This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Panoptic-aware Image-to-Image Translation

Liyun Zhang¹, Photchara Ratsamee^{1,2}, Bowen Wang¹, Zhaojie Luo¹, Yuki Uranishi¹, Manabu Higashida¹ and Haruo Takemura¹

¹Osaka University, Japan liyun.zhang@lab.ime.cmc.osaka-u.ac.jp ²Osaka Institute of Technology, Japan photchara@ime.cmc.osaka-u.ac.jp

Abstract

Despite remarkable progress in image translation, the complex scene with multiple discrepant objects remains a challenging problem. The translated images have low fidelity and tiny objects in fewer details causing unsatisfactory performance in object recognition. Without thorough object perception (i.e., bounding boxes, categories, and masks) of images as prior knowledge, the style transformation of each object will be difficult to track in translation. We propose panoptic-aware generative adversarial networks (PanopticGAN) for image-to-image translation together with a compact panoptic segmentation dataset. The panoptic perception (i.e., foreground instances and background semantics of the image scene) is extracted to achieve alignment between object content codes of the input domain and panoptic-level style codes sampled from the target style space, then refined by a proposed feature masking module for sharping object boundaries. The image-level combination between content and sampled style codes is also merged for higher fidelity image generation. Our proposed method was systematically compared with different competing methods and obtained significant improvement in both image quality and object recognition performance.

1. Introduction

Image-to-image (I2I) translation is a challenging problem in the computer vision field. It needs to combine the content information of input domain image and the style of target domain [9]. Initially, some image-level I2I translation models were proposed based on paired (*e.g.*, Pix2Pix [11] or BicycleGAN [43]) or unpaired datasets (*e.g.*, CycleGAN [42] or MUNIT [9]) and later translating high quality (realism, sharpness and diversity) images became a hot problem (*e.g.*, Pix2PixHD [38] or U-GAT-IT [17]). With the development of object-driven image synthesis (*e.g.*, SG2IM [15] and Layout2IM [40] synthesizes images from scenes and layout, respectively), semantics or object instance tends to promote synthesizing image with sharper objects. There-



Figure 1. Pipeline comparisons of image-level [43], instance-level [33] and proposed I2I translation methods. (a) uses randomly sampled image-level style codes for I2I translation; (b) uses instance-level (objects are only countable foreground instances 'thing', *e.g.*, car) and image-level style codes; Our approach (c) uses panoptic-level (objects are both 'thing' and uncountable background segments 'stuff', *e.g.*, road) and image-level style codes.

fore, some instance-level I2I translation methods based on object instance as a perception had been proposed, *e.g.*, INIT [33] achieved separate learning of instances/local and whole-background/global areas. They can generate highfidelity object instances. However, for the complex scene with multiple discrepant objects, the above methods cannot translate images to keep high fidelity and tiny objects in more detail for both foreground and background.

As illustrated in Fig. 1, the image-level I2I translation

method in Fig. 1 (a) extracts image representation as content codes to combine with image-level style codes, which are randomly sampled from the style space of the target domain for I2I translation. In contrast, the instance-level method in Fig. 1 (b) uses the pre-trained instance segmentation network [7] to extract instance perception from the input image, it provides bounding boxes, categories, and masks of 'thing' (foreground object instances). The 'thing' representations are extracted by Region of Interest Align (RoIAlign) [7] via bounding boxes from the image representation and then combine with the image representation as content codes. Sampled instance-level and imagelevel style codes are aligned with the corresponding content codes for an instance-aware I2I translation, which can refine the foreground object instance representation precisely but does not fully refine background semantic regions.

In this paper, our proposed panoptic-level method (PanopticGAN) in Fig. 1 (c) uses a pre-trained panoptic segmentation network [19] to extract panoptic perception, it provides bounding boxes, categories, and masks of 'thing' (foreground object instances) and 'stuff' (background semantic regions). The 'thing' and 'stuff' representations are extracted and then combine with image representation as whole content codes. Sampled panoptic-level and imagelevel style codes are aligned with the corresponding content codes for a panoptic-aware I2I translation, which can thoroughly refine each recognizable area in the image via panoptic perception for tracking the style transformation in translation and avoid losing too much information. Our main contributions are threefold:

- A novel GAN framework for panoptic-aware I2I translation: The proposed framework extracts panoptic perception to align object content codes of the input domain with sampled panoptic-level style codes of the target domain, the image-level combination between content and sampled style codes is also merged for panoptic-aware image translation, which has high fidelity and tiny objects in more detail.
- A feature masking module for sharping object **boundaries:** The style-aligned feature maps are further refined by feature masking to obtain sharp object boundaries for higher fidelity image generation.
- A compact panoptic segmentation thermal image dataset: We annotated a panoptic segmentation thermal image dataset on partial KAIST-MS dataset [10], augmented dataset can be used for training a panoptic segmentation model to extract panoptic perception of thermal images in I2I translation or other tasks.

2. Related work

Image-to-image translation. Image-to-image (I2I)

translation models transform the input domain image to target domain, it changes style but keeps content unchanged. Pix2Pix [11] achieved paired dataset learning, but it generates single-modal output. BicycleGAN [43] achieved a bijective mapping between latent and output spaces for multimode results. CycleGAN [42] uses cycle consistency loss for unpaired training. The disentangled representation models [9, 22, 25] combine input domain content and target domain style for unsupervised learning. Pix2PixHD [38] can translate high-resolution images by a multi-scale discriminator and coarse2fine generator. AGGAN [37] and U-GAT-IT [17] extracted attention regions as guidance to localize important content for high-quality results. TSIT [13] uses a two-stream model with feature transformations for coarseto-fine image synthesis. However, for the complex scene with multiple discrepant objects, the above methods cannot generate images in high fidelity.

Instance-level image-to-image translation. The Instance-level I2I translation is derived from the objectdriven image generation methods (e.g., synthesizing images from object scenes [1, 15] or generating images from layouts [35, 36, 40]), they use object perception (i.e., bounding boxes or masks) for generating sharp object boundaries. Instagan [29] incorporated a set of instance attributes for instance-aware I2I translation. DA-GAN [26] learned a deep attention encoder to consequently discover instancelevel correspondences. SCGAN [41] and SalG-GAN [12] regarded saliency maps as an object perception for image translation. Shen et al. [33], Su et al. [34] and Chen et al. [3] combined the instance-level feature maps with image-level feature maps for high quality instance-level I2I translation. However, they only use the instance-level objects 'thing' without considering specific background semantic regions 'stuff' in the image translation process.

Panoptic-level image-to-image translation. To the best of our knowledge, the panoptic-level I2I translation problem has not yet been investigated. From the theoretical perspective, instance-level I2I translation only considers foreground instances as objects for learning, it has certain disadvantages compared with panoptic-level I2I translation, which regards both foreground 'thing' and background 'stuff' as objects. Lin et al. [24] extracts image regions for the discriminator to improve the performance of GANs, Huang et al. [8] controls the output based on references semantically. Dundar et al. [5] proved panoptic perception makes generated images have higher fidelity and tiny objects in more detail. Panoptic segmentation [20] combines semantic segmentation and instance segmentation to define uncountable background semantics (e.g., sky) as 'stuff' and countable foreground instances (e.g., car) as 'thing'. We use a pre-trained panoptic segmentation network [19] to extract panoptic perception (covers 'thing' and 'stuff') to make sampled panoptic-level style codes and image-level style



Figure 2. Illustration of the training manner and pipeline for PanopticGAN. The red arrows in (a) corresponds to the process of (b).



Figure 3. Illustration of the architecture of our proposed PanopticGAN; the detailed notations are described in the Architecture section.

codes combine corresponding content codes for a higher fidelity panoptic-aware I2I translation.

3. The proposed method

3.1. Overview

We provide an overview of our proposed method from the training manner and pipeline, using the summer-towinter (transforming summer domain to winter domain) I2I translation as an example describes our framework details.

Training manner. In Fig. 2 (a), we use summer image s and winter image w from the Transient Attributes dataset [21] to extract content codes (summer: C_s , winter: C_w) and style codes (summer: S_s , winter: S_w). E_s and

 E_w are encoders of s and w; D_s and D_w are decoders. By combining C_s with S_s to feed into D_s , we can reconstruct summer image s2s. Similarly, w2w can be reconstructed by C_w with S_w . Style codes Z_w and Z_s are randomly sampled from a normal distribution. By hypothesizing that Z_w is from the winter style space and combining C_s with Z_w to feed into D_w , it can synthesize winter image s2w as indicated by the red arrows. Similarly, w2s can be translated by C_w and Z_s . The cross-domain (s2w and w2s) and withindomain (s2s and w2w) are trained together [9].

Pipeline. In Fig. 2 (b), we use a pre-trained panoptic segmentation network to obtain the panoptic perception of the input image scene, it provides panoptic-level bounding boxes, categories, and masks. The bounding boxes are pro-

vided to RoIAlign [7] and masks are provided to the proposed feature masking module. Firstly, Image-level representation is extracted, panoptic-level style codes Z_{obj} and image-level style codes Z_{img} are sampled from normal distribution, $Z_{obj} = \{Z_{obj_i}\}_{i=1}^{m}$ is processed by the proposed panoptic object style-align module, m is the number of objects perceived in the panoptic perception. Note that we treat both 'thing' and 'stuff' as objects in panoptic perception. Z_{obj} and Z_{img} will be aligned with corresponding panoptic-level and Image-level representations in the generator for panoptic-level image translation. The translated images are fed into the discriminator, where we use fusion hinge loss consisting of image-level and object-level adversarial hinge loss terms [23] for optimization.

3.2. Architecture

Our architecture, as illustrated in Fig. 3, is built upon a generator, discriminator and proposed novel modules (panoptic object style-align and feature masking). We deploy a generative adversarial learning setting via summer and winter domain images from the Transient Attributes dataset [21] for illustration of our architecture.

3.2.1 Generator

In generator, the input summer image s (e.g., $256 \times$ 256) is extracted by a backbone module consisting of down-sampling residual blocks for obtaining image content codes C_{imq} (size 32×32 , dimension 256). Let $P = \{(category_i, bbox_i, mask_i)_{i=1}^m\}$ be panoptic perception consisting of categories, bounding boxes, and masks, where m is the number of objects perceived from a pretrained panoptic segmentation network and $category_i \in$ CAT (CAT defines 134 categories in the COCO-Panoptic dataset [20], here 'thing' has 80 categories and 'stuff' has 54 categories). C_{img} is cropped by RoIAlign [7] through object bounding boxes of $P(bbox_i)_{i=1}^m$ into object content codes $C_{obj} = \{C_{obj_i}\}_{i=1}^{m}$ (size 8×8 , dimension 128). Define Z_{img} as image-level style codes (dimension 256) and $Z_{obj} = \{Z_{obj_i}\}_{i=1}^{\bar{m}}$ as panoptic-level style codes (dimension 64), which are randomly sampled from normal distribution. The goal of the generator in the summer-to-winter translation is to learn a generation function $G(\cdot)$, which is capable of translating summer image s to a generated winter image w' via a given (Z_{imq}, Z_{obj}) :

$$w' = G(s|Z_{img}, Z_{obj}; \Theta_G) \tag{1}$$

where Θ_G are the parameters of the generation function. **Panoptic object style-align.** We use a MLP network to process Z_{obj} to dynamically generate the parameters $y = (y_{\gamma}, y_{\beta})$ of Adaptive Instance Normalization (AdaIN) [16] layers, then C_{obj} are processed by the residual blocks with AdaIN layers. The parameters of AdaIN layers fuse panoptic-level style with content to translate the different objects in the target image.

2

$$AdaIN(x_i, y) = y_{\gamma, i} \left(\frac{x_i - \mu(x_i)}{\sigma(x_i)} \right) + y_{\beta, i} \qquad (2)$$

where x_i is each feature map of C_{obj} , which is normalized separately and then scaled and biased using the corresponding scalar components from style y. The μ and σ are channel-wise mean and standard deviation, γ and β are AdaIN parameters generated from Z_{obj} . This process achieves panoptic object style-align, we obtained the stylealigned object representation $O_{obj} = \{O_{obj_i}\}_{i=1}^{m}$.

$$O_{obj} = AdaIN(C_{obj}, Z_{obj}) \tag{3}$$

Similarly, image-level style codes Z_{img} are also processed by a MLP network to generate AdaIN parameters, which fuse the image-level style with image content codes C_{img} by the residual blocks with AdaIN layers to obtain a hidden representation H_{img} .

Feature masking. As illustrated in Fig. 3, Oobj contains *m* object feature maps $\{O_{obj_i}\}_{i=1}^m$. Since the object bounding boxes $P(bbox_i)_{i=1}^m$ define the size and location of each object in the original image, we firstly affine transform each object feature maps O_{obj_i} into its corresponding original bounding box, secondly we do zero padding outside each bounding box in the image to obtain new object feature maps $B_{obj} = \{B_{obj_i}\}_{i=1}^{m}$. To remove the redundant background information outside the object contour, we further refine B_{obj} via object masks $M = P(mask_i)_{i=1}^m$ for more precise object boundaries. Compared with the Convolutional Feature Masking (CFM) layer [4] using the pixel projection method, after affine transformation the size of each feature map in B_{obj} is the same as masks M, therefore we only need to align B_{obj} and M along the category sequence of $1 \sim m$ and multiply to mask the values outside of object contour. Finally, we can obtain finer object feature maps $F_{obj} = \{F_{obj_i}\}_{i=1}^{m}$.

$$F_{obj} = B_{obj} \cdot M \tag{4}$$

We feed F_{obj} into three layers convolutional Long-Short-Term Memory (cLSTM) module (see supplementary material) to integrate each object feature maps $\{F_{obj_i}\}_{i=1}^m$ along object sequence of $1 \sim m$ to obtain fused hidden representation H_{obj} . We concatenate H_{obj} with H_{img} as H, which is up-sampled by the decoder consisting of up-sampling residual blocks to generate translated winter image w'.

3.2.2 Discriminator

As illustrated in Fig. 3, our discriminator consists of imagelevel and object-level classifiers. Similar to generator, the translated image is encoded by the backbone as image content codes C_{img} , which is refined by RoIAlign [7] as object content codes $C_{obj} = \{C_{obj_i}\}_{i=1}^{m}$ via bounding boxes $P(bbox_i)_{i=1}^m$. The image-level classifier consists of a global average pooling and one-output fully connected (FC) layer to process C_{img} to obtain a scalar realness score s_{img} . The object-level classifier consists of a flatten layer and two FC layers. One FC layer processes C_{obj} to compute a realness score for each object, denoted by $s_{real} = \{s_{real_i}\}_{i=1}^m$. Another FC layer computes a category projection score [2, 28, 35] for each object, denoted by $s_{cls} = \{s_{cls_i}\}_{i=1}^m$, which is the inner product between category embedding (transforming each category of $P(category_i)_{i=1}^m$ to a corresponding latent vector sampled from normal distribution) and linear projection (using a FC layer) of down-sampled C_{obj} . Therefore, the overall object-level loss of an object is $s_{obj_i} = s_{real_i} + s_{cls_i}$. The discriminator will be denoted by $D(\cdot, \Theta_D)$ with parameters Θ_D .

$$(s_{img}, s_{obj_i}, \cdots, s_{obj_m}) = D(I; \Theta_D)$$
(5)

Given an image I (ground truth w or generated w'), the discriminator computes the prediction score for the image and the average scores for objects.

3.3. Loss function

The full objective comprises three loss functions:

Adversarial loss. We utilize image-level and objectlevel fusion hinge version [23] of standard adversarial loss [6] to train (Θ_G, Θ_D) in our PanopticGAN,

$$l_k(I) = \begin{cases} \min(0, -1 + s_k); & \text{if } I \text{ is ground truth } w \\ \min(0, -1 - s_k); & \text{if } I \text{ is generated } w' \end{cases}$$
(6)

where $k \in \{img, obj_i, \dots, obj_m\}$. The overall loss is $l(I) = \lambda \cdot l_{img}(I) + \frac{1}{m} \sum_{i=1}^{m} l_{obj_i}(I)$ with trade-off parameter λ (1.0 used in experiment) in fusion hinge losses between image-level and object-level. We define the losses for the discriminator and generator respectively [35],

$$L_{adv}(\Theta_D, \Theta_G) = - \underset{(I) \ p_{all}(I)}{\mathbb{E}} [l(I)]$$

$$L_{adv}(\Theta_G, \Theta_D) = - \underset{(I) \ p_{fake}(I)}{\mathbb{E}} [D(I; \Theta_D)]$$
(7)

where minimizing $L_{adv}(\Theta_D, \Theta_G)$ makes discriminator to distinguish ground truth and translated images; minimizing $L_{adv}(\Theta_G, \Theta_D)$ fools discriminator by translating finegrained images. $p_{all}(I)$ represents ground truth and translated images, $p_{fake}(I)$ represents translated images.

Image reconstruction loss. We penalize the L_1 difference by $L_1^{img} = \left\| w' - w \right\|_1$ between the translated image

w' and ground truth w, $\|_1$ calculates the L1 norm. Here, we mainly calculate the within-domain (s2s and w2w) way.

Perceptual loss. The $L_{\rm p}$ alleviates the problem that translated images are prone to producing distorted textures,

$$L_{\rm p} = \sum_{k} \frac{1}{C_k H_k W_k} \sum_{i=1}^{H_k} \sum_{j=1}^{W_k} \left\| \phi_k(w')_{i,j} - \phi_k(w)_{i,j} \right\|_1$$
(8)

where $\phi_k(\cdot)$ represents feature representations of the *k*th max-pooling layer in VGG-19 network [14], and $C_k H_k W_k$ represents the size of feature representations.

Full objective. The final loss function is defined as:

$$L_{\text{total}} = \lambda_1 L_{\text{adv}} + \lambda_2 L_1^{img} + \lambda_3 L_{\text{p}}$$
(9)

where λ_i are the parameters balancing different losses.

3.4. Implementation details

In L_{total} , the $\lambda_1 \sim \lambda_3$ were set to 0.1, 1 and 10. Model parameters were initialized using the Orthogonal Initialization method [32]. The spectral normalization [27] is to stabilize the training in both the generator and discriminator. We used leaky-ReLU with a slope of 0.2 for the activation function and Adam optimizer [18] with $\beta_1 = 0$ and $\beta_2 = 0.9$. The learning rates were set to 10^{-4} for the generator and 0.005 for the discriminator. We set 400,000 iterations for training on four NVIDIA V100 GPUs.

4. Experiments

We conducted extensive experiments to evaluate our method with state-of-the-art models to show superiority in aspects of image quality and object recognition performance. For competing methods, MUNIT [9], BicycleGAN [43] and TSIT [13] belong to image-level I2I translation. SCGAN [41] uses a saliency map as object perception for instance-level I2I translation. INIT [33] is an instance-level method, we also implemented it for a more fair evaluation comparison. To achieve an adequately fair comparison, we add panoptic perception to image-level competing methods, i.e., MUNIT, BicycleGAN and TSIT. The panoptic perception is extracted from a pre-trained panoptic segmentation network [19] and concatenated with image features as an additional feature channel for training, hence we call them MUNIT+Seg, BicycleGAN+Seg and TSIT+Seg. Based on specific metrics, we summarized the evaluation results as qualitative and quantitative aspects to discuss respectively. Note that the model efficiency, more experimental results and limitations are provided in supplementary material.

4.1. Datasets

We trained and evaluated our model on the Transient Attributes [21] and KAIST-MS [10] datasets for day-to-

Method	HP (%) ↑			IS ↑			FID ↓			DS ↑		
	t2c	d2n	s2w	t2c	d2n	s2w	t2c	d2n	s2w	t2c	d2n	s2w
MUNIT+Seg [9]	0.8	0.4	0.7	2.29	1.50	1.92	98.5	98.7	93.9	0.46	0.65	0.62
BicycleGAN+Seg [43]	3.0	2.2	1.7	2.61	1.86	1.81	98.8	97.9	92.2	0.47	0.60	0.61
TSIT+Seg [13]	12.3	12.4	17.1	2.64	1.78	1.96	95.3	80.8	81.3	0.43	0.67	0.64
SCGAN [41]	8.4	9.1	8.0	2.59	1.62	1.58	96.8	92.4	86.4	0.39	0.53	0.49
INIT [33]	34.1	36.3	32.0	2.70	1.22	1.84	83.2	76.7	78.9	0.37	0.65	0.57
Ours	41.2	39.4	40.1	2.85	1.93	2.01	72.7	69.4	71.1	0.54	0.72	0.69

Table 1. Human Preference (HP), Inception Score (IS), Fréchet Inception Distance (FID) and Diversity Score (DS) metrics evaluate image quality in thermal-to-color (t2c), day-to-night (d2n) and summer-to-winter (s2w) tasks. Higher HP, IS and DS, and lower FID are better.



Figure 4. Comparison of image quality for translated images. Top: summer \rightarrow winter; Middle: day \rightarrow night; Bottom: thermal \rightarrow color.

night, summer-to-winter and thermal-to-color I2I translation tasks respectively. In the day-to-night task, we used 17,823 images for training and 2,287 images for evaluation; in summer-to-winter task, the training set is 17,674 images and evaluation set is 2558 images; in thermal-to-color, training set is 11,610 images and evaluation set is 2,541 images. For panoptic perception in training and inference of day-tonight and summer-to-winter tasks, we use a Panoptic FPN model [19] pre-trained on COCO-Panoptic dataset [20] to perceive from input day and summer images respectively. For panoptic perception in the training of thermal-to-color task, we perceive it from the paired color images via pre-

Method	$PQ\uparrow$	$SQ\uparrow$	$RQ\uparrow$	$PQ^{Th}\uparrow$	$SQ^{Th}\uparrow$	$RQ^{Th}\uparrow$	$PQ^{St}\uparrow$	$SQ^{St}\uparrow$	$RQ^{St}\uparrow$
MUNIT+Seg [9]	3.3	12.1	4.2	0.6	9.6	0.8	9.0	17.5	11.3
BicycleGAN+Seg [43]	4.3	16.8	5.5	0.8	13.1	1.2	10.9	23.9	13.6
TSIT+Seg [13]	6.4	17.2	8.1	2.1	13.3	3.3	13.9	26.4	15.3
SCGAN [41]	5.6	15.2	7.4	1.7	11.8	2.6	13.6	22.5	17.4
INIT [33]	7.2	19.6	9.0	3.1	15.4	3.9	16.7	29.1	20.9
Ours	8.3	22.7	11.3	4.2	17.5	5.1	18.4	31.0	21.6

Table 2. The PQ, SQ and RQ series metrics (higher is better) evaluate the object recognition performance of translated images.

trained Panoptic FPN model on COCO-Panoptic dataset; in the inference, it is perceived from input images via pretrained Panoptic FPN model on a compact our contributed dataset (see supplementary material) of thermal panoptic segmentation, the source data are the pairs of thermal and color images from partial KAIST-MS [10] dataset.

4.2. Evaluation metrics

We use Human Preference (HP), Inception Score (IS) [31], Fréchet Inception distance (FID) [30] and Diversity Score (DS) metrics for image quality, and Panoptic Quality (PQ) [20] series metrics for object recognition performance. HP is a user perceptual study that compares the image quality of translated results from different methods including input and ground truth image, which are shown to twenty participants to select the best translated image corresponding to target domain (covering thermal-to-color, day-to-night and summer-to-winter tasks respectively). IS is a popular metric to measure image quality generated from GANs. FID improves IS by incorporating statistics from real images. We used LPIPS metric [39] to calculate DS, which measures the differences between two images via computing perceptual similarity [40]. PO contains segmentation quality (SO) and recognition quality (RQ) [20], it combines mean Intersection over Union (mIoU) in SQ and average precision (AP) in RQ for a more comprehensive score than instance segmentation and object detection. In addition, PQTh, SQTh, RQTh are only used on 'thing' (Th) categories; PQSt, SQSt, RQSt are only used on 'stuff' (St) categories. Detailed descriptions of metrics see supplementary material.

4.3. Qualitative results

For image quality, the human preference results in Table 1 show that our approach achieved significantly higher scores in human perceptual study of different tasks compared to other approaches. Fig. 4 demonstrates our PanopticGAN can translate higher fidelity and brightly colored images and have tiny objects in more details. In contrast, the results of other methods are more blurry, distorted and missing small objects. For translated objects, our results tend to have better sharpness, more natural color style and display diversity (*e.g.*, the appearance of cars). On the other methods side, object sharpness is not satisfactory, the style is far from ground truth and there is insufficient diversity.

For object recognition performance, we use panoptic segmentation results by a pre-trained Panoptic FPN model on the COCO-Panoptic dataset. We only show object recognition results on thermal-to-color task, because translated night images from day-to-night and winter images from summer-to-winter tasks have disadvantages of insignificant differences for object recognition comparison. Fig. 5 shows that our method can achieve better object recognition performance than other methods, *e.g.*, the number and boundaries of cars; the structure of sky, tree and road; and the areas where there are relatively fewer recognition failures. Also, our results are significantly better than the results of original thermal images, this verified the advantages of our method when adapted for image enhancement.

4.4. Quantitative results

For image quality, the scores of IS, FID and DS in Table 1 demonstrate that our approach achieved superiority in image quality of translated images compared to other approaches. Our method overall outperforms baselines since we avoid losing too much information in the translation. The higher IS and lower FID of our proposed approach demonstrated that the translated images from our method have higher fidelity and sharpened object information. The higher DS demonstrated our method can show more flexibility and high robustness when the scene is invariant, especially for the objects generated on the image. For object recognition performance, Table 2 shows that our method performed state of the art scores compared with other competing trained models on all PQ, SQ, RQ, PQTh, SQTh, RQTh, PQSt, SQSt, RQSt object recognition metrics. From the score difference, our results are uniformly higher than the state-of-the-art competing methods by a certain distance, which stated the superiority of our method.

4.5. Ablation study

We demonstrated the necessity of losses and modules $(L_{\rm obj}$: object-level hinge loss; $L_1^{\rm img}$: image reconstruction loss; L_p : perceptual loss; $M_{\rm msk}$: feature masking; $M_{\rm pano}$: panoptic object style-align; $M_{\rm clstm}$: cLSTM) of our model by comparing Inception Score (IS) [31], Fréchet Inception Distance (FID) [30] and Diversity Score (DS) [40]



Figure 5. Comparison of object recognition performance for translated images. Upper: translated images; Lower: panoptic segmentation.

for image quality; Panoptic Quality (PQ) [20], Segmentation Quality (SQ) and Recognition Quality (RQ) are for object recognition performance. The experiment results show several ablated versions of our model trained on the KAIST-MS [10] dataset for the thermal-to-color task. As shown in Table 3, removing any loss will decrease the overall performance. Removing $L_{\rm obj}$ and $L_1^{\rm img}$ have lower IS, DS, and higher FID due to generating low fidelity image and objects with fewer variations; PQ, SQ, and RQ are also decreased since L_{obj} compute the category projection scores for objects. Removing $L_{\rm p}$, the model produced distorted textures, this inevitably decreases image quality and object recognition performance. Removing any $M_{\rm msk}$, $M_{\rm pano}$ or $M_{\rm clstm}$ module decreased overall performance, it demonstrates their necessity. Because $M_{\rm msk}$ sharpens object boundaries and M_{clstm} sequentially integrates different objects back into image. Especially, removing M_{pano} destroyed the whole foundation of our proposed panopticlevel framework, and overall performance is significantly decreased. Therefore, the above study result for losses and modules shows the reasonability of our model design.

5. Conclusion

We propose a novel panoptic-aware image-to-image translation method (PanopticGAN) together with a compact panoptic segmentation dataset. The panoptic perception is extracted to achieve a panoptic-level combination between content and style codes, results are further refined

Method	IS \uparrow	$FID\downarrow$	$DS\uparrow$	$PQ\uparrow$	$SQ\uparrow$	$RQ\uparrow$
w/o $L_{ m obj}$	2.24	110.4	0.47	5.3	16.2	7.7
w/o $L_1^{ m img}$	2.64	104.3	0.45	6.4	20.4	10.1
w/o $L_{\rm p}$	2.66	97.1	0.42	6.7	19.4	10.5
w/o $M_{\rm msk}$	2.53	101.6	0.47	6.1	19.8	10.2
w/o $M_{ m pano}$	2.44	120.7	0.43	5.9	18.4	9.1
w/o $M_{\rm clstm}$	2.63	103.2	0.42	6.8	16.7	11.1
full model	2.85	72.7	0.54	8.3	22.7	11.3

Table 3. Ablation study. The performance removing losses and modules are compared in image quality (evaluated by IS, FID and DS) and object recognition (evaluated by PQ, SQ and RQ).

by our proposed feature masking module for sharp object boundary generation. The image-level combination of content and style codes is also merged for translating images with high fidelity and tiny objects in more detail in complex scenes with multiple discrepant objects. Extensive experiments demonstrated that our method obtained significant improvement in image quality and object recognition performance compared to different methods. Incorporating bounding boxes and mask regression losses in image translation model training will be the focus of our future work.

Acknowledgments: This work has been partly supported by the KAKENHI Fund for the Promotion of Joint International Research (fostering joint international research (B) No. 20KK0086) and the Mohamed Bin Zayed International Robotics Challenge (MBZIRC) Grant.

References

- Oron Ashual and Lior Wolf. Specifying object attributes and relations in interactive scene generation. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision, pages 4561–4569, 2019.
- [2] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [3] Tianlang Chen, Wei Xiong, Haitian Zheng, and Jiebo Luo. Image sentiment transfer. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4407–4415, 2020.
- [4] Jifeng Dai, Kaiming He, and Jian Sun. Convolutional feature masking for joint object and stuff segmentation. In *Proceed*ings of the IEEE conference on computer vision and pattern recognition, pages 3992–4000, 2015.
- [5] Aysegul Dundar, Karan Sapra, Guilin Liu, Andrew Tao, and Bryan Catanzaro. Panoptic-based image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8070–8079, 2020.
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- [7] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [8] Jialu Huang, Jing Liao, and Sam Kwong. Semantic example guided image-to-image translation. *IEEE Transactions on Multimedia*, 23:1654–1665, 2021.
- [9] 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.
- [10] Soonmin Hwang, Jaesik Park, Namil Kim, Yukyung Choi, and In So Kweon. Multispectral pedestrian detection: Benchmark dataset and baseline. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1037–1045, 2015.
- [11] 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*, pages 1125–1134, 2017.
- [12] Lai Jiang, Mai Xu, Xiaofei Wang, and Leonid Sigal. Saliency-guided image translation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 16509–16518, 2021.
- [13] Liming Jiang, Changxu Zhang, Mingyang Huang, Chunxiao Liu, Jianping Shi, and Chen Change Loy. Tsit: A simple and versatile framework for image-to-image translation. In *European Conference on Computer Vision*, pages 206–222. Springer, 2020.
- [14] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.

- [15] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1219–1228, 2018.
- [16] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4401–4410, 2019.
- [17] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwanghee Lee. U-gat-it: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-toimage translation. arXiv preprint arXiv:1907.10830, 2019.
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [19] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6399–6408, 2019.
- [20] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9404–9413, 2019.
- [21] Pierre-Yves Laffont, Zhile Ren, Xiaofeng Tao, Chao Qian, and James Hays. Transient attributes for high-level understanding and editing of outdoor scenes. ACM Transactions on graphics (TOG), 33(4):1–11, 2014.
- [22] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European conference on computer vision (ECCV)*, pages 35–51, 2018.
- [23] Jae Hyun Lim and Jong Chul Ye. Geometric gan. arXiv preprint arXiv:1705.02894, 2017.
- [24] Yu Lin, Yigong Wang, Yifan Li, Yang Gao, Zhuoyi Wang, and Latifur Khan. Attention-based spatial guidance for image-to-image translation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 816–825, 2021.
- [25] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In Advances in neural information processing systems, pages 700–708, 2017.
- [26] Shuang Ma, Jianlong Fu, Chang Wen Chen, and Tao Mei. Da-gan: Instance-level image translation by deep attention generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5657–5666, 2018.
- [27] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. arXiv preprint arXiv:1802.05957, 2018.
- [28] Takeru Miyato and Masanori Koyama. cgans with projection discriminator. arXiv preprint arXiv:1802.05637, 2018.
- [29] Sangwoo Mo, Minsu Cho, and Jinwoo Shin. Instagan: Instance-aware image-to-image translation. arXiv preprint arXiv:1812.10889, 2018.
- [30] Suman Ravuri and Oriol Vinyals. Classification accuracy score for conditional generative models. Advances in neural information processing systems, 32, 2019.

- [31] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing* systems, 29:2234–2242, 2016.
- [32] Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- [33] Zhiqiang Shen, Mingyang Huang, Jianping Shi, Xiangyang Xue, and Thomas S Huang. Towards instance-level imageto-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3683–3692, 2019.
- [34] Jheng-Wei Su, Hung-Kuo Chu, and Jia-Bin Huang. Instanceaware image colorization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7968–7977, 2020.
- [35] Wei Sun and Tianfu Wu. Image synthesis from reconfigurable layout and style. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10531– 10540, 2019.
- [36] Tristan Sylvain, Pengchuan Zhang, Yoshua Bengio, R Devon Hjelm, and Shikhar Sharma. Object-centric image generation from layouts. *arXiv preprint arXiv:2003.07449*, 1(2):4, 2020.
- [37] Hao Tang, Dan Xu, Nicu Sebe, and Yan Yan. Attentionguided generative adversarial networks for unsupervised image-to-image translation. In 2019 International Joint Conference on Neural Networks (IJCNN), pages 1–8. IEEE, 2019.
- [38] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 8798–8807, 2018.
- [39] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [40] Bo Zhao, Lili Meng, Weidong Yin, and Leonid Sigal. Image generation from layout. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8584–8593, 2019.
- [41] Yuzhi Zhao, Lai-Man Po, Kwok-Wai Cheung, Wing-Yin Yu, and Yasar Abbas Ur Rehman. Scgan: Saliency map-guided colorization with generative adversarial network. *IEEE Transactions on Circuits and Systems for Video Technology*, 2020.
- [42] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycleconsistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223– 2232, 2017.
- [43] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Multimodal image-to-image translation by enforcing bi-cycle con-

sistency. In Advances in neural information processing systems, pages 465–476, 2017.