Single-temporal Supervised Remote Change Detection for Domain Generalization

Qiangang Du¹^{*}, Jinlong Peng²^{*}, Xu Chen², Qingdong He², Liren He¹, Qiang Nie², Wenbing Zhu¹, Mingmin Chi¹, Yabiao Wang², and Chengjie Wang²

 $^1\,$ Fudan University, Shanghai, China $^2\,$ Tencent Youtu Lab

Abstract. Change detection is widely applied in remote sensing image analysis. Existing methods require training models separately for each dataset, which leads to poor domain generalization. Moreover, these methods rely heavily on large amounts of high-quality pair-labelled data for training, which is expensive and impractical. In this paper, we propose a multimodal contrastive learning (ChangeCLIP) based on visuallanguage pre-training for change detection domain generalization. Additionally, we propose a dynamic context optimization for prompt learning. Meanwhile, to address the data dependency issue of existing methods, we introduce a single-temporal and controllable AI-generated training strategy (SAIN). This allows us to train the model using a large number of single-temporal images without image pairs in the real world, achieving excellent generalization. Extensive experiments on series of real change detection datasets validate the superiority and strong generalization of ChangeCLIP, outperforming state-of-the-art change detection methods. Code will be available.

Keywords: domain generalization, remote change detection, single-temporal supervised learning

1 Introduction

Remote Sensing Change Detection (RSCD) detects changes in ground objects using bi-temporal remote sensing imagery. It is widely used and provides important reference value for urban planning, land monitoring, post-disaster assessment and construction.

The existing mainstream methods primarily rely on designing convolutional neural networks (CNNs) to improve prediction accuracy through network structure optimization. These methods can be categorized into supervised change detection [1,3,4,8,16,29] and unsupervised change detection [7,14,26]. These supervised methods often rely on training with large-scale and high-quality labeled bi-temporal image pairs, which are time-consuming and expensive to acquire. Unsupervised methods alleviate the dataset dependency and can be further divided into (a) self-supervised learning paradigm and (b) pretraining paradigm,

^{*} Equal contribution.



Fig. 1: (a) Self-supervised paradigm [7, 20, 25, 38], trained by contrastive learning or make pseudo image-pair by simple image augmentation. (b) Pretraining paradigm [26, 27, 30, 39], training encoder on single-temporal datasets and predicting by features analysis. (c) ChangeCLIP: Mutimodal contrastive learning based on single-temporal for domain adaptation RSCD.

as shown in Fig. 1. Self-supervised methods learn through image reconstruction or pseudo-image pair construction by basic image transformation. But it often lacks changes in non-identical objects through simple image transformations such as brightness augmentation. The pretraining paradigm trains the visual encoder on other large-scale dataset and relies on mathematical analysis to detect changes in the target data set. This way lacks training analysis of changes between image-pairs, and it is difficult to distinguish between real changes and some interfering factors by relying only on mathematical analyses such as PCA. This discrepancy from real-world scenarios can make them overly sensitive to noise in RSCD. Furthermore, both supervised and unsupervised methods require separate training on each dataset, resulting in poor domain generalization and significant performance variations across different datasets.

To address the above problems, we propose a unified model ChangeCLIP for change detection based on CLIP [23], as shown in Fig. 1 (c). Since CLIP is based on global text-image contrastive learning, it cannot directly be applied to pixel-dense recognition like RSCD. A number of pretraining approaches for dense prediction have emerged to alleviate this problem [24,34,41]. We propose a fine-tuning strategy by region alignment and dense pixel contrast learning to alleviate the lack of learning dense features in pre-training. Inspired by CoOp [43] and CoCoOp [42], we introduce a dynamic text-context optimization (DTCO) of prompt learning. It enhances the ability of the model to capture interdependencies between visual and textual information and model cross-learning. Additionally, to address the issue of data dependency, we propose a single-temporal image training strategy based on ControlNet [37]. This strategy allows us to train the change detection model on a large number of existing single-temporal images, further improving the model's generalization capabilities.

In summary, the main contributions of this paper are threefold:

- To address the issue of poor generalization in existing methods for change detection, we present a novel text-image alignment method ChangeCLIP. ChangeCLIP aims to effectively capture the interdependence between text and vision by a dynamic text-vision context representation learning DTCO.
- To address the problem of change detection relying on labeled pair images, we propose a reliable method for constructing pairwise images from singletemporal image, which follows natural distribution.
- Our approach achieves strong quantitative results, outperforming state-ofthe-art methods on multiple change detection datasets. Additionally, Change-CLIP shows excellent generalization across different datasets.

2 Related Work

2.1 Supervised Change Detection

RSCD detects the object changes by comparing the semantic features of pixels representing the same geographical region. Deep learning neural networks have demonstrated their excellence in change detection and these methods are based on siamese structure [21, 29, 33]. Existing supervised change detection methods often aim to improve model performance through network structure optimization or the incorporation of effective tricks, such as feature exchange [29,33], attention mechanism [6, 17, 18, 22], feature enhancement [17]. These supervised methods heavily rely on the availability of extensive labeled datasets and require training a separate model for each dataset. It results in overfitting issues, as the models become overly specialized to the dataset and the number of parameters rises. Consequently, the generalization performance of these models across different domains is significantly compromised.

Our method, based on the powerful prior knowledge of CLIP [23], demonstrates strong generalization on RSCD datasets through region alignment and pixel-wise text-vision contrastive learning.

2.2 Unsupervised Change Detection

Recently, unsupervised change detection based on deep learning have alleviated the dependence of supervised methods on high-quality annotated data. Some of the work use GANs and Diffusion technology to reconstruct images [14,25,38]. In the inference stage, mathematical methods such as principal component analysis (PCA) are used to predict [2,20,26,35]. Xu [31] et al. design proprietary pseudo-labelling generation structures to assist the above process.

However, the above unsupervised methods still face issues such as the diversity of changed objects, disturbing factors such as brightness and shadowing of different time, and data size limitations. In another hand, these models still need to be be trained separately for different datasets.

2.3 Single-Temporal Supervised Change Detection

Existing high-quality annotated change detection datasets [5, 13, 15] are very small, but there are already a large number of remote sensing single-temporal annotated datasets [11, 32]. ChangeStar [39] first proposed a single-temporal change detection training strategy STAR [39] which forms the pseudo-pair by randomly matching within the training batch. However, it lacks changes over long temporal span and leads to a significant number of changes that do not correspond to reality. It goes against the changes in the real world, significantly different from the true distribution. Moreover, it is difficult to perform shape analysis on the corresponding objects in the two randomly matched images. Changen [40] generates pseudo-image pairs based GAN [19], but its controllability and flexibility are limited.

In this paper, we propose a single-temporal training strategy based on Artificial Intelligence Generative Content (AIGC) with text-image mask controllable generation. By generating abundant trainable data through directed category editing, we mitigate data dependency. This approach constructs data closely resembling the natural real-world distribution, significantly enhancing model generalization.

2.4 Vision-Language Models

Computer vision has developed rapidly, driven by the pre-trained + finetuning paradigm. Pre-trained models such as Contrastive Language-Image Pre-training (CLIP) [23] and Imagen [28] are even more impressive for "text+vision" multimodal large language models. But these works are based on whole image classification. And the prompts provided by specialists are fixed. DenseCLIP [24] and MaskCLIP [41] propose CLIP-based methods to transfer CLIP to downstream dense prediction. CoOp [43] and CoCoOp [42] validate the effect of prompt engineering on pre-trained models and they adapt visual-language pre-trained models like CLIP for downstream image recognition.

In this paper, we propose the paradigm of pixel-wise text-vision contrastive learning and visual region contrastive learning, demonstrating superiority in dense pixel recognition and classification for RSCD. In addition, we propose a method of dynamic text-context optimization (DTCO), which ensures both the utilization of powerful prior knowledge and the generalization of the pre-training model.

3 Method

The overview of the proposed ChangeCLIP is illustrated in Fig. 2. We first revisit CLIP [23] and prompt engineering in Section 3.1. Then, we introduce contrastive learning in Section 3.2, dynamic context optimization in Section 3.3, and single-temporal training strategy in Section 3.4.



Fig. 2: Overview of ChangeCLIP. ChangeCLIP can be divided into two parts: visual processor and text processor. Firstly ChangeCLIP extracts the region visual-patch features \bar{f} and gets the dense visual features \hat{f} by FPN. Then a K-way classification is used for semantic segmentation. And the local visual-patch alignment loss \mathcal{L}_{lva} is used to alleviate the lack of large-scale pre-trained model on region feature learning. Then the text encoder extracts text embeddings and Fine-tuning CLIP for change detection of generalisability through visual-context alignment loss \mathcal{L}_{pca} .

3.1 Preliminary: CLIP and Prompt Engineering

CLIP [23] is trained on large-scale text-image pairs by whole-image level contrastive learning. CLIP consists of an image encoder (IE) and a text encoder (TE), and the image encoder provides two different backbones based on ResNet [12] and ViT [9]. The original CLIP text encoder input is a fixed text defined by specialists, such as "A photo of a dog", to generate the text embedding. CoOp [43] proposes independent learning prompt, replacing the text prompter with a shared learnable vector. CoCoOP [42] generates instance-conditional tokens through a lightweight network, and adds them to the text representation and the context input of the TE. Du et al. [10] and Rao et al. [24] demonstrate the effectiveness of this learnable prompt. The learnable prompt t_k is defined as:

$$t_k = TE([v_1, v_2 \dots v_{l-1}, c_i]) \tag{1}$$

where $\{v_1, v_2...v_{l-1}\}$ denotes the learnable context vectors, c_i is the word embeddings of k-th class word name, l denotes the context length.

3.2 Text-Vision Alignment

As shown in Fig. 2, ChangeCLIP consists of two alignments and two detections: local patch-visual alignment and pixel context alignment, semantic segmentation and object change detection. We utilize the CLIP image encoder to extract patchvisual embeddings (f_{po} and f_{pa}) from the bi-temporal images (I_{po} and I_{pa}), and extract the text embeddings (t) using CLIP text encoder.

Local patch-visual alignment. There is a large amount of pseudo-change noise between the bi-temporal images. We desire the image encoder to provide stable representations of image embedding. In the feature space, the features of the same category should be brought closer together. Conversely, the features of different categories should be pushed apart. Therefore, we employ a contrastive loss to align the local region features of the image patches, shown as in Fig. 2. Let \bar{f}_{po} and \bar{f}_{pa} be the normalized embeddings of f_{po} and f_{pa} , y_{po} and y_{pa} denotes the ground-truth of post-temporal and past-temporal image, and τ represents the temperature parameter. Based on this, we can obtain the similarity label: $\bar{y} = y_{po} \& y_{pa}$. The *i*-th pixel cosine similarity distance $d_i = \langle \bar{f}_{pa}, \bar{f}_{po} \rangle$. We define the local patch-visual alignment loss \mathcal{L}_{lva} mathematically as Equation 2.

$$\mathcal{L}_{lva} = -\frac{1}{N} \sum_{i}^{N} \bar{y} \log d_i + (1 - \bar{y}_i) \log(1 - d_i)$$
(2)

From an intuitive perspective, change detection can be seen as a parallel task to semantic segmentation. Existing pre-training is unaware of the regionlevel and pixel-level image representations with the text tokens. In this paper, we employ a Feature Pyramid Network (FPN) to upsample the region features f_{po} and f_{pa} of the bi-temporal patches from 1/32 to 1/4 of the original image size, shown as in Fig. 2. We represent the densely upsampled pixel-wise features as \hat{f}_{po} and \hat{f}_{pa} . Then we use a lightweight classification network to predict the semantic segmentation results for the bi-temporal images. Each pixel embedding $\hat{f}_i \in \mathbb{R}^D$ is fed into a K-way classification, where *i* is the pixel index. Assuming the pixel \hat{f}_i belongs to class *k*, the probability of f_i being classified as class *k* is computed as:

$$p(k|f_i) = \frac{\exp(w_k^T f_i)}{\sum_{k=1}^{K} \exp(w_k^T \hat{f}_i)}$$
(3)

where $\hat{f}_i \in {\{\hat{f}_{i,po}, \hat{f}_{i,pa}\}}$ denotes the *i*-th pixel, W denotes the learnable parameter of pixel-wise classification head in Fig. 2, $W = [w_1, ..., w_k] \in \mathbb{R}^{K \times D}$, $w_k \in \mathbb{R}^D$. The loss function is defined as follows:

$$\mathcal{L}_{seg} = -\log p(k|f_i) \tag{4}$$

$$= -\sum_{i=1}^{N} y_{i,po} log(p_{i,po}) - \sum_{i=1}^{N} y_{i,pa} log(p_{i,pa})$$
(5)

Pixel context alignment. With the pixel-level image representations $\hat{f}_{po}, \hat{f}_{pa} \in \mathbb{R}^{D \times N}$, and the context embeddings $t \in \mathbb{R}^{K \times D}$, we can align textual and visual

information through pixel-wise contrastive learning. This alignment process enables the alignment of textual and pixel-level visual information, facilitating effective cross-modal understanding and analysis. The category prediction of each pixel pair $f_i \in {\hat{f}_{po}, \hat{f}_{pa}}$ is achieved by a classification:

$$\hat{k}_{f_i} = \underset{(k)}{\operatorname{arg\,min}} \{ \langle f_i, t_k \rangle \}_{k=1}^K \tag{6}$$

We define the probability of pixel i over the class:

$$p(k|f_i) = \frac{\exp(-s_{i,k})}{\sum_{k'=1}^{K} \exp(-s_{i,k'})}$$
(7)

where $s_{i,k} = \langle f_i, t_k \rangle$ denotes the cosine similarity of f_i and text embeddings t_k of class k. Given the ground-truth y of bi-temporal image I_{pa}, I_{po} , i.e., $y_i \in \{1, ..., K\}$, the pixel-context alignment loss can be computed as:

$$\mathcal{L}_{pca,j}_{j \in \{pa,po\}} = -\frac{1}{N} \sum_{i=1}^{N} \log \frac{exp(-s_{i,k})}{exp(-s_{i,k}) + \sum_{k' \neq k_i} exp(-s_{i,k'})}$$
(8)

Change detection. With the visual features \hat{f}_{po} , \hat{f}_{pa} for dense pixels in the bi-temporal images, and the contrastive scoring map s between visual features and context embeddings, we employ a lightweight detection head to detect object changes. The visual information includes visual features such as object shapes, while the scoring map contains pixel-level class information. Therefore, we concatenate the scoring map with the visual feature map to form the input feature for the detection head.

$$\boldsymbol{p_k} = softmax(Convs([f_{po}, s_{po}, f_{pa}, s_{pa}])) \tag{9}$$

where $p_k \in [0, 1]$ denotes the predicted probability for positive class. This combined input allows the detection head to leverage both the visual and class information for accurate object change detection. With the ground-truth $\hat{y} \in \{0, 1\}$, we adopt binary cross-entropy loss for change detection loss:

$$\mathcal{L}_{cd}(\boldsymbol{p}_{\boldsymbol{k}}, \hat{y}) = -\hat{y}\log(\boldsymbol{p}_{\boldsymbol{k}}) + (1-\hat{y})\log(1-\boldsymbol{p}_{\boldsymbol{k}})$$
(10)

where $\hat{y} \in \{0, 1\}^{H \times W}$ specifies the ground-truth class.

Multi supervision. The overall loss function includes the losses of four components, as following:

$$\mathcal{L} = \mathcal{L}_{seg} + \mathcal{L}_{cd} + \alpha \mathcal{L}_{lva} + \beta (\mathcal{L}_{pca,po} + \mathcal{L}_{pca,pa})$$
(11)

Since the fundamental tasks are change detection and instance segmentation, the main components of the loss function are the semantic segmentation loss \mathcal{L}_{seg} and the change detection loss \mathcal{L}_{cd} . The local patch-visual alignment loss \mathcal{L}_{lva} and pixel-context alignment loss $\mathcal{L}_{pca,po}$, $\mathcal{L}_{pca,pa}$ serve as an auxiliary loss in this context. Therefore, we regulate the auxiliary losses weights by two parameters $\alpha = 0.1, \beta = 0.1$, and the main loss weight is always set to 1.

3.3 Dynamic Text-Vision Context Optimization

DenseCLIP [24] and CoCoOp [42] utilize lightweight networks to the embedding space of textual information. They then perform contrastive learning by directly adding these new contextual representations to the visual information. This way undermines the generalisability of CLIP. To address this issue, we propose a dynamic text context optimization (DTCO) to better adapt the visual features to context.

Initially, we utilize a transformer based on cross-attention to map visual features into the textual feature space. The multi-head cross-attention mechanism operates by attending to relevant textual features for each visual feature and vice versa. It allows the model to capture the interdependencies and correlations between visual and textual information, enabling a more comprehensive understanding of the context. Then, we employ two small networks, known as adapters, to separately map the visual information to the textual space and the textual information to the visual space. Finally, we use a learnable dynamic weight to combine the outputs of these adapters, resulting in a new contextual embedding. This dynamic weight allows the model to adaptively adjust the importance of visual and textual information, enhancing the fusion process.

Based on the aforementioned process, we can define the contextual representation as follows:

$$z = Transformer([f, g], t_k), with \ g = Pooling(f)$$
(12)

$$\hat{t}_k = \varphi(z) * t_k + \varsigma(t_k) * z + \sigma t_k \tag{13}$$

where φ, ς is a lightweight network (MLP-V and MLP-T in Fig 2), σ is the weight vector and $\sigma \in \mathbb{R}^L, \hat{t}_k \in \mathbb{R}^{K \times L}, t_k$ is calculated in Equation 1. By dynamically optimizing the textual context and focusing on relevant textual features for pixelwise visual feature, the model effectively captures the mutual dependencies and associations between visual and textual information.

3.4 Single-Temporal Controllable AI-Generated Training Strategy

Revisiting the essence of temporal image pairs for change detection, it is evident that pixels at the same location in these pairs represent two distinct states of the same geographical area at different times. These states may manifest as different physical conditions of the same object over time or as changes induced between different categories of objects at the same location. The establishment of correspondence between these diverse states is fundamental to train object change detection models based on single-temporal images. Acquiring high-quality instance annotations for such bi-temporal image pairs comes at a considerable cost. Nevertheless, a wealth of single-temporal remote sensing instance segmentation images is already available and highly relevant to change detection. ChangeStar [39] proposes a method of constructing image pairs by STAR [39]. However there are some shortcomings, the generated data deviate from practical significance, and violate the true changes observed in the natural world.



Fig. 3: Single-temporal Training Strategy (SAIN) Framework. We separate each class of objects and generates shape maskes by mask generation. During the training phase, we train the ControlNet by inpainting learning. In inference phase, we introduce deliberate category-based changes to each object in the image except buildings.

To address the aforementioned issue, we propose single-temporal controllable AI-generated training strategy (SAIN) based on AIGC for generating pseudoimage pairs as Fig. 3. Controlled image editing provides a promising avenue for addressing this issue. By selectively editing specific regions of an image, one can obtain pairs of images closely resembling each other, thus simulating different states over time.

In this paper, we implement this approach using ControlNet [37], denoted as \mathcal{G} . We begin by defining an image $\mathcal{I}_a = \{x_1, x_2, ..., x_i, ..., x_N\} \in \mathbb{R}^{N \times 3}$, composed of N objects belonging to K categories, where $x_i \in \{C_1, ..., C_K\}$, each with its segmentation annotation map $g \in 1, ..., K$. We define the process to generate pseudo-image as shown in Equation 14.

$$\mathcal{I}_b = \mathcal{G}(\mathcal{I}_a, M, P_t, n) \tag{14}$$

where I_a is the real image, M is a binary mask map to control the edited area, P_t is the prompt which specifies the category information of the object in the edited area, and n is the random noise.

During the training phase, we fine-tune the AIGC model by image inpainting approach to recognize objects in each category. In the inference phase, changes between categories or within the same category are generated through predefined priors. However, existing generative models perform poorly for small target objects and detailed parts. During the actual verification, we discovery that the generative model performs poorly in generating buildings. However, it excels in generating farmland, forests, and other large-scale targets. In real nature, changes in buildings, such as building-to-building changes, are often caused by variations in lighting, seasons, and other factors. Generating buildings by incorporating changes in image brightness and other relevant factors closely resembles

Table 1: Zero-shot change detection on LEVIR-CD [5] and WHU-CD [13] datasets. The model is trained on xView2, SECOND and SAIN 3.4 generated.

	xView2 [11]					SECOND [36]					SAIN Generated Data 3.4							
Model	LEV	VIR-CI) [<u>5</u>]	WH	IU-CD	[13]	LEV	VIR-CI) [5]	WH	IU-CD	[13]	LEV	VIR-CI) [5]	WH	IU-CD	[13]
	F1	IoU	OA	F1	IoU	OA	F1	IoU	OA	F1	IoU	OA	F1	IoU	OA	F1	IoU	OA
BiT [4]	66.12	49.35	97.05	62.47	45.42	96.90	70.23	53.23	97.42	66.94	50.31	97.35	80.00	66.67	98.20	69.74	53.55	96.70
ChangeFormer [1]	66.91	50.27	96.56	59.31	42.16	96.01	67.47	50.92	94.12	65.53	48.73	93.86	72.55	56.93	97.85	68.08	51.61	97.31
Changer [29]	62.26	45.20	96.74	37.96	23.43	94.05	57.13	39.99	94.86	62.30	45.24	95.71	68.83	52.47	96.95	65.59	48.79	96.76
SARAS-Net [3]	45.59	29.53	92.09	32.41	19.34	84.13	40.10	25.08	89.93	16.59	9.05	57.12	52.28	35.39	87.49	34.21	20.63	87.49
TinyCD [8]	56.65	39.52	94.97	29.52	17.32	88.18	67.43	50.87	96.07	60.32	43.18	95.55	68.21	51.76	95.83	74.28	59.08	97.62
USSFC-Net [16]	49.32	32.73	94.49	44.81	28.88	93.27	70.34	54.25	97.37	50.83	34.08	95.22	72.63	57.01	97.23	63.52	46.54	97.02
ChangeStar [39]	79.31	65.71	98.06	73.59	58.22	97.43	54.67	37.61	96.72	50.76	34.01	92.84	70.53	62.93	98.07	75.96	55.40	97.41
ChangeCLIP(Ours)	89.02	70.74	98.29	78.64	61.64	97.92	73.56	64.44	98.11	78.36	64.41	98.06	85.19	74.20	98.59	79.92	64.67	98.61

real-world scenarios. Therefore, we use the AIGC model for large targets like forest, while building generation is achieved through transformations in color, brightness, and other image properties. This strategy enables the creation of synthetic image pairs for training, offering a viable solution to the challenge of obtaining bi-temporal annotations in a cost-effective manner. The utilization of ControlNet [37] ensures the efficacy of our proposed method in generating realistic and informative pseudo-image pairs for enhancing the performance of change detection models.

4 Experiment

4.1 Experimental Setup

Dataset. Two high spatial resolution datasets (xView2 [11] and SECOND [36]) and the generated dataset by our proposed SAIN 3.4 using xView2 as a single time-phase data source are used to train models. Two RSCD datasets (LEVIR-CD [5] and WHU-CD [13]) are used to evaluate the performance.

xView2: xView2 [11] is a challenge set of xView, which is one of the largest publicly available datasets of overhead imagery, for assessing damaged buildings after a natural disaster. We have only pre-disaster images of "train" and "tier3" split available, with 9089 images of size 1024×1024 .

SECOND: SECOND [36] is a well-annotated semantic change detection dataset. There are 4662 pairs of aerial images from several platforms and sensors. Each image of these pairs has size 512×512 and is annotated at the pixel level. The SECOND reflects real distributions of land-cover categories when changes occur.

LEVIR-CD: LEVIR-CD [5] dataset is collected by Google Earth and contains 637 HRS bi-temporal remote sensing image pairs with size of 1024×1024 . Each image pair is annotated by a binary instance. LEVIR-CD is officially split into three parts, including train/val/test sets of samples 445/64/128. We merge all sets as the validation.

WHU-CD: WHU building change detection [13] dataset contains only one pair of images with size of 32507×15354 . We crop it into size of 512×512 and obtain 1920 pairs for verification.

4.2 Metrics and Implementation Details

Evaluation metric. In change detection, precision reflects resistance to pseudochange (e.g. color changes of buildings) noise, recall measures sensitivity to actual changes, and F1-score offers a comprehensive evaluation. Overall accuracy (OA) reflects the accuracy of the model in classification of the overall classes. Intersection over union (IoU) assesses shape analysis and change object localization. In this context, we adopt F1-score, IoU and OA as the primary focus centers.

Training: Our method is implemented on CLIP [23], utilizing ViT as the backbone. We initialize the image encoder and text encoder with pre-trained weights based on ViT-B. During training, the parameters of the text encoder are frozen. The image encoder is fine-tuned, with the gradient weight parameters set to 0.1 of the main learning rate. We use standard data augmentation techniques, including random flipping, random cropping, random color dithering, random rotation, and image pair swapping operations. We adopt AdamW as the parameter optimizer, with a learning rate set to 0.0001 and the learning rate is scheduled following the polynomial annealing policy.

In addition, we conduct model training experiments on three distinct datasets. The first dataset is xView2, a single-temporal dataset for building detection, where we construct pairs of data as ChangeStar [39]. The second dataset is SECOND, which is specifically designed for semantic change detection. Lastly, we create a dataset based on the generating strategy of Section 3.4, building upon the xView2 dataset.

Testing: We evaluate the performance of object change detection on the LEVIR-CD and WHU-CD datasets. We merge the training, validation, and testing sets of these two datasets separately to assess the performance on each dataset.

4.3 Comparison with State-of-the-Art Detectors

We conduct training of ChangeCLIP on the three aforementioned training datasets and subsequently validated its performance on two large-scale change detection benchmarks. Tab. 1 demonstrates that ChangeCLIP exhibits robust generalization and high accuracy, achieving excellent results on the LEVIR-CD [5] and WHU-CD [13] datasets even without the need for supervised learning. This underscores the model's ability to generalize well to diverse datasets, showcasing its potential for robust change detection applications.

Comparing horizontally in Tab. 1, contrasting models trained on the SEC-OND and xView2 datasets reveals that real-world variations present in SEC-OND contribute to improved generalization for BiT [4], TinyCD [8] and USSFC-Net [16]. However, due to its smaller data scale, performance is notably poor for Changer [29] and ChangeStar [39]. In contrast, leveraging our training strategy results in enhanced performance across all methods. In a vertical comparison, our method consistently outperforms other approaches, demonstrating that existing methods are prone to overfitting on the training dataset, with poor generalization and zero-shot capabilities. Specifically, when validating on the LEVIR-CD

Model	LEV	VIR-CI	D [5]	WHU-CD [13]			
Model	F1	IoU	OA	F1	IoU	OA	
BiT [4]	89.31	80.68	98.85	83.98	72.39	99.26	
ChangeFormer [1]	90.40	82.48	99.04	85.61	75.14	98.84	
SARAS-Net [3]	87.26	70.40	99.10	82.23	69.95	98.89	
TinyCD [8]	91.05	83.57	99.10	91.48	84.30	99.32	
USSFC-Net [16]	91.04	83.55	99.09	92.20	85.54	99.27	
ChangeStar [39]	90.82	83.19	99.02	86.25	75.98	95.26	
ChangeCLIP	91.14	83.72	99.11	92.18	85.50	99.33	

 Table 2: Supervised change detection.
 Comparison with the state-of-the-art change methods on the LEVIR-CD and WHU-CD test.

dataset trained on the xView2 dataset, our method outperforms the state-ofthe-art ChangeStar [39] by 5.03% in IoU and 9.71% in F1. Furthermore, when training on our single-temporal-phase generated dataset, our method achieves a 9.49% increase in IoU and a 5.88% improvement in F1 compared to ChangeStar on LEVIR-CD.

Supervised change detection results. Tab. 2 presents the supervised results on the LEIR-CD [5] and WHU-CD [13] datasets. We remove the two segmentation prediction branches of ChangeCLIP and fine-tune ChangeCLIP on the dataset separately. We found that change detection performance of Change-CLIP improves significantly after supervised fine-tuning on each dataset. This demonstrates the potential of ChangeCLIP as a unified model for change detection.

4.4 Ablation Study

In order to validate the effectiveness of each component, we conduct a series of ablation experiments, as shown in Tab. 3. Comparing with only use IE of CLIP (row 1), the method based on text-visual contrastive learning as CLIP (row 2) achieves 3.93% increase in IoU and 2.94% improving in F1. It demonstrates the feasibility of text-visual contrastive learning. From rows 3-5, we validate the effectiveness of text-visual alignment, DTCO, and single-temporal training strategy proposed in our approach. Comparing to the original CLIP (row 2), the improvements achieved through these components are significant. In row 6, we evaluate the combined effect of the alignment and DTCO, which results 5.2% increase in IoU and 3.66% increase in F1 comparing with row 2, further highlighting the effectiveness of our proposed ChangeCLIP approach. Finally, in row 7, we demonstrate the additional improvement achieved by SAIN.

Additionally, we conduct ablation experiments on the hyperparameter in Equation 3.2, as shown in Tab. 4. The close resemblance of the results demonstrates the robustness of our method.



 Table 3: Component effectiveness ablation experiment, test on LEVIR-CD.

Fig. 4: Visualization analysis for ChangeCLIP with benchmarks. The red regions indicate False Positives (FP), while the blue regions denote False Negatives (FN). Each column represents the respective performance of (a) Past-temporal, (b) post-temporal, (c) groundtruth, (d) Changer [29], (e) SARAS [3], (f) Tiny-CD [8], (g) USSFC [16], (h) ChangeStar [39], (i) ChangeCLIP.

4.5 Visualization Analysis

To illustrate the change detection performance of our method, we conducted a visual analysis of the detection results on the LEVIR-CD dataset. As depicted in Fig. 4, instances of false positives (FP) are denoted in red, while false negatives (FN) are represented in blue. The first row showcases the excellent performance of our method in detecting sparse, small-scale targets. The second row reflects the superior performance of our approach when dealing with scenarios involving both large-scale and small-scale objects, surpassing the capabilities of other methods. The third row demonstrate the outstanding performance of our method in detecting dense objects. This visual analysis provides a comprehensive representation of our method's versatility and effectiveness across various scenarios, highlighting its proficiency in handling sparse, small-scale, large-scale and dense objects.

Qualitative results. Fig. 5 illustrates the superiority of our single-temporal training strategy compared with single-temporal supervised learning (STAR) [39] used in ChangeStar [39]. It is evident that our generated images closely resem-

14 Qiangang Du et al.

Table 4: Different α and β in Equation 3.2.

α, β	0.1/0.1	0.1/1.0	1.0/0.1	0.5/0.5	1.0/1.0
F1	85.19	83.58	83.24	84.71	82.84
IoU	74.20	71.79	71.30	73.48	70.70



Fig. 5: Qualitative analysis result of single-temporal image in Section 3.4. The red area is the area where the buildings overlap in the STAR [39], which is impossible in real world. SAIN generates images that encompass both pseudo-change within the same category and changes between different categories.

ble with the changes in real world, including changes between different objects and color change between the same object. Additionally, our method avoids the unnatural changes, such as overlapping buildings that do not represent actual changes (as shown in the red region of the Fig. 5), which can occur in STAR [39]. This further validates the effectiveness of our approach.

5 Conclusion

In this paper, we propose a multimodal change detection network architecture, ChangeCLIP, to address the issue of poor generalization in existing change detection methods. To enhance the generalization of text-vision pretraining contrastive learning, we propose a dynamic context optimization method that effectively improves the model generalization ability. Furthermore, We alleviate the lack of CLIP for dense feature learning through local visual-patch alignment and visual-context alignment. Additionally, we introduce a single-temporal controllable AI-generated training strategy for change detection to overcome the reliance of existing methods on large-scale, high-quality annotated image pairs, which is expensive. Our network architecture is elegant, versatile, and exhibits outstanding performance via training-free on target datasets. Extensive experiments demonstrate that our approach is effective for single-temporal supervised training in change detection and offers insights into the application of pre-training in change detection.

References

- Bandara, W.G.C., Patel, V.M.: A transformer-based siamese network for change detection. In: IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium. pp. 207–210 (2022) 1, 10, 12
- Bergamasco, L., Saha, S., Bovolo, F., Bruzzone, L.: Unsupervised change detection using convolutional-autoencoder multiresolution features. IEEE Transactions on Geoscience and Remote Sensing 60, 1–19 (2022) 3
- Chen, C.P., Hsieh, J.W., Chen, P.Y., Hsieh, Y.K., Wang, B.S.: Saras-net: Scale and relation aware siamese network for change detection. Proceedings of the AAAI Conference on Artificial Intelligence 37(12), 14187–14195 (Jun 2023) 1, 10, 12, 13
- Chen, H., Qi, Z., Shi, Z.: Remote sensing image change detection with transformers. IEEE Transactions on Geoscience and Remote Sensing 60, 1–14 (2022) 1, 10, 11, 12
- Chen, H., Shi, Z.: A spatial-temporal attention-based method and a new dataset for remote sensing image change detection. Remote Sensing 12(10) (2020) 4, 10, 11, 12
- Chen, J., Yuan, Z., Peng, J., Chen, L., Huang, H., Zhu, J., Liu, Y., Li, H.: Dasnet: Dual attentive fully convolutional siamese networks for change detection in high-resolution satellite images. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 14, 1194–1206 (2021) 3
- Chen, Y., Bruzzone, L.: Self-supervised change detection in multiview remote sensing images. IEEE Transactions on Geoscience and Remote Sensing 60, 1–12 (2022) 1, 2
- Codegoni, A., Lombardi, G., Ferrari, A.: Tinycd: A (not so) deep learning model for change detection. arXiv preprint arXiv:2207.13159 (2022) 1, 10, 11, 12, 13
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. ICLR (2021) 5
- Du, Y., Wei, F., Zhang, Z., Shi, M., Gao, Y., Li, G.: Learning to prompt for open-vocabulary object detection with vision-language model. arXiv preprint arXiv:2203.14940 (2022) 5
- Gupta, R., Hosfelt, R., Sajeev, S., Patel, N., Goodman, B., Doshi, J., Heim, E., Choset, H., Gaston, M.: xbd: A dataset for assessing building damage from satellite imagery (2019) 4, 10
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition (2015) 5
- Ji, S., Wei, S., Lu, M.: 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 (2019) 4, 10, 11, 12
- de Jong, K.L., Sergeevna Bosman, A.: Unsupervised change detection in satellite images using convolutional neural networks. In: 2019 International Joint Conference on Neural Networks (IJCNN). pp. 1–8 (2019) 1, 3
- Lebedev, M.A., Vizilter, Y.V., Vygolov, O.V., Knyaz, V.A., Rubis, A.Y.: Change detection in remote sensing images using conditional adversarial networks. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLII-2, 565–571 (2018) 4
- 16. Lei, T., Geng, X., Ning, H., Lv, Z., Gong, M., Jin, Y., Nandi, A.K.: Ultralightweight spatial–spectral feature cooperation network for change detection in remote sensing

images. IEEE Transactions on Geoscience and Remote Sensing **61**, 1–14 (2023) 1, 10, 11, 12, 13

- Lei, T., Wang, J., Ning, H., Wang, X., Xue, D., Wang, Q., Nandi, A.K.: Difference enhancement and spatial–spectral nonlocal network for change detection in vhr remote sensing images. IEEE Transactions on Geoscience and Remote Sensing 60, 1–13 (2022) 3
- Li, Z., Yan, C., Sun, Y., Xin, Q.: A densely attentive refinement network for change detection based on very-high-resolution bitemporal remote sensing images. IEEE Transactions on Geoscience and Remote Sensing 60, 1–18 (2022) 3
- 19. Mirza, M., Osindero, S.: Conditional generative adversarial nets (2014) 4
- Noh, H., Ju, J., Seo, M., Park, J., Choi, D.G.: Unsupervised change detection based on image reconstruction loss. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1352–1361 (2022) 2, 3
- Peng, J., Jiang, Z., Gu, Y., Wu, Y., Wang, Y., Tai, Y., Wang, C., Lin, W.: Siamrcr: Reciprocal classification and regression for visual object tracking. In: International Joint Conference on Artificial Intelligence (IJCAI) (2021) 3
- Peng, J., Wang, C., Wan, F., Wu, Y., Wang, Y., Tai, Y., Wang, C., Li, J., Huang, F., Fu, Y.: Chained-tracker: Chaining paired attentive regression results for endto-end joint multiple-object detection and tracking. In: European Conference on Computer Vision (ECCV) (2020) 3
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision (2021) 2, 3, 4, 5, 11
- 24. Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., Zhou, J., Lu, J.: Denseclip: Language-guided dense prediction with context-aware prompting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022) 2, 4, 5, 8
- Ren, C., Wang, X., Gao, J., Zhou, X., Chen, H.: Unsupervised change detection in satellite images with generative adversarial network. IEEE Transactions on Geoscience and Remote Sensing 59(12), 10047–10061 (dec 2021) 2, 3
- Saha, S., Bovolo, F., Bruzzone, L.: Unsupervised deep change vector analysis for multiple-change detection in vhr images. IEEE Transactions on Geoscience and Remote Sensing 57(6), 3677–3693 (2019) 1, 2, 3
- Saha, S., Bovolo, F., Bruzzone, L.: Unsupervised multiple-change detection in vhr multisensor images via deep-learning based adaptation. In: IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium. pp. 5033–5036 (2019) 2
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Ayan, B.K., Mahdavi, S.S., Lopes, R.G., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding. CoRR abs/2205.11487 (2022) 4
- 29. Sheng Fang, Kaiyu Li, Z.L.: Changer feature interaction is what you need for change detection. Arxiv (2023) 1, 3, 10, 11, 13
- Sun, Y., Lei, L., Li, X., Tan, X., Kuang, G.: Structure consistency-based graph for unsupervised change detection with homogeneous and heterogeneous remote sensing images. IEEE Transactions on Geoscience and Remote Sensing 60, 1–21 (2022) 2
- Tang, X., Zhang, H., Mou, L., Liu, F., Zhang, X., Zhu, X.X., Jiao, L.: An unsupervised remote sensing change detection method based on multiscale graph convolutional network and metric learning. IEEE Transactions on Geoscience and Remote Sensing 60, 1–15 (2022) 3

- 32. Wang, J., Zheng, Z., Ma, A., Lu, X., Zhong, Y.: LoveDA: A remote sensing landcover dataset for domain adaptive semantic segmentation (Oct 2021) 4
- Wang, S., Li, Y., Xie, M., Chi, M., Wang, Y., Wang, C., Zhu, W.: Align, perturb and decouple: Toward better leverage of difference information for rsi change detection (2023) 3
- Wang, X., Zhang, R., Shen, C., Kong, T., Li, L.: Dense contrastive learning for selfsupervised visual pre-training. In: Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR) (2021) 2
- Wu, C., Chen, H., Du, B., Zhang, L.: Unsupervised change detection in multitemporal vhr images based on deep kernel pca convolutional mapping network. IEEE Transactions on Cybernetics 52(11), 12084–12098 (2022) 3
- Yang, K., Xia, G.S., Liu, Z., Du, B., Yang, W., Pelillo, M.: Asymmetric siamese networks for semantic change detection (2020) 10
- Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models (2023) 2, 9, 10
- Zhang, X., Tian, S., Wang, G., Zhou, H., Jiao, L.: Diffucd:unsupervised hyperspectral image change detection with semantic correlation diffusion model (2023) 2, 3
- 39. Zheng, Z., Ma, A., Zhang, L., Zhong, Y.: Change is everywhere: Single-temporal supervised object change detection in remote sensing imagery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 15193– 15202 (October 2021) 2, 4, 8, 10, 11, 12, 13, 14
- Zheng, Z., Tian, S., Ma, A., Zhang, L., Zhong, Y.: Scalable multi-temporal remote sensing change data generation via simulating stochastic change process. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21818–21827 (2023) 4
- 41. Zhou, C., Loy, C.C., Dai, B.: Extract free dense labels from clip. In: European Conference on Computer Vision (ECCV) (2022) 2, 4
- Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Conditional prompt learning for visionlanguage models. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022) 2, 4, 5, 8
- Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. International Journal of Computer Vision (IJCV) (2022) 2, 4, 5