hand, our proposed method has a significantly lower RMSE than TIR2Lab which indicates its robustness against outliers. On the FOI dataset, our proposed method marginally outperforms TIR2Lab with respect to all evaluation metrics.

Model Generalization To evaluate generalization capabilities of methods in comparison, we train them on the KASIT-MS dataset and evaluate on the FOI dataset. The former was captured in city and urban environment, while the latter was captured in natural environments and forests. As shown in Table 1, our proposed method maintains its performance to a big extent when evaluated on a different dataset. On the other hand, TIR2Lab model failed to generalize to unseen data.

4.4 Qualitative Analysis

The KAIST-MS dataset As shown in Figure 3, our proposed TIRcGAN produces much sharper and saturated images than the TIR2Lab model on the KAIST-MS dataset. In addition, TIRcGAN is more inclined to generate smooth lines and other objects as an attempt to make the transformed TIR image look more realistic. On the other hand, images from TIR2Lab are quite blurry and lacks details in some occasions. Pix2pix performs reasonably and produces sharp images, however, objects and lines are smeared out in some cases.




Fig. 3: Example images from the KAIST-MS dataset experiment. It is possible to note that the pix2pix and TIRcGAN produce sharper images than the TIR2Lab model. TIRcGAN model is able to distinguish between yellow and white lanes as seen in the third row. *Note:* in the third row, the corresponding frame for the pix2pix model was not available, so the closest frame was used.

The FOI dataset Figure 4 show the results for the TIR2Lab and our proposed TIRcGAN models on the FOI dataset. TIRcGAN consistently outperform the TIR2Lab model when it comes to producing perceptually realistic VIS images. TIR2Lab produces blurry images that lacks a proper amount of details, while TIRcGAN produces a significant amount of details that are very similar to the target VIS image.

Model Generalization Figure 5 show the TIR2Lab and our TIRcGAN ability to generalize from one dataset collected on one environment to a new dataset from a different environment. Both models struggle at generating accurate colors or perceptually realistic images since the two datasets have different colors distribution. However, TIRcGAN was able to predict objects in the image with a reasonable amount of details contrarily to TIR2Lab which completely failed.



Fig. 4: Example images from the FOI dataset experiment. We can see that the unpaired model produce much sharper and more realistic images.

4.5 Failure Cases

Figure 6 shows examples where different models fail on the KAIST-MS dataset. For all methods, predicting humans is quite troublesome. Road crossing lines are also challenging as they are not always visible in the TIR images. Figure 7 shows some failure case on the FOI dataset. Both models fails in predicting dense forests, side-roads and houses since they are not very common in the dataset.

5 Conclusions

In this paper, we addressed the problem of TIR to VIS spectrum transfer by employing unsupervised GAN model that train on unpaired data. Our method was able to handle misalignments in the KAIST-MS dataset and produced perceptually realistic and sharp VIS images compared to the supervised state-of-the-art methods. When our method was trained on the KAIST-MS dataset and eval-



Fig. 5: Examples for methods output when trained on the KAIST-MS dataset [4] and evaluated on the FOI dataset. Both models struggle with predicting the colors since the two datasets were captured in different environments. However, TIRcGAN can still predict the objects in the scene.



Fig. 6: Example images where different models fail on the KAIST-MS dataset. In the first row, we see that the model produces an inverse shadow, i.e., painting shadow only where there should not be shadow. In the second row we show that all the models struggle with producing perceptually realistic VIS images of humans.




Fig. 7: Example images where models fail on the FOI dataset experiment. Here we see that the models are not able to accurately colorize houses and some roads.

uated on the new FOI dataset, it maintained its performance to a big extent. This demonstrated the generalization capabilities of our proposed method.

References

1. Berg, A., Ahlberg, J., Felsberg, M.: Generating visible spectrum images from thermal infrared. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (June 2018)
2. Cao, Y., Zhou, Z., Zhang, W., Yu, Y.: Unsupervised diverse colorization via generative adversarial networks. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases. pp. 151–166. Springer (2017)
3. Cheng, Z., Yang, Q., Sheng, B.: Deep colorization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 415–423 (2015)
4. Hwang, S., Park, J., Kim, N., Choi, Y., So Kweon, I.: Multispectral pedestrian detection: Benchmark dataset and baseline. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1037–1045 (2015)
5. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Let there be Color!: Joint End-to-end Learning of Global and Local Image Priors for Automatic Image Colorization with Simultaneous Classification. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2016)* **35**(4) (2016)
6. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. *arXiv* (2016)
7. Kniaz, V.V., Gorbatshevich, V.S., Mizginov, V.A.: Thermalnet: a Deep Convolutional Network for Synthetic Thermal Image Generation. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* pp. 41–45 (May 2017). <https://doi.org/10.5194/isprs-archives-XLII-2-W4-41-2017>
8. Larsson, G., Maire, M., Shakhnarovich, G.: Colorization as a proxy task for visual understanding. In: CVPR. vol. 2, p. 8 (2017)
9. Levin, A., Lischinski, D., Weiss, Y.: Colorization using optimization. In: *ACM Transactions on Graphics (TOG)*. vol. 23, pp. 689–694. ACM (2004)
10. Limmer, M., Lensch, H.P.: Infrared colorization using deep convolutional neural networks. In: Machine Learning and Applications (ICMLA), 2016 15th IEEE International Conference on. pp. 61–68. IEEE (2016)
11. Liu, J., Zhang, S., Wang, S., Metaxas, D.N.: Multispectral deep neural networks for pedestrian detection. *arXiv preprint arXiv:1611.02644* (2016)
12. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: *Advances in Neural Information Processing Systems*. pp. 700–708 (2017)
13. Suárez, P.L., Sappa, A.D., Vintimilla, B.X.: Infrared image colorization based on a triplet dcgan architecture. In: *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. pp. 212–217. IEEE (2017)
14. Suárez, P.L., Sappa, A.D., Vintimilla, B.X.: Learning to colorize infrared images. In: *International Conference on Practical Applications of Agents and Multi-Agent Systems*. pp. 164–172. Springer (2017)
15. Surez, P.L., Sappa, A.D., Vintimilla, B.X.: Infrared image colorization based on a triplet dcgan architecture. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. pp. 212–217 (July 2017). <https://doi.org/10.1109/CVPRW.2017.32>
16. Welsh, T., Ashikhmin, M., Mueller, K.: Transferring color to greyscale images. In: *ACM Transactions on Graphics (TOG)*. vol. 21, pp. 277–280. ACM (2002)
17. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: *European Conference on Computer Vision*. pp. 649–666. Springer (2016)
18. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint* (2017)

19. Zhu, J.Y., Zhang, R., Pathak, D., Darrell, T., Efros, A.A., Wang, O., Shechtman, E.: Toward multimodal image-to-image translation. In: Advances in Neural Information Processing Systems. pp. 465–476 (2017)