

Inferring 3D change detection from bitemporal optical images

Valerio Marsocci^{a,*}, Virginia Coletta^b, Roberta Ravanelli^b,
Simone Scardapane^c, Mattia Crespi^{b,d}

^a*Department of Computer, Control and Management Engineering Antonio Ruberti (DIAG),
Sapienza University of Rome, 00185, Rome, Italy*

^b*Department of Civil, Constructional and Environmental Engineering (DICEA), Sapienza
University of Rome*

^c*Department of Information Engineering, Electronics and Telecommunication (DIET),
Sapienza University of Rome, 00184, Rome, Italy*

^d*Sapienza School for Advanced Studies (SSAS), Sapienza University of Rome, Rome, Italy*

*This is the post-print version of the article published in the "ISPRS Journal of
Photogrammetry and Remote Sensing", Volume 196, 2023, pages 325-339, doi:
10.1016/j.isprsjprs.2022.12.009*

Abstract

In recent years, change detection (CD) using deep learning (DL) algorithms has been a very active research topic in the field of remote sensing (RS). Nevertheless, the CD algorithms developed so far are mainly focused on generating two-dimensional (2D) change maps where the planimetric extent of the areas affected by changes is identified without providing any information on the corresponding elevation variations. The aim of this work is, hence, to establish the basis for the development of DL algorithms able to automatically generate an elevation (3D) CD map along with a standard 2D CD map, using only bitemporal optical images as input, and thus without the need to rely directly on elevation data during the inference phase. Specifically, our work proposes a novel network, capable of solving the 2D and 3D CD tasks simultaneously, and a modified version of the 3DCD dataset, a freely available dataset designed precisely for this twofold task. The proposed architecture consists of a Transformer network based on a

*Corresponding author: valerio.marsocci@uniroma1.it

Email addresses: virginia.coletta@uniroma1.it (Virginia Coletta),
roberta.ravanelli@uniroma1.it (Roberta Ravanelli), simone.scardapane@uniroma1.it (
Simone Scardapane), mattia.crespi@uniroma1.it (Mattia Crespi)

semantic tokenizer: the MultiTask Bitemporal Images Transformer (MTBIT). Encouraging results, obtained on the modified version of the 3DCD dataset by comparing the proposed architecture with other networks specifically designed to solve the 2D CD task, are shown. In particular, MTBIT achieves a metric accuracy (represented by the changed root mean squared error) equal to 6.46 m – the best performance among the compared architectures – with a limited number of parameters (13,1 M). The code and the 3DCD dataset are available at <https://sites.google.com/uniroma1.it/3dchangedetection/home-page>.

Keywords: 3D change detection, Remote Sensing, Deep Learning, Elevation Change Detection, Dataset

1. Introduction

Change detection (CD) is one of the most active research areas in Remote Sensing (RS) [1, 2]. Specifically, the term CD refers to the process of identifying areas of the Earth’s surface that have experienced changes through the joint analysis of two or more co-registered images captured at different epochs [3, 4, 1, 5].

CD algorithms, hence, estimate the changes that occurred at ground level and, for this reason, are applied to a wide range of applications, such as natural disaster management [6, 7], identification of urban changes [8, 9, 10, 11], and crops and forests management [12]. Nowadays, several algorithms can meet the increasingly stringent requirements of CD applications, driven by the unprecedented technological development of Earth Observation (EO) sensors. Despite these promising results, CD algorithms are generally focused exclusively on the generation of 2D change maps and thus they are only capable of identifying the areas affected by temporal variations of land use and land cover (LULC) [13, 14], estimating their planimetric (2D) extent. In this scenario, deep learning (DL) has overtaken traditional methods like, e.g., change vector analysis (CVA) [15] and image differencing, proving to be an effective solution to tasks such as semantic segmentation, and object and change detection [16, 17]. In particular,

several DL architectures are currently employed for solving CD tasks, such as CDNet [18], Siamese models [19], or ChangeFormer [20] networks. Nevertheless, the number of CD labelled RS datasets openly available is still limited, since the production of hundreds of annotated images requires a lot of time and specialised knowledge [21, 22]. The scarce availability of RS datasets represents thus an obstacle to the implementation of new methods and limits the development of viable solutions to the aforementioned CD tasks [4, 5]. Moreover, to the best of our knowledge [23, 16, 17], there is no freely available dataset providing pairs of bitemporal optical images and pixel-wise labels identifying both the 2D changes and the corresponding elevation (3D) variations. For this reason, we specifically designed and produced the 3DCD dataset¹ (Section 4), a freely available dataset [24] containing all the necessary information for solving both the 2D and the 3D CD tasks, starting from optical images captured at different epochs.

One interesting line of research of DL applied to RS is moving precisely towards this direction: to retrieve 3D information from the smallest possible amount of 2D information, usually extracted from optical images [25]. Indeed, 2D CD algorithms can only detect planimetric changes such as appearing/disappearing buildings/trees, shrinking/expanding structures; these results, though, are not sufficient to fulfil the requirements of applications needing also the volumetric/vertical information, such as quantitative estimation of changes in urban areas, forest biomasses and land morphology [23, 5]. The development of suitable 3D CD algorithms can hence offer numerous advantages. We provide a complete overview of the literature on 3D CD and possible applications later on in Section 2.

1.1. Contribution of the paper

The existing state-of-the-art DL CD methods rely mainly on deep CNNs thanks to their ability to extract powerful discriminative features [26]. Lately,

¹A preliminary version of the dataset was presented in [5].

though, the astounding results of Transformers in Natural Language Processing (NLP) has led researchers in applying Transformers also to computer vision tasks such as image classification and image segmentation, with architectures like Vision Transformer (ViT) [27] and SEgmentation Transformer (SETR) [28], and more recently, to 2D CD with the ChangeFormer [20] model. Our work goes one step further by proposing a novel Transformer-based network, trained on the modified version of the recently proposed 3DCD dataset, able to solve simultaneously the 2D and 3D CD tasks, as described in Section 3.

Indeed, the main contribution behind our work is to lay the foundations for the development of DL algorithms able to automatically produce 3D CD maps, i.e. maps containing the quantitative changes in elevation – together with standard 2D CD maps, i.e. maps identifying the planimetric extent of the areas affected by elevation changes –, using two optical images acquired in different epochs as input, without the need of 3D information during the inference step. In this contribution, we show how the proposed Transformer-based architecture performs very well when compared with the most recent networks currently available in the literature, as highlighted in Section 6.

1.2. Organization of the paper

The rest of this paper is organized as follows. Section 2 introduces the datasets and the methodologies currently available for the 2D and the 3D CD tasks. Sections 3 and 4 describe the proposed network and the modified version of the 3DCD dataset, respectively. Section 5 illustrates the experimental setup used to implement and train the proposed network. In Section 6, we discuss the results, comparing quantitatively and qualitatively the proposed network with different methods available in literature, designed specifically for 2D CD and adapted to solve also the 3D CD task, and with a baseline, that makes use of the IM2HEIGHT [25] architecture. In addition, Section 6 describes also the ablation studies performed to gain a better understanding of the proposed network. Finally, we draw conclusions and outlines for future developments (Section 7).

2. Related works

In this section we review both the openly available RS CD datasets (Section 2.1) and the different approaches developed so far to solve the 2D and 3D CD tasks using DL algorithms (Section 2.2).

2.1. Related datasets

In general, many datasets that contain bitemporal optical images and pixel-wise 2D CD maps are currently openly available, while only a small minority include 3D data, namely LiDAR/photogrammetric point cloud and/or Digital Surface Models (DSMs) [23, 5, 29, 30]. However, to our knowledge, there is no dataset including both bitemporal optical images, 2D CD maps and 3D CD maps (Table 1).

For this reason, we specifically designed and produced the 3DCD dataset [5], a freely available dataset providing all the necessary information for solving both the 2D and the 3D CD tasks, starting from bitemporal optical images. Some details about the production and the modification of the proposed dataset are given in Section 4.

In Table 1, the main features of the most used CD datasets freely available in literature are recalled.

Table 1: Main features of the most used CD datasets freely available in literature.

Name	N. images	Tile size	CD map	Notes
SZTAKI Air change [31, 32]	13	952x640	2D	
DSIFN-CD [33]	3600	512x512	2D	
SECOND [34]	4662	512x512	2D	6 classes
OSCD [35]	24	600x600	2D	multispectral
S2Looking [36]	5000	1024x1024	2D	
S2MTCP [37]	1520	600x600	2D	
SYSU-CD [38]	20000	256x256	2D	6 classes
DynamicEarthNet [39]	54750	1024x1024	2D	dense annotations
LEVIR-CD [40]	637	1024x1024	2D	
3DCD [5] (ours)	472	400x400	2D - 3D	

2.2. Related methodologies

CD is a challenging area of research in RS [41]. In this section we provide an overview of the available literature on 2D and 3D CD.

2.2.1. 2D Change Detection

In recent years, several approaches were developed to solve the 2D CD task using deep learning algorithms based on supervised learning, semi-supervised learning, weakly supervised learning, and unsupervised learning, including, among the others, CNNs, recurrent neural networks (RNNs), recurrent Boltzmann machines (RBM), Auto-Encoders (AEs), generative adversarial networks (GANs) and Transformer-based networks [41]. Moreover, different kinds of imagery (optical, synthetic-aperture-radar (SAR), multispectral, ...) have been used to detect 2D changes. Nevertheless, the majority of urban CD deep learning approaches currently available exclusively employ optical imagery [2]. Among these, the three models proposed in [19], i.e. Fully Convolutional Early Fusion (FC-EF), Fully Convolutional Siamese - Concatenation (FC-Siam-conc), Fully Convolutional Siamese - Difference (FC-Siam-diff), are among the first forerunners. Particularly, FC-EF processes the two bitemporal images together, concatenating them along their bands, with a U-Net [42] architecture. Fully Convolutional Siamese (FC-Siam) models, on the other hand, process the input images separately through the encoder. Then the features are concatenated and processed in a unique decoder. The difference between the two aforementioned models relies on the design of the skip connection layer: FC-siam-conc concatenates the two skip connections during the decoding steps, each one coming from one encoding stream, whereas FC-siam-diff concatenates the absolute value of their difference.

More recent algorithms propose novel approaches. For example, SNUNet-CD (a combination of Siamese network and NestedUNet) [43] optimizes the information transmission between encoder and decoder, and between decoder and decoder, through a compact densely connected siamese network and the Ensemble Channel Attention Module (ECAM).

An end-to-end superpixel-enhanced CD network (ESCNet) is proposed in [44]. Specifically, two weight-sharing superpixel sampling networks (SSNs) encode the features of the image pairs. Then, an UNet-based Siamese network is set as decoder. To enhance the performance, the architecture is integrated with an adaptive superpixel merging (ASM) module and a pixel-level refinement module.

Based on generative super resolution, [45] proposes a super-resolution-based CD network (SRCNet), in order to overcome the issue of different resolutions between bitemporal image pairs. Moreover, a stacked attention module (SAM) is presented to extract finer features.

Following a non-supervised approach, [46] proposes an unsupervised deep siamese kernel PCA convolutional mapping network (KPCA-MNet). In the KPCA-MNet, the feature encoder is a siamese network made of weight-shared KPCA convolutional layers. Then, the extracted features are mapped into a 2D polar domain. Finally, the CD results are generated by threshold segmentation and clustering, not requiring any labels.

Recently, the ChangeFormer [20] architecture was proposed. Based on a bitemporal image Transformer (BIT) [47], ChangeFormer combines a hierarchically structured Transformer encoder with a Multi-Layer Perception (MLP) decoder. Moreover, the same authors of [20], in [48], proposed a semi-supervised CD model in which they added an unsupervised loss term to the supervised one. The new term is computed constraining the output change probability map of the unlabeled bitemporal image pair to be consistent, while a small random perturbations is applied on the deep feature difference map.

Lately, several urban mapping studies have been started to integrate EO data fusion into supervised DL CD workflows [2]. SAR imagery, indeed, captures unique information about the physical properties of the Earth’s surface and is less affected by weather conditions than optical images. Nevertheless, the uni-modal application of SAR images in supervised multi-temporal CD analyses is limited by the difficulty of collecting pixel-wise labeled SAR imagery [49]. Moreover, network architectures which perform well on single sensor data may

not be suitable to process multi-sensor data [2]. Therefore, advanced network architectures should be considered to integrate optical and SAR images effectively [2]. In particular, individual streams based on the U-Net architecture were introduced in [2] to process SAR and optical images separately, before combining extracted features at a later decision stage. [49] proposed, instead, an unsupervised CD approach to train a cycle-consistent GAN (CycleGAN [50]) network that was used as a multi-temporal optical-like deep feature extractor from unlabeled bitemporal very-highspatial-resolution SAR images; then a deep CVA and fuzzy rules are applied to detect changed buildings. However, further studies are needed to investigate the use of DL algorithms for SAR and optical integration, before applying them at decision level for urban CD [2], but this is beyond the scope of this paper.

2.2.2. 3D Change Detection

Different kinds of imagery and approaches can be used to monitor 3D changes. However, differently from 2D CD applications for which SAR imagery has recently started to be integrated within the proposed pipelines, the few approaches developed so far for solving the 3D CD task are based exclusively on the use of optical images.

In particular, 3D CD applications can rely on two diverse approaches to detect 3D changes: the DSM difference or the cloud-to-cloud distance [23, 30]. The first approach computes a per pixel height difference between bitemporal DSMs in the format of a 2.5D map – this map is what we call 3D CD map in our workflow. Conversely, the second approach directly computes the 3D Euclidean distances between the points of one point cloud and the mesh built on the other point cloud used as reference. Due to its simplicity, the first approach is still the most used in EO applications, whereas only lately methods using the point clouds based approach have been developed [30, 51]. Indeed, 3D CD maps convey the elevation change information in the same format of 2D CD maps, thus they can be easily integrated in standard 2D CD pipelines, whereas point cloud data are unordered and irregularly distributed, making the extraction of

the information they carry more difficult.

Specifically, a critical review of the last developments and applications of 3D CD using RS and close-range data is given in [23]. More recently, [29] carried out a survey of the airborne LiDAR CD methods currently employed in Earth science applications. However, this survey did not address urban change detection [30], one of the most challenging CD scenarios due to the complex nature of urban spaces [2]. Finally, [52, 30, 51], are the only available studies where a comparative analysis for 3D CD was proposed [5]. In the first study [52], five methods were compared considering two criteria: the ability to detect demolished and new buildings, and the capability to provide information about the magnitude of the elevation changes. However, no quantitative results were provided, and the evaluation was carried out on a private dataset [52]. Moreover, this study did not employ any DL algorithms. On the other hand, in [30], an original simulator of multi-temporal aerial LiDAR urban point clouds was developed to automatically build a labelled 3D CD dataset consisting of pairs of 3D point clouds annotated according to the synthetic changes imposed by the authors. Six different 3D CD methods were assessed either by directly using the 3D point clouds or employing the DSMs generated from their rasterization [30]. In particular, the authors compared traditional methods, such as the application of different thresholding and filtering algorithms on the input DSMs, with both a machine learning algorithm (a random forest fed with hand-crafted features) and two DL networks (consisting in a feed-forward network and a Siamese network) [30]. In [51], a workflow for 3D CD was developed implementing an automated progressive bundle adjustment of multitemporal UAV-oblique images for the production of multitemporal point clouds where the areas affected by changes and the corresponding volumes are detected with a coarse-to-fine strategy. Nevertheless, all the aforementioned studies [29, 52, 30] employed exclusively 3D data as input and no optical imagery (in [51], the images are used to build the multitemporal point clouds and not to directly detect the volumetric changes), still leaving a strong barrier to a wide applicability of these methods to the 3D CD task.

Finally, the authors of [41] proposed a multi-path self-attentive hybrid coding network model (MAHNet) that fuses high-resolution satellite imagery and DSMs for CD of urban buildings. The model was tested on a private and self-produced dataset made of images captured by the Gaofen-7 stereo mapping satellite [41]. In this approach, however, the 3D information retrieved from the DSMs was only used to enhance the 2D change detection results, i.e the identification of the planimetric extent of the areas subjected to elevation changes, and no information about the corresponding elevation variation was returned in the output CD maps, making this contribution very different from ours.

Considering the aforementioned methods, to our knowledge, there is not yet an algorithm able to solve the 3D CD task starting from bitemporal optical images, neither as a single task or in a multitask fashion. In the following section, we hence propose a Transformer-based network to simultaneously solve the 2D and 3D CD tasks.

3. Methodology

The aim of this work is to develop a 2D-3D CD network able to generate two different CD maps starting from bitemporal optical images: a 2D CD map identifying the pixels affected by an elevation change and a 3D CD map quantifying the value of the elevation change for every pixel. Specifically, we propose a Transformer-based network, the MultiTask Bitemporal Images Transformer (MTBIT). MTBIT extends BIT [47], with the aim to solve both the 2D CD and the newly proposed 3D CD task simultaneously. The architecture is composed of four parts: i) a Siamese semantic tokenizer, which generates a set of tokens (i.e. visual words) for each optical input image; ii) a Transformer-based encoder, which encodes positional and context information in the token space; iii) a Transformer-based decoder, that extracts feature maps projecting the tokens in the pixel space; iv) a pair of prediction heads which allow double mapping from the refined feature maps to 2D and 3D prediction maps. Particularly, a Siamese CNN backbone is fed with two input optical images X_1 and

$X_2 \in \mathbb{R}^{H \times W \times B}$, where H and W are the height and the width of the input images, and B is the number of the input bands, to extract, respectively, two sets of feature maps Z_1 and $Z_2 \in \mathbb{R}^{H' \times W' \times C}$, where H' is the height and W' is the width and C are the feature maps channels. To get a vocabulary of semantic tokens (i.e. visual words), these feature maps pass through a two-step semantic tokenizer made of: a) spatial attention maps, that is a CNN layer, followed by a softmax, operated on the height H' and width W' of each feature map; b) a weighted average sum of pixels of the obtained maps. Formally:

$$\mathbf{T} = (\mathbf{A})^T \mathbf{X} \quad (1)$$

where $\mathbf{A} = \text{softmax}(\text{conv}(\mathbf{Z})) \in \mathbb{R}^{H \times W \times L}$ are the attention maps, where L is the length of the vocabulary.

The obtained two sets of semantic tokens \mathbf{T}^i (one for each image) are concatenated in one set, namely $\mathbf{T}^* \in \mathbb{R}^{2L \times C}$. Before passing them through the Transformer encoder, the tokens are used to learn positional encodings, useful to encode the position of the token in the space. Given the fact that the encoder consists of N layers of multihead self-attention (MSA) and MLP blocks, the positional encodings are also very useful to have shared representations across heads of the tokens. Particularly, the semantic tokens are the input of the self-attention matrices (the query $\mathbf{Q} = \mathbf{T}^{*(l-1)} \mathbf{W}^q$, the key $\mathbf{K} = \mathbf{T}^{*(l-1)} \mathbf{W}^k$ and the value $\mathbf{V} = \mathbf{T}^{*(l-1)} \mathbf{W}^v$. where $\mathbf{W}^q, \mathbf{W}^k, \mathbf{W}^v \in \mathbb{R}^{C \times d}$, where d is the channel dimension and l the layer). The single head-attention is defined as in [53]. MSA performs multiple independent attention heads in parallel. The outputs are, then, concatenated and projected. MSAs are shaped in order to capture different information in different representation spaces. From a formal point of view,

$$\text{MSA}(\mathbf{T}^{*(l-1)}) = \text{Concat}(\text{Att}_1, \dots, \text{Att}_H) \mathbf{W}^E \quad (2)$$

with

$$\text{Att}_h = \text{softmax} \left(\frac{\mathbf{Q}_h \mathbf{K}_h^T}{\sqrt{d}} \right) \mathbf{V}_h + \mathbf{PE}_h \quad (3)$$

where $h = 1, \dots, H$, with H denoting the number of attention heads, \mathbf{PE} are the learnable positional encodings of the encoder and $\mathbf{W}^E \in \mathbb{R}^{hd \times C}$ is the weight matrix. On the other hand, the MLP block is made of two linear projection layers. A Gaussian error linear unit (GELU) module [54] is inserted after the first linear transformation.

In this way, we obtained a new set of tokens \mathbf{S} that encodes the context information from the input images. We split it into two set \mathbf{S}^i with $i = 1, 2$ and we pass them through the decoder. The aim of the Transformer decoder is to project back to the pixel space the semantic information capable to correctly predict 2D and 3D changes. Practically, the Transformer decoder consists of M layers of multihead cross attention (MA) and MLP blocks. In MA, $\mathbf{Q} = \mathbf{Z}_i^{(l-1)} \mathbf{W}^q$, $\mathbf{K} = \mathbf{S}^{(l-1)} \mathbf{W}^k$ and $\mathbf{V} = \mathbf{S}^{(l-1)} \mathbf{W}^v$.

Finally, the output feature maps from the decoder are passed through the pair of prediction heads, each consisting in a Fully Convolutional Neural Network (FCN). Particularly, given two feature maps \mathbf{Y}_1 and $\mathbf{Y}_2 \in \mathbb{R}^{H'' \times W'' \times C}$, the final predictions maps, $\mathbf{M}_{2D} \in \mathbb{R}^{H \times W \times 2}$ and $\mathbf{M}_{3D} \in \mathbb{R}^{H \times W}$, are obtained in the following way:

$$\mathbf{M}_{2D} = \text{sigmoid}(f_1(\mathbf{Y}_2 - \mathbf{Y}_1)) \quad (4)$$

and

$$\mathbf{M}_{3D} = \text{tanh}(f_2(\mathbf{Y}_2 - \mathbf{Y}_1)) \quad (5)$$

where f_1 and f_2 are 3x3 convolutional layers.

As loss of the proposed network, we use a weighted combination of a 2D loss, taking as input \mathbf{M}_{2D} , and a 3D loss, taking as input \mathbf{M}_{3D} :

$$L = \alpha \cdot L_{2D} + \beta \cdot L_{3D} \quad (6)$$

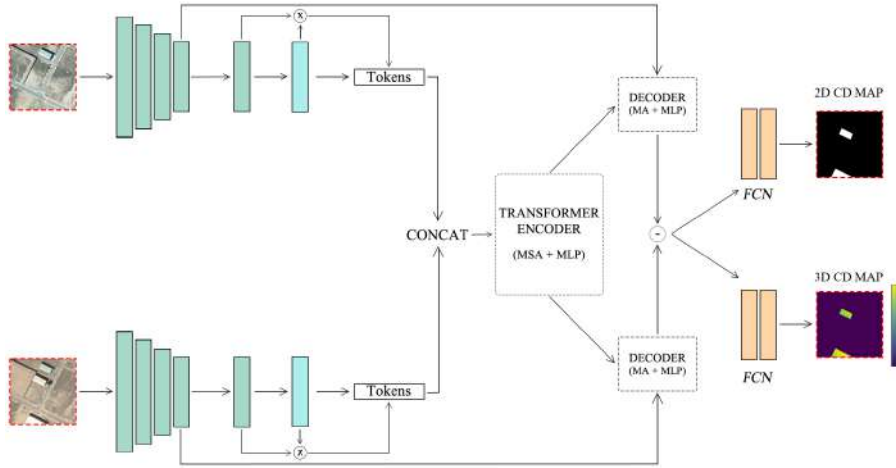


Figure 1: Schematic representation of the MultiTask Bitemporal Image Transformer (MTBIT) network architecture.

where α and β are fixed weights. In particular, we use a binary cross-entropy (BCE) for L_{2D} and the mean squared error (MSE) for L_{3D} .

The MTBIT architecture is summarized in Figure 1.

4. Dataset

The dataset used to evaluate the potentialities of the proposed 2D-3D CD network is a modified version of the previously presented 3DCD dataset [5], freely available at <https://sites.google.com/uniroma1.it/3dchangedetection/home-page>. The dataset, which can be classified as a top-view dataset according to the taxonomy of visualization scenarios proposed in [23], is composed of:

- 472 pairs of images cropped from optical orthophotos acquired through two different aerial surveys (performed respectively in 2010 and in 2017);
- 472 pairs of DSMs produced from the rasterization of point clouds acquired through two different LiDAR flights covering the same area and years

(2010 and 2017) of the orthophotos²;

- the corresponding 472 2D CD maps in raster format;
- the corresponding 472 3D CD maps – the elevation variation ($\Delta H = H_{2017} - H_{2010}$) maps – in raster format.

The original orthophotos and the LiDAR point clouds used to produce the dataset were made available by the *Centro de Descarga Nacional* [55]. The modified version of the dataset includes also the georeferencing information: the reference system is EPSG:3042 - ETRS89 / UTM zone 30N (N-E), i.e. the reference system of the original orthophotos and LiDAR point clouds.

In particular, the dataset covers the historical and urban centre of the city of Valladolid in Spain, including also the surrounding commercial areas. Agricultural zones were not considered since no significant elevation changes were observed in these areas. Furthermore, only the changes affecting artificial structures, such as the construction and the demolition of buildings/roads/bridges, were included in the 2D and 3D CD maps.

Two examples of the 3DCD dataset are shown in Figure 2. Each tile of the dataset (including one pair of images, one pair of DSMs, and the corresponding 2D and 3D CD maps) covers an area equal to 200 m \times 200 m.

The optical images (Figure 2-a and 2-b) and the 2D CD maps (Figure 2-e), which we produced from the orthophotos, have a size of 400 \times 400 pixels with a GSD of 0.5 m.

The 2D CD maps are binary maps, in which we annotated the pixels belonging to the areas affected by a change in elevation. Specifically, the 2D CD maps include the following two classes (Figure 2-e): (i) no change (null elevation change in the corresponding 3D CD map: $\Delta H = 0$); (ii) changes due to construction (positive elevation change in the corresponding 3D CD map: $\Delta H > 0$) or demolition (negative elevation change in the corresponding 3D CD map: $\Delta H < 0$) of anthropic artefacts (e.g. buildings, roads, bridges).

²The DSMs were not released in the previous version of the dataset.

The 3D CD maps (Figure 2-f) have a size of 200×200 pixels with a GSD of 1 m, as well as the DSMs (Figure 2-c and 2-d), and were obtained from the per-pixel differences between the two different LiDAR DSMs ($\Delta H = H_{2017} - H_{2010}$). In particular, through a visual comparison with the corresponding orthophoto pair, we only considered pixels affected by a real elevation change, while the ΔH of the pixels containing only noise was set to zero. In the 3D CD maps, the values of the elevation changes are represented on a continuous range varying from -30 m to -1 m and from 1 m to 35 m (Figure 3). Indeed, all the elevation changes lower than one meter in absolute value were considered as negligible with respect to the entity of elevation variations usually affecting man made artefacts, and for this reason their value was set to zero [5].

The main features of the data contained in the 3DCD dataset are recalled in Table 2.

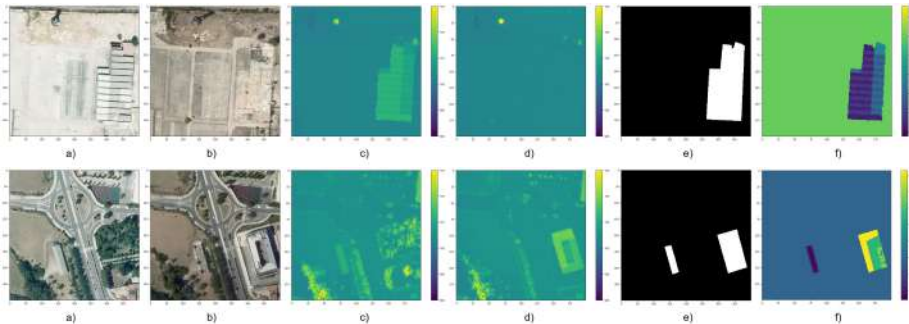


Figure 2: Two examples from the modified 3DCD dataset. a) 2010 optical image, b) 2017 optical image, c) 2010 DSM (the colour bar is expressed in meters), d) 2017 DSM (the colour bar is expressed in meters), e) 2D CD map (the pixels affected by a change in elevation are reported in white), f) 3D CD map (the colour bar is expressed in meters): first row: demolition ($\Delta H < 0$), second row: construction ($\Delta H > 0$).

Finally, the dataset was split into *train*, *validation* and *test* subsets to allow an easy and reproducible benchmarking. In particular, the train subset contains 320 images ($\sim 68\%$), the test subset 110 images ($\sim 23\%$) and the validation subset 42 images ($\sim 9\%$). The division was performed maintaining a similar distribution of changed pixels (2D CD maps) and of elevation changes (3D CD maps) among the three subsets (see Table 3 and Figures 3a, 3c, 3b for details),

verifying that the subsets did not share the same buildings in the modified version of the dataset, removing hence any spatial autocorrelation.

Table 2: Main features of the data contained in the modified version of the 3DCD dataset.

Image	Number of pixels	GSD
<i>Orthophotos from 2010</i>	400×400	0.50 m
<i>Orthophotos from 2017</i>	400×400	0.50 m
<i>DSMs from 2010</i>	200×200	1.0 m
<i>DSMs from 2017</i>	200×200	1.0 m
<i>2D CD maps</i>	400×400	0.50 m
<i>3D CD maps</i>	200×200	1.0 m

Table 3: Percentage of pixels characterised by ‘change’ and ‘no change’ for each subset.

Subset	Number of Images	Change	No change
<i>Train</i>	320	5%	95%
<i>Test</i>	110	4%	96%
<i>Val</i>	42	5%	95%

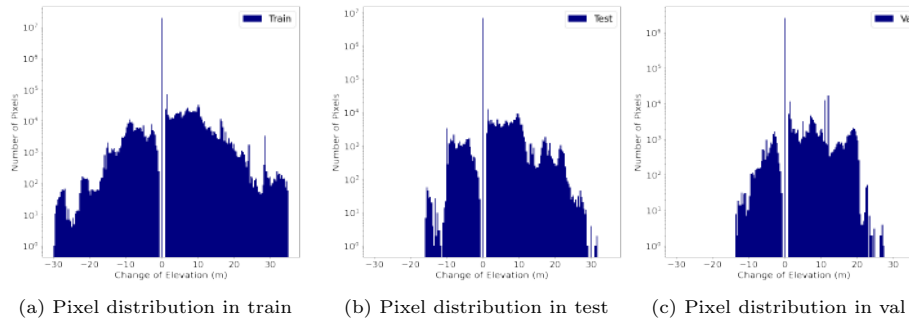


Figure 3: Distribution of the elevation change values among all the pixels of the 3D CD maps in each subset of the modified version of the 3DCD dataset.

5. Experimental setup

For the training phase, a single Tesla T4 16 GB GPU was used. We trained the proposed network – MTBIT – for 300 epochs with the AdamW optimizer and a learning rate of 0.0001. Moreover, we used a fixed batch size of 15 and we performed the following augmentation strategies on the input images: resizing to 256×256 , random horizontal flip, random geometric transformation

(i.e. shifting, scaling and rotating), addition of random gaussian noise, random radiometric transformation (i.e. brightness, sharpening, blurring, contrasting, saturation). Furthermore, we normalized the 3D CD maps between -1 and 1, using a TanH as last non-linear layer of the proposed network. For the 2D loss (i.e. BCE), we set the weights to 0.05 for no-change pixels and 0.95 for change ones. In addition, we set $\alpha = 1$ and $\beta = 3$ in Equation (6) (see Section 6.2.2 for further details).

To evaluate the performance of the proposed network, we opted for two 2D metrics, i.e. intersection over union (IoU) and F1-score (F1), referred to the change class (that is 1), and two 3D metrics, i.e. root mean squared error (RMSE) and changed root mean squared error (cRMSE). Particularly, cRMSE is a RMSE that considers only the error with respect to the ground truth pixels affected by an actual elevation change (i.e. $\Delta H \neq 0$). Formally, the metrics are expressed as follows:

$$\text{IoU} = \frac{\text{TP}}{\text{TP} + \text{FN} + \text{FP}} \quad (7)$$

$$\text{F1} = \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}} \quad (8)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\widehat{\Delta H}_i - \Delta H_i)^2} \quad (9)$$

$$\text{cRMSE} = \sqrt{\frac{1}{n_c} \sum_{i=1}^{n_c} (\widehat{\Delta H}_i^C - \Delta H_i^C)^2} \quad (10)$$

where TP stands for True Positives, TN for True Negatives, FP for False Positive and FN for False Negatives, $\widehat{\Delta H}_i$ for predicted pixels, ΔH_i for ground truth pixels, $\widehat{\Delta H}_i^C$ for predicted pixels, ΔH_i^C for respective ground truth changed pixels. Between the 3D CD metrics, cRMSE is preferable to standard RMSE because the latter is less sensible to changes (and errors), being most of the changes in the target map at 0 value and thus easy to be detected. In addition to several ablation studies to gain a better understanding of the behaviour of

the proposed architecture, we performed a comparison with different existing networks designed for the 2D CD, given the fact that, to our knowledge, no architecture able to solve the 3D CD task is currently available. Specifically, to make these existing networks able to predict the 3D CD maps, we changed the prediction head with a pair of FCNs, able to solve both the 2D and 3D CD tasks. We selected the following models for the comparison: ResNet18, ChangeFormer [20], FC-EF [19], FC-Siam-diff [19] and a modified version of FC-Siam-conc, that we named SUNet (Siamese ResU-Net).

Particularly, SUNet is made of: a) a Siamese encoder; b) a convolutional decoder; c) a pair of prediction heads. The Siamese encoder consists of two ResNet architectures with shared weights, each of which extracts the features from each input image. The choice of the ResNet is pivotal for the effective learning of the features, making possible to avoid the degradation problem. At the end of the ResNet encoder, the features are concatenated to be fed through a convolutional decoder, equipped with skip connections. Specifically, we opted for concatenating the skip connection of each layer of the encoder. This choice is justified by the structure of our problem, for which not all the land use changes corresponding to different features match a 2D change in the prediction maps. Finally the decoded feature maps are passed through a prediction head pair, that, as for the MTBIT, is made of two FCNs. The code is available at <https://sites.google.com/uniroma1.it/3dchangedetection/home-page>.

Moreover, we also provide a baseline obtained using IM2HEIGHT [25], a network designed to infer an elevation map from a single RGB image. We trained this model twice, firstly on the 2010 images and then on the 2017 images. After inferring the resulting elevation maps on the test images, we subtracted the 2010 elevation maps to the 2017 elevation maps, obtaining the baseline 3D CD maps. For this experiment, we started from the code available at <https://github.com/dettmar/im2height>.

6. Experimental Results

As described in Section 3, several models were trained with the final goal of automatically producing both the 2D and 3D CD maps, using a pair of bitemporal optical images as input. In the following sections, the best results and different ablation studies are described.

6.1. Comparison

We compared the proposed network, MTBIT, with the aforementioned architectures to evaluate its performance. In particular, in addition to the quantitative assessment of the results, performed evaluating the metrics described in Section 5 (Table 4), we performed also two qualitative comparisons: we visually compared the predicted 2D and 3D CD maps with the corresponding ground truth maps (Figure 4) and we compared the distribution of the predicted elevation changes with the corresponding distribution of the ground truth values (Figure 5).

Table 4: Validation results of the different networks applied to the modified version of the 3DCD dataset.

Architecture	Figure 4	Trainable parameters (M)	2D CD metrics		3D CD metrics	
			F1 (%)	IoU (%)	RMSE (m)	cRMSE (m)
ChangeFormer	a	1,8	42,89	27,30	1,31	7,09
FC-EF	b	1,3	46,32	30,14	1,41	7,04
FC-Siam-conc	c	1,5	38,30	23,69	1,42	10,66
SUNet	d	35,7	59,72	42,57	1,24	6,47
ResNet18	e	32,6	15,52	8,41	1,56	7,45
IM2HEIGHT	f	7,4	-	-	1,57	7,59
MTBIT (ours)	g	13,1	62,15	45,09	1,20	6,46

Table 4 shows that the ResNet18, ChangeFormer, FC-EF and FC-Siam-conc models provide the worst results when applied to the 3DCD dataset, for both the 2D and the 3D CD tasks. In fact, the F1 score reaches a value less than 50% for any of these four models and the IoU value is always less than 35%. Looking at the Figures 4a-a, 4a-b and 4a-c, relative to the 2D CD results, it is evident




Figure 4: Visualisation of 2D CD and 3D CD maps for the models described in Table 4.

that the 2D CD is not particularly accurate, both with respect to the detected areas and to their extent. The worst results are obtained with the ResNet18 architecture which cannot reconstruct correctly the 2D CD maps (Figure 4a-e). This problem, true also for the 3D CD task, is due to the absence of the skip connections, making the models unsuitable for the proposed tasks (Figure 4b-e). When adding skip connections and making the network deeper, the results are more accurate. This is the case of SUNet, which reaches results comparable to MTBIT. However, the number of its parameters (35,7 M) is higher than the counterparts. MTBIT, on the other hand, provides the best results for the 2D CD task, obtaining an F1 score of 62.15% and an IoU value of 45.09% (Table 4-g). From Figure 4-g, it can also be observed that changes are detected with high accuracy and the background noise is minimal.

Also for the 3D CD task, ResNet18, ChangeFormer, FC-EF and FC-Siam-conc networks do not achieve accurate results: all the four architectures have a cRMSE higher than 7 m, up to more than 10 m for the FC-Siam-conc model (Table 4). From the Figures 4b-a, 4b-b and 4b-c, we can see that these models in general show limitations in detecting real changes and have a tendency to underestimate elevation changes. In addition, we can observe that the 3D CD maps predicted by ChangeFormer (Figure 4b-a) and FC-Siam-conc (Figure 4b-c) are characterized by a background noise. The latter problem is highlighted by the histograms reported in Figures 5a, 5b and 5c; for all the three models, the distributions of the values of the predicted elevation changes are very different from the distribution of the Ground Truth. The reasons behind these failures vary depending on the considered network: ChangeFormer – a fully Transformer architecture, designed for 2D CD, which operates on few data – shows poor generalization capabilities, detecting changes in LULC that do not correspond to changes in elevation; conversely, FC-EF and FC-Siam-Conc are not able of learning effective features for the solution of the 2D-3D multitask (e.g., low number of parameters, lack of effective encoder). Furthermore, also the IM2HEIGHT-based baseline shows a poor performance (Table 4-f and Figure 4-f). This behaviour is mainly due to the fact that IM2HEIGHT is focused

on reconstructing the correct height of the single DSMs, and not on estimating the elevation changes occurred between two epochs (2010 and 2017). Moreover, any small error in reconstructing the DSMs for each of the two epochs is propagated through their difference.

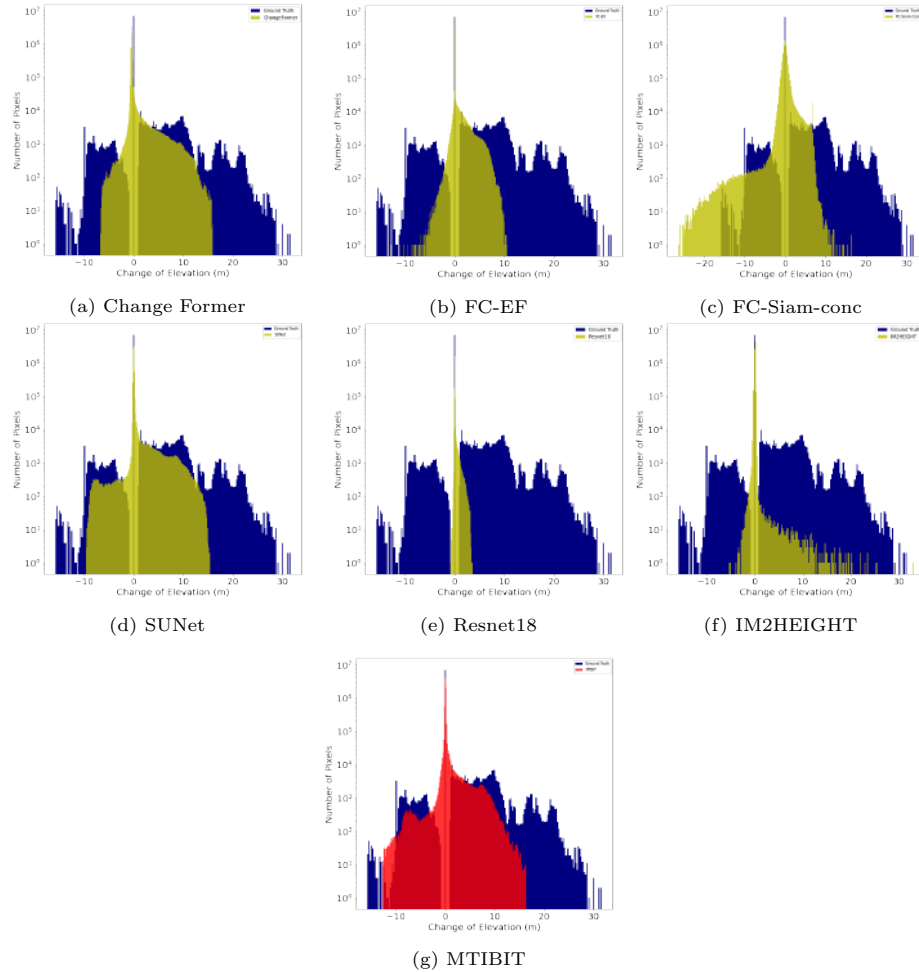


Figure 5: Comparison between the distribution of the elevation change values in the ground truth (in blue) and the distribution of the predicted elevation change values obtained with the different models reported in Table 4.

As far as the 3D CD task is concerned, MTBIT achieves the best results with a cRMSE of 6.46 m (Table 4-g), even if the 3D CD maps are characterised by a slight background noise (Figure 4b-g). Furthermore, the predicted

elevation changes (see Figure 4b-g) slightly underestimate the real ones. This phenomenon is more evident when analysing the histograms of the elevation changes: the predicted elevation changes follow the ground truth distribution well, but the network fails to reproduce the extreme values with a deviation of about 10 m (Figure 5g). Also SUNet provides optimal results for the 3D task, reaching a cRMSE value equal to 6.47 m (Table 4-e).

Focusing on the comparison among these two best models (SUNet and MTBIT), the 2D CD maps obtained using SUNet (Figure 4a-e) detect changes very well, but are slightly noisier than the 2D CD maps obtained with MTBIT (Figure 4a-d). The 3D CD maps predicted using the SUNet model are accurate: the areas are well detected, there is no background noise and the predicted elevation changes are in line with ground truth values (Figure 4b-d). Moreover, the elevation changes predicted by SUNet follow the distribution of the ground truth values very well (Figure 5d). However, also SUNet, like MTBIT, is not able to predict the extreme values of the elevation changes, but with a higher deviation than MTBIT (Figure 5d). The problem with the detection of extreme values is probably due to a deficiency in the training data: considering that the y-axis of the histograms is logarithmic, the values that the models cannot reproduce are present in less than 100 pixels. Also the 3D CD metrics of MTBIT (RMSE: 1.20 m and cRMSE 6.46 m) and SUNet (RMSE: 1.24 m and cRMSE: 6.47 m) are very similar. Nevertheless, MTBIT has a slightly better ability to identify areas affected by 3D changes, reaching lower RMSE and cRMSE values. SUNet instead does not perfectly identify the contours of the changed areas, increasing the RMSE, still providing however elevation change values very consistent with the Ground Truth.

In conclusion, MTBIT is more appropriate than SUNet for its lower number of parameters. Indeed, MTBIT obtains results comparable to the ones obtained by SUNet with a third of the parameters (13,1 M vs 35,7 M). This is even more true if we compare the number of parameters of these two models with the other ones used for comparison. In fact, except for ResNet18 (32,6 M), all the other models have about 1.5 M of parameters. IM2HEIGHT, designed for inferring

Table 5: Ablation studies assessing the behaviour of MTBIT with different sizes of the main components of the architecture.

Figure 7	Encoder depth	Token length	Decoder depth	Decoder dimension head	2D CD metrics		3D CD metrics	
					F1 (%)	IoU (%)	RMSE (m)	cRMSE (m)
a	1	4	8	8	63,02	46,00	1,22	6,88
b	2	4	8	8	62,05	44,98	1,17	6,52
c	1	4	16	8	63,81	46,86	1,23	6,71
d	1	8	8	8	60,98	43,86	1,19	6,60
e	2	8	16	16	59,82	42,68	1,22	6,84
f	4	16	32	32	61,61	44,52	1,22	6,67
g	1	4	8	16	62,15	45,09	1,20	6,46

Table 6: Ablation studies assessing the behaviour of MTBIT with different features.

Figure 8	Absolute value in feature difference	Concat	Learnable upsampling	2D CD metrics		3D CD metrics	
				F1 (%)	IoU (%)	RMSE (m)	cRMSE (m)
a	Yes	No	No	63,02	46,00	1,22	6,55
b	No	Yes	No	60,91	43,79	1,24	7,60
c	No	No	Yes	61,11	44,00	1,21	6,87
d	No	No	No	62,15	45,09	1,20	6,46

height from RGB images, has 7,4 M of parameters.

In the next sections, we describe in more detail the different ablation studies related to the proposed model.

6.2. Ablation Studies

6.2.1. Components of MTBIT

In these ablation studies, we evaluated under which configuration MTBIT can achieve the best performance. In particular, we tested different features such as: type of upsampling (learnable or bilinear), concatenation or difference of features in the decoder output, number of attention heads in both the encoder and the decoder. Specifically, we reported two sets of experiments: one concerning the sizes of the main components of the MTBIT architecture (see Table 5 and Figure 7) and one involving specific architecture features (see Table 6 and Figure 8)

In the first set of experiments, we started from a standard configuration [47] (see Table 5-a) and we tried to increase the sizes of some of the components of the MTBIT architecture, to assess their effect on the performance of the network. As we can argue from the results, the architecture is fairly stable with respect to these changes, considering that the dataset is rather small. Particularly, we can observe that larger components of the network, especially when combined (see Table 5-e and 5-f), do not lead to better performances. This behaviour can be noted also in Figures 7-e and 7-f. Moreover, Figure 6 shows the attention maps of two image examples. Once again, we can affirm that with a small dataset, $L = 4$ leads to a good performance of the network. However, among all the




Figure 6: Visualisation of the attention maps for two bitemporal examples. Particularly, in a) the considered images are shown, from b) to e) the respective four attention maps are presented.

tested configurations, we selected the one in which the decoder dimension head is 16 (Table 5-g). In fact, under these conditions, MTBIT reached the best

value of the cRMSE which is the most appropriate 3D CD metric, as previously described.

Concerning the modification of the network, we found that a learnable up-sampling, obtained through transposed convolutional modules, leads to a worse performance with a higher number of parameters. Once again this behaviour is mainly due to an overfitting of the training data (Table 6-c). Moreover, the intuition of concatenating the features, to give the model a better understanding of the elevation changes (with respect to LULC changes), was not confirmed by the experiments, even though the results are only slightly worse (Table 6-b). Nevertheless, we observed that to effectively understand the 3D changes and their sign (positive or negative), the network benefits by subtracting the features before the prediction heads, without considering their absolute values. Conversely, the absolute values work better for the 2D CD task, for which the sign is not representative of the phenomenon, as showed also in [47].

6.2.2. Loss weights

Table 7: Ablation studies assessing the behaviour of MTBIT with different weights.

Architecture	Figure 9	α	β	2D CD metrics		3D CD metrics	
				F1 (%)	IoU (%)	RMSE (m)	cRMSE (m)
MTBIT	a	0	1	5,76	2,98	1,24	6,93
	b	1	0	19,40	10,74	5,37	80,02
	c	1	1	62,03	44,96	1,22	6,75
	d	1	3	62,15	45,09	1,20	6,46
	e	3	1	62,77	45,74	1,21	6,80
	f	1	5	62,71	45,68	1,19	6,53
	g	5	1	63,65	46,68	1,24	6,83

This ablation study is aimed at analysing how MTBIT performs varying α and β , i.e. the weights of, respectively, the 2D loss and 3D loss (Equation (6)). In Table 7 and Figure 9, the quantitative and qualitative results are shown. Particularly, we can highlight how focusing exclusively on one task not only leads to a poor performance on the other task, but does not even make it possible to

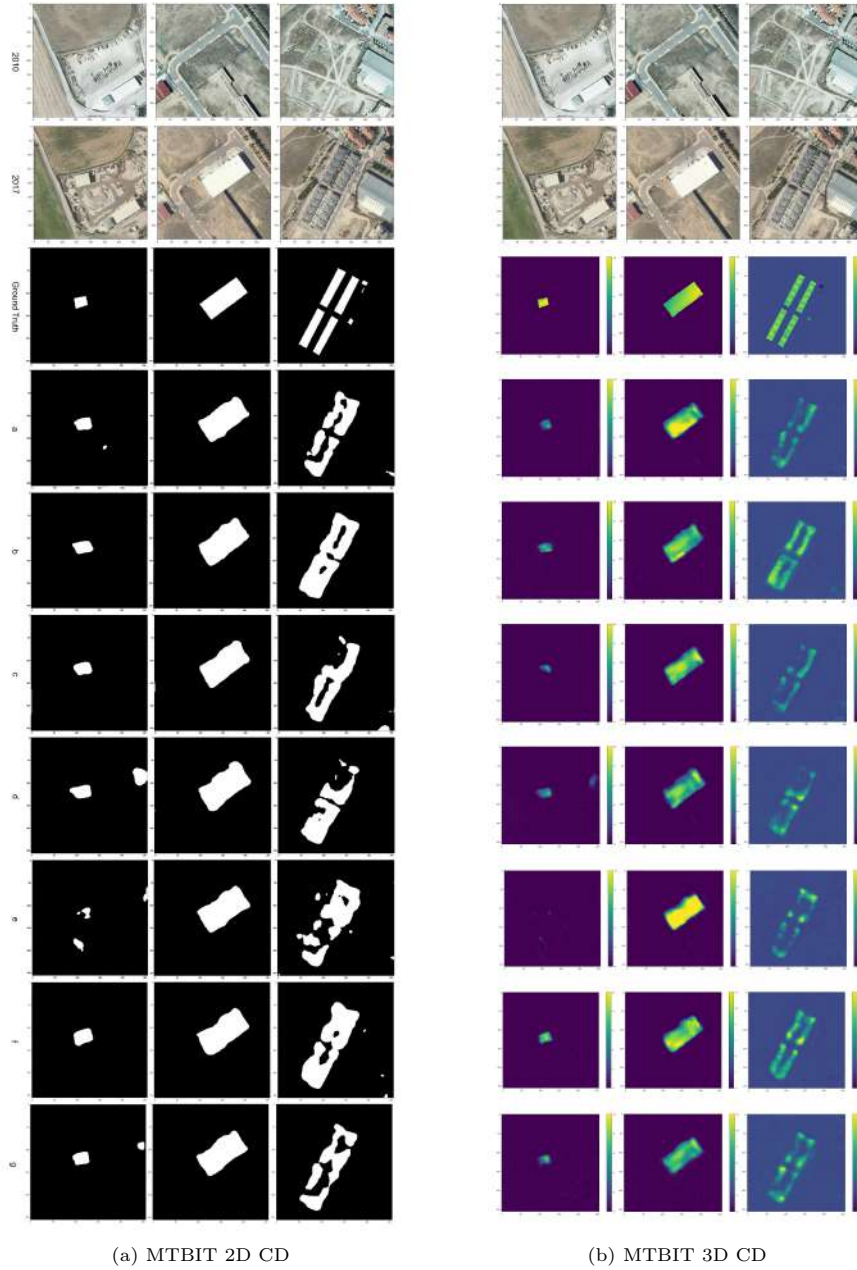


Figure 7: Visualisation of 2D CD and 3D CD maps obtained with the different models described in Table 5

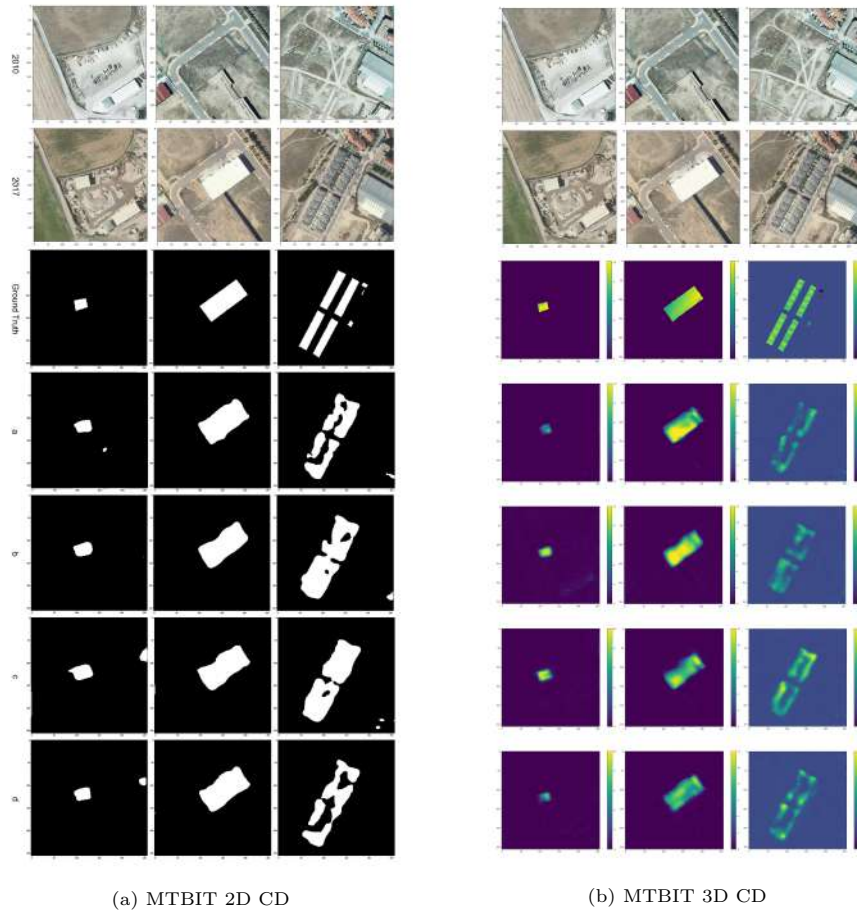


Figure 8: Visualisation of 2D CD and 3D CD maps obtained with the different models described in Table 6

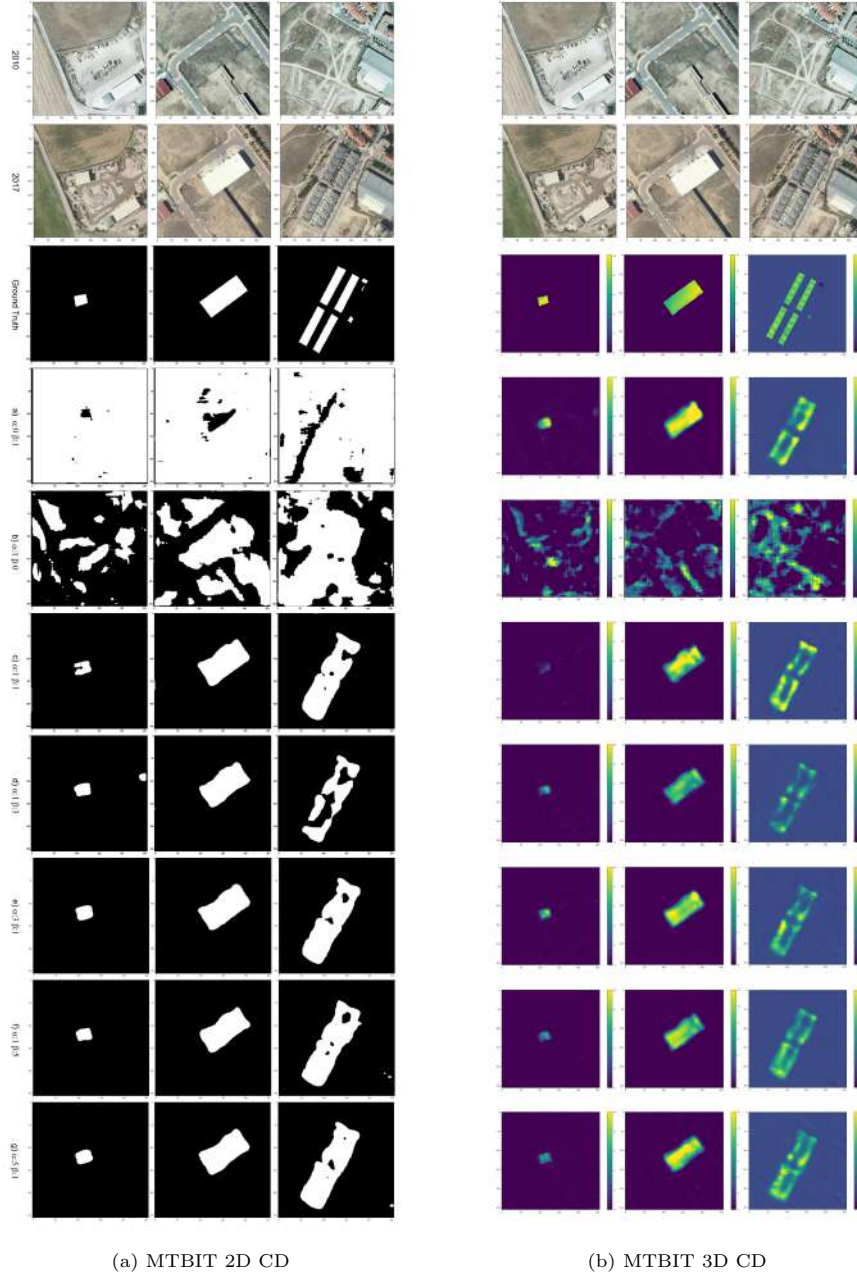


Figure 9: Visualisation of 2D CD and 3D CD maps obtained with the different models described in Table 7

Table 8: Ablation studies showing the behaviour of MTBIT under different encoders and strategies.

Figure: 10	Encoder	Frozen encoder	Pretraining strategy	Trainable parameters (M)	2D CD metrics		3D CD metrics	
					F1 (%)	IoU (%)	RMSE (m)	cRMSE (m)
a	ResNet18	No	Random	13,1	49,24	32,66	1,31	8,11
b	ResNet18	No	ImageNet	13,1	62,15	45,09	1,20	6,46
c	ResNet50	No	Random	42,9	46,11	29,96	1,40	8,16
d	ResNet50	No	ImageNet	42,9	63,50	46,52	1,15	6,35
e	ResNet18	Yes	ImageNet	1,4	49,36	32,77	1,28	7,11
f	ResNet50	Yes	ImageNet	17,3	56,72	39,58	1,22	6,96

reach the best performances on the considered task. This is mainly due to the fact that the representations learned by the network to solve one task are useful also to solve the other task and viceversa. So, we tested several combinations to find the optimal one. As we can argue from the Table 7 and Figures 9, there is not an evident best combination of the values of α and β which clearly overperforms the others. The main intuition behind preferring higher values of β is that the 2D CD task is easier to be solved and the 3D CD representations are more effective to solve also the 2D CD than viceversa. Concerning the magnitude of β , given $\alpha = 1$, the combination that reaches the most balanced results is with $\beta = 3$ (Figure 9-d). Indeed, with lower (e.g., the configuration with $\beta = 1$ reaches a good performance on 2D CD, but does not achieve good results with respect to cRMSE) and higher (e.g., the configuration with $\beta = 5$ reaches unstable performance, where neither the 2D CD or the 3D CD metrics are the best in their categories) values of β , the MTBIT performances start to decrease.

6.2.3. Encoders

In this set of ablation studies, we analyzed the behavior of MTBIT with different encoders and with different pretraining strategies. In particular, we tried to investigate how the performance of MTBIT changes with respect to the number of parameters and the initialization of the network. The results are shown in Figure 10 and in Table 8.

First of all, as we expected, ResNet50 overcomes the ResNet18 performances.

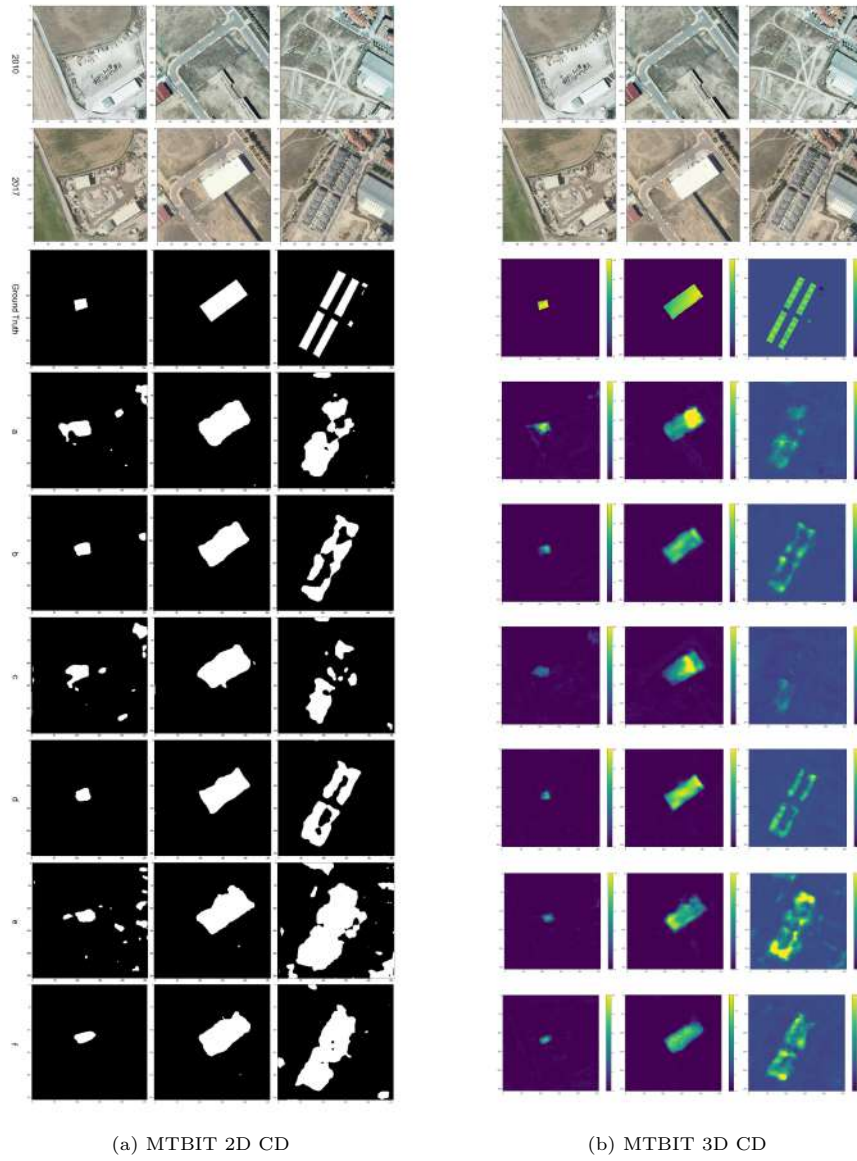


Figure 10: Visualisation of 2D CD and 3D CD maps obtained with the different models described in Table 8

However, with respect to the number of parameters, which is more than three times higher, MTBIT with ResNet50 encoder (42,9 M parameters) gains just 1.5% in F1 and 11 cm in cRMSE with respect to MTBIT with ResNet18 encoder (13,1 M parameters), as Tables 8-b and 8-d show. For this reason, we selected the ResNet18 encoder model. Moreover, we can observe how initializing the encoder with pretrained Imagenet weights leads to a good performance of the network, both with a ResNet18 and a ResNet50 encoder, as it is reasonable to expect. In fact, even though the ImageNet features are rather different from the features of remote sensing images, the 3DCD dataset is too small to make the network learn rich features. In addition, we can observe that ResNet18 encoder reaches better values of the metrics than the ResNet50 encoder (e.g. 49.24% vs 46.11% for F1; 8.11 m vs 8.16 m for cRMSE) in the random weight initialization scenario. This is mainly due to the poor generalization capabilities of the larger network (MTBIT with ResNet50 encoder), that cause an overfitting of the training data. This is confirmed by the Figures 10-a and 10-c. Regarding the experiments that involved freezing the encoder pretrained on ImageNet, we can observe that discrete performances are obtained. In general, we can affirm that freezing the encoder within such an architecture means to freeze the inputs of the semantic tokenizer. Since ImageNet features differ from the features of remote sensing images, we cannot expect good results. However, MTBIT with ResNet50 frozen encoder (17,3 M) loses 7% in F1 and 0.5 m in cRMSE, employing a number of parameters comparable to the one used by the benchmark MTBIT (13,1 M, see Table 4). Figures 10-e and 10-f show a distinct degradation, especially with respect to the 2D CD results. On the other hand, MTBIT with ResNet18 frozen encoder reaches similar performances to the ones obtained by the networks used for the comparison (Section 6.1) with a similar number of parameters (see Table 4), such as ChangeFormer (1,8 M) and FC-EF (1,3 M).

7. Conclusions

In this paper, we presented a task that, to our knowledge, had not yet been investigated: to infer elevation (3D) changes from bitemporal optical images. The final aim was indeed to establish the basis for the development of DL algorithms able to automatically generate an elevation 3D CD map together with a standard 2D CD map, using only optical images captured at different epochs as input. Specifically, we proposed MTBIT, a novel network belonging to the family of vision Transformers, one of most used DL architectures, and a modified version of the 3DCD dataset, a freely available dataset designed precisely for this twofold task.

In particular, the 3DCD dataset – consisting of more than 450 pairs of images and DSMs together with the corresponding 2D and 3D CD maps – made it possible to address the 3D CD problem in a direct and practical way.

Furthermore, the MTBIT model, which extends to the new 3D CD task some of the methods already proposed for solving the 2D CD task, lays a solid foundation towards the solution of the 3D CD problem. In details, MTBIT achieves a metric accuracy (represented by the cRMSE) equal to 6.46 m, a value comparable to the one obtained by SUNet – the most performant among the other tested architectures w.r.t. cRMSE – but using only approximately one third of its parameters (13,1 M vs 35,7 M).

Nevertheless, much remains to be investigated from both the data and algorithmic point of view. First of all, an expansion of the dataset is already planned. This step is considered essential to enable a more effective application of those architectures that require a large amount of data to perform at their best. Particularly, a fully-automated pipeline, able to annotate thousand of images, is currently under development. Moreover, it is also important to work further on the proposed network, which, despite the encouraging results, presents some limitations, such as the difficult inference of extreme values. In addition, other types of architectures, such as generative models (e.g. GANs or diffusion models), could be tested to solve more efficiently the 3D CD task in

the future stages of the research.

References

- [1] H. Si Salah, S. E. Goldin, A. Rezgoui, B. Nour El Islam, S. Ait-Aoudia, What is a remote sensing change detection technique? Towards a conceptual framework, *International Journal of Remote Sensing* 41 (5) (2020) 1788–1812. doi:10.1080/01431161.2019.1674463. 2
- [2] S. Hafner, A. Nascetti, H. Azizpour, Y. Ban, Sentinel-1 and Sentinel-2 Data Fusion for Urban Change Detection Using a Dual Stream U-Net, *IEEE Geoscience and Remote Sensing Letters* 19 (2022) 1–5. doi:10.1109/LGRS.2021.3119856. 2, 6, 7, 8, 9
- [3] L. Bruzzone, F. Bovolo, A novel framework for the design of change-detection systems for very-high-resolution remote sensing images, *Proceedings of the IEEE* 101 (3) (2013) 609–630. doi:10.1109/JPROC.2012.2197169. 2
- [4] R. C. Daudt, B. Le Saux, A. Boulch, Y. Gousseau, Multitask learning for large-scale semantic change detection, *Computer Vision and Image Understanding* 187 (2019) 102783. doi:10.1016/j.cviu.2019.07.003. 2, 3
- [5] V. Coletta, V. Marsocci, R. Ravanelli, 3DCD: A new dataset for 2D and 3D change detection using deep learning techniques, *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLIII-B3-2022* (2022) 1349–1354. doi:10.5194/isprs-archives-XLIII-B3-2022-1349-2022. 2, 3, 5, 9, 13, 15
- [6] Z. Chen, Y. Zhang, C. Ouyang, F. Zhang, J. Ma, Automated landslides detection for mountain cities using multi-temporal remote sensing imagery, *Sensors* 18 (3). doi:10.3390/s18030821. 2
- [7] Z. Zheng, Y. Zhong, J. Wang, A. Ma, L. Zhang, Building damage assessment for rapid disaster response with a deep object-based semantic change

- detection framework: From natural disasters to man-made disasters, *Remote Sensing of Environment* 265 (2021) 112636. doi:10.1016/j.rse.2021.112636. 2
- [8] H. Lyu, H. Lu, L. Mou, W. Li, J. Wright, X. Li, X. Li, X. X. Zhu, J. Wang, L. Yu, P. Gong, Long-Term Annual Mapping of Four Cities on Different Continents by Applying a Deep Information Learning Method to Landsat Data, *Remote Sensing* 10 (3). doi:10.3390/rs10030471. 2
- [9] X. Huang, L. Zhang, T. Zhu, Building change detection from multitemporal high-resolution remotely sensed images based on a morphological building index, *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 7 (1) (2013) 105–115. doi:10.1109/JSTARS.2013.2252423. 2
- [10] N. Wang, W. Li, R. Tao, Q. Du, Graph-based block-level urban change detection using sentinel-2 time series, *Remote Sensing of Environment* 274 (2022) 112993. doi:10.1016/j.rse.2022.112993. 2
- [11] X. Huang, Y. Cao, J. Li, An automatic change detection method for monitoring newly constructed building areas using time-series multi-view high-resolution optical satellite images, *Remote Sensing of Environment* 244 (2020) 111802. doi:10.1016/j.rse.2020.111802. 2
- [12] S. H. Khan, X. He, F. Porikli, M. Bennamoun, Forest change detection in incomplete satellite images with deep neural networks, *IEEE Transactions on Geoscience and Remote Sensing* 55 (9) (2017) 5407–5423. doi:10.1109/TGRS.2017.2707528. 2
- [13] T. Hermosilla, M. A. Wulder, J. C. White, N. C. Coops, Land cover classification in an era of big and open data: Optimizing localized implementation and training data selection to improve mapping outcomes, *Remote Sensing of Environment* 268 (2022) 112780. doi:10.1016/j.rse.2021.112780. 2

- [14] R. Goldblatt, M. F. Stuhlmacher, B. Tellman, N. Clinton, G. Hanson, M. Georgescu, C. Wang, F. Serrano-Candela, A. K. Khandelwal, W.-H. Cheng, R. C. Balling, Using Landsat and nighttime lights for supervised pixel-based image classification of urban land cover, *Remote Sensing of Environment* 205 (2018) 253–275. doi:10.1016/j.rse.2017.11.026. 2
- [15] R. D. Johnson, E. Kasischke, Change vector analysis: A technique for the multispectral monitoring of land cover and condition, *International journal of remote sensing* 19 (3) (1998) 411–426. doi:10.1080/014311698216062. 2
- [16] X. X. Zhu, D. Tuia, L. Mou, G.-S. Xia, L. Zhang, F. Xu, F. Fraundorfer, Deep learning in remote sensing: A comprehensive review and list of resources, *IEEE Geoscience and Remote Sensing Magazine* 5 (4) (2017) 8–36. doi:10.1109/MGRS.2017.2762307. 2, 3
- [17] L. Ma, Y. Liu, X. Zhang, Y. Ye, G. Yin, B. A. Johnson, Deep learning in remote sensing applications: A meta-analysis and review, *ISPRS journal of photogrammetry and remote sensing* 152 (2019) 166–177. doi:10.1016/j.isprsjprs.2019.04.015. 2, 3
- [18] P. F. Alcantarilla, S. Stent, G. Ros, R. Arroyo, R. Gherardi, Street-view change detection with deconvolutional networks, *Autonomous Robots* 42 (7) (2018) 1301–1322. doi:10.1007/s10514-018-9734-5. 3
- [19] R. C. Daudt, B. Le Saux, A. Boulch, Fully convolutional siamese networks for change detection, in: 2018 25th IEEE International Conference on Image Processing (ICIP), IEEE, 2018, pp. 4063–4067. doi:10.1109/ICIP.2018.8451652. 3, 6, 18
- [20] W. G. C. Bandara, V. M. Patel, A transformer-based siamese network for change detection, *CoRR abs/2201.01293*. arXiv:2201.01293. 3, 4, 7, 18
- [21] V. Marsocci, S. Scardapane, N. Komodakis, MARE: Self-supervised multi-

- attention REsu-net for semantic segmentation in remote sensing, *Remote Sensing* 13 (16) (2021) 3275. doi:10.3390/rs13163275. 3
- [22] H. Guo, Q. Shi, A. Marinoni, B. Du, L. Zhang, Deep building footprint update network: A semi-supervised method for updating existing building footprint from bi-temporal remote sensing images, *Remote Sensing of Environment* 264 (2021) 112589. doi:10.1016/j.rse.2021.112589. 3
- [23] R. Qin, J. Tian, P. Reinartz, 3D change detection—approaches and applications, *ISPRS Journal of Photogrammetry and Remote Sensing* 122 (2016) 41–56. doi:10.1016/j.isprsjprs.2016.09.013. 3, 5, 8, 9, 13
- [24] 3DCD Dataset, <https://sites.google.com/uniroma1.it/3dchangedetection/home-page> (2022). 3
- [25] L. Mou, X. X. Zhu, IM2HEIGHT: Height estimation from single monocular imagery via fully residual convolutional-deconvolutional network, arXiv preprint arXiv:1802.10249doi:10.48550/arXiv.1802.10249. 3, 4, 18
- [26] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778. 3
- [27] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby, An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale (2020). doi:10.48550/ARXIV.2010.11929. 4
- [28] S. Zheng, J. Lu, H. Zhao, X. Zhu, Z. Luo, Y. Wang, Y. Fu, J. Feng, T. Xiang, P. H. S. Torr, L. Zhang, Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers (2020). doi:10.48550/ARXIV.2012.15840. 4
- [29] U. Okyay, J. Telling, C. L. Glennie, W. E. Dietrich, Airborne lidar change detection: An overview of earth sciences applications, *Earth-Science Re-*

- views 198 (2019) 102929. doi:10.1016/j.earscirev.2019.102929. 5, 9
- [30] I. de Gélis, S. Lefèvre, T. Corpetti, Change detection in urban point clouds: An experimental comparison with simulated 3D datasets, *Remote Sensing* 13 (13). doi:10.3390/rs13132629. 5, 8, 9
- [31] C. Benedek, T. Szirányi, A mixed Markov model for change detection in aerial photos with large time differences, in: 2008 19th International Conference on Pattern Recognition, IEEE, 2008, pp. 1–4. doi:10.1109/ICPR.2008.4761658. 5
- [32] C. Benedek, T. Szirányi, Change detection in optical aerial images by a multilayer conditional mixed markov model, *IEEE Transactions on Geoscience and Remote Sensing* 47 (10) (2009) 3416–3430. doi:10.1109/TGRS.2009.2022633. 5
- [33] C. Zhang, P. Yue, D. Tapete, L. Jiang, B. Shangguan, L. Huang, G. 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 (2020) 183–200. doi:10.1016/j.isprsjprs.2020.06.003. 5
- [34] K. Yang, G.-S. Xia, Z. Liu, B. Du, W. Yang, M. Pelillo, L. Zhang, Asymmetric siamese networks for semantic change detection in aerial images, *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022) 1–18. doi:10.1109/TGRS.2021.3113912. 5
- [35] R. C. Daudt, B. Le Saux, A. Boulch, Y. Gousseau, Urban change detection for multispectral earth observation using convolutional neural networks, in: *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, 2018, pp. 2115–2118. doi:10.1109/IGARSS.2018.8518015. 5
- [36] L. Shen, Y. Lu, H. Chen, H. Wei, D. Xie, J. Yue, R. Chen, S. Lv, B. Jiang,

- S2Looking: A satellite side-looking dataset for building change detection, *Remote Sensing* 13 (24) (2021) 5094. doi:10.3390/rs13245094. 5
- [37] M. Leenstra, D. Marcos, F. Bovolo, D. Tuia, Self-supervised pre-training enhances change detection in Sentinel-2 imagery, in: *International Conference on Pattern Recognition*, Springer, 2021, pp. 578–590. doi:0.1007/978-3-030-68787-8_42. 5
- [38] Q. Shi, M. Liu, S. Li, X. Liu, F. Wang, L. Zhang, A deeply supervised attention metric-based network and an open aerial image dataset for remote sensing change detection, *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022) 1–16. doi:10.1109/TGRS.2021.3085870. 5
- [39] A. Toker, L. Kondmann, M. Weber, M. Eisenberger, A. Camero, J. Hu, A. P. Hoderlein, c. Şenaras, T. Davis, D. Cremers, G. Marchisio, X. X. Zhu, L. Leal-Taixé, *DynamiCEarthNet: Daily multi-spectral satellite dataset for semantic change segmentation* (2022). doi:10.48550/ARXIV.2203.12560. 5
- [40] H. Chen, Z. Shi, A spatial-temporal attention-based method and a new dataset for remote sensing image change detection, *Remote Sensing* 12 (10) (2020) 1662. doi:10.3390/rs12101662. 5
- [41] J. Pan, X. Li, Z. Cai, B. Sun, W. Cui, A self-attentive hybrid coding network for 3D change detection in high-resolution optical stereo images, *Remote Sensing* 14 (9). doi:10.3390/rs14092046. 6, 10
- [42] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: N. Navab, J. Hornegger, W. M. Wells, A. F. Frangi (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Springer International Publishing, 2015, pp. 234–241. doi:10.1007/978-3-319-24574-4_28. 6
- [43] S. Fang, K. Li, J. Shao, Z. Li, Snunet-cd: A densely connected siamese

- network for change detection of vhr images, *IEEE Geoscience and Remote Sensing Letters* 19 (2022) 1–5. doi:10.1109/LGRS.2021.3056416. 6
- [44] H. Zhang, M. Lin, G. Yang, L. Zhang, ESCNet: An end-to-end superpixel-enhanced change detection network for very-high-resolution remote sensing images, *IEEE Transactions on Neural Networks and Learning Systems* (2021) 1–15doi:10.1109/TNNLS.2021.3089332. 7
- [45] M. Liu, Q. Shi, A. Marinoni, D. He, X. Liu, L. Zhang, Super-resolution-based change detection network with stacked attention module for images with different resolutions, *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022) 1–18. doi:10.1109/TGRS.2021.3091758. 7
- [46] C. Wu, H. Chen, B. Du, L. Zhang, Unsupervised change detection in multitemporal VHR images based on deep kernel PCA convolutional mapping network, *IEEE Transactions on Cybernetics* 52 (11) (2022) 12084–12098. doi:10.1109/TCYB.2021.3086884. 7
- [47] H. Chen, Z. Qi, Z. Shi, Remote sensing image change detection with transformers, *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022) 1–14. doi:10.1109/TGRS.2021.3095166. 7, 10, 25, 26
- [48] W. G. C. Bandara, V. M. Patel, Revisiting consistency regularization for semi-supervised change detection in remote sensing images (2022). doi:10.48550/ARXIV.2204.08454. 7
- [49] S. Saha, F. Bovolo, L. Bruzzone, Building change detection in vhr sar images via unsupervised deep transcoding, *IEEE Transactions on Geoscience and Remote Sensing* 59 (3) (2021) 1917–1929. doi:10.1109/TGRS.2020.3000296. 7, 8
- [50] J.-Y. Zhu, T. Park, P. Isola, A. A. Efros, Unpaired image-to-image translation using cycle-consistent adversarial networks, in: 2017 IEEE International Conference on Computer Vision (ICCV), 2017, pp. 2242–2251. doi:10.1109/ICCV.2017.244. 8

- [51] N. Xu, D. Huang, S. Song, X. Ling, C. Strasbaugh, A. Yilmaz, H. Sezen, R. Qin, A volumetric change detection framework using uav oblique photogrammetry – a case study of ultra-high-resolution monitoring of progressive building collapse, *International Journal of Digital Earth* 14 (11) (2021) 1705–1720. doi:10.1080/17538947.2021.1966527. 8, 9
- [52] S. Shirowzhan, S. M. Sepasgozar, H. Li, J. Trinder, P. Tang, Comparative analysis of machine learning and point-based algorithms for detecting 3D changes in buildings over time using bi-temporal lidar data, *Automation in Construction* 105 (2019) 102841. doi:10.1016/j.autcon.2019.102841. 9
- [53] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Vol. 30, Curran Associates, Inc., 2017, pp. 5998–6008. 11
- [54] D. Hendrycks, K. Gimpel, Gaussian error linear units (GELUs), arXivdoi:10.48550/arXiv.1606.08415. 12
- [55] Organismo Autónomo Centro Nacional de Información Geográfica, Digital elevation models and maps in image format, <http://centrodedescargas.cnig.es/CentroDescargas/buscadorCatalogo.do?codFamilia=LIDAR#>. 14