

This ACCV 2020 workshop paper, provided here by the Computer Vision Foundation, is the author-create version. The content of this paper is identical to the content of the officially published ACCV 2020 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

A Weakly Supervised Convolutional Network for Change Segmentation and Classification

Philipp Andermatt^[0000-0003-3503-8693] and Radu Timofte^[0000-0002-1478-0402]</sup></sup>

Computer Vision Lab, ETH Zürich, Switzerland {anphilip,timofter}@ethz.ch

Abstract. Fully supervised change detection methods require difficult to procure pixel-level labels, while weakly supervised approaches can be trained with image-level labels. However, most of these approaches require a combination of changed and unchanged image pairs for training. Thus, these methods can not directly be used for datasets where only changed image pairs are available. We present W-CDNet, a novel weakly supervised change detection network that can be trained with image-level semantic labels. Additionally, W-CDNet can be trained with two different types of datasets, either containing changed image pairs only or a mixture of changed and unchanged image pairs. Since we use image-level semantic labels for training, we simultaneously create a change mask and label the changed object for single-label images. W-CDNet employs a Wshaped siamese U-net to extract feature maps from an image pair which then get compared in order to create a raw change mask. The core part of our model, the Change Segmentation and Classification (CSC) module, learns an accurate change mask at a hidden layer by using a custom Remapping Block and then segmenting the current input image with the change mask. The segmented image is used to predict the image-level semantic label. The correct label can only be predicted if the change mask actually marks relevant change. This forces the model to learn an accurate change mask. We demonstrate the segmentation and classification performance of our approach and achieve top results on AICD and HRSCD, two public aerial imaging change detection datasets as well as on a Food Waste change detection dataset. Our code is available at: https://github.com/PhiAbs/W-CDNet

1 Introduction

Change detection [1, 2] is an important computer vision task. It has applications in remote sensing [3, 4], video surveillance [5] and street view imaging [6, 7], amongst others. It is a challenging task since one has to distinguish relevant changes between two temporally different images from noise as well as from irrelevant semantic changes. An image pair can belong to one of two classes: *Changed* which means that there are relevant changes between the two images, or *unchanged* which means that there are either no changes at all or only irrelevant ones. The user has to define in advance what counts as relevant change.

Many data-driven change detection methods are trained with pixel-level labels [6–8]. However, creating a pixel-level change detection dataset is costly and time-consuming since two images have to be compared by hand and labeled on pixel-level. As a result, there is a need for change detection approaches which can be trained with simpler labels, for example image-level labels or bounding boxes. Additionally, many weakly supervised change detection approaches need both changed and unchanged images for training. This reduces their usability to datasets where unchanged images are actually available or can be generated with image augmentation.



Fig. 1. Overview of the proposed W-CDNet solution.

In this work, we propose W-CDNet, a data-driven change detection network that can be trained with image-level semantic labels. By image-level semantic labels we mean that the labels describe the changed object and not only the binary state of the image pair (*changed* or *unchanged*, as used in other works [9–11]). Thus, the label *changed* is further split up into several semantic labels like *apple* or *bread* in case of a Food Waste dataset.

In contrast to other approaches [9–11], our method can be trained in one of two ways: (i) With changed image pairs only, meaning that every image pair in the dataset contains at least some relevant change; (ii) With a mix of *changed* and *unchanged* image pairs.

Fig. 1 shows a high-level overview of our W-CDNet model. We use a W-shaped siamese network [12] based on two U-Nets [13] with shared weights for feature extraction. A group of custom Comparison Blocks compares the feature maps from the siamese network and creates a high-resolution feature map containing information about the differences between the two images. Our custom Change Segmentation and Classification (CSC) module remaps these features to create a meaningful change mask which is then used to segment the current input image. The segmented image gets encoded and, together with the feature vector from the Residual Block, gets fed to a classifier which predicts an image-level label.

If the change mask created within the CSC module does not contain relevant change, the classifier is not able to predict the correct image-level label. At the same time, the model can not simply focus on the complete image thanks to our remapping function (part of the CSC module). Thus, the model is forced to learn a change mask which marks relevant changes and suppresses irrelevant changes, with weak supervision. Moreover, we integrate a CRF-RNN layer [14] to refine the change mask created within the CSC module. Our approach works for single-label images. Since it is trained with image-level semantic labels, we can directly classify the changed object.

The key contributions of our work include:

- (i) W-CDNet, a novel change detection network which can be trained with weak supervision using image-level semantic labels. To the best of our knowledge we are the first ones to use image-level *semantic* labels instead of binary labels for weak supervision.
- (ii) CSC, a novel Change Segmentation and Classification module that creates a high-resolution change mask at a hidden layer and encodes the segmented current input image. The CSC module is crucial for W-CDNet's performance.
- (iii) We show that our W-CDNet improves the state-of-the-art for weakly supervised change detection on the AICD dataset [15]. Additionally, we are the first ones to report weakly supervised change detection results on the HRSCD [16] dataset. We also report performance on the Food Waste dataset.
- (iv) We make newly collected image-level semantic labels for the AICD dataset publicly available.

2 Related Work

Change detection plays an important role in remote sensing [2, 17, 18, 16] and in street view imagery [6, 19, 7]. Today, with the advent of deep learning [20, 21], data-driven approaches are used most often as they lead to better performance than previous approaches [22–24]. These related works can be categorized into fully supervised, weakly supervised and unsupervised change detection data-driven methods.

2.1 Fully Supervised Change Detection

Many methods are trained in full supervision with pixel-level labels to obtain a good segmentation performance. Depending on the exact task, either binary change masks or semantic change masks are used.

Guo *et al.* [19] proposed a siamese fully convolutional network. They subtracted the features from three hidden layers, used the result to create a change mask and proposed a new loss to punish noisy changes. Jiang *et al.* [25] introduced a pyramid feature-based attention-guided siamese network using a global co-attention mechanism to emphasize correlations between input feature pairs. Daudt *et al.* [8] presented three fully convolutional networks for change detection. Two use a siamese architecture with shared weights as an encoder. A decoder uses features from different hidden layers from the encoder to create a change mask. Lebedev *et al.* [18] created a conditional adversarial network for fully supervised change detection.

Bu *et al.* [26] proposed a method for change detection in unregistered images. A first module roughly predicts change areas and matching information for an image pair and a second one refines the change areas. Daudt *et al.* [16] proposed four methods to perform semantic change detection, all of which use the same encoder-decoder structure with residual connections. They either directly predict a semantic change mask or an individual change mask and semantic labels which are then combined in a second step.

2.2 Weakly Supervised and Unsupervised Change Detection

Sakurada *et al.* [27] proposed a siamese network which first creates a change mask and then performs semantic segmentation of the changed regions. Both model parts are trained with pixel-level labels. They use separate change detection and semantic segmentation datasets for training. Jiang *et al.* [10] used a siamese network in combination with a conditional random field (CRF) [28]. They only use image-level labels for training. Additionally, they proposed a weighted global average pooling layer (WGAP) which allows them to jointly predict pixel-level and image-level labels.

Yu *et al.* [17] used active learning and markov random fields to reduce the number of needed ground truth change masks. Jong *et al.* [29] trained their model on a standard segmentation dataset. They then compared extracted features of two images in order to create a difference image.

Minematsu *et al.* [11] proposed a convolutional neural network (CNN) which uses a simple pixel-wise subtraction of the two input images as an additional constraint. Their model is trained with image-level labels only. Chianucci *et al.* [30] used spatial transformer networks [31] and a sliding window to detect changes between two images. They did not perform change segmentation but rather simple change detection with bounding boxes. Khan *et al.* [9] used a CNN in combination with a CRF in order to simultaneously perform binary classification and create a change mask. They use a CNN to predict a first change mask and a CRF to refine the change mask. They trained their model with image-level binary labels.

In contrast to the existing weakly supervised change detection approaches, our method can be trained with image-level *semantic* labels which describe the changed object. Our proposed Change Segmentation and Classification (CSC) module allows us to create a change mask and simultaneously label the changed object for single-label images. As a result, our method can be trained with datasets which contain *changed* image pairs only or with datasets which contain a mix of *changed* and *unchanged* image pairs. Additionally, in contrast to other works [11, 10, 9], our model creates a change mask and segments the image containing change at a hidden layer. The resulting segmented image is then refined using a CRF-RNN [14] layer and used to predict an image-level semantic label. Akin to [25, 19, 8, 27, 10], we too employ a siamese network to extract feature maps from an image pair. However, we propose a W-shaped siamese network based on two U-Nets [13] and introduce a custom Comparison Block to compare these feature maps before our novel CSC module to jointly segment and classify the change.

3 Proposed Method

In this section we introduce our proposed W-CDNet. We start by providing an overview of the processing pipeline and architecture, then we provide the details for each of the major components. We also explain how the Change Segmentation and Classification (CSC) module lets the model learn an accurate change mask at a hidden layer.

3.1 W-CDNet overview

We introduce W-CDNet (see Fig. 2), a weakly supervised convolutional network which detects and segments relevant pixel-level changes in an image pair. The proposed method was trained with image-level semantic labels, as explained in introductory section 1. The model takes two temporally different, co-registered RGB images as input (called *previous image* and *current image*). If the image pairs were not yet co-registered, we aligned them in a pre-processing step.

A siamese model (shared weights), which consists of two U-Nets [13], extracts feature maps from the image pair. These feature maps get compared by a chain of Comparison Blocks. By comparing features from several hidden layers, we leveraged both high-level and low-level information which allowed us to filter out unwanted changes while still creating an accurate high-resolution change mask.

The output of the last Comparison Block is a raw change mask, which gets processed by the CSC module. The CSC module is the core of W-CDNet: It lets the model learn a change mask at a hidden layer. The CSC module and the individual blocks are described in more detail in the following sections.

3.2 Siamese Network

The Siamese Network is based on U-Net [13] which uses a VGG16 [32] model as feature extractor. An image pair is used as model input. The U-Nets extract high-level and low-level feature maps from the two images. The feature maps generated before every up-sampling layer are passed on to the Comparison Blocks.

3.3 Comparison Block

This block compares feature maps from the two U-Nets. It highlight differences of relevant features between the two images. Fig. 3 shows the different layers of a single Comparison Block. We used a total of six Comparison Blocks, each of which compares features from a different U-Net layer. This allows the model to use both high-level and low-level information to create a change mask. The feature maps created by one Comparison Block get up-sampled and are fed to the next Comparison Block. The feature maps from the last block get processed by three convolutional layers with ReLU [33] activation, which then output a single-channel raw change mask.



Fig. 2. W-CDNet model architecture.

3.4 CSC module

The CSC module is the core of W-CDNet. It is split up into two independent branches. One branch, consisting only of a Residual Block, extracts features which help to decide whether the image pair contains any relevant change at all.

The other branch further processes the raw change mask in the Remapping Block which assumes that the highest mask activations always correspond to relevant change, independent of whether there actually is relevant change in the image pair or not. The Remapping Block creates a change mask which gets refined by the CRF-RNN [14] layer. The resulting final change mask is used to segment the current input image. The segmented image then gets encoded by the VGG16 [32] model. The Fusion Block concatenates the feature vectors from the two branches and predicts an image-level label. The model has to highlight relevant changes, else the Fusion Block will not be able to predict the correct image-level label. We now describe the blocks from the CSC module in more detail.

Remapping Block. The Remapping Block receives the raw change mask from the Comparison Blocks. The change mask pixels with a high activation (relative to the other pixels) are assumed to mark relevant change. This is enforced by the following formulas:

$$\hat{X}_{map} = \left(\frac{\hat{X} - \min(\hat{X})}{\max(\hat{X}) - \min(\hat{X})} \cdot \alpha\right) - \frac{\alpha}{2} \tag{1}$$

$$\hat{Y} = sig\left(\hat{X}_{map}\right) \tag{2}$$

where \hat{X} is the raw change mask which can contain values from the interval $[0, \infty]$, α is a hyper-parameter, sig() is the sigmoid function and \hat{Y} is the intermediate predicted change mask which now contains values from the interval [0, 1].

Equation (1) shows the intermediate result of the remapping process. Independent of the input range, the values of \hat{X}_{map} lie in the interval $\left[-\frac{\alpha}{2}, \frac{\alpha}{2}\right]$ where $min(\hat{X})$ gets mapped to $-\frac{\alpha}{2}$ and $max(\hat{X})$ gets mapped to $\frac{\alpha}{2}$ (Exception: All features in \hat{X} have the exact same value). Since these values are then fed into a sigmoid function, the hyper-parameter α influences how strongly the values marking relevant change get separated from the values marking irrelevant change. The smaller α is, the smoother the distribution of the change mask values will be. We set α to 32 for training (without CRF-RNN layer) and to 16 for finetuning (including CRF-RNN layer).

The intermediate change mask can now be further refined with a CRF-RNN [14] layer. The final change mask then gets multiplied element-wise with the current input image, which results in an image where only regions with relevant change are visible. This step, combined with the remapping explained above, forces the model to learn an accurate change mask. The segmented image is then fed to a VGG16 [32] model which creates a feature vector describing the region marked by the change mask.

Residual Block. While the Remapping Block assumes that there is always *some* relevant change between the two input images, the Residual Block helps to decide if this is actually true or not. To this end the Residual Block takes the raw change mask as input and creates a feature vector which helps the model to distinguish between changed and unchanged image pairs. Without the Residual Block, the model is not able to do this distinction due to the Remapping Block. It can, however, still learn to segment changes when trained on changed images only. Fig. 3 shows the different layers of the Residual Block.

Fusion Block. The Fusion Block concatenates the feature vectors created by the VGG16 model and the Residual Block and classifies them into N classes, from which N - 1 describe the changed object and one is the unchanged class. Fig. 3 shows the different layers of the Fusion Block.

3.5 Model Variations

W-CDNet can be trained with a combination of changed and unchanged image pairs. By removing the Residual Block from the CSC module, the model can be trained with changed image pairs only. In this case, the Fusion Block has only one input and simply serves as a classifier. It is also possible to train our model with full supervision in a multitask learning setting [34]. In this case, one uses both the change mask and the image-level label as model outputs.



Fig. 3. Comparison Block, Residual Block and Fusion Block in more detail. The colors of the block inputs and outputs refer to the ones in Fig. 2.

3.6 Training

In order to speed up training, we leveraged transfer learning and used publicly available pre-trained weights for U-Net [35]. For the VGG16[32] model we used publicly available pre-trained weights from a training on ImageNet [36]. All other weights were initialized by sampling from a normal distribution.

The model has to be trained in two steps:

(1) Training: The CRF-RNN layer [14] is removed from the model. The model has to learn to correctly label the image pairs (image-level labels). During this process, it learns to create an unrefined change mask at a hidden layer. The boundaries of change objects are not accurately segmented yet, as one can see in Fig. 4.

(2) Finetuning: In a second step, the CRF-RNN layer is inserted into the model and the weights from step 1 for all other layers are loaded. The model is then trained again end-to-end, but with a reduced learning rate. After this training step, the changed objects get segmented more accurately thanks to the segmentation refinement by the CRF-RNN layer, as one can see in Fig. 4.



Fig. 4. Influence of CRF-RNN layer on change segmentation performance, shown for an image pair from the AICD dataset. From left to right: Previous image, Current image with change object, change mask without refinement, change mask with refinement by CRF-RNN layer, ground truth change mask.

Loss Function. We used categorical crossentropy loss to train our weakly supervised model. For fully supervised training we combined two losses: Categorical

crossentropy loss for image-level semantic labels \mathcal{L}_{IL} and a custom Conditional Loss \mathcal{L}_{CM} for the pixel-level change mask. \mathcal{L}_{CM} takes on the form of a standard binary crossentropy loss if max(Y) > 0, else it is set to 0. Y is the ground truth change mask. This Conditional Loss is needed since the model branch, which predicts the change mask, assumes that there is always at least some relevant change in the image pair. Thus, if the image pair does not contain any change, we do not want to induce any loss for the change mask, but only for the image-level semantic label. \mathcal{L}_{IL} and \mathcal{L}_{CM} are summed up to the final loss \mathcal{L} .

4 Experiments and Results

In this section we describe the experimental benchmark (datasets and performance measures), the conducted experiments, and discuss the achieved results in relation to the reported results from the literature. There is only a very limited number of weakly supervised change detection methods which report their results on publicly available datasets. Since the code for these methods is often not publicly available, we compare our results to the ones provided by the authors, by using the performance measures also used by them.

4.1 Datasets

Most change detection datasets contain pixel-level binary labels [37, 15, 38, 6]. In our work, however, we used image-level semantic labels. We could also extract them from pixel-level semantic labels. HRSCD [16], PSCD [27] and SCPA-WC [39] contain pixel-level labels for semantic change detection, but only HRSCD is publicly available. We refer to [1] for an extended review of datasets and related literature. In this work we validated our approach on HRSCD, AICD [15] for which we collected image-level semantic labels, and on the Food Waste dataset.

Food Waste Dataset. The Food Waste dataset contains 57,674 image pairs of which 3,447 pairs have the image-level label *unchanged*. The changed image pairs are spread over 20 classes, which results in a dataset with 21 classes, including the *unchanged* class. Since the Food Waste dataset is very new, most images only contain image-level labels. A total of 145 images also contain pixel-level binary labels. They were used to evaluate segmentation performance. The image-level labels of 1,050 image pairs was double-checked by hand. These pairs were used to evaluate semantic and binary classification performance (50 of which belonged to the *unchanged* class).

The images were resized from $3,280 \times 2,464$ pixels to 512×512 pixels for training and testing. The model was trained with image augmentation.

AICD Dataset. The Aerial Image Change Detection (AICD) dataset [15] is a synthetic dataset consisting of 1,000 image pairs (600×800 pixels), 500 of which include hard shadows. For our experiments we only used the image pairs with

hard shadows. For each image pair, there is only one object which is considered as relevant change. Since these objects are very small, we split up the images into 48 patches of size 122×122 pixels with minimal overlap, which results in 24,000 image pairs. The patches were resized to 128×128 pixels for training. We manually annotated the image pairs with image-level semantic labels. We will make the labels publicly available for further research.

HRSCD Dataset. The *High Resolution Semantic Change Detection* (HRSCD) dataset [16] is a very challenging dataset that contains 291 high-resolution aerial images pairs $(10,000 \times 10,000 \text{ pixels})$. The dataset contains pixel-level semantic labels for 5 classes (Artificial surfaces, Agricultural surfaces, Forests, Wetlands, Water). Additionally, there is a ground truth change mask available for every image pair. We extracted image-level semantic labels from pixel-level semantic labels for the changed regions.

We used 146 image pairs for training and 145 image pairs for testing. We work on image crops of size $1,000 \times 1,000$ pixels which were further resized to 512×512 pixels for training and testing. The original images were cropped without any overlap, which resulted in 29,100 image pairs.

Since our model only works with single-label image pairs but HRSCD contains multi-labels, we created one label for each possible combination of the 5 semantic labels. This resulted in 31 unique semantic labels, plus the additional unchanged label. We oversampled the underrepresented classes and made sure that every class had at least 600 training samples.



Fig. 5. Qualitative results of the predicted change masks. Row 1: Previous image. Row 2: Current image. Row 3: Predicted change mask. Row 4: Ground truth change mask. Columns 1-3: AICD dataset. Columns 4-6: Food Waste dataset. Columns 7-10: HRSCD dataset.

4.2 Performance Measures

For assessing the performance we use standard measures from the literature. For classification performance we employ:

average precision
$$(AP) = \sum_{n} (R_n - R_{n-1})P_n, \quad accuracy = \frac{T}{N}$$

where P_n and R_n are precision and recall at the n-th threshold, T is the number of all true predictions and N is the total number of samples.

For change segmentation performance we use the standard:

$$mIoU = mean(\frac{TP}{TP + FP + FN}), \qquad Kappa = \frac{p_o - p_e}{1 - p_e},$$
$$Dice = \frac{2 \cdot TP}{2 \cdot TP + FP + FN}, \qquad Total \ accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

where TP - true positive, TN - true negative, FP - false positive, FN - false negative, p_o - observed agreement between ground truth and predictions, p_e - expected agreement between ground truth and predictions given class distributions. Cohen's Kappa [40], the Sørensen-Dice coefficient [41] and Total accuracy are defined as in [16].

4.3 Influence of Residual and Comparison Blocks on Performance

We study the importance of the Residual Block in our CSC module by training our W-CDNet with and without that Block. Since the Residual Block is needed to decide whether an image pair contains any relevant change at all, the model without this block must be trained with changed images only. Thus, we compare the change detection and segmentation performance of the full model (dataset contains also unchanged image pairs) to the model without a Residual Block (trained with changed images only). Additionally, we studied the influence on performance of different numbers of filters in the Comparison Blocks. We compare two models where the number of filters was kept constant at either 64 or 256 and a model where the number of filters decreased (256, 128, 64, 32, 16, 16 filters, from first to last Comparison Block).

We show numerical results for tests on the Food Waste dataset in Tab. 1. The models without a Residual Block were trained on changed images only. The models with a Residual Block were trained on unchanged images plus the same changed images as the models without a Residual Block. We show that both models, with and without Residual Block, learn to create a change mask. They achieve similar segmentation performances. However, the influence on performance of the number of filters is very different for the two models. The model with a Residual Block shows the best segmentation results for 64 filters while the model without a Residual Block performs best with a decaying number of filters.

Additionally, the model with a Residual Block and 256 filters did not learn to segment any changes at all. We argue that while the Residual Block is crucial for the model to distinguish between changed and unchanged images, it increases the model complexity which made training unstable for too many filters in the Comparison Blocks.

We also show that change segmentation performance for the model without the Residual Block (and thus for a dataset without unchanged images) is on par with the performance for the model with the Residual Block. This shows that our model learns to distinguish relevant from irrelevant changes already from changed images only.

	Wit	h Re	sidual E	Block	Without Residual Block					
Number of filters	mIoU	AP	Top-1-	Top-5-	mIoII	ΔP	Top-1-	Top-5-		
			acc.	acc.	mioc	111	acc.	acc.		
64	46.5	74.9	68.2	91.7	45.5	77.7	72.2	93.0		
256	0.0	14.3	16.9	49.9	47.4	81.9	74.4	95.5		
Decaying	43.2	68.9	63.9	89.3	47.7	79.4	72.4	94.6		

Table 1. Results on Food Waste dataset for training a model with or without the Residual Block and with different numbers of filters in the Comparison Blocks.

4.4 Influence of CRF Refinement on Performance

In Tab. 2 we compare the change segmentation performance and the binary classification performance of our W-CDNet model without any refinement to a model with a standard CRF [42] as post-processing step and to a model with an integrated CRF-RNN layer [14] for segmentation refinement. The standard CRF was only applied after training the model without any refinement. The CRF-RNN layer was included in the training process.

We observe a significant segmentation performance increase for our model with CRF-RNN layer. The standard CRF (without parameter optimization) applied as post-processing shows only a marginal performance improvement on the AICD dataset and even significantly decreases performance on the Food Waste and HRSCD datasets. This was to be expected, as the CRF-RNN layer learns its parameters during training and adapts to each dataset individually.

4.5 Comparison Results on AICD Dataset

In Tab. 3 and Fig. 5 we report change segmentation and binary as well as semantic change classification results on the AICD dataset [15]. We directly compare our results to the ones reported by Khan *et al.* [9] and Chianucci *et al.* [30]. Additionally, we trained our W-CDNet model with full supervision in a multi-task learning setting as described in section 3.5.

	AICD			Food Waste			HRSCD		
	mIoU	AP	Acc.	mIoU	AP	Acc.	mIoU	AP	Acc.
No refinement	51.6	98.8	99.2	43.3	97.3	98.2	8.4	75.9	81.7
Post-processing with CRF	52.5	98.8	99.2	30.3	97.3	98.2	4.5	75.9	81.7
Integration of a CRF-RNN layer	66.2	98.7	99.2	46.5	89.8	92.9	9.5	75.1	81.0

Table 2. Influence on segmentation and classification performance of refining thechange mask with a standard CRF or with a CRF-RNN layer.

On AICD, we improved the state-of-the-art for weakly supervised change segmentation. However, the model from Khan *et al.* performs slightly better for full supervision. The qualitative results show that our model successfully suppresses irrelevant changes while highlighting relevant change.

Our model had the most difficulties with objects at the edge of the images. Also, our weakly supervised model interpreted hard shadows from changed objects as part of the object itself while the ground truth masks excluded these shadows.

Table 3. Change segmentation and classification results on the AICD dataset [15].

			Wea	ak sup	ervision	n Full supervision				
	Semantic classification			Bin classif	nary fication	Segmentation	Segmentation			
	AP	Top-1 acc.	Top-5 acc	AP	Acc.	mIoU	mIoU			
Khan et al. [9]	-	-	-	97.3	99.1	64.9	71.0			
Bu et al. [26]	-	-	-	-	-	-	64.9			
Chianucci et al. [30]	-	-	-	-	-	57.0	-			
W-CDNet (ours)	99.5	98.9	99.8	98.7	99.2	66.2	70.3			

4.6 Comparison Results on HRSCD Dataset

In Tab. 4 and Fig. 5 we report change segmentation and classification results on the HRSCD dataset [16]. We are the first ones to report weakly supervised change detection performance on HRSCD.

For full supervision we compare our results to the ones from Daudt *et al.* [16]. We compare to the performance of the change detector from their *strategy* 3 which is trained with pixel-level binary labels. Our fully supervised model shows better performance than *strategy* 3 and our weakly supervised model's performance is slightly better. Daudt *et al.* reported better results than us for their *strategy* 4.2, but it was additionally trained with pixel-level semantic labels and

thus was stronger supervised than our method. We refer to their work [16] for the results of all strategies trained by Daudt et al.

The change segmentation results on HRSCD are significantly lower than the ones on Food Waste or AICD. There are three main reasons for that: (i) The HRSCD ground truth change masks are generally too large: they not only contain the changed objects but they mark complete building zones; (ii) HRSCD is highly imbalanced, even more so than the AICD dataset; (iii) A part of the classes are visually very similar, which makes classification, in combination with the high class imbalance, very difficult.

Table 4. Change segmentation and classification results on the HRSCD dataset [16].

	Weak supervision									Full supervision		
	Semantic classification			Binary classification		Segmentation			Segmentation			
	AP	Top-1 acc.	Top-5 acc.	AP	Acc.	Kappa	Dice	Total acc.	Kappa	Dice	Total acc.	
Daudt <i>et al.</i> [16], strategy 3	-	-	-	-	-	-	-	-	12.5	13.8	94.7	
W-CDNet (ours)	83.0	79.9	98.0	75.1	81.0	13.3	14.1	98.4	15.9	16.8	98.0	

5 Conclusion

This paper presented W-CDNet, a novel data-driven convolutional network for change detection which can be trained with weak supervision using image-level semantic labels. The proposed Change Segmentation and Classification (CSC) module enables W-CDNet to learn a change mask at a hidden layer by using the custom Remapping Block and then segmenting the current input image with the predicted change mask. The segmented image is then used to predict the imagelevel semantic label which is used for training. In ablative studies we showed the importance of our Residual Block and of the integrated CRF-RNN layer and their impact on the W-CDNet overall performance.

We improved the state-of-the-art for weakly supervised change detection on AICD and are the first to report weakly supervised change detection results on the HRSCD dataset. On HRSCD the performance of our weakly supervised approach is on par with a related work using full supervision.

For future work we propose to try multi-label change detection. Since our model uses image-level semantic labels for training, one could predict an individual change mask for each label. One could also try to improve information sharing within the model to improve performance. Another interesting task would be to perform change detection on unregistered images. Additionally, we hope that the fact that our model can be trained with or without unchanged image pairs opens new applications for weakly supervised change detection methods.

References

- Shi, W., Min, Z., Zhang, R., Chen, S., Zhan, Z.: Change detection based on artificial intelligence: State-of-the-art and challenges. Remote Sensing 12 (2020) 1688
- Ban, Y., Yousif, O. In: Change Detection Techniques: A Review. Springer International Publishing, Cham (2016) 19–43
- Liu, Z., Li, G., Mercier, G., He, Y., Pan, Q.: Change detection in heterogenous remote sensing images via homogeneous pixel transformation. IEEE Transactions on Image Processing 27 (2018) 1822–1834
- Zhang, P., Gong, M., Su, L., Liu, J., Li, Z.: Change detection based on deep feature representation and mapping transformation for multi-spatial-resolution remote sensing images. ISPRS Journal of Photogrammetry and Remote Sensing 116 (2016) 24–41
- Wang, K., Gou, C.: M4cd: A robust change detection method for intelligent visual surveillance. IEEE Access 6 (2018) 15505 – 15520
- Sakurada, K., Okatani, T.: Change detection from a street image pair using cnn features and superpixel segmentation. In: BMVC. Volume 61. (2015) 1–12
- Alcantarilla, P.F., Stent, S., Ros, G., Arroyo, R., Gherardi, R.: Street-view change detection with deconvolutional networks. Autonomous Robots 42 (2018) 1301– 1322
- Daudt, R.C., Le Saux, B., Boulch, A.: Fully convolutional siamese networks for change detection. In: 2018 25th IEEE International Conference on Image Processing (ICIP), IEEE (2018) 4063–4067
- 9. Khan, S.H., He, X., Porikli, F.M., Bennamoun, M., Sohel, F., Togneri, R.: Learning deep structured network for weakly supervised change detection. In: IJCAI. (2017)
- Jiang, X., Tang, H.: Dense high-resolution siamese network for weakly-supervised change detection. In: 2019 6th International Conference on Systems and Informatics (ICSAI). (2019) 547–552
- Minematsu, T., Shimada, A., Taniguchi, R.i.: Simple background subtraction constraint for weakly supervised background subtraction network. In: 2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), IEEE (2019) 1–8
- Bromley, J., Bentz, J., Bottou, L., Guyon, I., Lecun, Y., Moore, C., Sackinger, E., Shah, R.: Signature verification using a "siamese" time delay neural network. International Journal of Pattern Recognition and Artificial Intelligence 7 (1993) 25
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F., eds.: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Cham, Springer International Publishing (2015) 234–241
- Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., Torr, P.: Conditional random fields as recurrent neural networks. In: International Conference on Computer Vision (ICCV). (2015)
- Bourdis, N., Marraud, D., Sahbi, H.: Constrained optical flow for aerial image change detection. In: 2011 IEEE International Geoscience and Remote Sensing Symposium, IEEE (2011) 4176–4179
- Daudt, R., Le Saux, B., Boulch, A., Gousseau, Y.: Multitask learning for largescale semantic change detection. Computer Vision and Image Understanding 187 (2019) 102783

- 16 P. Andermatt and R. Timofte
- Yu, H., Yang, W., Hua, G., Ru, H., Huang, P.: Change detection using high resolution remote sensing images based on active learning and markov random fields. Remote Sensing 9 (2017) 1233
- Lebedev, M., Vizilter, Y., Vygolov, O., Knyaz, V., Rubis, A.: Change detection in remote sensing images using conditional adversarial networks. ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLII-2 (2018) 565–571
- Guo, E., Fu, X., Zhu, J., Deng, M., Liu, Y., Zhu, Q., Li, H.: Learning to measure change: Fully convolutional siamese metric networks for scene change detection. CoRR abs/1810.09111 (2018)
- Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proceedings of the IEEE 86 (1998) 2278–2324
- 21. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature 521 (2015) 436-444
- Muchoney, D.M., Haack, B.N.: Change detection for monitoring forest defoliation. Photogrammetric engineering and remote sensing 60 (1994) 1243–1252
- Lambin, E.F.: Change detection at multiple temporal scales: seasonal and annual variations in landscape variables. Photogrammetric engineering and remote sensing 62 (1996) 931–938
- Collins, J.B., Woodcock, C.E.: Change detection using the gramm-schmidt transformation applied to mapping forest mortality. Remote Sensing of Environment; (United States) 50:3 (1994)
- Jiang, H., Hu, X., Li, K., Zhang, J., Gong, J., Zhang, M.: Pga-siamnet: Pyramid feature-based attention-guided siamese network for remote sensing orthoimagery building change detection. Remote Sensing 12 (2020) 484
- Bu, S., Li, Q., Han, P., Leng, P., Li, K.: Mask-cdnet: A mask based pixel change detection network. Neurocomputing 378 (2019)
- 27. Sakurada, K., Shibuya, M., Weimin, W.: Weakly supervised silhouette-based semantic scene change detection. Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2020)
- Lafferty, J.D., McCallum, A., Pereira, F.C.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: Proceedings of the Eighteenth International Conference on Machine Learning. (2001) 282–289
- de Jong, K.L., Bosman, A.S.: Unsupervised change detection in satellite images using convolutional neural networks. In: 2019 International Joint Conference on Neural Networks (IJCNN), IEEE (2019) 1–8
- Chianucci, D., Savakis, A.: Unsupervised change detection using spatial transformer networks. In: 2016 IEEE Western New York Image and Signal Processing Workshop (WNYISPW), IEEE (2016) 1–5
- Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. Advances in Neural Information Processing Systems 28 (NIPS 2015) (2015)
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR abs/1409.1556 (2014)
- Hahnloser, R., Sarpeshkar, R., Mahowald, M., Douglas, R., Seung, H.: Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. Nature 405 (2000) 947–51
- 34. Caruana, R. In: Multitask Learning. Springer US, Boston, MA (1998) 95-133
- 35. Yakubovskiy, P.: Segmentation models. https://github.com/qubvel/ segmentation_models (2019)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR09. (2009)

- 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 (2019) 574–586
- Wang, Y., Jodoin, P.M., Porikli, F., Konrad, J., Benezeth, Y., Ishwar, P.: Cdnet 2014: An expanded change detection benchmark dataset. IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (2014)
- Cheng, W., Zhang, Y., Lei, X., Yang, W., Xia, G.: Semantic change pattern analysis (2020)
- Cohen, J.: A coefficient of agreement for nominal scales. Educational and Psychological Measurement 20 (1960) 37–46
- Sørensen, T.: A Method of Establishing Groups of Equal Amplitude in Plant Sociology Based on Similarity of Species Content and Its Application to Analyses of the Vegetation on Danish Commons. Biologiske skrifter. I kommission hos E. Munksgaard (1948)
- Krähenbühl, P., Koltun, V.: Efficient inference in fully connected crfs with gaussian edge potentials. In Shawe-Taylor, J., Zemel, R.S., Bartlett, P.L., Pereira, F., Weinberger, K.Q., eds.: Advances in Neural Information Processing Systems 24. Curran Associates, Inc. (2011) 109–117