

Scalable Multi-Temporal Remote Sensing Change Data Generation via Simulating Stochastic Change Process

Zhuo Zheng^{1,2}, Shiqi Tian¹, Ailong Ma¹, Liangpei Zhang¹, Yanfei Zhong^{1,*}
¹Wuhan University ²Stanford University

<https://github.com/Z-Zheng/Changen>




Figure 1. This village is disappearing, which is a stochastic change process simulated by our generative probabilistic change model (GPCM). This is also synthetic multi-temporal change data produced by our change generator (Changen), a GAN-based GPCM, enabling object change data generation with controllable object property (*e.g.*, scale, position, orientation), and change event.

Abstract

*Understanding the temporal dynamics of Earth’s surface is a mission of multi-temporal remote sensing image analysis, significantly promoted by deep vision models with its fuel—labeled multi-temporal images. However, collecting, preprocessing, and annotating multi-temporal remote sensing images at scale is non-trivial since it is expensive and knowledge-intensive. In this paper, we present a scalable multi-temporal remote sensing change data generator via generative modeling, which is cheap and automatic, alleviating these problems. Our main idea is to simulate a stochastic change process over time. We consider the stochastic change process as a probabilistic semantic state transition, namely generative probabilistic change model (GPCM), which decouples the complex simulation problem into two more trackable sub-problems, *i.e.*, change event simulation and semantic change synthesis. To solve these two problems, we present the change generator (Changen), a GAN-based GPCM, enabling controllable object change data generation, including customizable object property, and change event. The extensive experiments suggest that our Changen has superior generation capability, and the change detectors with Changen pre-training exhibit excellent transferability to real-world change datasets.*

* corresponding author (zhongyanfei@whu.edu.cn)

1. Introduction

Change detection is one of the most fundamental Earth vision tasks to understand the temporal dynamics of Earth’s surface. Tremendous progress in change detection has been achieved by joint efforts of remote sensing and computer vision communities. Deep change detection models [7,39,40,42] represented by Siamese networks [3] have dominated in recent years. The key behind their success lies in large-scale training datasets [6,27,31,32]. However, building a large-scale remote sensing change detection dataset is difficult and expensive, because collecting, preprocessing, annotating remote sensing images needs more expertise and efforts.

Synthetic data, as an alternative, is a promising direction to alleviate the data hungry. There are currently two main paradigms for change detection data synthesis in the remote sensing domain, *i.e.*, graphics-based [2,19] and data augmentation-based [4] approaches. The graphics-based methods synthesize images by rendering manual constructed 3D models. IAug [4], as a data augmentation-based approach, synthesizes new image pairs by pasting object instances into existing bitemporal image pairs.

However, the scalability and diversity of these conventional synthetic change datasets remain limited. This is because 3D modeling and rendering need expertise and tremendous efforts for graphics-based approaches, and an existing

change dataset is needed for data augmentation-based approaches. Besides, the synthetic data is aimed to assist change detection model in improving the performance on real-world data. However, the relationship between the quality of synthetic data and the transferability of features learned from synthetic data is still unclear due to scale-limitation of graphics-based approaches and the coupling of synthetic and real-world data of data-augmentation based approaches.

In this paper, we present a scalable multi-temporal change data generator via generative modeling. Our data generator is aimed to generate realistic and diverse multi-temporal change-detection images from a single-temporal image and its semantic segmentation mask, by *simulating the change process*.

To this end, we first describe the stochastic change process as a probabilistic graphical model, namely *generative probabilistic change model* (GPCM), considering each image and semantic mask as random variables, as shown in Fig. 2. The change is always driven by the event, therefore we make a Markov assumption that the image and its semantic mask at time $t + 1$ only depends on the image and its semantic mask at time t , to simplify this modeling problem. Based on this condition, the whole problem can be decoupled to two sub-problems, *i.e.*, the change event simulation at semantic mask-level and the semantic change synthesis at image-level.

To solve above two sub-problems, we propose a generative model called *Changen*, which is a GPCM parameterized with generative adversarial networks (GANs). Our *Changen* creates or removes objects in the semantic mask at time t as a stochastic change event, to generate a new semantic mask at time $t + 1$, and synthesizes a post-event image at time $t + 1$, by progressively applying simulated semantic changes to the pre-event image at time t .

As demonstration, our *Changen* generates a large-scale building change detection dataset with diverse object properties (e.g., scale, position, orientation) and two change events. The change detector pre-trained on this dataset has the superior transferability on the real-world building change detection datasets, significantly outperforming commonly used ImageNet [8] pre-training, and additionally possessing zero-shot prediction capability. More importantly, based on our *Changen*, we find that the temporal diversity of synthetic change data is a key factor in ensuring transferability after model pre-training. The main contributions of this paper are summarized as follows:

- **Generative change modeling** decouples the complex stochastic change process simulation to more tractable change event simulation and semantic change synthesis.
- **Change generator**, *i.e.*, *Changen*, enables object change generation with controllable object property (e.g., scale, position, orientation), and change event.
- **Our synthetic change data pre-training** empowers the change detectors with better transferability and zero-shot prediction capability.




Figure 2. Generative Probabilistic Change Model.

2. Related Work

Change Data Synthesis. Based on the computer graphics, the early studies mainly use game engines to generate realistic remote sensing images from existing assets. For example, AICD dataset [2], consists of 1,000 pairs of 800×600 images that are automatically generated by the built-in rendering engine of a computer game. However, due to limited assets and an underdeveloped rendering engine, this dataset is of low diversity and graphics quality. To improve the diversity and realism of synthesized images, a semi-automatic data generation pipeline [19] is proposed to synthesize a change detection dataset, which adopts cartographic data for manual 3D scene modeling and the professional game engine for rendering. Another way is based on data augmentation, especially Copy-Paste [10], *e.g.*, IAUG [4] adopts SPADE [23] to generate object instances and randomly pastes them into an existing change detection dataset to increase training samples. Our work provides a new perspective, *i.e.*, generative modeling for change detection data synthesis. Benefiting from a deep generative model, our dataset generator is automatic and relies only on mono-temporal data, ensuring scalability and economy of synthetic data.

Semantic Image Synthesis. Recent studies [23, 26, 28, 29, 35, 45] mainly focus on translating the semantic segmentation label to the image based on different conditional GANs [21], pix2pix [15] and pix2pixHD [35], as early representative methods, directly use semantic mask to generate the image. SPADE [23] reveals that normalization layers tend to “wash away” semantic information, and propose a spatially adaptive conditional normalization layer to incorporate semantic information and spatial layout. The following studies further improve SPADE on region-level control ability [45], diversity [29], better trade-off between fidelity and diversity [28]. These methods strongly depend on VGG-based perceptual loss [9, 17] to guarantee image quality, however, this increases the complexity of the whole pipeline. To simplify GAN models for semantic image synthesis, OASIS [26] reveals that the segmentation-based discriminator is the key to

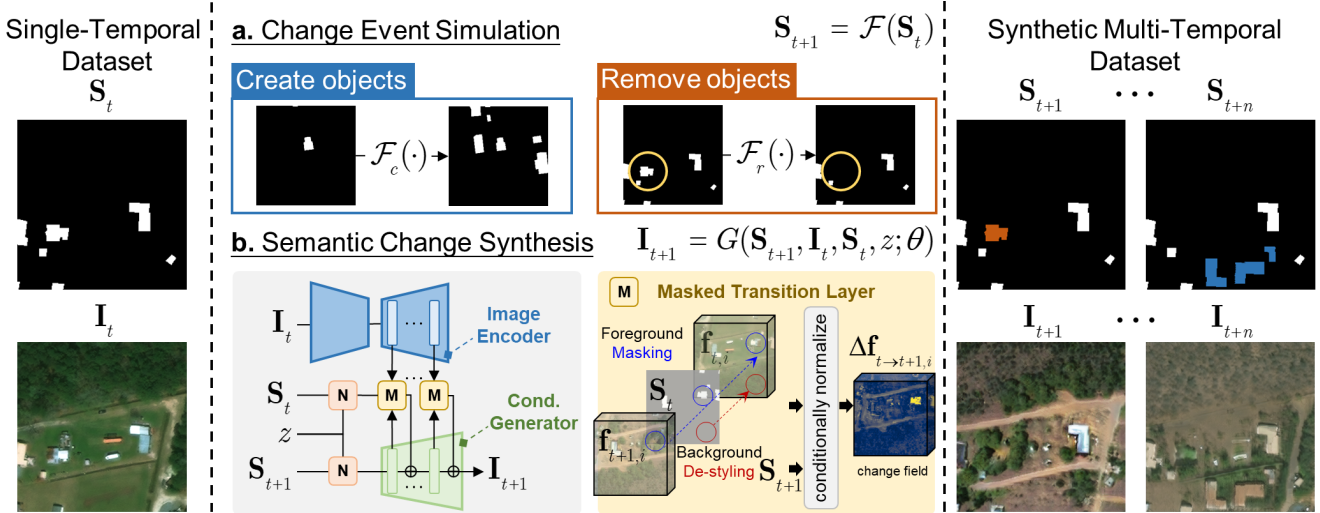


Figure 3. **Our Chagen architecture.** The change event simulation enables adding or removing objects in the semantic mask at time t to customize new semantic masks at times $t + 1$ to n . For the semantic change synthesis, the new images at times $t + 1$ to n will be synthesized by progressively accumulating estimated change fields on the image at time t . Our Chagen can generate the multi-temporal dataset with controllable scene layout, object property (e.g., scale, position, orientation, see \mathbf{I}_{t+n}), and change event. Legend: **Create**; **Remove**.

enable GAN model to synthesize high-quality images with only adversarial supervision. Our work extends conventional semantic image synthesis in temporal dimensionality, enabling semantic time-series synthesis from single-temporal data and with better consistency on the spatial layout.

3. Generative Probabilistic Change Model

The main idea of our dataset generator is to simulate stochastic change process started from each single-temporal image and its semantic mask. To this end, we frame the stochastic change process as a probabilistic graphical model shown in Fig. 2, to describe the relationship between the distributions of variables (i.e., images, semantic masks) over time. In this way, we can parameterize several smaller factors instead of directly parameterizing high-dimensionality joint distribution over all variables, simplifying this simulation problem. Based on this modeling framework, we provide a parameterization with a deep generative model, to synthesize the multi-temporal change detection dataset from single-temporal data.

3.1. Generative Change Modeling Framework

The image \mathbf{I}_t has a pre-defined pixel-wise semantics \mathbf{S}_t at time t . Given a stochastic change event occurred in the time period of t to $t + 1$, the image \mathbf{I}_t gradually evolves into the image \mathbf{I}_{t+1} with a semantic transition from \mathbf{S}_t to \mathbf{S}_{t+1} . The joint distribution of this stochastic change process is denoted as $P_{\text{scp}} := P(\mathbf{S}_{t+1}, \mathbf{S}_t, \mathbf{I}_{t+1}, \mathbf{I}_t)$. Here, we make a Markov assumption that the post-event image-semantics pair $(\mathbf{I}_{t+1}, \mathbf{S}_{t+1})$ only depends on the pre-event

image-semantics pair $(\mathbf{I}_t, \mathbf{S}_t)$, because the change is always event-driven. Based on this assumption, P_{scp} can be factorized as $P(\mathbf{S}_{t+1}, \mathbf{I}_{t+1} | \mathbf{S}_t, \mathbf{I}_t) P(\mathbf{S}_t, \mathbf{I}_t)$. Since the simulation is a generative task, i.e., semantics-to-images, the graph structure of the set of these four random variables can be described as shown in Fig. 2. The following factorization can be obtained:

$$P_{\text{scp}} = P(\mathbf{I}_{t+1} | \mathbf{S}_{t+1}, \mathbf{I}_t) P(\mathbf{S}_{t+1} | \mathbf{S}_t) P(\mathbf{I}_t | \mathbf{S}_t) P(\mathbf{S}_t) \quad (1)$$

where the semantics distribution $P(\mathbf{S}_t)$ and corresponding conditional image distribution $P(\mathbf{I}_t | \mathbf{S}_t)$ can be approximately seen as two known distributions because the single-temporal data is given for this simulation problem. This means that the semantic transition distribution $P(\mathbf{S}_{t+1} | \mathbf{S}_t)$ and the conditional image distribution $P(\mathbf{I}_{t+1} | \mathbf{S}_{t+1}, \mathbf{I}_t)$ need to be further estimated, so as to sample from P_{scp} .

Through above modeling, the whole simulation problem can be decoupled to two sub-problems, i.e., the *change event simulation* to approximate $P(\mathbf{S}_{t+1} | \mathbf{S}_t)$ and the *semantic change synthesis* to approximate $P(\mathbf{I}_{t+1} | \mathbf{S}_{t+1}, \mathbf{I}_t)$.

3.2. Chagen: GPCM parameterized with GANs

To solve above two sub-problems, we present a parameterization for GPCM with the learning-free function $\mathcal{F}(\cdot)$ for change event simulation and the well-designed deep generative model G for semantic change synthesis, as follows:

$$\mathbf{S}_{t+1} = \mathcal{F}(\mathbf{S}_t) \quad (2)$$

$$\mathbf{I}_{t+1} = G(\mathbf{S}_{t+1}, \mathbf{I}_t, \mathbf{S}_t, z; \theta), z \sim \mathcal{N}(\mathbf{0}, \mathbf{1}) \quad (3)$$

Fig. 3 illustrates the whole architecture, introduced next.

3.2.1 Change Event Simulation

Sampling from $P(\mathbf{S}_{t+1}|\mathbf{S}_t)$ is to obtain \mathbf{S}_{t+1} as a semantic guidance for the subsequent semantic change synthesis. We consider two common change events, *i.e.*, constructing new objects and destroying existing objects in real-world. From the perspective of image editing, we refer to these two events as creating objects and removing objects. To simulate these two events in semantic level, we design two learning-free functions, *i.e.*, $\mathcal{F}_c(\cdot) : R^{h \times w} \rightarrow R^{h \times w}$ for object creation and $\mathcal{F}_r(\cdot) : R^{h \times w} \rightarrow R^{h \times w}$ for object removal. In detail, we first randomly select some instances from semantic segmentation labels. The selected instances are either pasted on the rest area of the semantic segmentation label to simulate object creation or directly assigned the background label to simulate object removal, as shown in Fig. 3a.

3.2.2 Semantic Change Synthesis

Sampling from $P(\mathbf{I}_{t+1}|\mathbf{S}_{t+1}, \mathbf{I}_t)$ is to obtain the post-event image \mathbf{I}_{t+1} . To this ends, we design a deep generative model as shown in Fig. 3b, which has an image encoder that maps the pre-event image to multi-scale latent representations and a conditional decoder that progressively generates the post-event image from the latent representations, the post-event semantic representation, and Gaussian noise. Besides, the masked transition layers inserted between the image encoder and the conditional decoder, is to reassemble their feature maps for avoiding mode collapse caused by a trivial solution. **Image Encoder.** We adopt the U-Net [24] with ResNet-18 [13] as the image encoder to obtain multi-scale latent representations $\{\mathbf{f}_{t,i}\}_{i=0}^5$ of the pre-event image, where $\mathbf{f}_{t,i}$ has a size of $\frac{512}{2^i} \times \frac{h}{2^{(5-i)}} \times \frac{w}{2^{(5-i)}}$, and h, w are the image size. Spectral normalization [22] is applied to all conv layers.

Conditional Decoder. Our decoder is stacked by six residual blocks with spatially-adaptive group normalization [23, 36] for consistent training and inference. The decoder takes one-hot encoded post-event semantic mask concatenated with a 3D Gaussian noise map [26] in channel axis as input and yields multi-scale post-event feature maps $\{\mathbf{f}_{t+1,i}\}_{i=0}^5$ at time $t + 1$. The first feature map $\mathbf{f}_{t+1,0}$ is obtained by computing a 3×3 conv layer over $32 \times$ downsampled input. The remaining post-event feature map $\mathbf{f}_{t+1,i+1}$ is computed as follows:

$$\Delta \mathbf{f}_{t \rightarrow t+1,i} = \mathbf{M}_i(\mathbf{f}_{t,i}, \mathbf{f}_{t+1,i}, \mathbf{S}_t, \mathbf{S}_{t+1}) \quad (4)$$

$$\mathbf{f}_{t+1,i+1} = G_i(\mathbf{f}_{t+1,i} + \Delta \mathbf{f}_{t \rightarrow t+1,i}, \mathbf{S}_{t+1}, z) \quad (5)$$

where the change field $\Delta \mathbf{f}_{t \rightarrow t+1}$ is first computed by the masked transition layer \mathbf{M} (described next in detail). The each spatial position of the change field is a change vector encoding change information from the pre-event image to the post-event image. These change fields are progressively applied to intermediate feature maps of our decoder, and i -th result is fed to the residual block G_i for further conditional

Algorithm 1 Bitemporal Adversarial Learning

Require:

- $\mathcal{D}_t \sim P(\mathbf{I}_t, \mathbf{S}_t)$: Single-temporal dataset
 - N : Number of iterations
 - 1: **for** $n = 1$ **to** N **do**
 - ▷ random mini-batch sampling
 - 2: $\mathbf{I}_t, \mathbf{S}_t \sim \mathcal{D}_t$
 - ▷ stochastic change event simulation
 - 3: $\mathbf{S}_{t+1} \leftarrow \mathcal{F}(\mathbf{S}_t)$
 - ▷ semantic change synthesis
 - 4: $\tilde{\mathbf{I}}_{t+1} \leftarrow G(\mathbf{S}_{t+1}, \mathbf{I}_t, \mathbf{S}_t, z; \theta), z \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$
 - ▷ update the generator once
 - 5: $\min_{\theta} \mathbb{E}_{\mathbf{v}_g} [D(\tilde{\mathbf{I}}_{t+1}, \mathbf{S}_{t+1}; \phi)]$
 - ▷ update the discriminator once (ours)
 - 6: $\max_{\phi} \mathbb{E}_{\mathbf{v}_g} [D(\tilde{\mathbf{I}}_{t+1}, \mathbf{S}_{t+1}; \phi)] + \mathbb{E}_{\mathbf{I}_t} [D(\mathbf{I}_t, \mathbf{S}_t; \phi)]$
 - ▷ update the discriminator once (classical, unavailable)
 - $\max_{\phi} \mathbb{E}_{\mathbf{v}_g} [D(\tilde{\mathbf{I}}_{t+1}, \mathbf{S}_{t+1}; \phi)] + \mathbb{E}_{\mathbf{I}_{t+1}} [D(\mathbf{I}_{t+1}, \mathbf{S}_{t+1}; \phi)]$
 - 7: **return** θ
-

$\mathbf{v}_g := (z, \mathbf{S}_{t+1}, \mathbf{I}_t)$. For simplicity, we ignore the notation of the classification score converter for the discriminator D , which is identical to [26].

synthesis. The tanh activation is applied to the final feature map to obtain the synthesized post-event image. To stable training, the spectral normalization is also applied in our decoder. The layer \mathbf{N} is the noise injection via `concat`.

Masked Transition Layer. With the guidance of the pre-event image, the synthesis of the post-event image can be easier. However, this also causes a trivial solution, *i.e.*, the generative model directly copies the content from the pre-event image rather than synthesizing visual content, yielding completely incorrect change information. To avoid this “feature leakage”, we propose the masked transition layer, as illustrated in Fig. 3b, to yield a change field to progressively update the post-event feature map rather than to produce it directly.

Our masked transition layer first masks¹ pre-event foreground region features with the post-event feature maps and de-styles² the pre-event background region features for explicitly weakening over-emphasized pre-event features. We further inject the post-event semantic information into the weakened pre-event features via a spatially-adaptive normalization conditioned by the post-event semantic mask, yielding a change field.

Bitemporal Adversarial Learning. Since there is no ground-truth for the image to be generated, we propose the bitemporal adversarial learning to train our *Changen*, as

¹Pseudo code of *Masking* in a PyTorch-like style:
`mask = (S_t == fg); return torch.where(mask, f_{t+1,i}, f_{t,i})`

²For simplicity, we adopt a 1×1 conv followed by instance normalization [33] and a Leaky ReLU for De-styling.

illustrated in Algorithm 1. Like all adversarial learning for semantic image synthesis, we have a conditional discriminator to determine whether the image-mask pair is real or fake. Following OASIS [26], we adopt a segmentation-based discriminator D with spectral normalization and train our Changen with only conditional adversarial losses. Unlike classical conditional adversarial losses to optimize the discriminator, we use the pair $(\mathbf{I}_t, \mathbf{S}_t)$ as the real target instead of the pair $(\mathbf{I}_{t+1}, \mathbf{S}_{t+1})$ because there is no ground truth for the image \mathbf{I}_{t+1} , as presented in the line 6 of Algorithm 1.

4. Experiments

4.1. Experimental Setting

Single-temporal dataset. We use the pre-disaster part of xView2 dataset [11] as our single-temporal dataset, called *xView2 pre-disaster*. The `train` and `tier3` splits are used to train generative models. The `hold` split, which is cropped into 256×256 non-overlapped patches, is used to evaluate the quality of synthesized images.

Training generative models. We use Adam [18] with $\beta_1 = 0$ and $\beta_2 = 0.999$ as our optimizer for both the generator and the discriminator. The learning rate is 0.0001 for the generator and 0.0004 for the discriminator. We use a mini-batch size of 32 in 8 GPUs. We train all generative models for 100k iterations with the same data and settings. Random flip, rotation, transpose, scale jitter, cropping into 256×256 are adopted for training data augmentation.

Change dataset generation. We sample pre-event image-mask pairs from the xView2 pre-disaster, resulting $\sim 90k$ pairs. Changen is applied to these pre-event pairs for generating post-event pairs, resulting a change dataset. Each pre-event mask can be mapped to $n \in \{1, 2, \dots\}$ post-event masks, therefore, we refer to the change dataset with n as `Changen- $\{n \times 90\}k$` , which means that this dataset consists of $n \times 90k$ synthetic bitemporal training sample pairs. We use `Changen-90k` as default for fast experiments.

Pre-training on synthetic dataset. We use ChangeStar [39], a simple multi-task change detection architecture, with an optimized configuration of FarSeg [41, 43] with 96 channels, ChangeMixin with $N = 1, d_c = 96$, and ResNet-18 backbone, which denotes ChangeStar (1×96), as default for simplicity. We pre-train the change detector for 20 epochs on each synthetic dataset if not specified. SGD with a momentum of 0.9 and a weight decay of 0.0001 is used as our optimizer. We use a mini-batch size of 64 with an initial learning rate of 0.03 and multiplied by $(1 - \frac{\gamma}{N})^\gamma$ with $\gamma = 0.9$. Random flip, rotation, transpose, color jitter are adopted for training data augmentation.

4.2. Experiments on Change Data Generation

Setup. We train all deep generative models on xView2 pre-disaster and evaluate the quality of synthetic images with FID score [14] and Inception score (IS) [25] following common

Table 1. Comparison with other change data generation approaches on the image quality.

Method	Modeling	FID↓	IS↑
SPADE [23] + GPCM	$P(\mathbf{I}_t \mathbf{S}_t)$	204.01	3.41
OASIS [26] + GPCM		45.13	4.95
Changen (ours)	$P(\mathbf{I}_{t+1} \mathbf{S}_{t+1}, \mathbf{I}_t)$	34.74	5.41

Table 2. Comparison with other change data generation approaches on the transferability. The reported IoU and F_1 are obtained by pre-training ChangeStar with synthetic data, fine-tuning it on 5% LEVIR-CD^{train} and testing it on LEVIR-CD^{test}, thereby measuring the transferability.

Method	Modeling	IoU↑	F_1 ↑
SPADE [23] + GPCM	$P(\mathbf{I}_t \mathbf{S}_t)$	-	-
OASIS [26] + GPCM		75.9	86.3
Copy-Paste	$P(\mathbf{I}_{t+1} \mathbf{S}_{t+1}, \mathbf{I}_t)$	74.3	85.2
Inpainting		74.6	85.5
Changen (ours)		(ours)	79.3

practices. We do supervised pre-training on each synthetic dataset and then fine-tune the change detection model on 5% training set of LEVIR-CD [6] to measure the transferability.

Main Result. Table 1 and 2 presents the comparison with other deep generative models and data-augmentation-based synthesis methods. For deep generative models, we compare with classical SPADE [23] and state-of-the-art OASIS [26]. SPADE is hard to generate high-quality images in our settings due to complex global disaster scenarios in xView2, while OASIS yields a FID score of 45.13 and a fine-tuned F_1 score of 86.3%. Thanks to our GPCM modeling, our Changen outperforms OASIS both on the image quality and the transferability, yielding a FID score of 34.74 and 88.4% F_1 score. Besides, OASIS yields stochastic spatial layout due to the absence of temporal information, while our Changen yields more coherent and realistic spatial layout with the pre-event image, as shown in Fig. 4. This suggests that the pre-event image guidance is important for the image quality and the transferability if carefully exploited.

For data augmentation-based synthesis approaches, we extend Copy-Paste as [4, 10] in our framework for a fair comparison. Besides, similar to Copy-Paste (creating objects), we also extend an inpainting methods [30, 38] (removing objects) as strong baselines, called Inpainting and DeepFillv2. Since there is only single-temporal data available, above two methods yield bitemporal images with highly consistent background pixels, as shown in Fig. 4a-b, only generating new and effective positive samples for change detection. Our Changen produces not only more realistic positive samples but also diverse negative samples via carefully designed masked transition layer to handle the pre-event image guidance, significantly surpassing them by large margins of $\sim 3\%$ fine-tuned F_1 score, respectively.

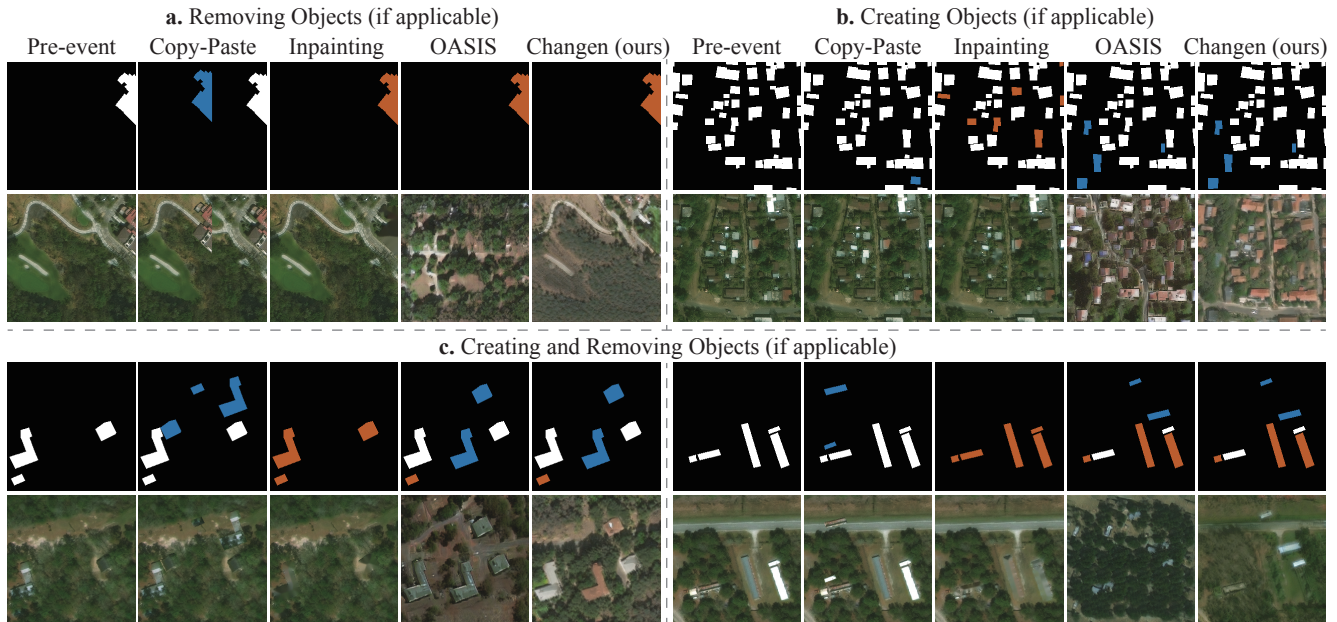


Figure 4. **Visual comparison of change data synthesis results.** Each pre-event column presents a real pre-event image-mask pair. Other columns present a synthesized post-event image-mask pair for each method. **a.** and **b.** simulate object removal and object creation, respectively. **c.** simulates more complex scenarios that include object removal and object creation simultaneously. Legend: **Create**; **Remove**.

Table 3. **Ablation study on Masked Transition Layer**, investigating the impact of each component on the image quality and the transferability.

Method	FID↓	IS↑	IoU↑	F ₁ ↑
Masked Transition Layer	34.74	5.41	79.3	88.4
w/o De-styling	37.49	5.20	79.1	88.3
w/o De-styling and Masking	20.38	5.00	74.3	85.2

Ablation: Masked Transition Layer. To yield an accurate change field, *Masking* and *De-styling* operations are two core designs for foreground and background feature weakening, respectively. Table 3 presents that without *De-styling*, the image quality degenerates slightly (+2.75 FID), but the transferability is hardly impacted (-0.1% F₁). This suggests that (i) the original style of the pre-event image background region may perturb the image generation, and (ii) the new image style generated from noise can guarantee diverse background regions in the synthesized image for sufficient negative samples of change.

Removing *Masking* further, a serious drop of the transferability (-3.2% F₁ score) and an abnormal image quality improvement (-14.36 FID score) are observed because of the trivial solution (“feature leakage”). This suggests that (i) *Masking* is an indispensable component in the masked transition layer, and (ii) correct positive samples of change is essential to guarantee transferability. This is because the trivial solution results in two approximately identical images with different semantic masks, yielding incorrect positive samples of change, while *Masking* explicitly corrects these positive samples, as shown in Fig. 5.




Figure 5. **Ablation study on the Masked Transition Layer.** “M” denotes the masked transition layer. Without M, the model trends to copy everything from the pre-event image, failing to remove objects, while with M, the model can successfully remove objects and maintain the spatial layout of the pre-event image.

4.3. Experiments on Transferability

Setup. We investigate the transferability using two protocols: (i) zero-shot change evaluation, *i.e.*, the change detector is first pre-trained on synthetic data and then directly tested on the real-world data, and (ii) end-to-end fine-tuning. The transferability is measured by the model accuracy (*e.g.*, F₁ score) for each synthetic data. We use ChangeStar (FarSeg) with ResNet-18 as the change detector.

Zero-shot Change Evaluation. Table 4 presents the results on the entire LEVIR-CD and WHU-CD [16] datasets, respectively, where these two datasets include building object creation and removal events, resulting in two change types. Pre-training on the synthetic change dataset generated by Copy-Paste yields collapsing zero-shot performance, which suggests that there is an insuperable domain gap between the

Table 4. **Zero-shot Change Evaluation for Object Change Detection** results on LEVIR-CD^{all} and WHU-CD^{all}. “C” and “R” denote creating objects and removing objects, respectively.

Pre-train from	Supported event type	LEVIR-CD ^{all}				WHU-CD ^{all}			
		IoU	F ₁	Prec.	Rec.	IoU	F ₁	Prec.	Rec.
Copy-Paste	C	0.7	1.4	1.9	1.2	1.9	3.8	2.4	9.0
Inpainting	R	10.8	19.5	11.9	54.0	12.2	21.8	13.3	59.3
OASIS + GPCM	C&R	39.7	56.8	45.5	75.7	12.7	22.6	13.8	61.3
Changen (ours)	C&R	45.8	62.8	49.3	86.4	15.3	26.6	15.7	87.1

Table 5. **Transfer Learning on LEVIR-CD**. Fine-tuning on LEVIR-CD^{train} with different training sample ratios and evaluating on LEVIR-CD^{test}. “C” and “R” denote creating objects and removing objects, respectively.

Pre-train from	supported event type	5%LEVIR-CD ^{train}				100%LEVIR-CD ^{train}			
		IoU	F ₁	Prec.	Rec.	IoU	F ₁	Prec.	Rec.
scratch	-	65.4	79.1	80.3	77.9	79.4	88.5	89.8	87.3
ImageNet sup.	-	74.3	85.2	85.9	84.5	82.7	90.5	91.9	89.2
Copy-Paste	C	74.3	85.2	86.7	83.8	83.2	90.8	92.2	89.5
Inpainting	R	74.6	85.5	87.7	83.4	83.2	90.8	92.0	89.6
DeepFillv2	R	74.8	85.6	87.5	83.7	82.7	90.5	91.9	89.2
Changen (ours)	C&R	79.3	88.4	90.2	86.7	83.7	91.1	92.0	90.2

changes synthesized by Copy-Paste and real-world changes in these two cases at least. We argue that objects pasted at random positions of the image can be considered as perturbations to this high-dimensional data point, which, if not properly constrained, can move this data point away from the local manifold of the natural image. Different from Copy-Paste, Inpainting synthesizes post-event images via removing objects and then filling removed regions with its own local statistics. The perturbations from removal are reduced by the constraints from self-filling, thereby yielding non-collapsing zero-shot performance.

Our Changen significantly outperforms the above two methods by large margins. Because the perturbations are overcome by the prior from generative models as a constraint. This prior guarantees that the synthesized image is in training data distribution as much as possible.

Fine-tuning on LEVIR-CD. We fine-tune ChangeStar (1×96) for 200 epochs following common practices [5, 6, 39] (see supplement for training details). Table 5 shows that our Changen yields the synthetic change dataset with superior transferability over other approaches under both 5% and 100% training set. Compared with commonly used ImageNet pre-training, our Changen pre-training significantly improves transferability by 3.2%/0.6% F₁ under 5% and full training set, respectively. This property is very beneficial for those application scenarios where real-world data is limited. Compared with Copy-Paste, Inpainting, and DeepFillv2, which can be seen the approaches with perfect temporal coherence, our Changen produces more realistic changes and diverse background, *i.e.*, a better trade-off between temporal diversity and temporal coherence. This suggests that temporal diversity is a key factor in ensuring the transferability of Changen pre-training.

Table 6. Comparison with the state-of-the-art change detectors on LEVIR-CD^{test}. “R-18”: ResNet-18. † indicates that the backbone is pre-trained on ImageNet and then ADE20K [44]. * indicates their modified backbone. The amount of floating point operations (Flops) was computed with a float32 tensor of shape [2,256,256,3].

Method	Pre-train from	Backbone	F ₁ †	#Params.	#Flops
• ConvNet-based					
FC-Siam-Diff [7]	-	-	63.0	1.3M	4.7G
STANet [6]	ImageNet-1K	R-18	87.3	16.9M	19.6G
CDNet+IAug [4]	ImageNet-1K	R-18	89.0	14.3M	-
BiT [5]	ImageNet-1K	R-18	89.3	11.9M	8.6G
ChangeStar [39]	ImageNet-1K	R-18	90.2	19.3M	22.3G
ChangeStar [39]	ImageNet-1K	R-50	90.8	33.9M	29.0G
ChangeStar [39]	ImageNet-1K	RX-101	91.2	52.5M	39.2G
ChangeStar (1×96)	ImageNet-1K	R-18	90.5	16.4M	16.3G
+ self-supervised	SeCo-1M [20]	R-18	89.9	16.4M	16.3G
+ seg. supervised	xView2 pre-disaster	R-18	90.6	16.4M	16.3G
+ synthetic data	OASIS-90k	R-18	90.6	16.4M	16.3G
Ours	Changen-90k	R-18	91.1(†0.6)		+0
◦ Transformer-based					
ChangeFormer [1]	IN-1K, ADE20K†	MiT-B2*	90.4	41.0M	203.1G
BiT [5, 34]	ImageNet-1K	ViTAEv2-S	91.2	19.6M	15.7G
ChangeStar (1×96)	ImageNet-1K	MiT-B1	90.0	18.4M	16.0G
Ours	Changen-90k	MiT-B1	91.5(†1.5)		+0

We also compare ChangeStar (1×96) with our Changen pre-training to the state-of-the-art change detection methods and three pre-training baselines in Table 6. For ConvNet-based change detectors, with R-18 backbone, our Changen pushes ChangeStar further, outperforming other methods with similar complexity, on par with vanilla ChangeStar with RX-101 backbone. For Transformer-based change detectors, our method with MiT-B1 backbone outperforms BiT with ViTAEv2, setting a new record of 91.5% F₁. Besides, it is observed that Transformer backbones (*e.g.*, MiT [37]) benefits more from our synthetic Changen-90k, improving by 1.5% F₁. Compared to three strong pre-training baselines (OASIS synthetic data pre-training, segmentation data (xView2 pre-disaster) pre-training, and self-supervised pre-training), our Changen pre-training is superior since it considers temporal dynamics and has zero pretext task gap.

Fine-tuning on S2Looking. This dataset is much more challenging because of large geometry offset, radiation differences, tiny change regions, domain gaps between geographic regions [27]. We fine-tune ChangeStar (1×96) for 60k iterations (see supplement for training details). Table 7 shows a consistent observation that Changen pre-training enables the change detector to learn stronger transferable representation. Changen pre-training outperforms ImageNet pre-training by 3.0%/0.8% F₁ under 5% and full training set. This suggests that the gain from Changen pre-training is hardly affected by the complexity of real-world data, which confirms the robustness of Changen pre-training.

We further compare ChangeStar (1×96) with our

Table 7. **Transfer Learning on S2Looking**. Fine-tuning on S2Looking^{train} with different training sample ratios and evaluating on S2Looking^{test}. “C” and “R” denote creating objects and removing objects, respectively.

Pre-train from	supported event type	5%S2Looking ^{train}				100%S2Looking ^{train}			
		IoU	F ₁	Prec.	Rec.	IoU	F ₁	Prec.	Rec.
scratch	-	27.8	43.5	65.5	32.6	48.3	65.2	69.9	61.0
ImageNet sup.	-	33.2	49.9	66.1	40.0	49.6	66.3	70.9	62.2
Copy-Paste	C	34.0	50.8	68.3	40.4	50.0	66.7	69.4	64.1
Inpainting	R	34.0	50.8	69.1	40.1	49.6	66.3	69.6	63.4
DeepFillv2	R	34.2	51.0	66.8	41.2	50.0	66.7	71.6	62.4
Changen (ours)	C&R	36.0	52.9	65.9	44.6	50.5	67.1	70.1	64.3

Table 8. Comparison with the state-of-the-art change detectors on the S2Looking^{test} set. “R-18”: ResNet-18. The amount of floating point operations (Flops) was computed with a float32 tensor of shape [2,512,512,3] as input.

Method	Backbone	F ₁ ↑	Prec.	Rec.	#Params.	#Flops
• ConvNet-based						
FC-Siam-Diff [7]	-	13.1	83.2	15.7	1.3M	18.7G
STANet [6]	R-18	45.9	38.7	56.4	16.9M	156.7G
CDNet [4]	R-18	60.5	67.4	54.9	14.3M	-
BiT [5]	R-18	61.8	72.6	53.8	11.9M	34.7G
ChangeStar (1×96)	R-18	66.3	70.9	62.2	16.4M	65.3G
+ Changen-90k	R-18	67.1	70.1	64.3	16.4M	65.3G
		↑0.8	↓0.8	↑2.1		+0
◦ Transformer-based						
ChangeStar (1×96)	MiT-B1	64.3	69.3	59.9	18.4M	67.3G
+ Changen-90k	MiT-B1	67.9	70.3	65.7	18.4M	67.3G
		↑3.6	↑1.0	↑5.8		+0

Changen pre-training to the state-of-the-art methods in change detection in Table 8. We can see that ChangeStar (1×96) with R-18 or MiT-B1, as our strong baselines pre-trained on ImageNet, have outperformed other methods by large margins. Nevertheless, our Changen pre-training still surpasses the ImageNet pre-training by 0.8%/3.6% F₁ for R-18 and MiT-B1 backbones, respectively.

There are two interesting observations: (i) Transformer-based ChangeStar with ImageNet pre-training is inferior to ConvNet-based when they have similar complexity. With Changen pre-training, Transformer-based ChangeStar obtains a significant improvement (64.3→67.9), overtaking ConvNet-based one; (ii) The improvements from Changen pre-training are mainly from higher recall gains for both ConvNet-based and Transformer-based ChangeStar.

Based on above two observations and consistent results on LEVIR-CD, we argue that (i) large model capacity guarantees the transferability of Changen pre-training. This is because even though these two models have similar complexity, the model capacity of the transformer is considered to be larger [12] due to the lack of inductive bias. (ii) temporal diversity is a key factor in ensuring transferability since it brings rich negative samples that are also important to learn a generalizable change representation, in addition to necessary positive samples.

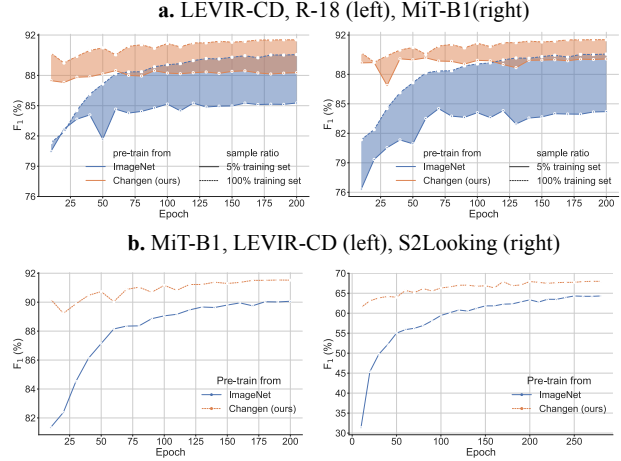


Figure 6. **Learning curves** of ChangeStar (1×96) with different pre-training datasets and backbones, respectively. **a.** the accuracy gap is represented by the filled region between solid and dashed lines. **b.** Changen pre-training brings much faster convergence than ImageNet pre-training.

Understand Changen Pre-training. To understand why the synthetic change dataset yielded by Changen works well on the paradigm of pre-training and fine-tuning, we plot several learning curves in Fig. 6 to investigate the behavior of each model. From Fig. 6a, we can observe that the accuracy gap between the models trained with 5% and 100% training set is significantly reduced by our Changen pre-training, compared to ImageNet pre-training. Besides, our Changen pre-training yields higher average and peak accuracies in different backbones and training sample ratios. To our surprise, ChangeStar pre-trained on Changen-90k and fine-tuned on 5% training set is on-par with it pre-trained on ImageNet and fine-tuned on 100% when using MiT-B1 backbone. From Fig. 6b, the model with Changen pre-training has much faster convergence than with ImageNet pre-training. These results point out that Changen pre-training can find task-specific starting points for deep change detectors, resulting in faster and better convergence. Transformer-based change detectors can benefit more from Changen pre-training.

5. Conclusion

We present a scalable multi-temporal remote sensing change data generator to alleviate “data hungry” problem for deep change detectors. Our work enables continuous and controllable multi-temporal change data generation via simulating stochastic change process. This data generation is easy to extend and scale up via new object properties, and change events. Our large-scale synthetic change dataset, Changen-90k, is confirmed its superior transferability, making conventional deep change detectors great again, e.g., for the ChangeStar with MiT-B1 backbone, Changen pre-training significantly outperforms ImageNet pre-training and achieves new records on LEVIR-CD and S2Looking datasets.

We also find that temporal diversity is a key factor in ensuring the transferability of Changen pre-training. These two properties also bring two open problems: the trade-off between fidelity and diversity in image generation, and the trade-off between zero-shot capability and transferability in model generalization, which deserve to be explored in the future.

Acknowledgements

This work was supported by National Key Research and Development Program of China under Grant No. 2022YFB3903404, National Natural Science Foundation of China under Grant No. 42325105, 42071350, and LIES-MARS Special Research Funding.

References

- [1] Wele Gedara Chaminda Bandara and Vishal M. Patel. A transformer-based siamese network for change detection. In *IGARSS*, pages 207–210, 2022.
- [2] Nicolas Bourdis, Denis Marraud, and Hichem Sahbi. Constrained optical flow for aerial image change detection. In *2011 IEEE International Geoscience and Remote Sensing Symposium*, pages 4176–4179. IEEE, 2011.
- [3] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a” siamese” time delay neural network. In *NeurIPS*, 1994.
- [4] Hao Chen, Wenyuan Li, and Zhenwei Shi. Adversarial instance augmentation for building change detection in remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–16, 2021.
- [5] Hao Chen, Zipeng Qi, and Zhenwei Shi. Remote sensing image change detection with transformers. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2021.
- [6] Hao Chen and Zhenwei Shi. A spatial-temporal attention-based method and a new dataset for remote sensing image change detection. *Remote Sensing*, 12(10):1662, 2020.
- [7] Rodrigo Caye Daudt, Bertr Le Saux, and Alexandre Boulch. Fully convolutional siamese networks for change detection. In *ICIP*, pages 4063–4067. IEEE, 2018.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009.
- [9] Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. In *NeurIPS*, volume 29, 2016.
- [10] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, pages 2918–2928, 2021.
- [11] Ritwik Gupta, Richard Hosfelt, Sandra Sajeev, Nirav Patel, Bryce Goodman, Jigar Doshi, Eric Heim, Howie Choset, and Matthew Gaston. xbd: A dataset for assessing building damage from satellite imagery. *arXiv preprint arXiv:1911.09296*, 2019.
- [12] Kai Han, Yunhe Wang, Hanting Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chun-jing Xu, Yixing Xu, et al. A survey on vision transformer. *IEEE TPAMI*, 2022.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [14] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, volume 30, 2017.
- [15] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, pages 1125–1134, 2017.
- [16] Shunping Ji, Shiqing Wei, and Meng Lu. Fully convolutional networks for multisource building extraction from an open aerial and satellite imagery data set. *IEEE Transactions on Geoscience and Remote Sensing*, 57(1):574–586, 2018.
- [17] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711. Springer, 2016.
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2014.
- [19] Maria Kolos, Anton Marin, Alexey Artemov, and Evgeny Burnaev. Procedural synthesis of remote sensing images for robust change detection with neural networks. In *International Symposium on Neural Networks*, pages 371–387. Springer, 2019.
- [20] Oscar Manas, Alexandre Lacoste, Xavier Giró-i Nieto, David Vazquez, and Pau Rodriguez. Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data. In *ICCV*, pages 9414–9423, 2021.
- [21] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [22] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.
- [23] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, pages 2337–2346, 2019.
- [24] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241. Springer, 2015.
- [25] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NeurIPS*, volume 29, 2016.
- [26] Edgar Schönfeld, Vadim Sushko, Dan Zhang, Juergen Gall, Bernt Schiele, and Anna Khoreva. You only need adversarial supervision for semantic image synthesis. In *ICLR*, 2021.
- [27] Li Shen, Yao Lu, Hao Chen, Hao Wei, Donghai Xie, Jiabao Yue, Rui Chen, Shouye Lv, and Bitao Jiang. S2looking: A satellite side-looking dataset for building change detection. *Remote Sensing*, 13(24):5094, 2021.
- [28] Yupeng Shi, Xiao Liu, Yuxiang Wei, Zhongqin Wu, and Wangmeng Zuo. Retrieval-based spatially adaptive normalization for semantic image synthesis. In *CVPR*, pages 11224–11233, 2022.

- [29] Zhentao Tan, Menglei Chai, Dongdong Chen, Jing Liao, Qi Chu, Bin Liu, Gang Hua, and Nenghai Yu. Diverse semantic image synthesis via probability distribution modeling. In *CVPR*, pages 7962–7971, 2021.
- [30] Alexandru Telea. An image inpainting technique based on the fast marching method. *Journal of graphics tools*, 9(1):23–34, 2004.
- [31] Shiqi Tian, Yanfei Zhong, Zhuo Zheng, Ailong Ma, Xicheng Tan, and Liangpei Zhang. Large-scale deep learning based binary and semantic change detection in ultra high resolution remote sensing imagery: From benchmark datasets to urban application. *ISPRS Journal of Photogrammetry and Remote Sensing*, 193:164–186, 2022.
- [32] Aysim Toker, Lukas Kondmann, Mark Weber, Marvin Eisenberger, Andrés Camero, Jingliang Hu, Ariadna Pregel Hoderlein, Çağlar Şenaras, Timothy Davis, Daniel Cremers, et al. DynamicEarthNet: Daily multi-spectral satellite dataset for semantic change segmentation. In *CVPR*, pages 21158–21167, 2022.
- [33] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- [34] Di Wang, Jing Zhang, Bo Du, Gui-Song Xia, and Dacheng Tao. An empirical study of remote sensing pretraining. *IEEE Transactions on Geoscience and Remote Sensing*, 2022.
- [35] 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 *CVPR*, pages 8798–8807, 2018.
- [36] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, pages 3–19, 2018.
- [37] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. SegFormer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, volume 34, pages 12077–12090, 2021.
- [38] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *ICCV*, pages 4471–4480, 2019.
- [39] Zhuo Zheng, Ailong Ma, Liangpei Zhang, and Yanfei Zhong. Change is Everywhere: Single-temporal supervised object change detection in remote sensing imagery. In *ICCV*, pages 15193–15202, 2021.
- [40] Zhuo Zheng, Yanfei Zhong, Shiqi Tian, Ailong Ma, and Liangpei Zhang. ChangeMask: Deep multi-task encoder-transformer-decoder architecture for semantic change detection. *ISPRS Journal of Photogrammetry and Remote Sensing*, 183:228–239, 2022.
- [41] Zhuo Zheng, Yanfei Zhong, Junjue Wang, and Ailong Ma. Foreground-aware relation network for geospatial object segmentation in high spatial resolution remote sensing imagery. In *CVPR*, pages 4096–4105, 2020.
- [42] Zhuo Zheng, Yanfei Zhong, Junjue Wang, Ailong Ma, and Liangpei Zhang. Building damage assessment for rapid disaster response with a deep object-based semantic change detection framework: From natural disasters to man-made disasters. *Remote Sensing of Environment*, 265:112636, 2021.
- [43] Zhuo Zheng, Yanfei Zhong, Junjue Wang, Ailong Ma, and Liangpei Zhang. FarSeg++: Foreground-aware relation network for geospatial object segmentation in high spatial resolution remote sensing imagery. *IEEE TPAMI*, 2023.
- [44] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *IJCV*, 127(3):302–321, 2019.
- [45] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. SEAN: Image synthesis with semantic region-adaptive normalization. In *CVPR*, pages 5104–5113, 2020.