

Communication

# SS-CPGAN: Self-Supervised Cut-and-Pasting Generative Adversarial Network for Object Segmentation

Kunal Chaturvedi, Ali Braytee, Jun Li and Mukesh Prasad \* 

School of Computer Science, FEIT, University of Technology Sydney, Sydney, NSW 2007, Australia

\* Correspondence: mukesh.prasad@uts.edu.au

**Abstract:** This paper proposes a novel self-supervised based Cut-and-Paste GAN to perform foreground object segmentation and generate realistic composite images without manual annotations. We accomplish this goal by a simple yet effective self-supervised approach coupled with the U-Net discriminator. The proposed method extends the ability of the standard discriminators to learn not only the global data representations via classification (real/fake) but also learn semantic and structural information through pseudo labels created using the self-supervised task. The proposed method empowers the generator to create meaningful masks by forcing it to learn informative per-pixel and global image feedback from the discriminator. Our experiments demonstrate that our proposed method significantly outperforms the state-of-the-art methods on the standard benchmark datasets.

**Keywords:** generative adversarial networks; self-supervised learning; cut-and-paste; segmentation

## 1. Introduction

Generative adversarial networks (GANs) [1] have become a popular class of image synthesis methods due to their demonstrated ability to create high-dimensional samples with desired data distribution. The primary objective of GANs is to generate diverse, high-quality images while also ensuring the stability of GAN training [2,3]. GAN consists of generator and discriminator networks trained in an adversarial manner. The generator attempts to synthesize the real data distribution to fool the discriminator, whereas the discriminator's goal is to distinguish between the generator's real and fake data. In image segmentation, several compositional generative models have been proposed [4–9], where the generator creates a synthesized composite image by copying the object from one image and pasting it in another to fool the discriminator into thinking the synthesized composite image is real. However, the generator may not perform any segmentation, and the background may look realistic. Therefore, for effective training, the discriminator must provide the generator with informative learning signals by learning relevant semantics and structures of the data that may result in more effective generators. However, the current state-of-the-art GANs [9–13] employ discriminators based on the classification network, which learn only a single discriminative signal such as the difference between real and fake images. In such a non-stationary environment, the generator becomes prone to catastrophic forgetting and may lead to training instability or mode collapse [14].

To address the issues above, additional discriminatory signals are required to guide the training mechanism and assist the generator in producing high-quality images. This can be accomplished by increasing the capacity of the discriminator with auxiliary tasks and signals. These auxiliary tasks on the labeled datasets resist the forgetting issues and improve the training stability of GANs, but they suffer with unlabeled datasets. Recently, self-supervised learning has been explored on numerous GANs methods [14–17]. The self-supervised tasks provide the learning environment with additional guidance to the standard training mechanism. Most of the recent self-supervision methods on GANs use auxiliary tasks on transformation. For example, SS-GAN developed by Chen et al. [14] uses



**Citation:** Chaturvedi, K.; Braytee, A.; Li, J.; Prasad, M. SS-CPGAN: Self-Supervised Cut-and-Pasting Generative Adversarial Network for Object Segmentation. *Sensors* **2023**, *23*, 3649. <https://doi.org/10.3390/s23073649>

Academic Editors: Shengzong Zhou and Jingsha He

Received: 27 February 2023

Revised: 16 March 2023

Accepted: 29 March 2023

Published: 31 March 2023



**Copyright:** © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

rotation prediction as an auxiliary task. In FX-GAN, Huang et al. [16] use the pretext task of prediction on corrupted real images, and in LT-GAN [15], the authors use distinguishing GAN-induced transformation as a pretext task. However, the goals of these self-supervised transformation tasks need to be consistent with the GAN's goal of mimicking the real data distribution. Moreover, this problem amplifies when the generator's task is to construct segmentation masks from the foreground images.

Recent self-supervised learning methods have demonstrated remarkable promise in global tasks such as image classification by training simple classifiers on features learned through instance discrimination. However, the pre-text tasks of these global feature-learning approaches do not explicitly retain spatial information, making them unsuitable for object segmentation [18–21]. To maintain an enriched real data representation and improve the quality of generated segmentation masks, we propose a Self-Supervised Cut-and-Paste GAN(SS-CPGAN) using U-net architecture [22], which unifies cut-and-paste adversarial training with a self-supervised task. It allows the discriminator to learn local and global differences between real and fake data. In contrast to the existing transformation self-supervision methods, our self-supervision learning method creates pseudo labels using unsupervised segmentation methods. Then, it simultaneously forces the discriminator to provide the generator with global feedback (real or fake) and the per-pixel feedback of the synthesized images with the help of pseudo labels.

To sum up, the contributions of this paper are as follows:

- This paper proposes a novel Self-Supervised Cut-and-Paste GAN (SS-CPGAN), which unifies cut-and-paste adversarial training with a segmentation self-supervised task. SS-CPGAN leverages unlabeled data to maximize segmentation performance and generates highly realistic composite images.
- The proposed self-supervised task in SS-CPGAN improves the discriminator's representation ability by enhancing structure learning with global and local feedback. This enables the generator with additional discriminatory signals to achieve superior results and stabilize the training process.
- This paper comprehensively analyzes the benchmark datasets and compares the proposed method with the baseline methods.

## 2. Related Works

### 2.1. Unsupervised Object Segmentation via GANs

Unsupervised segmentation using GANs is an important topic in research. Several works [4–7] investigate the use of compositional generative models to obtain high-quality segmentation masks. Copy-pasting GAN [7] performs unsupervised object discovery by extracting foreground objects and then copying and pasting them onto different backgrounds. Similarly, PerturbGAN [5] generates a foreground mask along with a background and foreground image in an adversarial manner. Recently, Abdal et al. (2021) [6] proposed the use of an alpha network that includes two pre-trained generators and a discriminator on the StyleGAN to generate high-quality masks. These methods learn object segmentation without needing to use annotations. However, they are prone to degenerate solutions or other trivial cases. For example, the generator may not perform any segmentation, and the background looks realistic, or the generator may segment foreground masks consisting of all-ones. To avoid such problems, special care must be taken while training the compositional generative models. Copy-pasting GAN uses anti-shortcut, border-zeroing, blur, and grounded fakes to prevent trivial solutions [7]. PerturbGAN avoids such solutions by randomly shifting object segments relative to the background [5]. However, Abdal et al. (2021) [6] make several changes to the original StyleGAN and use a truncation trick along with regularization to avoid degenerate solutions. While these methods achieve object segmentation of foreground objects [23], the generated segmentation masks are often inferior in quality. Furthermore, due to such non-trivial procedures, the training of GANs becomes very challenging and complex [24].

## 2.2. Self-Supervised Learning

Self-supervised learning belongs to unsupervised learning, which learns useful feature representations from unlabeled data with the help of pretext tasks. It helps reduce the enormous data collection and annotation cost [25,26]. The traditional way to do this is to give the model some pretext tasks to solve. In this way, the networks learn good feature representations with the help of pseudo labels created by the pretext tasks [27]. Recently, many pretext tasks, and adversarial training, have been introduced [14–16,28,29]. The motivation for using self-supervised learning in GANs is to (1) prevent discriminator forgetting [30]; (2) improve training stability [31]; (3) and ensure the high quality of image generated [32]. The self-supervision techniques rely on pretext tasks on geometric transformations (e.g., prediction on rotated images [14], corrupted images [16], GAN-induced transformations [15], clustering representations [28], or a deshuffling task that predicts the shuffled orders [29]) to increase the discriminator's representation power. These self-supervised tasks may not work well for segmentation due to the inherent differences between the classification and segmentation tasks. In addition to this, image generation for segmentation requires GANs to capture contextual information between the foreground object and the background, which can be complicated in the absence of relevant visual representations. Unlike the aforementioned methods, we incorporate segmentation using self-supervised learning coupled with the Cut-and-Paste GAN to obtain high-quality segmentation masks. Most importantly, with our self-supervised approach, no extra care is needed to deal with the trivial solutions prevalent in compositional generative models.

## 3. Method

In this section, we first present the standard terminology of adversarial training and the encoder–decoder discriminator. We then introduce our SS-CPGAN method built upon the cut-and-paste adversarial training. The unified framework with the segmentation using self-supervised task encourages the generator to emphasize local and global structures while synthesizing masks.

### 3.1. Adversarial Training

We build a generative model in which the generator takes the foreground image as the input and generates a composite image using a combination of the predicted mask, the source foreground image, and the background image to fool the discriminator. Formally, we define the input foreground source image as  $I_f \in P_{data}$  and the background image as  $I_b \in P_{data}$ , where  $P_{data}$  denotes the set of input images. Now, we define a generator ( $G$ ) that is trained in an adversarial manner against the discriminator ( $D$ ). During the training process, the generator predicts a segmentation mask defined by  $m_g(I_f) \in [0, 1]$ . Then, using the predicted mask:  $m_g(I_f)$ ; the foreground source image:  $I_f$ ; and the resized background image:  $I_b$ , we define the composite image as follows:

$$I_C = m_g(I_f)I_f + (1 - m_g(I_f))I_b \quad (1)$$

The discriminator's objective is to classify the composite image as real or fake. As a result, the standard objective of the discriminator and the generator of the CPGAN is defined as follows:

$$\mathcal{L}_D = E[\log D(I_f) + \log(1 - D(I_C))] \quad (2)$$

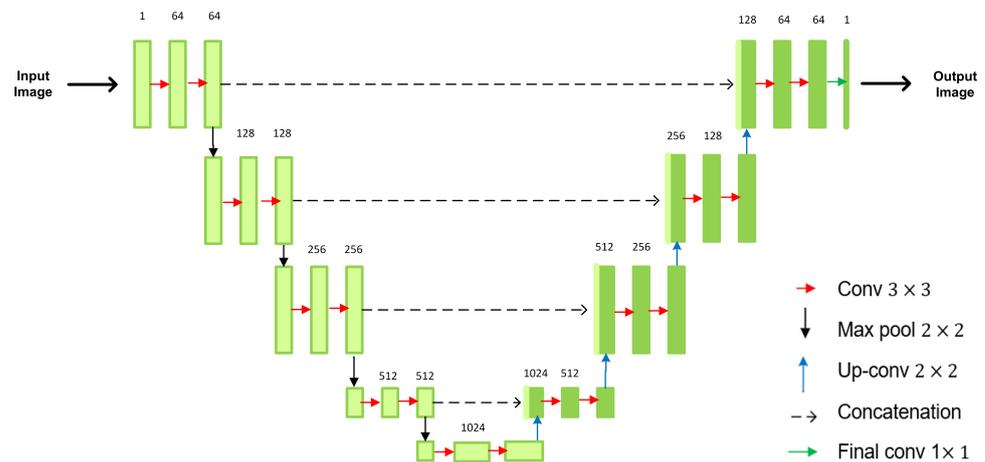
$$\mathcal{L}_G = -E[\log D(I_C)] \quad (3)$$

The discriminator works as a classification network restricted to learning only through the discriminative differences between the real and fake samples. Thus, the discriminator fails to provide any useful information to the generator. Therefore, we use an encoder–decoder discriminator network with self-supervised learning to mitigate this problem.

### 3.2. Encoder–Decoder Discriminator

In this work, we replace the standard classification discriminator with the U-net based discriminator. The U-net is an encoder–decoder architecture that consists of a network of convolutional layers, and skip connections for semantic segmentation [33–35]. It was initially proposed for biomedical image segmentation, which achieved precise segmentation results with few training images. Further, it demonstrates good results in other applications, including geo-sciences [36], remote sensing [37], and others. Its architecture (see Figure 1) is symmetric and consists of two paths: an encoder that extracts spatial features from the input image (downscaling process), and a decoder that constructs the segmentation map from the extracted feature maps (upscaling process). The use of U-net architecture in the proposed model adds major advantages: (1) it enables the simultaneous use of the global location and context while predicting masks; (2) it retains the full context of the input images, which is a significant advantage over patch segmentation approaches [38].

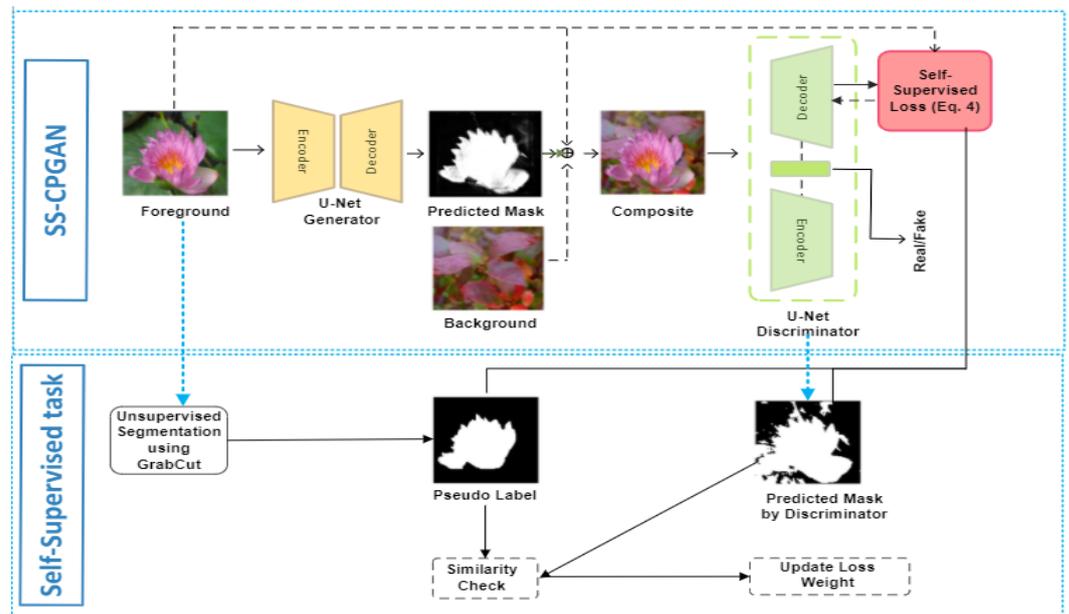
We use the encoder part of the U-net as the standard classification discriminator that performs the binary decision on real/fake composite images. Additionally, the decoder part of the U-net architecture is utilized by the self-supervised task to give per-pixel feedback on the synthesized images with the help of pseudo labels. This allows the discriminator to learn relevant local and global differences between real and fake images.



**Figure 1.** An overview of U-net architecture. The different arrows denote the different operations used in the encoder–decoder architecture.

### 3.3. Self-Supervised Cut-and-Paste GAN (SS-CPGAN)

To improve the representation learning ability of the CPGAN, the discriminator must learn semantic and structural information from the synthesized images. Therefore, we use self-supervised learning to build comprehensive representations for the CPGAN. In this work, we employ a segmentation self-supervised task to enable the discriminator with enhanced learned features that ultimately empower the generator to create consistent and structurally coherent masks. The pseudo segmentation masks  $m_{US}(I_f) \in [0, 1]$  are created using a graph unsupervised segmentation algorithm [39]. These masks obtained by the GrabCut technique act as a suitable prior for the U-net discriminator (see, Figure 2 (top)). Here, the discriminator performs two important tasks, i.e., (1) classification of real/fake compositing images; and (2) performing per-pixel classification on  $I_f \in P_{data}$  to generate segmentation masks. Given the self-supervised pseudo labels, we train the discriminator for accurate pixel-level prediction. Integrating self-supervisory signals empowers the discriminator by enhancing its localization ability and forces it to learn useful semantic representations. This mechanism enables the generator to achieve optimized results and makes the training process more stabilized.



**Figure 2.** The proposed self-supervised cut-and-paste GAN (SS-CPGAN).

Formally, we define  $I_f \in P_{data}$  as the source image containing the foreground object, and  $P_{data}$  denotes the set of input images. Further, we create a pseudo label denoted by  $m_{US}(I_f) \in [0, 1]$ , using an unsupervised segmentation algorithm. Then, we define  $m_w(I_C) \in [0, 1]$  as the pixel-wise segmentation mask produced by the decoder of the discriminator. Hereafter, we optimize the overall discriminator loss function (Equation (5)) by augmenting a new self-supervision loss (Equation (4))

$$\mathcal{L}_{self-supervised} = \mathcal{L}(m_w(I_C), 1 - m_{US}(I_f)) \quad (4)$$

$$\mathcal{L}'_D = \mathcal{L}_D + \lambda \mathcal{L}_{self-supervised} \quad (5)$$

where  $\mathcal{L}$  is the cross-entropy loss, and  $\lambda$  denotes the loss weight for the self-supervision loss. This hyperparameter is updated according to the comparison between  $m_{US}(I_f)$  and  $m_w(I_C)$ , using intersection-over-union (IoU). The details of the hyperparameter chosen are explained in the implementation details section. The framework of self-supervised learning is shown in Figure 2 (bottom).

#### 4. Experimentation

This section discusses the implementation details of the proposed method and an extensive set of experiments on various datasets.

##### 4.1. Datasets

We utilize five different datasets for the foreground and background set to train our SS-CPGAN as described below:

- Caltech-UCSD Birds (CUB) 200-2011 is a frequently used benchmark for unsupervised image segmentation. It consists of 11,788 images from 200 bird species.
- Oxford 102 Flowers consists of 8189 images from 102 flower classes.
- FGVC Aircraft (Airplanes) contains 102 different aircraft model variants with 100 images of each. This dataset was used initially for fine-grained visual categorization.
- MIT Places2 is a scene-centric dataset with more than 10 million images consisting of over 400 unique scene classes. However, in the experiments, we use the classes rain-forest, forest, sky, and swamp as a background set for the Caltech-UCSD Birds dataset, and in the Oxford 102 Flowers, we use the class: herb garden as a background set.

- Singapore Whole-sky IMaging CATEGORIES (SWIMCAT) contains 784 images of five categories: patterned clouds, clear sky, thick dark clouds, veil clouds, and thick white clouds. We use the SWIMCAT dataset as a background set for the FGCV dataset.

We chose background datasets similar to the background of the images from the foreground dataset. For the foreground datasets, we use Caltech-UCSD Birds (CUB) 200-2011 [40], Oxford 102 Flowers [41], and FGCV Aircraft (Airplanes) [42]. During the training, we do not utilize the masks available with datasets, Caltech-UCSD Birds, and Flowers-102. For the background datasets, we use MIT Places2 [43], and SWIMCAT [44].

#### 4.2. Experimental Settings

Our implementation used the PyTorch framework. For training our models, we deploy a batch size of 16 and the Adam optimizer with an initial learning rate of  $2 \times 10^{-4}$ . The training images are reshaped to  $64 \times 64$ ,  $128 \times 128$ , and  $256 \times 256$ . For the self-supervision task, we use the GrabCut technique [39] as the unsupervised segmentation algorithm.

#### 4.3. Hyper-Parameter Range

SS-CPGAN presents a new self-supervised loss, i.e.,  $\mathcal{L}_{self-supervised}$ , which needs to be validated. As shown in Table 1, we present Structural similarity (SSIM) scores according to different values of  $\lambda$ . The SSIM scores vary between the range of [0,1], with lower values indicating the lower quality of generated images. During the experiments, we find that the optimal values of the hyper-parameter can vary depending on the intersection-over-union (IoU) score between the pseudo label (mask) and the predicted mask by the discriminator. Initially, when  $\text{IoU} < 0.2$ , the hyperparameter value is set to 0.5 to boost the model's ability to learn useful representations from the pseudo label. When the  $0.2 < \text{IoU} < 0.8$ , we refine the predicted mask using the hyperparameter value  $\lambda$  of 0.1. To avoid the pseudo labels compromising the predicted masks, we restrict the value  $\lambda$  to 0 when the  $\text{IoU} > 0.8$ .

**Table 1.** Validation of hyper-parameter choices for  $\lambda$  in the self-supervised loss.

	SSIM $\uparrow$				
lambda	0.1	0.5	1	10	100
SSIM	0.657	0.834	0.450	0.421	0.125

#### 4.4. Results

We utilized the Fréchet inception distance (FID) score and mean Intersection over Union (mIoU) metric for the quantitative evaluation of our methods. In this work, we use the FID score on the datasets CUB2011, Oxford 102 Flowers, and FGCV Aircraft (see Table 2) to compare the SS-CPGAN model with the CPGAN model images spatially scaled to  $64 \times 64$ ,  $128 \times 128$ , and  $256 \times 256$ . For the datasets with available ground truth masks, including CUB2011, and Oxford 102 Flowers, we use the mIoU metric as shown in Table 3.

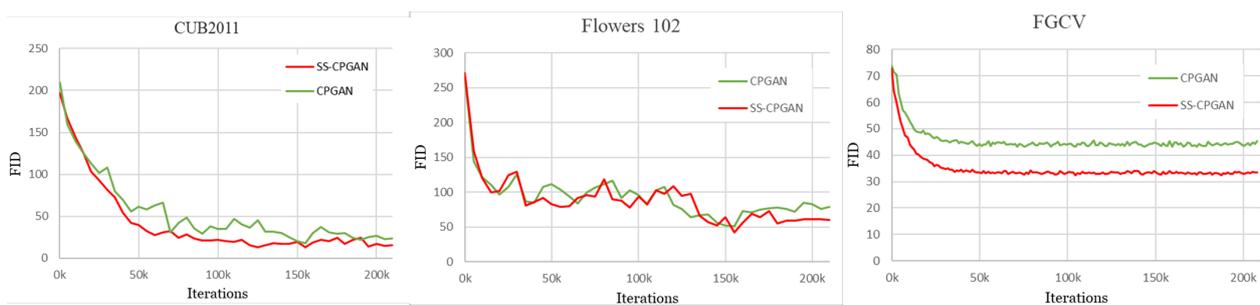
**Table 2.** FID comparison of the proposed method with the baseline CPGAN model.

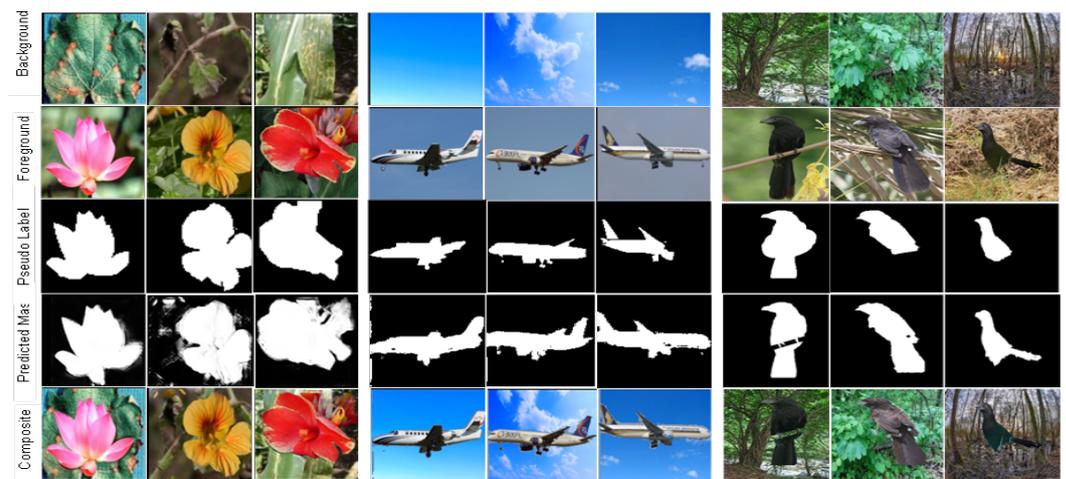
	FID $\downarrow$			
Methods	Image Size	Caltech UCSD-Bird 200	FGCV-Aircraft	Oxford 102 Flowers
CPGAN	$64 \times 64$	26.724	43.353	81.724
	$128 \times 128$	23.002	39.674	44.825
	$256 \times 256$	21.346	44.825	51.218
SS-CPGAN	$64 \times 64$	<b>22.342</b>	<b>39.578</b>	<b>63.343</b>
	$128 \times 128$	<b>15.634</b>	<b>37.756</b>	<b>54.982</b>
	$256 \times 256$	<b>13.113</b>	<b>33.149</b>	<b>49.181</b>

**Table 3.** mIoU comparison of the proposed method with the baseline CPGAN model.

mIoU ↑			
Methods	Image Size	Caltech UCSD-Bird 200	Oxford 102 Flowers
w/o Self-Supervision	$64 \times 64$	0.537	0.632
	$128 \times 128$	0.492	0.674
	$256 \times 256$	0.484	0.779
<b>Self-Supervision</b>	$64 \times 64$	<b>0.571</b>	<b>0.625</b>
	$128 \times 128$	<b>0.543</b>	<b>0.719</b>
	$256 \times 256$	<b>0.518</b>	<b>0.791</b>

In Figure 3, we report the FID scores over the training iterations. We show that our method stabilizes GAN training across all of the datasets by allowing GAN training to converge faster and consistently improve performance throughout the training. According to Figure 3, our method, SS-CPGAN, utilizing self-supervision outperforms the baseline method, CPGAN, on each dataset used. Furthermore, as shown in Figure 4, the generated masks and composite images of our proposed SS-CPGAN are of superior quality. The standard classification discriminator of CPGAN does not provide effective guidance to the generator. During the training, the standard discriminator is not encouraged to learn more robust data representation. The classification task learns only the representation based on the discriminative differences between real/fake images and fails to give information on why the synthesized image looks fake. Notably, our self-supervision task assigned to the U-net discriminator provides the generator with global feedback (real or fake) and per-pixel feedback of the masks with the help of pseudo labels. The self-supervisory signals prevent the two scenarios for the generator, which the standard discriminator fails to do, i.e., creating constant masks of only all-zeros pixel values or all-ones pixel values. The enhanced discriminator of SS-CPGAN influences the generator to create high quality masks that are devoid of any such anomalies. As shown in Figure 4, the qualitative analysis of the proposed SS-CPGAN shows that the generated masks and composite images are of superior quality.

**Figure 3.** FID training curves for CPGAN and SS-CPGAN on the datasets: CUB2011, Flowers 102, and FGCV).



**Figure 4.** Visualization results with the proposed SS-CPGAN on the datasets: Oxford 102 Flowers (**left**), FGVC Aircraft (**center**), and Caltech-UCSD Birds (CUB) 200-2011 (**right**).

#### 4.5. Comparison with the State-of-the-Art

We compare our self-supervision-based Cut-and-Paste GAN (SS-CPGAN) with state-of-the-art. As shown in Table 4, we report and compare the FID score on the Caltech UCSD-Bird 200 dataset. Specifically, the FID scores of StackGANv2 [45], OneGAN [46], LR-GAN [47], ELGAN [48], and FineGAN [49] are listed. The results in Table 4 show that our method delivers better performance and outperforms the existing methods. LR-GAN [47] performed the worst, followed by the other methods. The low performance of layer-wise GANs [47,48] is attributed to the fact that these methods are prone to degenerate during the training phase, with all the pixels being assigned as one component. In Table 5, we compare the performance of our method to the recent methods using the mIoU metric on Caltech UCSD-Bird 200 and Oxford flowers-102 respectively. In comparison to PerturbGAN [5], ContraCAM [50], ReDO [4], UISB [51], and IIC-seg [52], our method outperforms by a large margin on the Caltech UCSD-Bird 200 dataset. On the Oxford flowers-102 dataset, we perform better than the methods ReDO [4], Kyriazi et. al [53] and Voynov et. al. [54]. Here, ReDO and Kyriazi et. al (2021) are unsupervised approaches, whereas Voynov et. al (2021) is a weakly supervised approach to create segmentation maps. The ability to leverage pseudo labels in the training of Cut-and-Paste GAN assists in creating foreground masks of superior quality.

**Table 4.** FID comparison of our proposed method SS-CPGAN with the state-of-art on Caltech UCSD-Bird 200 dataset.

Method	FID
StackGANv2	21.4
FineGAN	23.0
OneGAN	20.5
LR-GAN	34.91
ELGAN	15.7
<b>SS-CPGAN</b>	<b>13.11</b>

**Table 5.** Quantitative comparison of the segmentation performance of our method SS-CPGAN with the state-of-art.

Dataset	Method	mIoU
Caltech UCSD-Bird 200	PerturbGAN	0.380
	ContraCAM	0.460
	ReDO	0.426
	UIB	0.442
	IIC-seg	0.365
	<b>SS-CPGAN</b>	<b>0.571</b>
Oxford 102 flowers	ReDO	0.764
	Kyriazi et. al.	0.541
	Voynov et al.	0.540
	<b>SS-CPGAN</b>	<b>0.791</b>

## 5. Conclusions

In this work, we proposed a novel Self-Supervised Cut-and-Paste GAN method to learn object segmentation. Specifically, we unified the cut-and-paste adversarial training with the proposed segmentation based self-supervision learning. Unlike the existing transformation self-supervised methods, our method improves the discriminator's representation ability by enhancing structure learning with global and local feedback from the synthesized masks. Furthermore, SS-CPGAN overcomes the issue of unwanted trivial solutions (generating constant masks of only all-zeros or all-ones pixel values) that plagues the generator. The experimental results show that our approach generates superior quality images and achieves promising results on the benchmark datasets.

**Author Contributions:** Conceptualization, K.C., A.B. and M.P.; methodology, K.C. and A.B.; software, K.C.; validation, K.C., A.B., J.L. and M.P.; formal analysis, K.C. and A.B.; investigation, K.C. and J.L.; resources, A.B. and M.P.; data curation, K.C. and A.B.; writing—original draft preparation, K.C.; writing—review and editing, K.C., A.B., J.L. and M.P.; visualization, A.B. and J.L.; supervision, A.B. and M.P.; project administration, K.C. and A.B.; and funding acquisition, J.L. and M.P. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The datasets presented in this study are openly available in Caltech-UCSD Birds (CUB) 200-2011 [40], Oxford 102 Flowers [41], FGCV Aircraft (Airplanes) [42], MIT Places2 [43], and SWIMCAT [44].

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative adversarial nets. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, USA, 8–13 December 2014; Volume 27.
- Brock, A.; Donahue, J.; Simonyan, K. Large scale GAN training for high fidelity natural image synthesis. *arXiv* **2018**, arXiv:1809.11096.
- Chaturvedi, K.; Braytee, A.; Vishwakarma, D.K.; Saqib, M.; Mery, D.; Prasad, M. Automated Threat Objects Detection with Synthetic Data for Real-Time X-ray Baggage Inspection. In Proceedings of the 2021 International Joint Conference on Neural Networks (IJCNN), Shenzhen, China, 18–22 July 2021; pp. 1–8.
- Chen, M.; Artières, T.; Denoyer, L. Unsupervised object segmentation by redrawing. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2019; Volume 32.
- Bielski, A.; Favaro, P. Emergence of object segmentation in perturbed generative models. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2019; Volume 32.
- Abdal, R.; Zhu, P.; Mitra, N.J.; Wonka, P. Labels4free: Unsupervised segmentation using stylegan. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, BC, Canada, 10–17 October 2021; pp. 13970–13979.
- Arandjelović, R.; Zisserman, A. Object discovery with a copy-pasting gan. *arXiv* **2019**, arXiv:1905.11369.

8. Zhang, Q.; Ge, L.; Hensley, S.; Isabel Metternicht, G.; Liu, C.; Zhang, R. PolGAN: A deep-learning-based unsupervised forest height estimation based on the synergy of PolInSAR and LiDAR data. *ISPRS J. Photogramm. Remote Sens.* **2022**, *186*, 123–139. [[CrossRef](#)]
9. Zhan, C.; Dai, Z.; Samper, J.; Yin, S.; Ershadnia, R.; Zhang, X.; Wang, Y.; Yang, Z.; Luan, X.; Soltanian, M.R. An integrated inversion framework for heterogeneous aquifer structure identification with single-sample generative adversarial network. *J. Hydrol.* **2022**, *610*, 127844. [[CrossRef](#)]
10. Zhou, G.; Song, B.; Liang, P.; Xu, J.; Yue, T. Voids Filling of DEM with Multiattention Generative Adversarial Network Model. *Remote Sens.* **2022**, *14*, 1206. [[CrossRef](#)]
11. Li, W.; Tang, Y.M.; Yu, K.M.; To, S. SLC-GAN: An automated myocardial infarction detection model based on generative adversarial networks and convolutional neural networks with single-lead electrocardiogram synthesis. *Inf. Sci.* **2022**, *589*, 738–750. [[CrossRef](#)]
12. Fu, L.; Li, J.; Zhou, L.; Ma, Z.; Liu, S.; Lin, Z.; Prasad, M. Utilizing Information from Task-Independent Aspects via GAN-Assisted Knowledge Transfer. In Proceedings of the 2018 International Joint Conference on Neural Networks (IJCNN), Rio de Janeiro, Brazil, 8–13 July 2018; pp. 1–6. [[CrossRef](#)]
13. Zhang, L.; Li, J.; Huang, T.; Ma, Z.; Lin, Z.; Prasad, M. GAN2C: Information Completion GAN with Dual Consistency Constraints. In Proceedings of the 2018 International Joint Conference on Neural Networks (IJCNN), Rio de Janeiro, Brazil, 8–13 July 2018; pp. 1–8. [[CrossRef](#)]
14. Chen, T.; Zhai, X.; Ritter, M.; Lucic, M.; Houlby, N. Self-supervised gans via auxiliary rotation loss. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 16–17 June 2019; pp. 12154–12163.
15. Patel, P.; Kumari, N.; Singh, M.; Krishnamurthy, B. Lt-gan: Self-supervised gan with latent transformation detection. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Virtual, 5–9 January 2021; pp. 3189–3198.
16. Huang, R.; Xu, W.; Lee, T.Y.; Cherian, A.; Wang, Y.; Marks, T. Fx-gan: Self-supervised gan learning via feature exchange. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Snowmass Village, CO, USA, 1–5 March 2020; pp. 3194–3202.
17. Hou, L.; Shen, H.; Cao, Q.; Cheng, X. Self-Supervised GANs with Label Augmentation. In Proceedings of the Advances in Neural Information Processing Systems, Online, 6–12 December 2021; Volume 34.
18. Shi, Y.; Xu, X.; Xi, J.; Hu, X.; Hu, D.; Xu, K. Learning to Detect 3D Symmetry From Single-View RGB-D Images With Weak Supervision. *IEEE Trans. Pattern Anal. Mach. Intell.* **2023**, *45*, 4882–4896. [[CrossRef](#)] [[PubMed](#)]
19. Li, Y.; Che, P.; Liu, C.; Wu, D.; Du, Y. Cross-scene pavement distress detection by a novel transfer learning framework. *Comput.-Aided Civ. Infrastruct. Eng.* **2021**, *36*, 1398–1415. [[CrossRef](#)]
20. Liu, Y.; Zhang, Z.; Liu, X.; Wang, L.; Xia, X. Efficient image segmentation based on deep learning for mineral image classification. *Adv. Powder Technol.* **2021**, *32*, 3885–3903. [[CrossRef](#)]
21. Dong, C.; Li, Y.; Gong, H.; Chen, M.; Li, J.; Shen, Y.; Yang, M. A Survey of Natural Language Generation. *ACM Comput. Surv.* **2022**, *55*, 173. [[CrossRef](#)]
22. Ronneberger, O.; Fischer, P.; Brox, T. U-net: Convolutional networks for biomedical image segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Munich, Germany, 5–9 October 2015; Springer: Berlin/Heidelberg, Germany, 2015; pp. 234–241.
23. Zhang, H.; Luo, G.; Li, J.; Wang, F.Y. C2FDA: Coarse-to-Fine Domain Adaptation for Traffic Object Detection. *IEEE Trans. Intell. Transp. Syst.* **2022**, *23*, 12633–12647. [[CrossRef](#)]
24. Yang, B.; Gu, S.; Zhang, B.; Zhang, T.; Chen, X.; Sun, X.; Chen, D.; Wen, F. Paint by Example: Exemplar-based Image Editing with Diffusion Models. *arXiv* **2022**, arXiv:2211.13227.
25. Xie, B.; Li, S.; Lv, F.; Liu, C.H.; Wang, G.; Wu, D. A Collaborative Alignment Framework of Transferable Knowledge Extraction for Unsupervised Domain Adaptation. *IEEE Trans. Knowl. Data Eng.* **2022**, *Early Access*. [[CrossRef](#)]
26. Dang, W.; Guo, J.; Liu, M.; Liu, S.; Yang, B.; Yin, L.; Zheng, W. A Semi-Supervised Extreme Learning Machine Algorithm Based on the New Weighted Kernel for Machine Smell. *Appl. Sci.* **2022**, *12*, 9213. [[CrossRef](#)]
27. Ericsson, L.; Gouk, H.; Loy, C.C.; Hospedales, T.M. Self-Supervised Representation Learning: Introduction, advances, and challenges. *IEEE Signal Process. Mag.* **2022**, *39*, 42–62. [[CrossRef](#)]
28. Feng, J.; Zhao, N.; Shang, R.; Zhang, X.; Jiao, L. Self-Supervised Divide-and-Conquer Generative Adversarial Network for Classification of Hyperspectral Images. *IEEE Trans. Geosci. Remote. Sens.* **2022**, *60*, 5536517. [[CrossRef](#)]
29. Baykal, G.; Unal, G. DeshuffleGAN: A self-supervised gan to improve structure learning. In Proceedings of the 2020 IEEE International Conference on Image Processing (ICIP), Abu Dhabi, United Arab Emirates, 25–28 October 2020; pp. 708–712.
30. Thanh-Tung, H.; Tran, T. Catastrophic forgetting and mode collapse in GANs. In Proceedings of the 2020 International Joint Conference on Neural Networks (IJCNN), Glasgow, UK, 19–24 July 2020; pp. 1–10.
31. Mao, Q.; Lee, H.Y.; Tseng, H.Y.; Ma, S.; Yang, M.H. Mode seeking generative adversarial networks for diverse image synthesis. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 1429–1437.
32. Tran, N.T.; Tran, V.H.; Nguyen, B.N.; Yang, L.; Cheung, N.M.M. Self-supervised gan: Analysis and improvement with multi-class minimax game. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, USA, 8–13 December 2019; Volume 32.

33. Xie, B.; Li, S.; Li, M.; Liu, C.; Huang, G.; Wang, G. SePiCo: Semantic-Guided Pixel Contrast for Domain Adaptive Semantic Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2023**, 1–17. [[CrossRef](#)]
34. Yang, D.; Zhu, T.; Wang, S.; Wang, S.; Xiong, Z. LFRSNet: A robust light field semantic segmentation network combining contextual and geometric features. *Front. Environ. Sci.* **2022**, *10*, 1443. [[CrossRef](#)]
35. Sheng, H.; Cong, R.; Yang, D.; Chen, R.; Wang, S.; Cui, Z. UrbanLF: A Comprehensive Light Field Dataset for Semantic Segmentation of Urban Scenes. *IEEE Trans. Circuits Syst. Video Technol.* **2022**, *32*, 7880–7893. [[CrossRef](#)]
36. Chen, Y.; Wei, Y.; Wang, Q.; Chen, F.; Lu, C.; Lei, S. Mapping post-earthquake landslide susceptibility: A U-Net like approach. *Remote Sens.* **2020**, *12*, 2767. [[CrossRef](#)]
37. Tran, L.A.; Le, M.H. Robust U-Net-based road lane markings detection for autonomous driving. In Proceedings of the 2019 International Conference on System Science and Engineering (ICSSE), Dong Hoi, Vietnam, 20–21 July 2019; pp. 62–66.
38. Huang, H.; Lin, L.; Tong, R.; Hu, H.; Zhang, Q.; Iwamoto, Y.; Han, X.; Chen, Y.W.; Wu, J. UNet 3+: A Full-Scale Connected UNet for Medical Image Segmentation. In Proceedings of the ICASSP 2020—2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Barcelona, Spain, 4–8 May 2020; pp. 1055–1059. [[CrossRef](#)]
39. Rother, C.; Kolmogorov, V.; Blake, A. “GrabCut” interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph. (TOG)* **2004**, *23*, 309–314. [[CrossRef](#)]
40. Welinder, P.; Branson, S.; Mita, T.; Wah, C.; Schroff, F.; Belongie, S.; Perona, P. Caltech-UCSD Birds 200. 2010. Available online: [https://www.vision.caltech.edu/datasets/cub\\_200\\_2011/](https://www.vision.caltech.edu/datasets/cub_200_2011/) (accessed on 4 November 2020).
41. Nilsback, M.E.; Zisserman, A. Automated Flower Classification over a Large Number of Classes. 2008. Available online: <https://www.robots.ox.ac.uk/~vgg/data/flowers/102/> (accessed on 4 November 2020).
42. Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M.; Vedaldi, A. Fine-Grained Visual Classification of Aircraft. 2013. Available online: <https://www.robots.ox.ac.uk/~vgg/data/fgvc-aircraft/> (accessed on 4 November 2020).
43. Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; Torralba, A. Places: A 10 Million Image Database for Scene Recognition. 2017. Available online: <http://places2.csail.mit.edu/download.html> (accessed on 4 November 2020).
44. Dev, S.; Lee, Y.H.; Winkler, S. Categorization of Cloud Image Patches Using an Improved Texton-Based Approach. 2015. Available online: <https://stefan.winkler.site/Publications/icip2015cat.pdf> (accessed on 4 November 2020).
45. Zhang, H.; Xu, T.; Li, H.; Zhang, S.; Wang, X.; Huang, X.; Metaxas, D.N. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *41*, 1947–1962. [[CrossRef](#)]
46. Benny, Y.; Wolf, L. Onegan: Simultaneous unsupervised learning of conditional image generation, foreground segmentation, and fine-grained clustering. In *Proceedings of the European Conference on Computer Vision, Glasgow, UK, 23–28 August 2020*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 514–530.
47. Yang, J.; Kannan, A.; Batra, D.; Parikh, D. Lr-gan: Layered recursive generative adversarial networks for image generation. *arXiv* **2017**, arXiv:1703.01560.
48. Yang, Y.; Bilen, H.; Zou, Q.; Cheung, W.Y.; Ji, X. Unsupervised Foreground-Background Segmentation with Equivariant Layered GANs. *arXiv* **2021**, arXiv:2104.00483.
49. Singh, K.K.; Ojha, U.; Lee, Y.J. Finegan: Unsupervised hierarchical disentanglement for fine-grained object generation and discovery. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–19 June 2019; pp. 6490–6499.
50. Mo, S.; Kang, H.; Sohn, K.; Li, C.L.; Shin, J. Object-aware contrastive learning for debiased scene representation. In Proceedings of the Advances in Neural Information Processing Systems, Online, 6–12 December 2021; Volume 34.
51. Kim, W.; Kanezaki, A.; Tanaka, M. Unsupervised learning of image segmentation based on differentiable feature clustering. *IEEE Trans. Image Process.* **2020**, *29*, 8055–8068. [[CrossRef](#)]
52. Ji, X.; Henriques, J.F.; Vedaldi, A. Invariant information clustering for unsupervised image classification and segmentation. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Republic of Korea, 27 October–2 November 2019; pp. 9865–9874.
53. Melas-Kyriazi, L.; Rupprecht, C.; Laina, I.; Vedaldi, A. Finding an unsupervised image segmenter in each of your deep generative models. *arXiv* **2021**, arXiv:2105.08127.
54. Voynov, A.; Morozov, S.; Babenko, A. Object segmentation without labels with large-scale generative models. In Proceedings of the International Conference on Machine Learning, PMLR, Virtual, 18–24 July 2021; pp. 10596–10606.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.