

A Mamba-based Siamese Network for Remote Sensing Change Detection

Jay N. Paranjape
 Johns Hopkins University, Baltimore
 jparanj1@jhu.edu

Celso de Melo
 DEVCOM Army Research Laboratory, Adelphi
 celso.m.demelo.civ@army.mil

Vishal M. Patel
 Johns Hopkins University, Baltimore
 vpatel36@jhu.edu

Abstract

Change detection in remote sensing images is an essential tool for analyzing a region at different times. It finds varied applications in monitoring environmental changes, man-made changes as well as corresponding decision-making and prediction of future trends. Deep learning methods like Convolutional Neural Networks (CNNs) and Transformers have achieved remarkable success in detecting significant changes, given two images at different times. In this paper, we propose a **Mamba-based Change Detector (M-CD)** that segments out the regions of interest even better. Mamba-based architectures demonstrate linear-time training capabilities and an improved receptive field over transformers. Our experiments on four widely used change detection datasets demonstrate significant improvements over existing state-of-the-art (SOTA) methods. Code: <https://github.com/JayParanjape/M-CD>

1. Introduction

In remote sensing, Change Detection (CD) refers to the task of detecting significant changes to our planet's surface across time. It is generally addressed by processing satellite images of a region, taken at different instances in time [2, 38]. The task of CD is non-trivial since the change of interest can vary across different applications, ranging from man-made structures, natural vegetation, or effects of climate change. Subsequently, CD finds its applications in monitoring and assessing natural disasters and climate shifts [26, 28, 43, 50, 69], policy planning [7, 55], surveying land and farmland cover [16, 37, 61], as well as military applications [36, 62]. Remote sensing images over different points in time can have multiple data disparities, including illumination changes [42, 63], varying resolutions [5], noise [1, 39], and errors in registration [6, 31] among others.

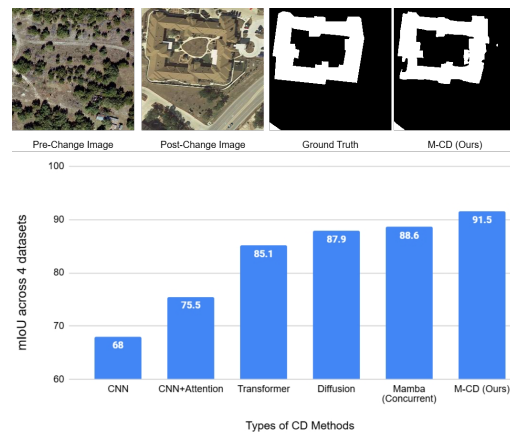


Figure 1. (Top) Example result of our M-CD. (Bottom) Average performance of M-CD with respect to existing types of CD models

Thus, making CD systems that are robust to such discrepancies is a complex and challenging task.

With the advent of deep learning, various systems have been developed for CD that employ deep networks and show substantial gains on multiple datasets [3, 4, 11, 21, 44, 71]. These methods include Convolutional Neural Networks [8, 21, 44], Attention-based networks [4, 11, 12] as well as diffusion-based networks [3]. Deep learning-based methods for CD generally take in a pair, consisting of a pre-change image and a post-change image as input, and output a segmentation mask locating the regions where the change of interest has occurred.

Recently, selective state space models (a.k.a Mamba) [22] were proposed that have a more global receptive field than transformers and are more scalable. While these were introduced for text inputs, VMamba [45] adapted this model for representing images and showed improvements in general vision tasks like classification [45] and segmentation

[64]. In this paper, we propose a Mamba-based Change Detector (M-CD) that outperforms all existing SOTA CD methods, as shown in Fig. 1. In particular, we develop a VMamba-based Siamese image encoder that is used to generate rich features from the two input images. We then develop a concatenation-based difference module that takes the features for the two images as input and combines them at multiple scales to create a joint feature vector. This is passed to a VMamba-based decoder that performs an aggregation over channels to produce the final output mask. In summary, the main contributions of our work are as follows:

1. We develop a novel Mamba-based Change Detector, called M-CD, that generates precise change masks, given two images.
2. To this extent, we develop a difference module that combines the features from the pair of images. This is provided to the mask decoder, which is modified to be more aware of the channel information along with the temporal and spatial information.
3. We evaluate our method on four widely used public datasets. We demonstrate consistent improvements over existing SOTA methods.

2. Related Work

Classical Change Detection. Traditionally, CD has been approached using algebraic, transformation-based or classification-based methods. Algebraic methods generally perform predefined mathematical operations to find a difference measure between the pre-change and post-change images. This is followed by using a pre-defined threshold to detect changed areas in the image. Some examples of these methods include ImageDiff [17], ImageRegr [47], ImageRatio [17], and Change Vector Analysis (CVA) [58]. While these methods are simple to implement, they are not sufficient for modeling the complexity of the problem and are heavily reliant on the threshold.

Transformation-based methods generally apply various transforms to the images and then compute a “difference” image based on the transformed input images. Examples of these transforms include Principal Component Analysis [19, 34], Karhunen-Loeve Transform (KT) [17], Gramm-Schmidt (GS) [17], Multivariate Alteration Detection (MAD) [54], Re-Weighted Multivariate Alteration Detection (IRMAD) [53], and Chi-Square [17]. The common principle behind these methods is that they aim to condense the important information present in the input pair of images using transforms. With just the relevant information remaining after the transform, they propose that a difference operation along with a threshold can give the required segmentation map. However, these methods are also limited by the selection of threshold as well as poor performance.

Finally, classification-based methods like Expectation-Maximization CD [17] and spectral-temporal CD [17] use machine learning techniques to provide per-pixel classification and perform additional post-processing to generate the final mask. However, as compared to modern-day deep networks, these techniques have an inferior quality.

Change Detection Using Deep Learning. Deep learning techniques produce rich embeddings for images, allowing them to surpass traditional methods for CD. Convolutional Neural Networks paved the way for utilization of DL for processing the pre and post-change images and then fusing them to produce the predicted mask. FC-EF [8] and FC-Siam-Conc [8] were one of the first methods to do so. In FC-EF, the pre-change and post-change image are concatenated along the channel dimension and a fully convolutional encoder-decoder structure is applied on the joint image. In FC-Siam-Conc, the same encoder operates on both the input images and the resulting embeddings are concatenated before passing them through the decoder. The latter approach was shown to be more effective in modeling the temporal difference between the two images. Later, with the introduction of deeper backbones like VGG [59], Resnet [27] and DenseNet [30], multiple CD techniques have evolved which improve the performance of FC-Siam-Conc. Some examples include DT-SCN [44] and SNUNet [21]. DT-SCN combines three subnetworks, one for change detection and two for semantic segmentation, which adopt a ResNet backbone. SNUNet combines dense connections and nested UNet [75] with Siamese networks for CD.

The introduction of Transformer networks for image analysis tasks in ViT [20] was instrumental in capturing the long-range dependencies in images, which was previously missing from CNN architectures. Since CD involves analyzing the relation between two different images, transformer-based approaches like BIT [11] and ChangeFormer [4] were found to be more effective than CNN methods [4, 11]. BIT uses a combination of ResNet, followed by a transformer encoder, that allows it to extract spatial and comparative features. Its successor, the Changeformer uses a purely transformer architecture and showed significant improvements over BIT. It comprises of a transformer encoder for the images as well as a transformer decoder, which is more effective than CNNs for feature representations.

One recent line of work uses self-supervised learning to pretrain the models with vast amounts of unannotated data. These may be obtained from datasets like ImageNet 1k as well as publicly available data on the internet. Some examples of this approach include SiamSiam [15], MoCo-v2 [52], DenseCL [66], CMC [9], SeCo [49] and SaDL-CD [10], among others. These methods usually pretrain their encoders on self-supervision tasks, thereby vastly im-

proving their representative nature. For instance, SaDL-CD performs background-swapping during training and applies data augmentations to produce three different views of a given image. Next, it optimizes a loss function that appropriately measures the difference between the views. SeCo, on the other hand, enforces consistencies between features of the seasonally varying images.

Another recent work proposes the use of a generative modeling framework called Denoising Diffusion Probabilistic Models (DDPM) [29] as a pretraining strategy for learning powerful encoders [3]. Once learnt, they can be used as feature extractors to train a light-weight model for CD. Hence, these can also be considered as a method that requires a large amount of pretraining. The proposed approach, called DDPM-CD, outperforms all the traditional and deep learning techniques mentioned above and is the current SOTA method for CD.

In this paper, we develop a Mamba-based mechanism that outperforms DDPM-CD as well as other methods on four widely used public datasets for CD, showing the effectiveness of selective state space modeling as a fundamental building block for CD over CNNs and transformers.

Mamba in Vision Tasks. Various recent works propose the use of State Space Modeling (S4) as an effective method of representing visual information [32, 33, 51, 65]. Trans4mer [33] and Vis4mer [32] were used to create representations of movie clips using S4, while S4ND [51] uses S4 to effectively model multidimensional signals in 1D, 2D and 3D. On the other hand, S5 [65] proposes a selective mechanism for capturing long-range dependencies in long videos, which became the motivation for Mamba [22]. Mamba has been adapted for a multitude of computer vision tasks recently, which shows its effectiveness in modeling visual representations. Vision Mamba [76] proposed using the selective scan mechanism in Mamba from two directions for an image, while VMamba [45] showed that scanning an image from four directions has a better representation. In our work, we follow the four-directional strategy of VMamba in the Siamese Image Encoder. Mamba has also been explored for various applications in Computer Vision including medical image segmentation [48, 57, 67, 70], multi-modal segmentation [64], point cloud computations [41], and image restoration [25]. However, most of these methods involve direct replacement of transformer blocks with Mamba blocks, without adapting for the particular task. Unlike this plug-and-play approach, we develop a Mamba-based difference module and decoder that is tailored for the task of CD. Some methods like Sigma [64], which are used for multimodal analysis introduce an additional fusion module using Mamba. This method is similar to our method in using Mamba blocks for encoding and decoding, along with an added fusion

strategy. However, this method uses cross-modal fusions which are not suited for CD according to our experiments.

Concurrent Works in Mamba for CD. A couple of recent CD methods have proposed using Mamba as the backbone architecture. While working on the proposed method, we recognize these recent methods as concurrent works and compare our method with them. ChangeMamba [14] uses three different designs in the decoder for modeling the spatio-temporal relationship between the pre-change and post-change images. RSMamba [74], on the other hand increases the number of scanning directions in VMamba from four to eight in order to capture relations from all directions. In contrast, our method uses a single decoder design and the standard four directions for scanning. Instead, we design the decoder to use features at multiple scales for spatial learning and design a separate difference module to learn temporal relations. CDMamba [73] comes close to our proposed method in that it uses multi-scale features in the decoder. However, our method also additionally uses a difference module to allow the model to learn temporal relations between the two images better. Experimental results on multiple datasets show that our method outperforms concurrent Mamba-based methods showing the significance of our architecture design for CD.

3. Proposed Method

3.1. Preliminaries

State Space Modeling. State Space Models (SSM) are a class of sequence-to-sequence modeling algorithms, which follow the linear time invariance (LTI) property [23, 24, 60]. These systems work by maintaining a hidden state that changes with time. The output of the system at any point in time is dependent on the hidden state as well as the input. More specifically, given input $x(t) \in \mathbb{R}$, the following equations define the output $y(t) \in \mathbb{R}$ and the hidden state $h(t) \in \mathbb{R}^N$, where N is the size of the hidden state:

$$y(t) = Ch(t) + Dx(t), \quad \dot{h}(t) = Ah(t) + Bx(t), \quad (1)$$

where $\dot{h}(t)$ denotes the first derivative with respect to time t . The matrices $A \in \mathbb{R}^{N \times N}$, $B \in \mathbb{R}^{N \times 1}$, $C \in \mathbb{R}^{1 \times N}$, and $D \in \mathbb{R}$ are characteristic to the system and are learnt during training. Further, to model discrete sequences like images, SSMs utilize a predefined discretization parameter Δ , which maps the parameters A and B to the discrete space. This technique is known as Zero Order Hold Discretization [23]. More specifically, the discretization process is defined as follows:

$$y_k = \bar{C}h_k + \bar{D}x_k, \quad h_k = \bar{A}h_{k-1} + \bar{B}x_k, \quad (2)$$

where $\{x_1, x_2, \dots, x_k\}$ denotes the discrete input sequence, $\{y_1, y_2, \dots, y_k\}$ denotes the output sequence, and the SSM

matrices $\bar{A}, \bar{B}, \bar{C}, \bar{D}$ are defined as follows:

$$\begin{aligned} \bar{A} &= \exp(\Delta A), \quad \bar{B} = (\Delta A)^{-1}(\exp(A) - I)(\Delta B) \\ \bar{C} &= C, \quad \bar{D} = D. \end{aligned} \quad (3)$$

Selective Scan Modification. The LTI property of SSMs causes them to learn a general set of parameters that are expected to work well for all kinds of inputs. Mamba (Selective State Space Model) [22] addresses this concern by making the parameters input-dependent and shows that this allows the model to perform significantly better. More specifically, the parameters B, C and Δ are dependent on the input x during training, making Mamba capable of effectively learning the complex relations in longer sequences. In addition, Mamba also approximates $\bar{B} \simeq \Delta B$ using Taylor approximation.

3.2. M-CD Architecture

An overview of the proposed M-CD architecture is shown in Figure Fig. 2. M-CD consists of three main components - the Siamese Image Encoder (SIE), the Difference Module (DM) and the Mask Decoder (MD). Given two images, they are passed through the two branches of the encoder to generate image features. The two branches work on the same modality of images and so weights between them are shared. This also reduces the computational complexity. The SIE is responsible for extracting features on multiple scales, facilitated by the cascading of four Visual State Space (VSS) blocks and downsampling operations. The DM is responsible for analyzing features from both images together at different scales and generating combined multi-scale features. These are further transformed by the Mask Decoder using Channel-Averaged VSS blocks and upsampling operations. These transformed features are finally passed through a classifier that segments out the regions involving significant change. In the subsequent subsections, we describe each of the three components of M-CD in greater detail.

3.3. Siamese Image Encoder (SIE)

The SIE has two branches with common weights, each of which takes an RGB image. One of these is the pre-change image and the other is the post-change image. These images go through a stem module which is a series of 2D convolution layers to extract preliminary features, similar to Vision Transformer (ViT) [20]. This is followed by the VSS block for further processing. As seen in Fig. 3, the VSS block consists of a linear layer, followed by a Depth-wise Convolution layer, similar to the original Mamba [22]. This is followed by the Selective Scan 2D (SS2D) Module, similar to VMamba [45] and MambaIR [25]. As shown in Fig. 3, The SS2D module flattens the input and processes it from four different directions - top-left to bottom-right, bottom-right to top-left, top-right to bottom-left, and bottom-left to

top-right. This is done to extract long-range dependencies from multiple directions, as suggested by VMamba [45]. For each of the directions, the selective scan state model is learnt by learning the parameters A, B, C, D, Δ , and the results are merged to produce the final output of the SS2D block. Lastly, the VSS block consists of a layernorm operation followed by a linear operator. The SIE consists of four such VSS blocks. After the first block, each of the VSS blocks are preceded by a downsampling module. Thus, the SIE produces deep multi-scale features as the outputs of each of the VSS blocks for both the input images.

3.4. Difference Module (DM)

As seen in Fig. 4, the difference module takes in features corresponding to the two images (pre-change and post-change), and outputs one feature vector to be used by the mask decoder. Each of the input features passes through a linear and a depth-wise convolution operation, followed by a Joint Selective Scan (JSS) module. The JSS module concatenates the outputs from the previous layer in two ways - $Pre; Post$ and $Post; Pre$, where Pre denotes the pre-change image, $Post$ denotes the post-change image, and ; represents concatenation. The selective scan operations are performed on these two concatenated features and the outputs are added. Unlike transformers which use self-attention to attend to each token, the DM uses a scanning operation. Hence, it is important to concatenate in both directions to maintain symmetry of the network and aid the training process. The resulting vector is then split into the pre-image and post-image parts and passed through a layernorm operation. The linear and convolution operations are used to refine the incoming features, while the JSS module is used for learning to identify the significant differences between the two images. In addition, a residual connection is added between the features before the convolution and after the normalization layer to aid training by reducing the effect of vanishing gradients. The resulting vectors are concatenated and passed through a final linear layer, before passing the result to the mask decoder. For capturing the difference information between images at different scales, we employ four separate DMs corresponding to four different scaled inputs from the image encoder.

3.5. Mask Decoder (MD)

As shown in Fig. 2, the Mask Decoder is responsible for generating the output mask. The Mask decoder consists of a series of Channel-Averaged VSS (CAVSS) blocks and upsampling operations. The architecture of a CAVSS block is shown in Fig. 5. It has a VSS block in the beginning that processes the output from the Difference Module. However, while the VSS block can extract global context well, it struggles to learn the inter-channel dependencies [64]. Hence, the CAVSS additionally contains average-

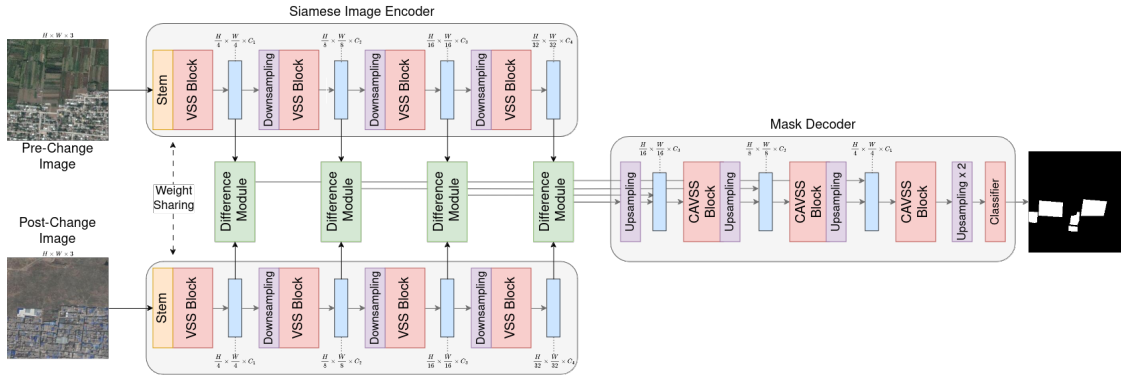


Figure 2. An overview of the M-CD Architecture. Given a pre-change and post-change image, they are passed through a Mamba-based encoder with shared weights (Siamese) and features at multiple scales are extracted. The Difference Module combines them before sending to the Mask Decoder, which uses a Mamba-based Decoder to generate the predicted change map.

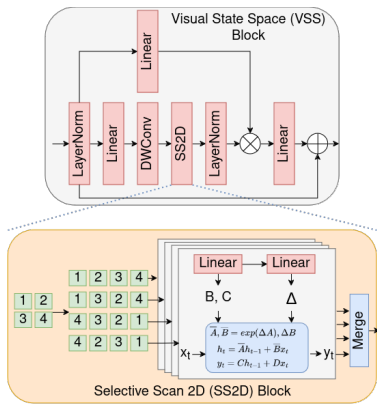


Figure 3. Architecture of Visual State Space Block and Selective Scan 2D block.

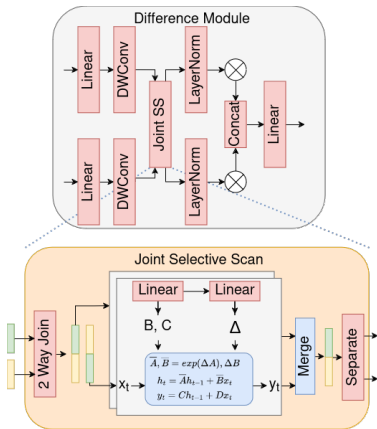


Figure 4. Architecture of the Difference Module and Joint Selective Scan.

pool and max-pool operations along the channel dimension, that cater to the channel information, similar to CBAM [68]. Thus, the VSS and the Channel-Averaged Pooling layers

learn the spatial as well as the channel context. The Mask Decoder follows a UNet [56] structure, where the output of a Difference Module passes through the CAVSS block, is upsampled and added to the output of the previous Difference Module. Adding these skip connections allows the model to benefit from a wider context and reduce the effect of vanishing gradients. Finally, the output from the last CAVSS block is upsampled to the original image size and fed into a classifier that classifies each pixel as “Change” or “No Change” classes.




Figure 5. Architecture of Channel Averaged VSS Block

4. Experiments and Results

Datasets: We compare our method with current state-of-the-art (SOTA) methods for change detection on four widely used public datasets: WHU-CD [35], DSIFN-CD [72], LEVIR-CD [13], and CDD [40]. For all the datasets, we use the same preprocessing and data splits as [4]. WHU-CD consists of paired aerial images from 2012 and 2016. It represents an earthquake-affected area, with many changes including rebuilt and new buildings. The training set consists of 5947 image pairs, while the validation set and test set have 743 and 744 pairs respectively. DSIFN-CD [72] has images from six cities in China from different times. The training, validation and test sets have 14400, 1360, and 192 image pairs respectively. LEVIR-CD [13] consists of remote sensing images, collected from

Google Earth and majorly covers building appearances and disappearances. It has 7120 image pairs in training set, 1024 in validation set and 2048 in the testing set. CDD [40] consists of Google Earth remote sensing images from varying seasons. It caters to changes in large architectural structures, cars, and seasonal changes in natural vegetation. All datasets consist of high-resolution images and the data splits were created using 256×256 crops from the original images, and obtained from [4].

Experimental Setup: We use the AdamW optimizer [46] with initial learning rate $6e^{-5}$ and weight decay 0.01. The batch size during training is kept at 8 and the model is trained for 150 epochs. For initialization of the SIE, we use the ImageNet 1k pretrained model VMamba-Small [45]. The training is divided among four Nvidia RTX A5000 GPUs, each taking 6GB memory, while just one of these is used during evaluation. Following [3], we use three metrics for evaluation - F1 score, Intersection-Over-Union (IoU) and Overall Accuracy (OA).

Results: The quantitative results comparing M-CD with existing SOTA methods are tabulated in Tab. 1. The first three rows in the table represent methods that employ CNNs for CD. Among these, SNUNet [21] and FC-Siam-conc [8] start training with random weight initialization, while DT-SCN [44] and IFNet [71] use pretrained ImageNet 1k weights for initializing the models. While convolution operations are useful in capturing the spatial context of an image, they lack in comparing the temporal differences between the two images in CD. In contrast, STANet [12] uses spatial and temporal attention, while BIT [11] and ChangeFormer [4] adapt transformers and employ cross-attention operations between the two images, which significantly improves results over convolution-based methods.

Other recent methods like SiamSiam [15] and SaDL-CD [10] require a large amount of pertaining and hence, we only present the results which were mentioned in their original papers, namely the F1 score and IoU score on the WHU-CD and LEVIR-CD datasets. We find that our method is able to significantly outperform all the self-supervision-based approaches without requiring a large amount of data. This shows the effectiveness of the Mamaba-based modeling over strong pretrained encoders. DDPM-CD [3] instead learns a diffusion model to learn the structure of remote sensing images and uses its encoder to generate strong predictions, thus being the current SOTA for CD.

On the other hand, it was shown that Selective State Space Models like Mamba are capable of learning through a larger receptive field as compared to convolutions or transformer-based mechanisms [22, 45]. This can be seen through the on-par performance of Mamba-based methods like RSMamba [74], CDMamba [73] and ChangeMamba

[14]. However, we see that they do not improve the performance significantly. In contrast, by using multi-scale features and a dedicated module to model the temporal difference, we see that our method M-CD significantly improves over existing methods. M-CD is able to outperform CNN-based methods by around 7-10% in IoU score, and outperform transformer-based methods by around 5%. In addition, it also performs better than DDPM-CD and Mamba-based methods, with a rise of around 3% in IoU score. At the same time, our method is also initialized with ImageNet 1k weights and so it does not require a large amount of pretraining with remote sensing image data, unlike DDPM-CD. All results with M-CD have a p-value of at most 10^{-5} compared to other methods, indicating statistical significance.

We present the qualitative results of our method in Fig. 6. We see that while DDPM-CD generates some artefacts in the output masks, M-CD generates cleaner segmentation predictions, with more precise boundaries.

5. Ablation Studies

Varying Mamba Backbones: In our main experiments, we use the VMamba-Small backbone in the SIE. VMamba [45] provides three different model settings, namely Vmamba Tiny, Small and Base in the increasing order of complexity. We measure the effect of the choice of the backbone for our task by comparing the results on LEVIR-CD in Tab. 2. We find that while there is no significant difference in the performance, VMamba-Small performs slightly better. One reason for this could be that the number of parameters from VMamba-Small to Base increase by 70%, which can be more prone to overfitting. Hence, there is no significant improvement, when going from Small to Tiny. A similar argument can also hold while going from Tiny to Small.

Varying Difference Module: We measure the effect of the design choice of the Difference Module (DM) on model performance for LEVIR-CD and DSIFN-CD datasets in Tab. 3. We evaluate three different designs for the Difference Module. In the first row, “Difference” signifies that the DM simply subtracts the post-change image from the pre-change image and sends the difference to the decoder. While this approach seems naive, this allows the decoder to learn from the residuals between the image features and performs well on the Change Detection task. In the second row, we use the strategy of combining of Cross Mamba and Concat Mamba blocks used by Sigma [64] to generate segmentation masks using multiple modalities. This technique performs well when multiple modalities guide the model towards a single segmentation results. However, for the case of Change Detection, we find that interchanging the parameters for the two images during training, as proposed by Sigma leads to a lower model performance. This could be because an interchange encourages the model to treat the

Method	Extra training data	WHU-CD [35]			DSIFN-CD [72]			LEVIR-CD [13]			CDD [40]		
		F1 (↑)	IoU (↑)	OA (↑)	F1 (↑)	IoU (↑)	OA (↑)	F1 (↑)	IoU (↑)	OA (↑)	F1 (↑)	IoU (↑)	OA (↑)
CNN-based Methods:													
FC-Siam-conc [8]	None	0.798	0.665	98.5	0.597	0.426	87.6	0.837	0.720	98.5	0.751	0.601	94.9
SUNet [21]	None	0.835	0.717	98.7	0.662	0.495	87.3	0.882	0.788	98.8	0.839	0.721	96.2
IFNet [71]	IN1k	0.834	0.715	98.8	0.601	0.430	87.8	0.881	0.788	98.9	0.840	0.719	96.03
CNN + Attention based Methods:													
DT-SCN [44]	IN1k	0.914	0.842	99.3	0.706	0.545	82.9	0.877	0.781	98.8	0.921	0.853	98.2
STANet [12]	IN1k	0.823	0.700	98.5	0.645	0.478	88.5	0.873	0.774	98.7	0.841	0.722	96.1
Transformer-based Methods:													
BIT [11]	IN1k	0.905	0.834	99.3	0.876	0.780	92.3	0.893	0.807	98.92	0.889	0.800	97.5
ChangeFormer [4]	None	0.886	0.795	99.12	0.947	0.887	93.2	0.904	0.825	99.0	0.946	0.898	98.7
Methods with self supervised pretraining													
SiamSiam [15]	IN1k, IBSD, Google Earth	0.847	0.734	-	-	-	-	0.880	0.786	-	-	-	-
MoCo-v2 [52]	IN1k, IBSD, Google Earth	0.882	0.789	-	-	-	-	0.879	0.784	-	-	-	-
DenseCL [66]	IN1k, IBSD, Google Earth	0.867	0.765	-	-	-	-	0.869	0.780	-	-	-	-
CMC [9]	IN1k, IBSD, Google Earth	0.886	0.795	-	-	-	-	0.877	0.780	-	-	-	-
SeCo [49]	IN1k, IBSD, Google Earth	0.883	0.790	-	-	-	-	0.881	0.787	-	-	-	-
SaDL-CD [10]	IN1k, IBSD, Google Earth	0.909	0.833	-	-	-	-	0.899	0.818	-	-	-	-
Mamba-based Methods													
RSMamba [74]	IN1k	0.927	0.865	99.4	0.965	0.913	97.0	0.897	0.814	98.9	0.943	0.902	98.8
ChangeMamba [14]	IN1k	0.925	0.861	99.4	0.875	0.778	95.8	90.16	0.821	99.0	0.944	0.920	99.0
CDMamba [73]	IN1k	<u>0.937</u>	<u>0.882</u>	<u>99.5</u>	0.966	<u>0.914</u>	97.0	0.907	0.831	99.0	<u>0.960</u>	0.919	<u>99.1</u>
Diffusion-based Methods													
DDPM-CD [3]	Google Earth	0.927	0.863	99.4	<u>0.967</u>	0.913	<u>97.1</u>	<u>0.909</u>	<u>0.833</u>	<u>99.1</u>	0.956	0.916	99.0
M-CD (Ours)	IN1k	0.953	0.911	99.6	0.970	0.935	98.9	0.921	0.850	99.2	0.982	0.963	99.5

Table 1. Comparison of M-CD with respect to SOTA CD methods. F1 denotes the F1 metric, IoU denotes the Intersection-Over-Union metric and OA denotes the overall pixel accuracy. IN1k denotes training data from ImageNet 1k dataset [18]. The best result is indicated in **bold** and the second-best result is underlined. Our method outperforms existing methods for all datasets.




Figure 6. Qualitative results on four public datasets. White represents true positives, black represents true negatives, green represents false positives and red represents false negatives.

post-change image as being similar to the pre-change image semantically. This proves to be beneficial when fusing two modalities like in Sigma, but is harmful for the task of Change Detection. Hence, we see a lower performance. Finally, the last row represents our method, which uses the concatenation of the feature vectors from the two images. This allows the DM to combine the two features with more

freedom than simple difference, while not enforcing similarity. Hence, it achieves a higher performance than the other two strategies.

Computational Complexity Analysis: We calculate the compute required for running our method using three metrics, namely the number of trainable parameters the model has, amount of floating point operations required for a for-

Backbone	Number of Parameters	LEVIR-CD IoU
Tiny	36.5 M	0.848
Small	58.1 M	0.850
Base	100.7 M	0.845

Table 2. Ablation Experiments with varying encoder backbones. While the number of parameters increase significantly, there is no significant change in model performance on LEVIR-CD.

Difference Module Design Strategy	LEVIR-CD IoU	DSIFN-CD IoU
Difference	0.840	0.924
Cross+Concatenation [64]	0.821	0.909
Concatenation (Ours)	0.850	0.935

Table 3. Ablation Experiments with varying DM designs. Our method (concatenation) shows higher performance on the LEVIR-CD and DSIFN-CD datasets over other strategies like difference of features, or combining parameter crossing and concatenation, as proposed in [64].

Method	Trainable Parameters (million)	GFLOPS	Average Inference Time Per Image Pair (ms)
SiamSiam	12.49	4.76	1.04
MoCo-v2	11.24	4.76	1.92
DenseCL	11.69	4.76	2.66
CMC	22.48	4.66	1.55
ScCo	12.16	9.52	3.62
DDPM-CD [3]	46.41	2175.46	88.10
M-CD (Ours)	69.80	29.58	160

Table 4. Computational Complexity Analysis of M-CD. Our method has greater trainable parameters and inference time. The number of GFLOPs are comparable. However, while this is a limitation of M-CD, we believe the improvement in performance without additional training data outweighs the shortcomings.

ward pass, in gigaflops (GFLOPS) and the average inference time required for one pair of pre and post-change image inputs. The image size for these experiments is 256X256 and one Nvidia RTX A5000 GPU is used for the calculation of these metrics. The results are tabulated in Tab. 4. We see that our method has a greater number of trainable parameters as well as has a higher inference time as compared to other SOTA methods. However, GFLOPS in Mamba-based architectures have a linear scaling relation with the number of parameters [22]. Hence, the number of floating point operations are comparable to other methods and significantly lower than DDPM-CD. For rest of the methods, while the computational complexity of our method is greater, the gain in performance, especially in the absence of large amounts of pertaining, outweighs this limitation in our opinion.

Effective Receptive Field (ERF): We utilize the code provided by VMamba [45] to generate the ERF for our method (Mamba-based) in comparison to ChangeFormer [4] (transformer-based) and SNUNet [21] (CNN-based). This represents the importance of every input image pixel for generating the center of the predicted mask. As shown in Fig. 7, using Mamba architecture includes information




Figure 7. Effective Receptive Field (ERF) of SNUNet v.s. ChangeFormer v.s. M-CD for WHU dataset. The top row indicates ERF with respect to the pre-change image (A) and the bottom row with respect to the post-change image (B). Darker shade of green indicates higher dependence on the input pixel. M-CD and Transformer have a more global ERF than SNUNet. M-CD has a more structured ERF than the uniform ERF of ChangeFormer

from both the input images at a more global level compared to CNNs, as seen from the darker shades of green at non-centric parts of the ERF. On the other hand, the transformer has a uniformly distributed attention map unlike Mamba, which is more structured. This shows that Mamba can achieve a wider receptive field across both the pre and post-change images, and can find and prioritize certain areas of the inputs which are more relevant for CD. These results are consistent with similar studies conducted in [45].

6. Conclusion

In this work, we introduce M-CD, a model for the task of Change Detection, that is based upon the Selective State Space (a.k.a. Mamba) architecture. M-CD comprises of a Siamese Image Encoder to featurize the input images, a Difference Module that performs multi-scale analysis of the change of interest, and finally, a Mask Decoder to generate the segmentation masks. Extensive experiments on four widely used datasets clearly show that M-CD beats the existing SOTA methods. At the same time, it requires much lesser pretraining than the recently proposed self-supervised or diffusion methods, while outperforming them significantly, thus making M-CD the new SOTA in this field.

Acknowledgments: This Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-23-2-0008. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- [1] Salem Saleh Al-amri, N. V. Kalyankar, and S. D. Khamitkar. A comparative study of removal noise from remote sensing image, 2010. [1](#)
- [2] Ting Bai, Le Wang, Dameng Yin, Kaimin Sun, Yepei Chen, Wenzhuo Li, and Deren Li. Deep learning for change detection in remote sensing: a review. *Geo-spatial Information Science*, 26(3):262–288, 2023. [1](#)
- [3] Wele Gedara Chaminda Bandara, Nithin Gopalakrishnan Nair, and Vishal M. Patel. Ddpm-cd: Denoising diffusion probabilistic models as feature extractors for change detection, 2024. [1](#), [3](#), [6](#), [7](#), [8](#)
- [4] Wele Gedara Chaminda Bandara and Vishal M. Patel. A transformer-based siamese network for change detection. In *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium*, pages 207–210, 2022. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [5] Jon Atli Benediktsson, Jocelyn Chanussot, and Wooil M. Moon. Very high-resolution remote sensing: Challenges and opportunities [point of view]. *Proceedings of the IEEE*, 100(6):1907–1910, 2012. [1](#)
- [6] Y. Bentoutou, N. Taleb, K. Kpalma, and J. Ronsin. An automatic image registration for applications in remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 43(9):2127–2137, 2005. [1](#)
- [7] Jose Bolorinos, Newsha K. Ajami, and Ram Rajagopal. Consumption change detection for urban planning: Monitoring and segmenting water customers during drought. *Water Resources Research*, 56(3):e2019WR025812, 2020. [1](#)
- [8] Rodrigo Caye Daudt, Bertr Le Saux, and Alexandre Boulch. Fully convolutional siamese networks for change detection. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 4063–4067, 2018. [1](#), [2](#), [6](#), [7](#)
- [9] Keumgang Cha, Junghoon Seo, and Yeji Choi. Contrastive multiview coding with electro-optics for sar semantic segmentation. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022. [2](#), [7](#)
- [10] Hao Chen, Wenyuan Li, Song Chen, and Zhenwei Shi. Semantic-aware dense representation learning for remote sensing image change detection. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–18, 2022. [2](#), [6](#), [7](#)
- [11] Hao Chen, Zipeng Qi, and Zhenwei Shi. Remote sensing image change detection with transformers. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2022. [1](#), [2](#), [6](#), [7](#)
- [12] 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), 2020. [1](#), [6](#), [7](#)
- [13] 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), 2020. [5](#), [7](#)
- [14] Hongruixuan Chen, Jian Song, Chengxi Han, Junshi Xia, and Naoto Yokoya. Changemamba: Remote sensing change detection with spatio-temporal state space model, 2024. [3](#), [6](#), [7](#)
- [15] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML’20*. JMLR.org, 2020. [2](#), [6](#), [7](#)
- [16] Nicholas Coops, Piotr Tompalski, Tristan Goodbody, Alexis Achim, and Christopher Mulverhill. Framework for near real-time forest inventory using multi source remote sensing data. *Forestry*, pages 1–19, 05 2022. [1](#)
- [17] E. Brondízio D. Lu Corresponding author, P. Mausel and E. Moran. Change detection techniques. *International Journal of Remote Sensing*, 25(12):2365–2401, 2004. [2](#)
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [7](#)
- [19] M. Dharani and G. Sreenivasulu. Land use and land cover change detection by using principal component analysis and morphological operations in remote sensing applications. *International Journal of Computers and Applications*, 43(5):462–471, 2021. [2](#)
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [2](#), [4](#)
- [21] Sheng Fang, Kaiyu Li, Jinyuan Shao, and Zhe Li. Snunet-cd: A densely connected siamese network for change detection of vhr images. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022. [1](#), [2](#), [6](#), [7](#), [8](#)
- [22] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces, 2023. [1](#), [3](#), [4](#), [6](#), [8](#)
- [23] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces, 2022. [3](#)
- [24] Albert Gu, Isys Johnson, Karan Goel, Khaled Saab, Tri Dao, Atri Rudra, and Christopher Ré. Combining recurrent, convolutional, and continuous-time models with linear state space layers. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 572–585. Curran Associates, Inc., 2021. [3](#)
- [25] Hang Guo, Jinmin Li, Tao Dai, Zhihao Ouyang, Xudong Ren, and Shu-Tao Xia. Mambair: A simple baseline for image restoration with state-space model, 2024. [3](#), [4](#)
- [26] Ebrahim Hamidi, Brad G. Peter, David F. Muñoz, Hamed Moftakhari, and Hamid Moradkhani. Fast flood extent monitoring with sar change detection using google earth engine. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–19, 2023. [1](#)
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [2](#)
- [28] Xiaoning He, Shuangcheng Zhang, Bowei Xue, Tong Zhao, and Tong Wu. Cross-modal change detection flood extraction based on convolutional neural network. *International Journal of Applied Earth Observation and Geoinformation*, 117:103197, 2023. [1](#)

- [29] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA, 2020. Curran Associates Inc. **3**
- [30] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269, 2017. **2**
- [31] Jordi Inglada, Vincent Muron, Damien Pichard, and Thomas Feuvrier. Analysis of artifacts in subpixel remote sensing image registration. *IEEE Transactions on Geoscience and Remote Sensing*, 45(1):254–264, 2007. **1**
- [32] Md Islam and Gedas Bertasius. *Long Movie Clip Classification with State-Space Video Models*, pages 87–104. 11 2022. **3**
- [33] Md. Mohaiminul Islam, Mahmudul Hasan, Kishan Athrey, Tony Braskich, and Gedas Bertasius. Efficient movie scene detection using state-space transformers. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18749–18758, 2023. **3**
- [34] Y. H. Deng J. S. Deng, K. Wang and G. J. Qi. Pca-based land-use change detection and analysis using multitemporal and multisensor satellite data. *International Journal of Remote Sensing*, 29(16):4823–4838, 2008. **2**
- [35] 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, 2019. **5, 7**
- [36] Pongpun Juntakut, Peerasak Aemlaor, and Yaowaret Jantakut. Detecting change in military sites using satellite imagery case study: Russia-ukraine war in 2022. *Defence Technology Academic Journal*, 4(10):48–57, Oct. 2022. **1**
- [37] Ravneet Kaur, Reet Tiwari, Raman Maini, and Sartajvir Singh. A framework for crop yield estimation and change detection using image fusion of microwave and optical satellite dataset. *Quaternary*, 6:28, 04 2023. **1**
- [38] Lazhar Khelifi and Max Mignotte. Deep learning for change detection in remote sensing images: Comprehensive review and meta-analysis. *IEEE Access*, 8:126385–126400, 2020. **1**
- [39] David A. Landgrebe and Erick Malaret. Noise in remote-sensing systems: The effect on classification error. *IEEE Transactions on Geoscience and Remote Sensing*, GE-24(2):294–300, 1986. **1**
- [40] M. A. Lebedev, Y. V. Vizilter, O. V. Vygolov, V. A. Knyaz, and A. Y. Rubis. 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. **5, 6, 7**
- [41] Ding Kang Liang, Xin Zhou, Wei Xu, Xingkui Zhu, Zhikang Zou, Xiaoqing Ye, Xiao Tan, and Xiang Bai. Pointmamba: A simple state space model for point cloud analysis, 2024. **3**
- [42] Jun Liu, Xing Wang, Min Chen, Shuguang Liu, Zhenfeng Shao, Xiran Zhou, and Ping Liu. Illumination and contrast balancing for remote sensing images. *Remote. Sens.*, 6:1102–1123, 2014. **1**
- [43] Qunqun Liu, Shiquan Wan, and Bin Gu. A review of the detection methods for climate regime shifts. *Discrete Dynamics in Nature and Society*, 2016:1–10, 01 2016. **1**
- [44] Yi Liu, Chao Pang, Zongqian Zhan, Xiaomeng Zhang, and Xue Yang. Building change detection for remote sensing images using a dual-task constrained deep siamese convolutional network model. *IEEE Geoscience and Remote Sensing Letters*, 18(5):811–815, 2021. **1, 2, 6, 7**
- [45] Yue Liu, Yunjie Tian, Yuzhong Zhao, Hongtian Yu, Lingxi Xie, Yaowei Wang, Qixiang Ye, and Yunfan Liu. Vmamba: Visual state space model, 2024. **1, 3, 4, 6, 8**
- [46] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. **6**
- [47] Luigi Tommaso Luppino, Filippo Maria Bianchi, Gabriele Moser, and Stian Normann Anfinsen. Unsupervised image regression for heterogeneous change detection. *IEEE Transactions on Geoscience and Remote Sensing*, 57(12):9960–9975, 2019. **2**
- [48] Jun Ma, Feifei Li, and Bo Wang. U-mamba: Enhancing long-range dependency for biomedical image segmentation, 2024. **3**
- [49] Oscar Mañas, Alexandre Lacoste, Xavier Giró-i Nieto, David Vazquez, and Pau Rodríguez. Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9414–9423, October 2021. **2, 7**
- [50] Anannya Mazumder Md Ariful Haque, Sharmin Shishir and Mehedi Iqbal. Change detection of jamuna river and its impact on the local settlements. *Physical Geography*, 44(2):186–206, 2023. **1**
- [51] Eric Nguyen, Karan Goel, Albert Gu, Gordon Downs, Preey Shah, Tri Dao, Stephen Baccus, and Christopher Ré. S4nd: Modeling images and videos as multidimensional signals with state spaces. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 2846–2861. Curran Associates, Inc., 2022. **3**
- [52] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8162–8171. PMLR, 18–24 Jul 2021. **2, 7**
- [53] Allan Aasbjerg Nielsen. The regularized iteratively reweighted mad method for change detection in multi- and hyperspectral data. *IEEE Transactions on Image Processing*, 16(2):463–478, 2007. **2**
- [54] Allan A. Nielsen, Knut Conradsen, and James J. Simpson. Multivariate alteration detection (mad) and maf postprocessing in multispectral, bitemporal image data: New approaches to change detection studies. *Remote Sensing of Environment*, 64(1):1–19, 1998. **2**
- [55] Merrill K Ridd and Jiajun Liu. A comparison of four algorithms for change detection in an urban environment. *Remote Sensing of Environment*, 63(2):95–100, 1998. **1**

- [56] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241. Cham, 2015. Springer International Publishing. 5
- [57] Jiacheng Ruan and Suncheng Xiang. Vm-unet: Vision mamba unet for medical image segmentation, 2024. 3
- [58] Sudipan Saha, Francesca Bovolo, and Lorenzo Bruzzone. Unsupervised deep change vector analysis for multiple-change detection in vhr images. *IEEE Transactions on Geoscience and Remote Sensing*, 57(6):3677–3693, 2019. 2
- [59] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015. 2
- [60] Jimmy T. H. Smith, Andrew Warrington, and Scott W. Linderman. Simplified state space layers for sequence modeling, 2023. 3
- [61] Bui Bao Thien, Vu Phuong, and Do Thi Viet Huong. Detection and assessment of the spatio-temporal land use/cover change in the thai binh province of vietnam’s red river delta using remote sensing and gis. *Modeling Earth Systems and Environment*, 9:1–12, 12 2022. 1
- [62] Paul Tueller, Robert Ramsey, Thomas Frank, Robert Washington-Allen, and Scott Tweddale. Emerging and contemporary technologies in remote sensing for ecosystem assessment and change detection on military reservations. page 74, 12 1998. 1
- [63] Xue Wan, Jianguo Liu, Shengyang Li, John Dawson, and Hongshi Yan. An illumination-invariant change detection method based on disparity saliency map for multitemporal optical remotely sensed images. *IEEE Transactions on Geoscience and Remote Sensing*, 57(3):1311–1324, 2019. 1
- [64] Zifu Wan, Yuhao Wang, Silong Yong, Pingping Zhang, Simon Stepputtis, Katia Sycara, and Yaqi Xie. Sigma: Siamese mamba network for multi-modal semantic segmentation, 2024. 2, 3, 4, 6, 8
- [65] Jue Wang, Wentao Zhu, Pichao Wang, Xiang Yu, Linda Liu, Mohamed Omar, and Raffay Hamid. Selective structured state-spaces for long-form video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6387–6397, June 2023. 3
- [66] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3024–3033, June 2021. 2, 7
- [67] Ziyang Wang and Chao Ma. Weak-mamba-unet: Visual mamba makes cnn and vit work better for scribble-based medical image segmentation, 2024. 3
- [68] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 5
- [69] Lei Wu, Rui Liu, Gulin Li, Jingsong Gou, and Yuzhu Lei. Landslide detection methods based on deep learning in remote sensing images. In *2022 29th International Conference on Geoinformatics*, pages 1–4, 2022. 1
- [70] Zhaohu Xing, Tian Ye, Yijun Yang, Guang Liu, and Lei Zhu. Segmamba: Long-range sequential modeling mamba for 3d medical image segmentation, 2024. 3
- [71] Chenxiao Zhang, Peng Yue, Deodato Tapete, Liangcun Jiang, Boyi Shangguan, Li Huang, and Guangchao Liu. A deeply supervised image fusion network for change detection in high resolution bi-temporal remote sensing images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 166:183–200, 2020. 1, 6, 7
- [72] Chenxiao Zhang, Peng Yue, Deodato Tapete, Liangcun Jiang, Boyi Shangguan, Li Huang, and Guangchao Liu. A deeply supervised image fusion network for change detection in high resolution bi-temporal remote sensing images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 166:183–200, 2020. 5, 7
- [73] Haotian Zhang, Keyan Chen, Chenyang Liu, Hao Chen, Zhengxia Zou, and Zhenwei Shi. Cdmamba: Remote sensing image change detection with mamba, 2024. 3, 6, 7
- [74] Sijie Zhao, Hao Chen, Xueliang Zhang, Pengfeng Xiao, Lei Bai, and Wanli Ouyang. Rs-mamba for large remote sensing image dense prediction, 2024. 3, 6, 7
- [75] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In Danail Stoyanov, Zeike Taylor, Gustavo Carneiro, Tanveer Syeda-Mahmood, Anne Martel, Lena Maier-Hein, João Manuel R.S. Tavares, Andrew Bradley, João Paulo Papa, Vasileios Belagiannis, Jacinto C. Nascimento, Zhi Lu, Sailesh Conjeti, Mehdi Moradi, Hayit Greenspan, and Anant Madabhushi, editors, *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 3–11, Cham, 2018. Springer International Publishing. 2
- [76] Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggong Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model, 2024. 3