Parallel Generative Adversarial Network for Third-person to First-person Image Generation

Gaowen Liu, Hugo Latapie, Ozkan Kilic, Adam Lawrence
Emerging Technologies and Incubation, Cisco Research, USA
{gaoliu, hlatapie, okilic, adamlawr}@cisco.com

Abstract

Cross-view image generation has been recently proposed to generate images of one view from another dramatically different view. In this paper, we investigate third-person (exocentric) view to first-person (egocentric) view image generation. This is a challenging task since egocentric view sometimes is remarkably different from exocentric view. Thus, transforming the appearances across the two views is a non-trivial task. To this end, we propose a novel Parallel Generative Adversarial Network (P-GAN) with a novel cross-cycle loss to learn the shared information for generating egocentric images from exocentric view. We also incorporate a novel contextual feature loss in the learning procedure to capture the contextual information in images. Extensive experiments on the Exo-Ego datasets [5] show that our model outperforms the state-of-the-art approaches.

1. INTRODUCTION

Wearable cameras, also known as first-person cameras, nowadays are widely used in our daily lives since the appearance of low price but high quality wearable products such as GoPro cameras. Meanwhile, egocentric (first-person) vision is also becoming a critical research topic in the field. As we know, egocentric view have some unique properties other than exocentric (third-person) view. Traditional exocentric cameras usually give a wide and global view of the high-level appearances happened in a video. However, egocentric cameras can capture the objects and people at a much finer level of granularity. In the early egocentric vision studies, researchers [13] found that people perform different activities or interacting with objects from a first-person egocentric perspective and seamlessly transfer knowledge between egocentric and exocentric perspective. Therefore, analyzing the relationship between egocentric and exocentric perspectives is an extremely useful and interesting topic for image and video understanding. However, there is few research to address this important problem in literature.

Recently, Generative Adversarial Networks (GANs) [11] have been shown effectively in image generation tasks. Isola *et al.* [12] propose Pix2Pix adversarial learning framework on paired image generation, which is a supervised model and uses a conditional GAN framework to learn a translation function from input to output image domain. Zhu *et al.* [45] introduce CycleGAN which develops cycle-consistency constraint to deal with unpaired image generation. However, these existing works consider an application scenario in which the objects and the scenes have a large degree of overlapping in appearance and view. Recently, some works investigate cross-view image generation problems to generate a novel scene which is drastically different from a given scene image. This is a more challenging task since different views share little overlap information. To tackle this problem, Regmi and Borji [26] propose X-Fork and X-Seq GAN-based architecture using an extra semantic segmentation map to facilitate the generation. Moreover, Tang *et al.* [35] propose a multi-channel attention selection module within a GAN framework for cross-view image generation. However, these methods are not able to generate satisfactory results due to the drastically differences between exocentric and egocentric views.

To bridge egocentric and exocentric analaysis, in this paper we propose a novel Parallel GAN (P-GAN) to generate exocentric images from egocentric view. P-GAN framework is able to automatically learn the shared information between two parallel generation tasks via a novel cross-cycle loss and hard-sharing of network layers. We also utilize a novel contextual loss in our objective function to capture texture information over the entire images. To the best of our knowledge, we are the first to attempt to incorporate a parallel generative network for exocentric to egocentric image translation. The proposed P-GAN is related to CoGAN [18] and DualGAN [42]. However, CoGAN and DualGAN have limited ability in generating image pairs with dramatically different viewpoints. As shown in Fig. 1, our architecture is designed in a bi-directional parallel fashion to discover the shared information between egocentric and exocentric images. Two parallel GANs are trained simulta-

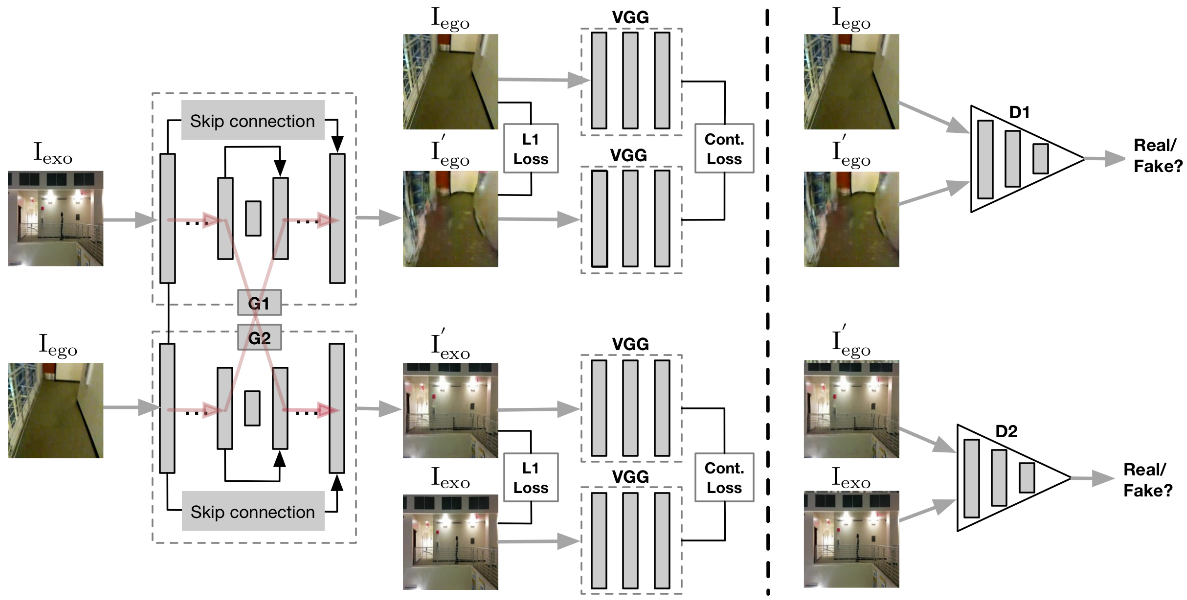

Figure 1. The pipeline of our P-GAN model. It consists of two parallel generators (G_1, G_2) and two discriminators (D_1, D_2). The total loss contains pairs of $L1$ loss, contextual loss and adversarial loss.

neously with hard-sharing of certain layers.

In summary, our contributions can be highlighted as follows:

- A novel P-GAN is proposed to learn the shared information between different views simultaneously via a novel cross-cycle loss.
- A novel contextual feature loss is incorporated in the training to capture the contextual information.
- Experiments on Exo-Ego dataset show the effectiveness of our hard-sharing of network layers in multi-directional parallel generative models. The early version of this paper appeared in [17].

2. Related Work

In this section, we review related work about Generative Adversarial Networks, Egocentric Vision and Synthesis.

2.1. Generative Adversarial Networks (GANs)

Over the last few years, GANs [10] have been shown effectively in many image generation tasks [12,32,36,45]. For example, Isola et al. [12] propose Pix2Pix adversarial learning framework for paired image generation. Zhu et al. [45] introduce CycleGAN which developed cycle-consistency constraint to deal with unpaired image generation. However, these works aim to generate images which have a large degree of overlapping in the appearance and view with input images. Synthesis is much more challenging when the generation is conditioned on images with drastically different views. Recently, researchers investigate cross-view image generation problems [27]. This is a more challenging task since different views share little overlap information. To tackle this problem, Krishna et al. [26] propose X-Fork and X-Seq GAN-based architecture using an extra semantic map to facilitate generation. Tang et al. [35] propose a semantic-guided multi-channel attention selection module within a GAN framework for cross-view image generation. However, these methods are limited to cross-view image generation task, they are not able to generate satisfactory results for cross-view video generation.

2.2. Egocentric Vision

Egocentric vision has been recently explored in the computer vision field [1, 7, 8, 21, 24, 25, 37, 43]. Aghazadeh et al. [1] propose an approach for discovering anomalous events from videos captured from a small camera attached to a person's chest. Fathi et al. [8] introduce a method for individuating social interactions in first-person videos collected during social events. Some recent works [7, 21, 24, 29, 37] have focused on activity analysis considering different scenarios (e.g., kitchen, office, home). Xu et al. [41] propose a semi-Siamese CNN architecture to address the person-level correspondences across first- and third-person videos. They formulate the problem as learning a joint embedding space for first- and third-person videos that considers both spatial- and motion-domain cues.

2.3. Synthesis

There is few recent work investigate image and video generation problem [3,6,9,16,30]. TGAN [30] directly generate video clips from noise by using generative adversarial networks. MoCoGAN [39] employ unsupervised adversarial training to decompose motion and content to control the image-to-video generation. Pan et al. [22] work on video-to-video translation to generate a sequence of frames from a sequence of aligned semantic representations. Some recent works such as RecycleGAN [2] and Vid2Vid [40] learn mapping between different videos and transferred motion between faces and from poses to body, respectively. Frameworks [18, 23, 38, 44] propose image generation networks for 3D view synthesis.

However, existing frameworks on image/video generation require the input and output image/video scenes sharing the similar architecture, which were insufficient for cross-view image and video generation. Particularly, exocentric to egocentric cross-view image/video generation has not yet been studied in literature yet. Our method investigates both cross-view generation and image/video generation in the exocentric to egocentric perspective setting, which is more challenging than various image/video generation problems. To the best of our knowledge, this is the first attempt in literature.

3. Parallel GANs

In this section, we address the problem of generating images across two drastically different views, namely top-ego view and side-ego view. We show the details of our network architectures which captured the shared information in two different viewpoints.

3.1. Network Architecture

Cross-view exocentric to egocentric image synthesis is a challenging task, because these two views have little overlapping in image appearance. Most existing works on cross-view image synthesis are based on GANs. A traditional GAN consists of a generative model and a discriminative model. The objective of the generative model is to synthesize images resembling real images, while the objective of the discriminative model is to distinguish real images from synthesized ones. Both the generative and discriminative models are realized as multi-layer perceptrons. Since there will be some shared high-level concept information in a pair of corresponding images between exocentric and egocentric views, we propose a P-GAN with two GANs in parallel which is able to learn the shared high-level semantic information among different views. Fig. 1 shows our framework which contains two generators and two discriminators. A set number of layers from two generators are shared across P-GAN. We force the first three layers of two generators to have the identical structure and share the weights, and the rest layers are task-specific. The experiments show that sharing three layers of generators yield the best performance.

Particularly, we employ U-Net [28] as the architecture of our generators G_1 and G_2. We impose skip connection strategy from down-sampling path to up-sampling path to avoid vanishing gradient problem. To learn the shared information between exocentric and egocentric view, we perform hard-sharing in the first three layers of down-sampling path. We adopt PatchGAN [12] for the discriminator D_1 and D_2. The feature maps for contextual loss are extracted by the VGG-19 network pretrained on ImageNet.

3.2. Overall Optimization Objective

The training objective can be decomposed into four main components which are contextual loss, adversarial loss, cross-cycle loss and reconstruction loss.

3.2.1 Contextual loss.

Different from the commonly used $L1$ loss function which compares pixels at the same spatial coordinates between the generated image and the target image, we incorporate contextual loss in our P-GAN learning framework. The key idea is to measure similarity between images at the high-level feature space.

Given a generated fake image I'_{ego} and a real image I_{ego} in egocentric view, we obtain a list of VGG-19 [33] features as $I_{ego}{=}\{I_i\}$ and $I'_{ego}{=}\{I'_j\}$, where $I_i{=}\psi^i(I_{ego})$, $I'_j{=}\psi^j(I'_{ego})$, ψ means VGG-19 feature. i, j are i-th and j-th layer in the network ψ. The similarity between the generated image I'_{ego} and the real image I_{ego} in egocentric view can be defined as follows,

$$\mathcal{S}_{I_i,I'_j} = \exp\left(1 - \frac{1-d_{ij}}{\min_k d_{ik} + \zeta}\right)/h \tag{1}$$

where d_{ij} is the cosine distance between I_{ego} and I'_{ego}. We define $\zeta{=}1e^{-5}$, $h{=}0.5$ in our experiments. The similarity can be normalized as,

$$\bar{S}_{ij} = \frac{\mathcal{S}_{I_i,I'_j}}{\sum_k \mathcal{S}_{I_i,I'_k}} \tag{2}$$

Then the contextual loss is formulated as follows,

$$\mathcal{L}_{cont}(I_i, I'_j) = \frac{1}{\max(\mid I_{ego} \mid, \mid I'_{ego} \mid)} \sum_j \max \bar{S}_{ij} \tag{3}$$

where $|\cdot|$ denotes the numbers of feature maps.

3.2.2 Cross-cycle loss.

As shown in Fig. 1, we employ U-Net [28] as our generators G_1 and G_2. Each U-Net contains a down-sampling encoder

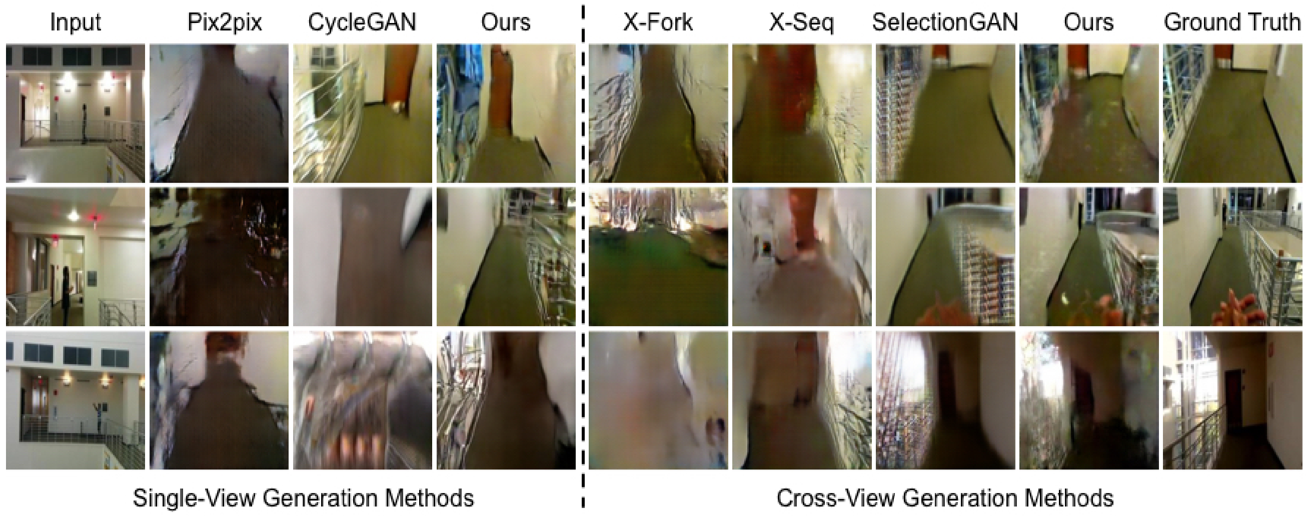

Figure 2. Results generated by different methods on Side2Ego dataset. These samples were randomly selected for visualization purposes. Columns from left to right are: Input, Pix2pix [12], CycleGAN [45], P-GAN (ours), X-Fork [26], X-Seq [26], SelectionGAN [35], P-GAN + Segmentation map (ours), Ground Truth.

EN which is a feature contracting path, and an up-sampling decoder DE which is a feature expanding path. Inspired by the U-net properties, we design a novel cross-cycle loss as follows,

$$\mathcal{L}_X(G_1, G_2) = \mathbb{E}_{I_{exo}, I'_{exo}}[\|I_{exo} - DE_2(EN_1(I_{exo}))\|_1] + \lambda_1 \mathbb{E}_{I_{ego}, I'_{ego}}[\|I_{ego} - DE_1(EN_2(I_{ego}))\|_1] \tag{4}$$

3.2.3 Adversarial loss.

Recent works [4, 11, 19, 31, 34] have shown that one can learn a mapping function by tuning a generator and a discriminator in an adversarial way. Assuming we target to learn a mapping $G\colon I_{exo} \to I_{ego}$ from input exocentric image I_{exo} to output egocentric image I_{ego}. The generator G is trained to produce outputs to fool the discriminator D. The adversarial loss can be expressed as,

$$\mathcal{L}_{GAN_1}(G_1, D_1) = \mathbb{E}_{I_{exo}, I_{ego}}[\log D_1(I_{exo}, I_{ego})] + \mathbb{E}_{I_{exo}, I'_{ego}}[\log(1 - D_1(I_{exo}, G_1(I_{exo})))] \tag{5}$$

$$\mathcal{L}_{GAN_2}(G_2, D_2) = \mathbb{E}_{I_{ego}, I_{exo}}[\log D_2(I_{ego}, I_{exo})] + \mathbb{E}_{I_{ego}, I'_{exo}}[\log(1 - D_2(I_{ego}, G_2(I_{ego})))] \tag{6}$$

The adversarial loss is the sum of Eq. (5) and Eq. (6).

$$\mathcal{L}_{GAN} = \mathcal{L}_{GAN_1}(G_1, D_1) + \lambda_2 \mathcal{L}_{GAN_2}(G_2, D_2) \tag{7}$$

3.2.4 Reconstruction loss.

The task of the generator is to reconstruct an image as close as the target image. We use $L1$ distance in the reconstruction loss,

$$\mathcal{L}_{re}(G_1, G_2) = \mathbb{E}_{I_{exo}, I'_{ego}}[\|I_{ego} - DE_1(EN_1(I_{exo}))\|_1] + \lambda_3 \mathbb{E}_{I_{ego}, I'_{exo}}[\|I_{exo} - DE_2(EN_2(I_{ego}))\|_1] \tag{8}$$

3.2.5 Overall loss.

The total optimization loss is a weighted sum of the above losses. Generators G_1, G_2 and discriminators D_1, D_2 are trained in an end-to-end fashion to optimize the following objective function,

$$\mathcal{L} = \mathcal{L}_{GAN} + \lambda_4 \mathcal{L}_X + \lambda_5 \mathcal{L}_{re} + \lambda_6 \mathcal{L}_{cont} \tag{9}$$

where λ_i's are the regularization parameters.

4. Experimental Results

In this section, we provide detailed experiment setup and results.

4.1. Datasets

To explore the effectiveness of our proposed P-GAN model, we compare our model with the state-of-the-art methods on Exo-Ego dataset [5] which contains two different viewpoint subsets (Side2Ego and Top2Ego). This dataset is challenging due to two reasons. First, it contains dramatically different indoor and outdoor scenes. Second, the dataset is collected simultaneously by an exocentric camera (side and top view) and an egocentric body-worn wearable camera. It includes a huge amount of blurred images for egocentric view. For Side2Ego subset, there are

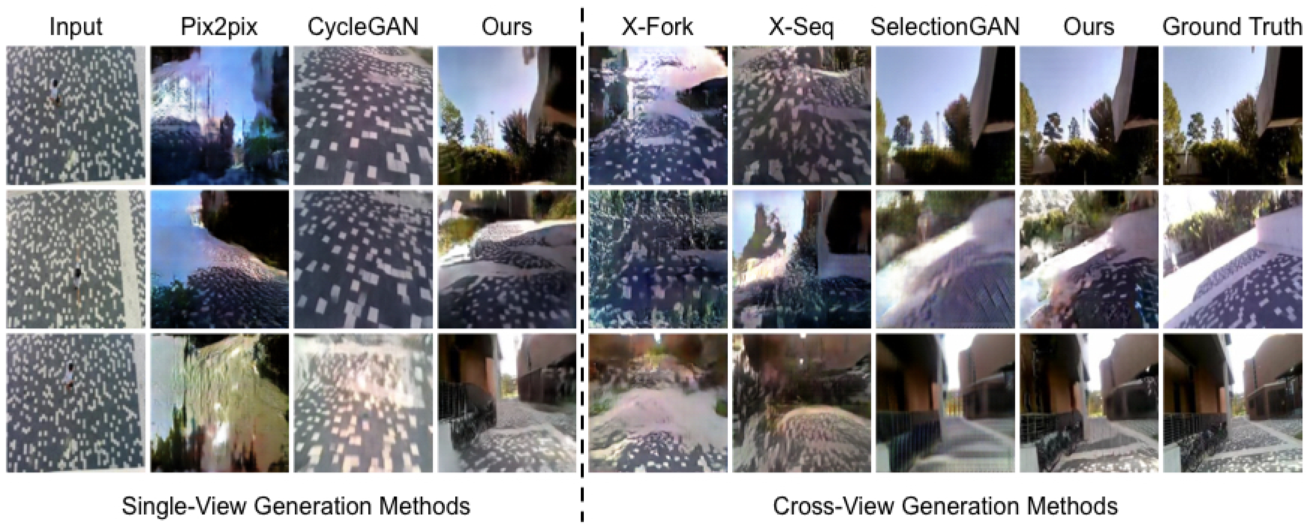

Figure 3. Results generated by different methods on Top2Ego dataset. These samples were randomly selected for visualization purposes. Columns from left to right are: Input, Pix2pix [12], CycleGAN [45], P-GAN (ours), X-Fork [26], X-Seq [26], SelectionGAN [35], P-GAN + Segmentation map (ours), Ground Truth.

Table 1. SSIM, PSNR, Sharpness Difference (SD), KL score (KL) and Accuracy of different single-view image generation methods. For these metrics except KL score, higher is better.

2*Dataset	2*Method	2*SSIM	2*PSNR	2*SD	2*KL	Top-1 Accuracy (%)		Top-5 Accuracy (%)	
3*Top2Ego	Pix2pix [12]	0.2514	15.0532	18.1002	62.74 ± 1.78	1.24	1.22	4.21	4.35
	CycleGAN [45]	0.2806	15.5486	18.5678	52.09 ± 1.69	**2.10**	0.99	5.37	2.72
	Ours	**0.3098**	**17.0236**	**18.6043**	**31.46 ± 1.74**	1.81	**5.90**	**5.74**	**9.17**
3*Side2Ego	Pix2pix [12]	0.3946	16.0716	19.8664	75.27 ± 2.01	3.20	5.18	8.41	13.30
	CycleGAN [45]	0.4017	15.9678	19.7533	62.41 ± 2.41	4.18	7.60	15.62	21.45
	Ours	**0.4908**	**17.995**1	**20.6521**	**13.92 ± 1.53**	**16.21**	**30.80**	**27.57**	**46.51**

Table 2. SSIM, PSNR, Sharpness Difference (SD), KL score (KL) and Accuracy of different cross-view image generation methods. For these metrics except KL score, higher is better.

2*Dataset	2*Method	2*SSIM	2*PSNR	2*SD	2*KL	Top-1 Accuracy (%)		Top-5 Accuracy (%)	
4*Top2Ego	X-Fork [26]	0.2952	15.8849	18.7349	63.96±1.74	0.8	1.22	3.16	4.08
	X-Seq [26]	0.3522	16.9439	19.2733	54.91 ± 1.81	1.07	1.77	4.29	6.94
	SelectionGAN [35]	0.5047	22.0244	19.1976	**10.07 ± 1.29**	8.85	16.55	24.32	33.90
	Ours	**0.5287**	**22.2891**	**19.2389**	12.07 ± 1.69	**9.76**	**29.67**	**24.80**	**51.79**
4*Side2Ego	X-Fork [26]	0.4499	17.0743	20.4443	51.20 ± 1.94	4.49	9.76	11.63	19.44
	X-Seq [26]	0.4763	17.1462	20.7468	45.10 ± 1.95	6.51	12.70	11.97	19.36
	SelectionGAN [35]	0.5128	18.3021	20.9426	**7.26 ± 1.27**	20.84	37.49	42.51	65.22
	Ours	**0.5205**	**19.4521**	**20.9684**	25.25 ± 1.88	**20.96**	**39.08**	**42.58**	**66.00**

26,764 pairs of images for training and 13,788 pairs for testing. For Top2Ego subset, there are 28,408 pairs for training and 14,064 pairs for testing. All images are in high-resolution 1280×720 pixels.

4.2. Experimental Setup

We compare our P-GAN with both single-view image generation methods [12, 45] and cross-view image generation methods [26, 35]. We adopt the same experimental setup as in [12, 26, 35]. All images are scaled to 256×256. We enable image flipping and random crops for data augmentation. To compute contextual loss, we follow [20] and use the VGG-19 network to extract image feature maps pretrained on ImageNet. We train 35 epochs with the batch size of 4. In our experiments, we set λ_1=10, λ_2=10, λ_3=100, λ_4=10, λ_5=1, λ_6=1 in Eq. (4), (7), (8) and (9), respectively. The state-of-the-art cross-view generation methods, *i.e.*, X-Fork [26], X-Seq [26] and SelectionGAN [35] utilize segmentation map to facilitate target view image generation. To compare with these cross-view methods, we adopt RefineNet [14, 15] to generate segmentation maps on Side2Ego and Top2Ego subsets as in [26, 35]. The generated segmentation maps are used as the conditional input of G_1 and G_2. The proposed P-GAN is implemented using PyTorch.

4.3. Evaluation Metrics

We apply metrics such as top-k prediction accuracy and KL score for evaluations as in [26, 35]. We also employ pixel-level similarity metrics, *i.e.*, Structural-Similarity (SSIM), Peak Signal-to-Noise Ratio (PSNR) and Sharpness Difference (SD). These metrics evaluate the generated images in a high-level feature space.

4.4. Quantitative Results

The quantitative results are presented in Tables 1 and 2. We observe that our P-GAN achieves better results than state-of-the-art methods in most cases. Compared with single-view image generation methods, our P-GAN outperforms Pix2pix [12] and CycleGAN [45]. On the other hand, we also achieve better results than other cross-view image generation methods in most metrics while incorporating semantic segmentation map as in the SelectionGAN [35].

4.5. Qualitative Results

Qualitative results are shown in Fig. 2 and Fig. 3. The results confirm that the proposed P-GAN network has the ability to transfer the image representations from exocentric to egocentric view, *i.e.,* objects are in the correct positions for generated egocentric images. Results show that egocentric images generated by P-GAN are visually much better compared with other baselines.

5. Conclusions

In this paper, we introduce a novel P-GAN which is able to learn shared information between cross-view images via a novel cross-cycle loss for a challenging exocentric to egocentric view image generation task. Moreover, we incorporate a novel contextual feature loss to capture the contextual information in images. Experimental results demonstrate that the hard-sharing of network layers in multi-directional parallel generative models can be used to increase the performance of cross-view image generation.

References

[1] O. Aghazadeh, J. Sullivan, and S. Carlsso. Novelty detection from an egocentric perspective. In *CVPR*, 2011. 2

[2] A. Bansal, S. Ma, D. Ramanan, and Y. Sheikh. Recycle-gan: Unsupervised video retargeting. In *ECCV*, 2018. 3

[3] C. Chan, S. Ginosar, T. Zhou, and A. A. Efros. Everybody dance now. In *ICCV*, 2019. 3

[4] Tong Che, Yanran Li, Athul Paul Jacob, Yoshua Bengio, and Wenjie Li. Mode regularized generative adversarial networks. In *ICLR*, 2017. 4

[5] Mohamed Elfeki, Krishna Regmi, Shervin Ardeshir, and Ali Borji. From third person to first person: Dataset and baselines for synthesis and retrieval. In *CVPR*, 2019. 1, 4

[6] D. Emily and F. Rob. Stochastic video generation with a learned prior. In *ICML*, 2018. 3

[7] A. Fathi, A. Farhadi, and J. M. Rehg. Understanding egocentric activities. In *ICCV*, 2011. 2

[8] A. Fathi, J. K. Hodgins, and J. M. Rehg. Social interactions: A first-person perspective. In *CVPR*, 2012. 2

[9] A. Fathi, Y. Li, and J. M. Rehg. Learning to recognize daily actions using gaze. In *ECCV*, 2012. 3

[10] I. Goodfellow, J.Pouget-Abadie, M. Mirza, B. Xu, D. W. Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 2

[11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 1, 4

[12] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 1, 2, 3, 4, 5, 6

[13] Takeo Kanade and Martial Hebert. First-person vision. *Proceedings of the IEEE*, 100(8):2442–2453, 2012. 1

[14] Guosheng Lin, Fayao Liu, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for dense prediction. *IEEE TPAMI*, 2019. 6

[15] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. RefineNet: Multi-path refinement networks for high-resolution semantic segmentation. In *CVPR*, 2017. 6

[16] Gaowen Liu, Hao Tang, Hugo Latapie, Jason Corso, and Yan Yan. Cross-view exocentric to egocentric video synthesis. In *ACM MM*, 2021. 3

[17] Gaowen Liu, Hao Tang, Hugo Latapie, and Yan Yan. Exocentric to egocentric image generation via parallel generative adversarial network. In *ICASSP*, 2020. 2

[18] Ming-Yu Liu and Oncel Tuzel. Coupled generative adversarial networks. In *NeurIPS*, 2016. 1, 3

[19] Shaohui Liu, Yi Wei, Jiwen Lu, and Jie Zhou. An improved evaluation framework for generative adversarial networks. *arXiv preprint arXiv:1803.07474*, 2018. 4

[20] Roey Mechrez, Itamar Talmi, and Lihi Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *ECCV*, 2018. 6

[21] K. Ogaki, K. M. Kitani, Y. Sugano, and Y. Sato. Coupling eye-motion and ego-motion features for first-person activity recognition. In *CVPR Workshops*, 2012. 2

[22] J. Pan, C. Wang, X. Jia, J. Shao, L. Sheng, J. Yan, and X.Wang. Video generation from single semantic label map. In *CVPR*, 2019. 3

[23] E. Park, J. Yang, E. Yumer, D. Ceylan, and A. C. Berg. Transformation-grounded image generation network for novel 3d view synthesis. In *CVPR*, 2017. 3

[24] H. Pirsiavash and D. Ramanan. Detecting activities of daily living in first-person camera views. In *CVPR*, 2012. 2

[25] Yair Poleg, Chetan Arora, and Shmuel Peleg. Temporal segmentation of egocentric videos. In *CVPR*, 2014. 2

[26] Krishna Regmi and Ali Borji. Cross-view image synthesis using conditional gans. In *CVPR*, 2018. 1, 2, 4, 5, 6

[27] K. Regmi and A. Borji. Cross-view image synthesis using geometry-guided conditional gans. *CVIU*, 2019. 2

[28] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *CVPR*, 2015. 3

[29] M. S. Ryoo and L. Matthies. First-person activity recognition: What are they doing to me? In *CVPR*, 2013. 2

[30] M. Saito, E. Matsumoto, and S. Saito. Temporal generative adversarial nets with singular value clipping. In *ICCV*, 2017. 3

[31] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NeurIPS*, 2016. 4

[32] T. R. Shaham, T. Dekel, and T. Michaeli. Singan: Learning a generative model from a single natural image. In *ICCV*, 2019. 2

[33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 3

[34] Hao Tang, Dan Xu, Gaowen Liu, Wei Wang, Nicu Sebe, and Yan Yan. Cycle in cycle generative adversarial networks for keypoint-guided image generation. In *ACM MM*, 2019. 4

[35] Hao Tang, Dan Xu, Nicu Sebe, Yanzhi Wang, Jason J. Corso, and Yan Yan. Multi-channel attention selection gan with cascaded semantic guidancefor cross-view image translation. In *CVPR*, 2019. 1, 2, 4, 5, 6

[36] Hao Tang, Dan Xu, Yan Yan, Philip HS Torr, and Nicu Sebe. Local class-specific and global image-level generative adversarial networks for semantic-guided scene generation. In *CVPR*, 2020. 2

[37] E. Taralova, F. De la Torre, and M. Hebert. Source constrained clustering. In *ICCV*, 2011. 2

[38] M. Tatarchenko, A. Dosovitskiy, and T. Brox. Multi-view 3d models from single images with a convolutional network. In *ECCV*, 2016. 3

[39] S. Tulyakov, Ming-Yu Liu, X. Yang, and J. Kautz. MoCoGAN: Decomposing motion and content for video generation. In *CVPR*, 2018. 3

[40] T. Wang, M. Liu, J. Zhu, G. Liu, A. Tao, J. Kautz, and B. Catanzaro. Video-to-video synthesis. In *NeurIPS*, 2018. 3

[41] M. Xu, C. Fan, Y. Wang, M. S Ryoo, and D. J Crandall. Identifying first-person camera wearers in third-person videos. In *CVPR*, 2017. 2

[42] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *ICCV*, 2017. 1

[43] R. Yonetani, K. M. Kitani, and Y. Sato. Visual motif discovery via first-person vision. In *ECCV*, 2016. 2

[44] T. Zhou, S. Tulsiani, W. Sun, J. Malik, and A. A. Efros. View synthesis by appearance flow. In *ECCV*, 2016. 3

[45] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networkss. In *ICCV*, 2017. 1, 2, 4, 5, 6