

Semi-Supervised Building Footprint Generation With Feature and Output Consistency Training

Qingyu Li, *Student Member, IEEE*, Yilei Shi, *Member, IEEE*, and Xiao Xiang Zhu^{id}, *Fellow, IEEE*

Abstract—Accurate and reliable building footprint maps are vital to urban planning and monitoring, and most existing approaches fall back on convolutional neural networks (CNNs) for building footprint generation. However, one limitation of these methods is that they require strong supervisory information from massive annotated samples for network learning. State-of-the-art semi-supervised semantic segmentation networks with consistency training can help deal with this issue by leveraging a large amount of unlabeled data, which encourages the consistency of model output on data perturbation. Considering that rich information is also encoded in feature maps, we propose to integrate the consistency of both features and outputs in the end-to-end network training of unlabeled samples, enabling to impose additional constraints. Prior semi-supervised semantic segmentation networks have established cluster assumption, in which the decision boundary should lie in the vicinity of low sample density. In this work, we observe that for building footprint generation, low-density regions are more apparent at the intermediate feature representations within the encoder than the encoder’s input or output. Therefore, we propose an instruction to assign the perturbation to the intermediate feature representations within the encoder, which considers the spatial resolution of input remote sensing imagery and the mean size of individual buildings in the study area. The proposed method is evaluated on three datasets with different resolutions: Planet dataset (3 m/pixel), Massachusetts dataset (1 m/pixel), and Inria dataset (0.3 m/pixel). Experimental results show that the proposed approach can well

extract more complete building structures and alleviate omission errors.

Index Terms—Building footprint, consistency training, semantic segmentation, semi-supervised.

I. INTRODUCTION

BUILDING footprint generation is a hot topic in the community of remote sensing, which involves numerous applications such as identifying undocumented buildings and assessing building damage after natural disasters. Remote sensing imagery that offers potential for meaningful geospatial target extraction on a large scale becomes a fundamental data source for building footprint generation. However, obtaining accurate and reliable building footprint maps from remote sensing imagery is still challenging due to several reasons. On one hand, complex and heterogeneous appearance of buildings leads to internal variability. On the other hand, mixed backgrounds and other objects with similar spectral signatures further limit the class separability.

Nowadays, convolutional neural networks (CNNs) have been widely used for remote sensing tasks [1]–[3], as they surpass conventional methods in terms of accuracy of efficiency. CNNs are capable of directly learning hierarchical contextual features from the original input, which have greater generalization capabilities for building footprint generation from remote sensing imagery. Although the existing CNNs are able to deliver very promising results [2], [4]–[6], there remains a challenge for extracting building footprints on a large scale. This challenge arises from that CNNs require massive annotated data to obtain strong supervisory information. However, manual annotation of reference data is a time-consuming and costly process.

To address this issue, a straightforward idea is to use semi-supervised learning, which can leverage a large amount of unlabeled data and alleviate the need for labeled examples. In general, the semi-supervised semantic segmentation methods are summarized into three types: weakly supervised training-based, adversarial training-based, and consistency training-based. Nevertheless, weakly supervised training-based methods need additional annotations, e.g., image-level labels or region-level labels. The adversarial training-based methods are able to make use of unlabeled data but are difficult to train. The consistency training-based approaches not only are simple to implement but also require no additional weakly labeled examples. The core idea of consistency training-based

Manuscript received September 1, 2021; revised January 5, 2022 and March 15, 2022; accepted May 2, 2022. Date of publication May 12, 2022; date of current version May 26, 2022. This work was supported in part by the European Research Council (ERC) through the European Union’s Horizon 2020 Research and Innovation Program under Agreement ERC-2016-StG-714087, Acronym: *So2Sat*; in part by the Helmholtz Association through the Framework of Helmholtz AI under Grant ZT-I-PF-5-01–Local Unit “Munich Unit @Aeronautics, Space and Transport (MASTr);” in part by the Helmholtz Excellent Professorship “Data Science in Earth Observation–Big Data Fusion for Urban Research” under Grant W2-W3-100; in part by the German Federal Ministry of Education and Research (BMBF) in the Framework of the International Future AI Laboratory “AI4EO–Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond” under Grant 01DD20001; in part by the German Federal Ministry of Economics and Technology in the Framework of the “National Center of Excellence ML4Earth” under Grant 50EE2201C; in part by the “Investigation of Building Cases Using AI” Project funded by the Bavarian State Ministry of Finance and Regional Identity (StMFH); and in part by the Bavarian Agency for Digitization, High-Speed Internet and Surveying. (*Corresponding author: Xiao Xiang Zhu.*)

Qingyu Li and Xiao Xiang Zhu are with the Chair of Data Science in Earth Observation, Technische Universität München (TUM), 80333 Munich, Germany, and also with the Remote Sensing Technology Institute (IMF), German Aerospace Center (DLR), 82234 Wessling, Germany (e-mail: qingyu.li@tum.de; xiaoxiang.zhu@dlr.de).

Yilei Shi is with the Chair of Remote Sensing Technology, Technische Universität München (TUM), 80333 Munich, Germany (e-mail: yilei.shi@tum.de).

Digital Object Identifier 10.1109/TGRS.2022.3174636

methods is to encourage the network to give consistent outputs for unlabeled inputs that are perturbed in various ways, thus improving the generalization of the network [7].

The state-of-the-art consistency training-based methods exploit the teacher–student framework [8]. Specifically, a student model is applied to the unlabeled sample, while a teacher model is applied to a perturbed version of the same sample. Afterward, the consistency is imposed between the outputs of two models to improve the performance of the student model [8]. However, there is still a certain gap in performance between these two models when the outputs are not completely correct during training. Inspired by [9] that feature maps can capture more discriminative contextual information, we further improve the performance of consistency training by proposing a new consistency loss that measures the discrepancy between both feature maps and outputs of the student model and those of the teacher model. By doing so, it can offer a strong constraint to regularize the learning of the network.

The effectiveness of the consistency training-based approaches depends heavily on the behavior of data distribution, i.e., cluster assumption, where the classes must be separated by low-density regions. However, the low-density regions separating the classes are not within the inputs, which offers an explanation for why semi-supervised is a challenging problem for semantic segmentation [10]. [11] observes that for natural images, low-density regions separating the classes are present at the encoder’s output, thus proposing to assign the perturbation at this position. However, for remote sensing imagery with low spatial resolution, we observe that the presence of low-density regions separating the classes is within the intermediate feature representations in the encoder rather than the encoder’s input or output. Motivated by this observation, in this work, we propose to enforce consistency over perturbation applied to feature representations at a certain depth within the encoder, where this depth should be in line with the spatial resolution of remote sensing imagery and the mean size of individual buildings in the study area.

Specifically, we consider a shared encoder and a main decoder that are trained together using labeled examples. To leverage unlabeled data, we then consider an auxiliary decoder whose inputs are perturbed versions of the shared encoder’s output. The consistency is imposed between outputs and feature maps of the main decoder and those of the auxiliary decoder. By doing so, the shared encoder’s representation is enhanced using the additional training signal extracted from the unlabeled data.

This work’s contributions are threefold.

1) We propose a semi-supervised network for building footprint generation, which has not been adequately addressed in the current literature. When the annotated samples are insufficient, the proposed method can leverage a large amount of unlabeled data to improve the performance of a model.

2) Our proposed method integrates consistency training of features and outputs into a unified objective function, which formulates an efficient end-to-end training framework. Compared with other competitors, our approach gains significant improvements.

3) Observing that the low-density regions separating the classes are within the intermediate feature representations in the encoder, we propose an instruction, in which the perturbation is applied on the feature representations at a certain depth within the encoder according to the spatial resolution of input remote sensing imagery and the mean size of individual buildings in the study area.

The remainder of this article is organized as follows. Related work is reviewed in Section II. Section III details the proposed network for building footprint generation. The experiments are described in Section IV. The results and Discussions are provided in Sections V and VI, respectively. Eventually, Section VII summarizes this work.

II. RELATED WORK

A. Building Footprint Generation

A tremendous amount of remote sensing imagery can be collected with recent technological advances, providing huge potential for mapping buildings. A variety of methods have been proposed to generate building footprints from remote sensing imagery.

Early studies can be categorized into four types: geometrical primitive-based, index-based, segmentation-based, and classification-based methods. The geometrical primitive-based methods [12] first extract geometric primitives (e.g., building edges and corners) and then group them to form building hypotheses. In the index-based methods [13], an index is designed to discriminate buildings from other objects. Afterward, buildings are extracted by selecting an empirical threshold. Using over-segmentation algorithms, the segmentation-based methods [14] aim at partitioning an image into different segments, so-called superpixels, and identify those belonging to buildings. In the classification-based methods [15], spectral and/or spatial features of each pixel are taken as input of classifiers to differentiate building from other classes. Nonetheless, a general limitation of these methods is that they rely heavily on manually defined rules and handcrafted features, resulting in a decrease in accuracy and efficiency.

In the past few years, deep-learning-based methods have shown remarkable performance on this task, as discriminative features from raw images can be automatically and adaptively learned. Early methods [16], [17] use a patch-wise classification framework and assign the label to each pixel according to the class of its enclosing patch. However, the large overlap among patches leads to redundant operation and low efficiency. Therefore, semantic segmentation networks that can efficiently perform pixel-wise segmentation have become more popular in the task of building footprint generation [5], [6], [18]–[25]. The commonly used network architectures involve fully convolutional networks (FCNs) [26] and encoder–decoder based architectures (e.g., DeepLabv3+ [27] Efficient-UNet [28], FC-DenseNet [29]). To take the characteristics of buildings in remote sensing imagery into account, some methods (e.g., efficient separable factorized network (ESFNet) [30], multiscale aggregation (MA)-FCN [31], holistically-nested attention (HA) U-Net [32], and

Multitask [33]) have made some specific adaptations to these network architectures, e.g., attention block and multiscale feature aggregation. More recently, instance segmentation networks are exploited to delineate individual building instances in several novel studies [34], [35]. Instance segmentation networks can not only assign a semantic label to each pixel with the class of its enclosing object but also distinguish different instances. The commonly used instance segmentation architecture for this task is Mask R-CNN [36].

B. Semi-Supervised Semantic Segmentation

The deep learning methods require strong supervisory information for network training; however, collection of large volumes of annotated data is time-consuming and costly. Especially for the task of semantic segmentation, acquisition of pixel-level labels is more expensive and laborious. Therefore, semi-supervised learning is favored in this task, and it can leverage a large amount of unlabeled data to compensate for limited supervisory information. In general, the semi-supervised semantic segmentation methods are summarized into three types: weakly supervised training-based, adversarial training-based, and consistency training-based.

Weakly supervised training-based methods [37]–[40] integrate weakly supervised learning in their approaches. Apart from the limited pixel-level labels, they still require weaker labels that can be regarded as supervisory information for network training. For the application of building footprint generation, weaker labels include image-level labels, bounding boxes, and point labels. The image-level label has two classes, where “building” refers to the images occupying building pixels more than a certain amount of the total pixels, and “nonbuilding” corresponds to images without building pixels [41], [42]. In [43], bounding box annotations are used to generate probabilistic masks using bivariate Gaussian distribution for every image. Point labels (two points inside and outside each small building, respectively) are used in [44], which are helpful to detect small buildings. Nevertheless, weakly supervised training-based methods fail to take advantage of massive unlabeled data. The adversarial training-based methods [45], [46] are able to exploit unlabeled samples, which adapt generative adversarial networks (GANs) [47] for semi-supervised semantic segmentation. Both the generator and the discriminator are first trained by labeled samples. Afterward, the generator outputs the segmentation masks of unlabeled images, while the discriminator distinguishes trustworthy regions in their predicted results to provide additional supervisory signals. Considering that the adversarial training strategy may be insufficient to guide network training, pseudo labels are generated by selecting high-confident segmentation predictions for unlabeled images [48]. Afterward, pseudo-building masks are incorporated to expand the training data and the generator is retrained. However, the adversarial training-based methods are very hard to train due to the instability of GANs [49]. In contrast, consistency training-based methods not only can leverage unlabeled images to improve the performance of the segmentation network but also are simple and efficient to implement. The goal of consistency training is to enforce the consistency of model’s predictions for unlabeled inputs that are

applied by small perturbations. By doing so, the robustness of the learned model will be enhanced.

Recently, several consistency training-based methods are proposed for the task of semi-supervised semantic segmentation, e.g., CutMix [10] and cross-consistency training (CCT) [11]. CutMix [10] applies perturbations to raw input and uses MixUp [50] to enforce the consistency between mixed outputs and outputs from mixed inputs. CCT [11] imposes an invariance of the model’s outputs over small perturbations applied to the encoder’s output. In the remote sensing community, two consistency training-based methods have been proposed for the application of building footprint generation, i.e., consistency regularization (CR) [51] and PiCoCo [52]. Color jitter and random noise are chosen as the perturbation for CR [51] and are applied to raw input. Then, the consistency of their outputs is enforced. PiCoCo [52] is also an input perturbation method, which augments the input images randomly and imposes the consistency constraint between the predictions of augmented images. In addition, it implements contrast learning on labeled images, which can regularize the compactness of intra- and interclass latent representation spaces [52].

However, these consistency training-based methods still have two limitations. On one hand, these methods ignore the rich information encoded in feature maps and generally impose consistency only over the outputs of the models. On the other hand, they add perturbations over the raw input or encoder’s output for all types of data, failing to take the characteristics of target objects into consideration when selecting the optimal position to apply perturbations.

III. METHODOLOGY

In this section, the consistency training-based methods are first introduced. Afterward, the proposed framework in the end-to-end network learning procedure is described. Finally, we propose an instruction to assign perturbation for the task of building footprint generation, which is based on our observation and analysis of cluster assumption.

A. Consistency Training-Based Methods

Given a small set of n input–target pairs $S_l = \{(x_1^l, y_1), \dots, (x_n^l, y_n)\}$ sampled from an unknown joint distribution $\beta(x, y)$, the goal of supervised learning (SL) is to derive a prediction function $f_\theta(x)$ parametrized by θ , and this prediction function is able to assign the correct target y to an unseen sample from $\beta(x)$. In semi-supervised learning, a larger set of m unlabeled examples $S_u = \{x_1^u, \dots, x_m^u\}$ is additionally provided. Semi-supervised learning aims to derive a more accurate prediction function than what is obtained by only using S_l . For instance, additional structure about input distribution $\beta(x)$ can be learned from S_u to produce an estimate of the decision boundary, which makes a better separation of samples into different classes [53].

The consistency training-based methods follow an intuitive goal to perform semi-supervised learning: when a perturbation is assigned to the data points $x \in S_u$ as \hat{x} , the output of $f_\theta(x)$ should not be significantly changed. Therefore, the objective of the consistency training-based methods is to minimize the

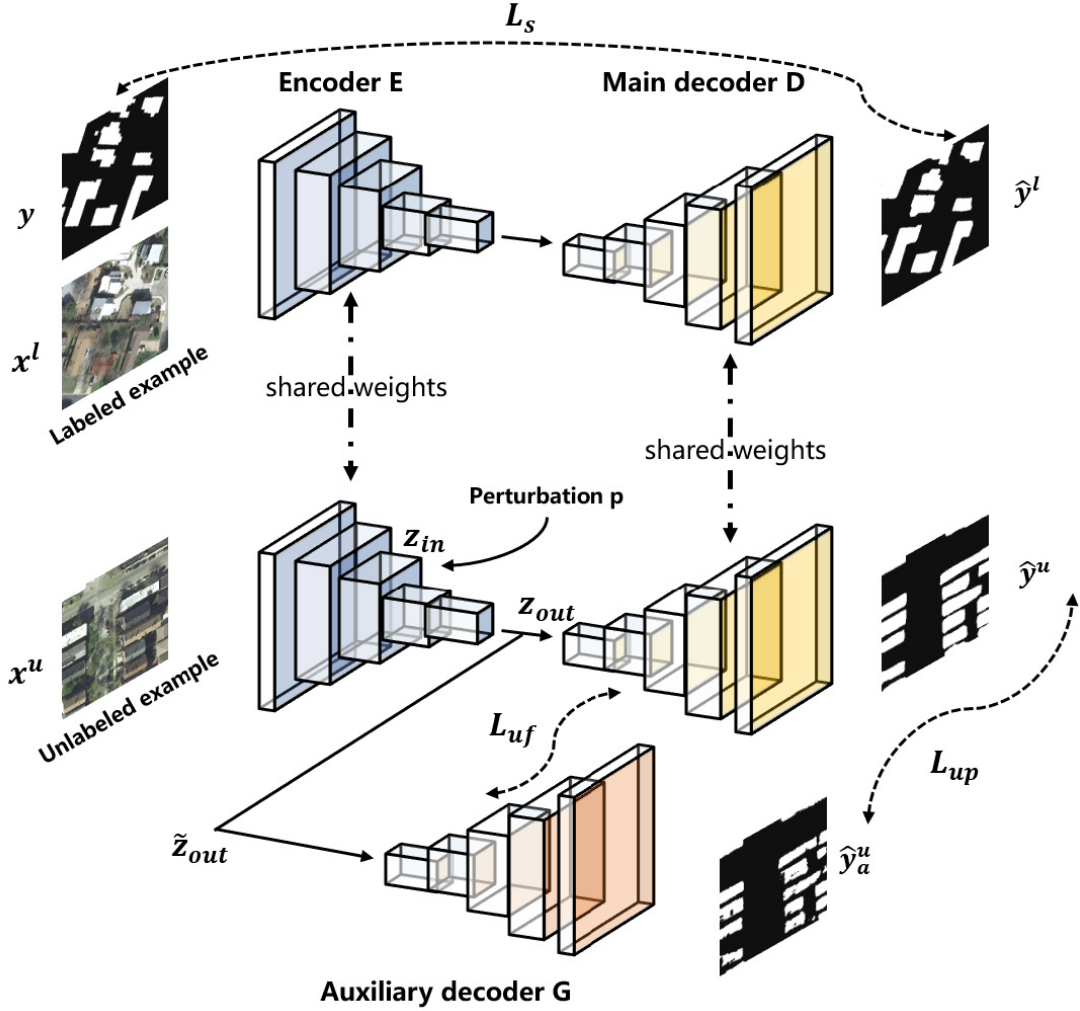


Fig. 1. Overview of the proposed semi-supervised building footprint generation network.

following loss function:

$$L = L_s + \lambda_u \cdot L_{cons} \quad (1)$$

where L_s is a supervised loss on labeled data. λ_u is a weighting function to control the importance of a consistency loss term L_{cons} which is formalized as

$$L_{cons} = \mathbf{T}(f_\theta(x), f_\theta(\hat{x})) \quad (2)$$

where $\mathbf{T}(\cdot, \cdot)$ measures a discrepancy between the outputs of prediction functions. In this regard, the unlabeled data can be leveraged to find a smooth manifold where the dataset lies [54].

Different settings in assigning perturbation or minimizing the L_{cons} lead to a wide variety of approaches for semi-supervised classification, e.g., virtual adversarial training (VAT) [55] and interpolation consistency training (ICT) [56], and those from semi-supervised semantic segmentation, e.g., CutMix [10], CCT [11], CR [51], and PiCoCo [52]. These methods are conducted in teacher–student frameworks, where a teacher model is first constructed from data perturbation, and then the output of the teacher model on unlabeled data is used to supervise a student model [8]. However, they have

not fully leveraged the information of the teacher model. This is because they fail to use intermediate feature maps of the teacher model that can also be regarded as knowledge to guide the learning of the student model. Therefore, a more precise consistency toward the underlying invariance of features and outputs between the student model and the teacher model is preferable in our research.

B. Proposed Framework in End-to-End Network Learning

Recently, the perceptual mechanism has achieved promising results for image reconstruction [9], and they make use of extracted high-level feature maps to improve the network performance. Inspired by it, we propose to impose consistency on both features and predictions for the training of unlabeled data, which is capable of fully harnessing information in deep features and output predictions. As a consequence, our network can guarantee that deep feature maps are consistent, alleviating the loss of detailed information during network training.

As shown in Fig. 1, the proposed framework is composed of a shared encoder E , a main decoder D , and an auxiliary decoder G . The segmentation network F is constituted as

$F = E \circ D$ and is trained on the labeled set in a fully supervised manner. The auxiliary network $A = E \circ G$ is trained on the unlabeled examples by enforcing the consistency of both features and outputs between D and G . D takes as input the encoder’s output \mathbf{z}_{out} , but G is fed with its perturbed version $\tilde{\mathbf{z}}_{\text{out}}$, in which the perturbation p is applied to feature representations \mathbf{z}_{in} at a certain depth within E . By doing so, the representation learning of E can be further improved by unlabeled examples, and subsequently, that of the segmentation network F .

For each iteration of training, a labeled input image x^l and its label y are sampled together with an unlabeled image x^u . Both x^l and x^u are passed through E and D , obtaining two main predictions \hat{y}^l and \hat{y}^u , respectively. The supervised loss L_s is computed with y and \hat{y}^l . For x^u , the perturbation p is applied to \mathbf{z}_{in} with \mathbf{z}_{in} being its feature representation within E and its output from E is $\tilde{\mathbf{z}}_{\text{out}}$. Afterward, an auxiliary prediction \hat{y}_a^u is generated from G using $\tilde{\mathbf{z}}_{\text{out}}$. The consistency loss L_{cons} consists of two parts L_{uf} and L_{up} , where L_{uf} is computed between the features of G and those of D , and L_{up} is computed between the outputs of G and that of D .

In the proposed approach, S_l and S_u are jointly trained by minimizing a global loss function L as (1). Following [57], λ_u is set to ramp up starting from zero along a Gaussian curve up to a fixed weight α , which can avoid the use of initial noisy output from the main encoder. The total loss L is derived and backpropagated to train the segmentation network F and the auxiliary network A . Note that L_{cons} is not backpropagated through D , and D is trained only by labeled examples. By doing so, D is only trained on the original input data. This is helpful from two aspects. On one hand, it can avoid collapsing solutions. If L_{cons} is backpropagated through both main decoder D and auxiliary decoder G , the main decoder D will collapse since L_{cons} will be minimized if predictions of both D and G are zeros. On the other hand, the method can be better adapted to the test stage since no perturbation is applied to test images.

For the labeled set, a supervised loss L_s is exploited to train the segmentation network F . To avoid overfitting, an annealed version of the bootstrapped cross-entropy loss [11] is chosen to compute the supervised loss L_s , and it is denoted as

$$L_s = \frac{1}{|S_l|} \sum_{x_i^l, y_i \in S_l} \{F(x_i^l) < \eta\} \mathbf{H}(y_i, F(x_i)) \quad (3)$$

where $F(x_i)$ is the output probability from F for a labeled example x_i , y_i is its ground reference label, and $\mathbf{H}(\cdot, \cdot)$ is the cross-entropy-based loss. In semi-supervised learning, the model is often overfitted to the limited amount of labeled data while being underfit to unlabeled data. To address this issue, a labeled example is used only if the model’s confidence in it is lower than a predefined threshold η . In other words, L_s is computed only over pixels with a probability less than the threshold η that serves as a ceiling to prevent overtraining on easy labeled data [58]. Following [11], we gradually increase η from 0.5 to 0.9 during the beginning of training.

For an unlabeled example x_i^u , \mathbf{z}_{out} is derived as the output from the shared encoder E . One contribution in our approach

is to apply perturbation to feature representation \mathbf{z}_{in} for x^u within the encoder E according to our proposed instruction. Afterward, the perturbed feature representations $\tilde{\mathbf{z}}_{\text{in}}$ will be fed to the subsequent layers in the encoder to generate the perturbed encoder’s output $\tilde{\mathbf{z}}_{\text{out}}$. Finally, \mathbf{z}_{out} and $\tilde{\mathbf{z}}_{\text{out}}$ are taken as input for D and G , respectively.

The training objective of the unlabeled set is to minimize a consistency loss L_{cons} , which is defined as

$$L_{\text{cons}} = L_{\text{up}} + \omega_u \cdot L_{\text{uf}} \quad (4)$$

where L_{uf} and L_{up} measure the discrepancy between the features and outputs of D and those of G , respectively. ω_u is a hyperparameter to introduce a weight to model the relative importance of two losses. More specifically, L_{up} is defined as

$$L_{\text{up}} = \frac{1}{|S_u|} \sum_{x_i^u \in S_u} \mathbf{T}(D(\mathbf{z}_{\text{out}}), G(\tilde{\mathbf{z}}_{\text{out}})) \quad (5)$$

where $\mathbf{T}(\cdot, \cdot)$ is the mean-squared-error-based loss.

Note that a contribution of our approach is that a loss term L_{uf} is introduced into the proposed network by imposing the consistency on features between the main decoder and the auxiliary decoder, which is able to harness detailed information in feature maps. Let $\phi_j(q)$ be the activations of the j th layer of the network ϕ when processing the input q . For D and G , $D_j(\mathbf{z}_{\text{out}})$ and $G_j(\tilde{\mathbf{z}}_{\text{out}})$ will be the corresponding feature maps at j th depth in the decoder. Here, j represents the position where upsampling operations are applied in the decoder. Then, L_{uf} is denoted as

$$L_{\text{uf}} = \frac{1}{|S_u|} \sum_{x_i^u \in S_u} \sum_{j=1}^J \mathbf{T}(D_j(\mathbf{z}_{\text{out}}), G_j(\tilde{\mathbf{z}}_{\text{out}})) \quad (6)$$

where J is the total number of depth in the decoder. In other words, J represents how many upsampling operations are applied in the decoder.

The proposed semi-supervised method can be summarized by the following Algorithm 1:

C. Instruction to Assign Perturbation for the Task of Building Footprint Generation

The effectiveness of consistency training-based methods relies on cluster assumption, i.e., two samples belonging to the same cluster in the input distribution are likely to have the same label [59]. In this case, the decision boundary should lie in low-density regions [60]. In other words, if a decision boundary crosses a high-density region, it will divide a cluster into two different classes, which violates cluster assumption. From the formal analysis, the expected value of L_{cons} is proportional to the squared magnitude of the Jacobian of the network’s outputs with respect to its inputs [7]. Therefore, minimizing L_{cons} indicates that the decision function in the regions of unsupervised samples will be flattened, and the decision boundary will be moved into the vicinity of low sample density [10].

The cluster assumption has inspired many recent consistency training-based methods for semi-supervised semantic segmentation [10], [11] which propose to assign perturbation

Algorithm 1 Algorithm for Feature and Output Consistency Training

Input: Labeled image x^l and pixel-level label y , as well as unlabeled image x^u

Require: Shared encoder E , main decoder D with the total depth number J , and auxiliary decoder G

- 1: Forward x^l through E and D : $\hat{y}^l = D(E(x^l))$
 - 2: Forward x^u through E : $\mathbf{z}_{out} = E(x^u)$
 - 3: Generate the main decoder's feature maps for \mathbf{z}_{out} :
 - 4: **For** $j = 1$ to J **do**
 Derive $D_j(\mathbf{z}_{out})$
 - 5: **End for**
 - 6: Generate the main decoder's output for \mathbf{z}_{out} :
 Derive $D(\mathbf{z}_{out})$
 - 7: Forward x^u through E and apply a noise perturbation \mathbf{N} to feature representations \mathbf{z}_{in} : $\tilde{\mathbf{z}}_{in} = (\mathbf{z}_{in} \odot \mathbf{N}) + \mathbf{z}_{in}$
 - 8: Forward $\tilde{\mathbf{z}}_{in}$ through subsequent layers in E to generate the perturbed encoder's output $\tilde{\mathbf{z}}_{out}$
 - 9: Generate the auxiliary decoder's feature maps for $\tilde{\mathbf{z}}_{out}$:
 - 10: **For** $j = 1$ to J **do**
 Derive $G_j(\tilde{\mathbf{z}}_{out})$
 - 11: **End for**
 - 12: Generate the auxiliary decoder's output for $\tilde{\mathbf{z}}_{out}$:
 Derive $G(\tilde{\mathbf{z}}_{out})$
 - 13: Training the network.
 $L_s = \{\hat{y}^l < \eta\} \mathbf{H}(y, \hat{y}^l)$
 $L_{up} = \mathbf{T}(D(\mathbf{z}_{out}), G(\tilde{\mathbf{z}}_{out}))$
 $L_{uf} = \sum_{j=1}^J \mathbf{T}(D_j(\mathbf{z}_{out}), G_j(\tilde{\mathbf{z}}_{out}))$
 Update the network by $L = L_s + \lambda_u \cdot (L_{up} + \omega_u \cdot L_{uf})$
-

to the raw input or encoder's output. However, they are not suitable for the task of building footprint generation, as the characteristics of both building objects and remote sensing imagery have not been taken into account. Therefore, we propose an instruction to assign perturbation for this task, which is inspired by the observation and analysis of cluster assumption in building footprint generation from remote sensing imagery. To examine cluster assumption, the local variations at an encoder depth d are measured between the value of each pixel and its local neighbors, and local variations with high values depict the presence of low-density regions [10]. Here, d represents the position where many downsampling operations are applied in the encoder. For instance, when $d = 1$, the spatial size (i.e., height and width) of feature representation is half of that of the raw input. Similarly, when $d = 2$, the spatial size (i.e., height and width) of feature representation is 1/4 of that of the raw input. Following [11], the average Euclidean distance at each spatial location and its eight intermediate neighbors is computed for the encoder's input ($d = 0$), and the feature representations of both the intermediate layer ($d = 2$) and encoder's output ($d = 5$). Both feature representations are first resampled to the input size, and then the average distance between the neighboring activations is calculated. Fig. 2 illustrates the example results for Planet satellite imagery (3 m/pixel). The feature representations from intermediate layer and encoder's output are 24-dimensional

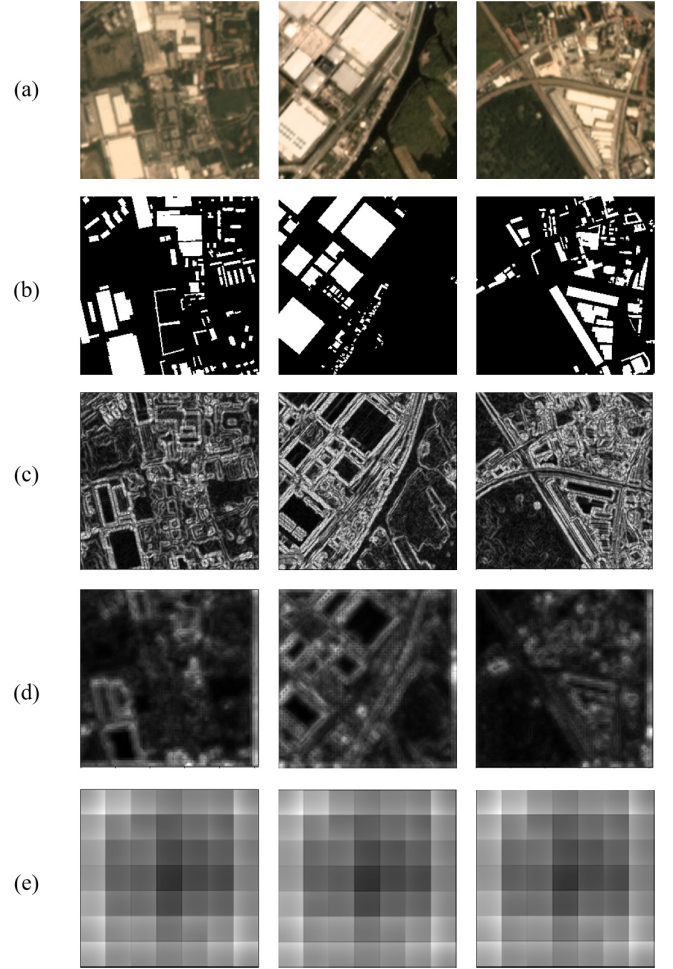


Fig. 2. Cluster assumption in consistency training-based methods for building footprint generation. Examples from (a) Planet satellite imagery (3 m/pixel), (b) pixel-level labels, and local variations at (c) encoder's input, (d) intermediate layer in the encoder, and (e) encoder's output. Bright regions indicate large variation.

and 1280-dimensional feature vectors learned from Efficient-UNet [28], respectively. It can be observed that low-density regions are not aligned with the class boundaries at the encoder's input or encoder's output, where cluster assumption is violated. In contrast, the cluster assumption is maintained at the intermediate layer, given that the class boundaries with high average distance coincide with low-density regions. This observation may be related to the receptive field of the network. The receptive field will be enlarged when the depth increases within the encoder, but when the receptive field exceeds a certain value that is much beyond the size of target objects, it might introduce more noise for network learning [61]. Furthermore, for remote sensing imagery with varying resolutions, the receptive fields of the network are various at the same depth within the encoder, when the unit is meters.

Based on the above observation and analysis, we propose an instruction to assign perturbation. The perturbation should be added to feature presentations at depth d within the encoder according to the spatial resolution of remote sensing imagery and the mean size of individual buildings in the study area.



Fig. 3. Satellite imagery of Lisbon in the Planet dataset (spatial resolution: 3 m/pixel) and three zoomed-in areas.

More specifically, d is computed as

$$d = \left\lfloor \log_2 \left(\frac{l_{\min} + l_{\max}}{2r} \right) \right\rfloor \quad (7)$$

where r is the spatial resolution of the remote sensing imagery, and l_{\min} and l_{\max} are the mean values of max and min lengths that are derived from the ground reference of individual buildings in the study area. $\lfloor \cdot \rfloor$ is the rounding down function, which aims to get the largest integer that does not exceed the original value.

A noise tensor $\mathbf{N} \sim \mu(-0.3, 0.3)$ of the same size as feature presentations \mathbf{z}_{in} is uniformly sampled as perturbation p . It is first multiplied with \mathbf{z}_{in} to adjust its amplitude, and then injected into \mathbf{z}_{in} to get perturbed feature maps $\tilde{\mathbf{z}}_{\text{in}}$

$$\tilde{\mathbf{z}}_{\text{in}} = (\mathbf{z}_{\text{in}} \odot \mathbf{N}) + \mathbf{z}_{\text{in}} \quad (8)$$

where \odot denotes element-wise multiplication. Afterward, it will be fed to subsequent layers in the encoder to generate perturbed intermediate representation $\tilde{\mathbf{z}}_{\text{out}}$ of the unlabeled input sample x^u .

IV. EXPERIMENT

A. Dataset

The effectiveness of the proposed method is validated on three datasets with different spatial resolutions, i.e., Planet dataset [62], Massachusetts dataset [16], and Inria dataset [18].

1) Planet dataset: In this research, PlanetScope satellite imagery is collected from eight European cities (Amsterdam, Berlin, Lisbon, Madrid, London, Paris, Milan, and Zurich) to create a Planet dataset. The PlanetScope satellite images have three bands (i.e., red, green, and blue) at a spatial resolution of 3 m/pixel. The corresponding building footprints that are stored as vector files are acquired from OpenStreetMap. Fig. 3 presents example imagery of Lisbon.

2) Massachusetts dataset: The Massachusetts dataset is composed of 151 tiles of aerial imagery over the city of Boston. Each aerial imagery has three bands (i.e., red, green, and blue) at a spatial resolution of 1 m/pixel, and its size is 1500 × 1500 pixels. A sample aerial image is illustrated



Fig. 4. Aerial image in the Massachusetts dataset (spatial resolution: 1 m/pixel) and three zoomed-in areas.

TABLE I
STATISTICS OF THE SELECTED DATASETS USED IN THIS RESEARCH

Dataset	City	The number of patches		
		train	validation	test
Planet dataset	Amsterdam	4800	1600	2400
	Berlin			
	Lisbon			
	Madrid			
	London			
	Paris			
	Milan			
	Zurich			
Massachusetts dataset	Boston	3424	100	250
Inria dataset	Austin	39852	6044	6044
	Chicago			
	Kitsap County			
	Western Tyrol			
	Vienna			

in Fig. 4. The corresponding ground reference building masks are also included in this benchmark dataset.

3) Inria dataset: The Inria dataset is a benchmark dataset consisting of 360 large-scale aerial images, in which each image is of the size 5000 × 5000 and has three bands (i.e., red, green, and blue) at a spatial resolution of 0.3 m/pixel. A sample aerial image is shown in Fig. 5. The ground reference building masks of this dataset are only publicly released for five cities (Austin, Chicago, Kitsap County, Western Tyrol, and Vienna).

For all three datasets, all the remote sensing images and ground-truth building masks are cut into small patches with the size of 256 × 256 pixels. For the Planet dataset, we have manually selected 1100 pairs of proper patches for each of eight European cities. The selected pairs are then separated into three parts, and the ratio of the train, validation, and




Fig. 5. Aerial image in the Inria dataset (spatial resolution: 0.3 m/pixel) and three zoomed-in areas.

test sets is 6:2:3. The data split in the Inria dataset are according to the setup in [18], [33]. More specifically, for each city, images with ids 1–5 are used for validation, and the remaining 31 images are for training. The statistics are derived from the validation set. The training/validation/test split of the Massachusetts dataset follows [16], where 137 tiles are used for training, 4 tiles are for validation, and the remaining 10 tiles are used to test models. The numbers of patches collected from each dataset for network training, validation, and test are reported in Table I.

B. Experiment Setup

Since the semantic segmentation network is an essential part of our approach, we first investigate which CNN modelx (i.e., Efficient-UNet [28], FC-DenseNet [29], DeepLabv3+ [27], ESFNet [30], MA-FCN [31], HA U-Net [32], and Multitask [33]) has a better performance for the task of building footprint generation. The CNN model achieving the best results under the fully supervised setting is selected as the backbone. Afterward, for each dataset, we randomly split the training data into two parts, which are labeled set and unlabeled set, and pixel-level annotations are excluded from the unlabeled set. Under the semi-supervised setting, the ratios of labeled data to unlabeled data are set as three different ratios (e.g., 1:2, 1:5, and 1:10). To validate the superiority of the proposed method, we make a comparison with other competitors, including SL, supervised learning + data augmentation (SL + DA), ICT [56], VAT [55], CutMix [10], CCT [11], CR [51], and PiCoCo [52]. The settings of λ_u and ω_u being the weights of consistency loss term and feature consistency loss term, respectively, and the position of the assigned perturbation in different methods

are shown in Table II for a better understanding of their differences. Furthermore, the effectiveness of our proposed feature and output consistency, being imposed between the main decoder and the auxiliary decoder, is analyzed. The position within the encoder to apply perturbation is also carefully investigated for different datasets. Finally, we explore whether the auxiliary decoder is able to improve the performance of the proposed method.

C. Training Details

Our experiments are conducted within a Pytorch framework on an NVIDIA Tesla with 16 GB of memory. For all the methods, the optimizer is stochastic gradient descent (SGD) with a learning rate of 0.1 and a momentum of 0.9, and the training batch size is set as 4. Detailed configurations of all the methods included in our experiments are listed as follows:

1) Efficient-UNet [28]: EfficientNet [63] is adopted as the encoder to learn feature maps. The decoder comprises five transposed convolutional layers that upsample the convolved image to predict segmentation masks.

2) DeepLabv3+ [27]: The feature extractor in DeepLabv3+ is the Xception model [64].

3) FC-DenseNet [29]: Both the encoder and the decoder in FC-DenseNet are composed of five dense blocks, and each dense block has five convolutional layers.

4) ESFNet [30]: This method uses separable factorized residual block (SFRB) as the core module. The encoder is composed of 16 blocks, where 3 blocks are downsampling blocks and 13 blocks are SFRB. The decoder consists of seven blocks for transposed convolutions and SFRB.

5) MA-FCN [31]: This approach has proposed a feature fusion structure to aggregate multiscale feature maps. It uses a feature pyramid network (FPN) [65]-based structure as the backbone where the encoder is a four-layer visual geometry group (VGG)-16 [66] architecture, and a corresponding decoder implements lateral connections between them.

6) HA U-Net [32]: The encoder of this network adopts ResNet34 [67]. The decoder comprises four modules that include up-sampling module, attention module, overall nesting module, and auxiliary loss module.

7) Multitask [33]: This method is based on SegNet [68]. It first adds one convolutional layer after the decoder to learn the distance to the border of buildings. Afterward, this learned distance mask and feature maps produced by the decoder are concatenated and fed into another convolutional layer to learn the final building masks.

8) Proposed method: The hyperparameter α in the unsupervised loss weighting function λ_u is set as 0.6. The loss term weighting parameter of feature consistency ω_u is chosen as 0.2. The network architectures of F and A are the same as that of the backbone.

9) SL: The backbone is learned from labeled samples. Note that unlabeled samples are not considered during training.

10) SL + DA: Following [69], data augmentation is first performed by randomly horizontally or vertically flipping, or rotating the image patches before training. Afterward, the backbone is trained on labeled samples.

TABLE II
SETTINGS OF ALL METHODS USED IN THIS RESEARCH. λ_u AND ω_u REPRESENT THE WEIGHTS OF
CONSISTENCY LOSS TERM AND FEATURE CONSISTENCY LOSS TERM, RESPECTIVELY

Method	λ_u	ω_u	The position of the assigned perturbation
Supervised Learning (SL)	= 0	= 0	-
Supervised Learning + Data Augmentation (SL+DA)	= 0	= 0	-
ICT [56]	> 0	= 0	encoder's input
VAT [55]	> 0	= 0	encoder's input
CutMix[10]	> 0	= 0	encoder's input
CCT [11]	> 0	= 0	encoder's output
CR [51]	> 0	= 0	encoder's input
PiCoCo [52]	> 0	= 0	encoder's input
Propose method	> 0	> 0	encoder's intermediate feature representations

TABLE III
ACCURACIES OF DIFFERENT SEMANTIC SEGMENTATION NETWORKS FOR SL ON THREE DATASETS. (%)

Method	Planet dataset (3 m/pixel) 4800 labeled		Massachusetts dataset (1 m/pixel) 3424 labeled		INRIA dataset (0.3 m/pixel) 39852 labeled	
	F1 score	IoU	F1 score	IoU	F1 score	IoU
Efficient-UNet [28]	59.03	41.87	68.70	52.32	85.34	74.42
DeepLabv3+ [27]	45.99	29.86	65.96	49.21	80.67	67.61
FC-DenseNet [29]	55.78	38.68	68.17	51.87	84.66	73.41
ESFNet [30]	56.55	39.42	67.37	50.80	83.65	71.90
MA-FCN [31]	58.40	41.35	68.95	52.62	85.03	74.27
HA U-Net [32]	53.70	36.70	64.87	48.00	84.28	72.82
Multi-task [33]	48.05	31.62	65.43	48.63	84.56	73.26

11) ICT [56] and VAT [55]: Following [10], we adapt these two semi-supervised classification methods for the task of semantic segmentation. The CNN model is the same as the backbone in our proposed method.

12) CutMix [10], CCT [11], CR [51], and PiCoCo [52]: For a fair comparison, we replace the CNN model with the same backbone in our proposed method.

D. Evaluation Metrics

The performance of models is evaluated by two metrics: F1 score and intersection over union (IoU). They can be computed as follows:

$$\text{F1 score} = \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}} \quad (9)$$

$$\text{IoU} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}} \quad (10)$$

$$\text{precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (11)$$

$$\text{recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (12)$$

where TP indicates the number of true positives, FN is the number of false negatives, and FP is the number of false positives. F1 score realizes a harmonic mean between precision and recall.

V. RESULTS

A. Results of Different Semantic Segmentation Networks for SL

The comparisons among different semantic segmentation networks for SL are presented in this section. Their respective performance is evaluated according to both quantitative (cf. Table III) and qualitative results (cf. Figs. 6–8) on three datasets, respectively. The goal of this comparison is to select

the best semantic segmentation network as the backbone for different learning methods in further experiments. In this case, we can avoid potential impacts due to convolutional layers and architectural differences.

Among these semantic segmentation networks, Efficient-UNet [28] performs better than DeepLabv3+ [27], FC-DenseNet [29], ESFNet [30], HA U-Net [32], and Multitask [33] on all three datasets. Especially for the Planet dataset that has a relatively low spatial resolution, Efficient-UNet [28] obtains increments of 13.04% and 12.01% in F1 score and IoU, respectively, when compared with DeepLabv3+ [27]. Although MA-FCN [31] is superior to Efficient-UNet [28] on the Massachusetts dataset, Efficient-UNet surpasses it by about 0.5% in IoU on both the Planet and Inria datasets. Fig. 8 presents a visual comparison among different methods on three datasets. For the Inria dataset with relatively high spatial resolution, some nonbuilding objects are wrongly identified as buildings by other methods. On the contrary, Efficient-UNet [28] is able to avoid such false alarms. The superiority of Efficient-UNet [28] on different resolution data can be attributed to its capability of systematically improving performance with all compound coefficients of the architecture (width, depth, and image resolution) balanced [28]. Thus, we take Efficient-UNet [28] as the backbone in both SL and semi-supervised learning approaches for further comparisons.

B. Comparison With Other Competitors

Furthermore, we make comparisons among the proposed method, SL, SL + DA, ICT [56], VAT [55], CutMix [10], CCT [11], CR [51] and PiCoCo [52]. Here, the ratios of labeled data to unlabeled data are designed as 1:2, 1:5, and 1:10, respectively. SL is regarded as the baseline method that




Fig. 6. Results obtained from (a) Efficient-UNet [28], (b) DeepLabv3+ [27], (c) FC-DenseNet [29], (d) ESFNet [30], (e) MA-FCN [31], (f) HA U-Net [32], and (g) Multitask [33]. (h) Satellite imagery from the Planet dataset (spatial resolution: 3 m/pixel). Pixel-based true positives, false positives, and false negatives are marked in white, green, and red, respectively.




Fig. 7. Results obtained from (a) Efficient-UNet [28], (b) DeepLabv3+ [27], (c) FC-DenseNet [29], (d) ESFNet [30], (e) MA-FCN [31], (f) HA U-Net [32], and (g) Multitask [33]. (h) Satellite imagery from the Massachusetts dataset (spatial resolution: 1 m/pixel). Pixel-based true positives, false positives, and false negatives are marked in white, green, and red, respectively.

TABLE IV
ACCURACIES OF DIFFERENT METHODS ON THE PLANET DATASET (3 m/pixel). (%)

Method	labeled:unlabeled $\approx 1 : 2$ (1600 labeled, 3200 unlabeled)		labeled:unlabeled $\approx 1 : 5$ (800 labeled, 4000 unlabeled)		labeled:unlabeled $\approx 1 : 10$ (400 labeled, 4400 unlabeled)	
	F1 score	IoU	F1 score	IoU	F1 score	IoU
Proposed method	59.35	42.20	56.19	39.07	53.78	36.78
SL	53.15	36.19	51.80	34.95	48.03	31.60
SL + DA	53.84	36.83	52.87	35.93	48.53	32.04
ICT [56]	54.23	37.20	52.20	35.32	49.87	33.22
VAT [55]	36.25	22.13	34.25	20.67	33.77	20.32
CutMix [10]	54.10	37.08	52.43	35.53	49.86	33.21
CCT [11]	56.09	38.97	53.10	36.15	50.81	34.06
CR [51]	47.17	30.86	44.60	28.77	41.38	26.08
PiCoCo [52]	54.12	37.10	52.43	35.54	46.94	30.67

is only trained with the labeled data, while SL + DA is trained on the labeled data that are already augmented. Labeled and unlabeled data are jointly trained for the proposed method,

ICT [56], VAT [55], CutMix [10], CCT [11], CR [51], and PiCoCo [52]. Their performance is evaluated from quantitative (cf. Tables IV–VI) perspectives. As an example, experiments




Fig. 8. Results obtained from (a) Efficient-UNet [28], (b) DeepLabv3+ [27], (c) FC-DenseNet [29], (d) ESFNet [30], (e) MA-FCN [31], (f) HA U-Net [32], and (g) Multitask [33]. (h) Satellite imagery from the Inria dataset (spatial resolution: 0.3 m/pixel). Pixel-based true positives, false positives, and false negatives are marked in white, green, and red, respectively.

TABLE V
ACCURACIES OF DIFFERENT METHODS ON THE MASSACHUSETTS DATASET (1 m/pixel). (%)

Method	labeled:unlabeled $\approx 1 : 2$ 1100 labeled, 2324 unlabeled		labeled:unlabeled $\approx 1 : 5$ 560 labeled, 2864 unlabeled		labeled:unlabeled $\approx 1 : 10$ 300 labeled, 3124 unlabeled	
	F1 score	IoU	F1 score	IoU	F1 score	IoU
Proposed method	70.26 \pm 0.57	54.15 \pm 0.68	68.59	52.20	67.69	51.16
SL	66.31 \pm 0.40	49.65 \pm 0.47	62.75	45.72	57.91	40.76
SL + DA	66.56 \pm 0.44	49.85 \pm 0.52	63.26	46.26	58.76	41.60
ICT [56]	67.03 \pm 0.19	50.42 \pm 0.23	63.33	46.34	60.19	43.05
VAT [55]	66.10 \pm 0.70	49.40 \pm 0.84	64.45	47.55	60.77	43.65
CutMix [10]	66.84 \pm 0.32	50.22 \pm 0.38	63.13	46.13	59.41	42.26
CCT [11]	67.79 \pm 0.33	51.30 \pm 0.39	64.54	47.64	60.70	43.58
CR [51]	65.83 \pm 0.55	51.08 \pm 0.66	63.91	46.96	60.68	43.56
PiCoCo [52]	68.76 \pm 0.52	52.39 \pm 0.62	65.73	48.96	64.76	47.88

TABLE VI
ACCURACIES OF DIFFERENT METHODS ON THE INRIA DATASET (0.3 m/pixel). (%)

Method	labeled:unlabeled $\approx 1 : 2$ 13000 labeled, 26852 unlabeled		labeled:unlabeled $\approx 1 : 5$ 6000 labeled, 33852 unlabeled		labeled:unlabeled $\approx 1 : 10$ 3600 labeled, 36252 unlabeled	
	F1 score	IoU	F1 score	IoU	F1 score	IoU
Proposed method	85.86	75.22	84.65	73.39	83.74	72.03
SL	81.94	69.41	80.38	67.61	77.87	64.12
SL + DA	82.40	70.07	81.01	68.08	78.56	64.69
ICT [56]	82.53	70.26	81.10	68.21	78.60	64.75
VAT [55]	82.79	70.63	81.42	68.66	78.48	64.58
CutMix [10]	82.89	70.77	81.34	68.55	78.59	64.74
CCT [11]	85.21	74.23	83.74	72.02	83.00	70.93
CR [51]	82.38	70.04	81.00	68.05	78.27	64.30
PiCoCo [52]	84.59	73.29	83.65	71.90	80.91	67.94

are carried out for five runs on the Massachusetts dataset where the ratio of labeled data to unlabeled data is 1:2. This provides a fair comparison, and the corresponding F1 score and IoU are shown as mean and variance. Figs. 9–11 illustrate visual results obtained by different methods for the ratio 1:10.

It can be seen from the statistics of three datasets that the proposed approach significantly boosts performance in F1 score and IoU when compared with other methods. The challenge induced by the ratio of 1:10 is the limited

data representation for buildings; however, we note that the proposed method still manages to perform better on three datasets when compared with its competitors. Our method gains improvements of 5.18%, 10.40%, and 7.91% in IoU than SL for the Planet, Massachusetts, and Inria datasets, respectively. In particular, on the Massachusetts dataset, the IoU of the proposed approach is improved by more than 7% when compared with other methods. When the ratio of labeled data to unlabeled data is 1:2, the number of labeled samples




Fig. 9. Results obtained from (a) proposed method, (b) SL, (c) SL + DA, (d) ICT [56], (e) VAT [55], (f) CutMix [10], (g) CCT [11], (h) CR [51], and (i) PiCoCo [52]. In this experiment, the ratio of labeled data to unlabeled data is 1:10 (400 labeled, 4400 unlabeled). (j) Satellite imagery from the Planet dataset (spatial resolution: 3 m/pixel). Pixel-based true positives, false positives, and false negatives are marked in white, green, and red, respectively.




Fig. 10. Results obtained from (a) proposed method, (b) SL, (c) SL + DA, (d) ICT [56], (e) VAT [55], (f) CutMix [10], (g) CCT [11], (h) CR [51], and (i) PiCoCo [52]. In this experiment, the ratio of labeled data to unlabeled data is 1:10 (300 labeled, 3124 unlabeled). (j) Aerial imagery from the Massachusetts dataset (spatial resolution: 1 m/pixel). Pixel-based true positives, false positives, and false negatives are marked in white, green, and red, respectively.

is already sufficient for SL, but our method still provides advantages over it. Note that the proposed approach performs even better than the other semantic segmentation networks (cf. Table III) that are trained on the full labeled sets. This proves that the effectiveness and robustness of the proposed approach for the task of building footprint generation.

The accuracy metric of IoU obtained by our method for the ratio of 1:2 is higher than that for the ratio of 1:10. This suggests that using more labeled samples increases the overall performances (42.20% versus 36.78% in the Planet dataset, $54.15 \pm 0.68\%$ versus 51.16% in the Massachusetts dataset, 75.22% versus 72.03% in the Inria dataset). It should be

mentioned that the proposed approach is capable of reducing the gap between different ratios. For instance, Table V shows that the IoU produced by our method, which is trained on the data of ratio of 1:10, only drops 1% than that of ratio of 1:5. This demonstrates that our method can obtain reliable segmentation results even when there is only a small number of annotated samples.

The visual results on the Planet dataset are illustrated in Fig. 9. There is a lot of missed detection in results provided by SL, VAT [55], CCT [11], CR [51], and PiCoCo [52], as the number of labeled samples is insufficient. On the contrary, our method can extract more building structures. Fig. 11 presents




Fig. 11. Results obtained from (a) proposed method, (b) SL, (c) SL + DA, (d) ICT [56], (e) VAT [55], (f) CutMix [10], (g) CCT [11], (h) CR [51], and (i) PiCoCo [52]. In this experiment, the ratio of labeled data to unlabeled data is 1:10 (3600 labeled, 36252 unlabeled). (j) Aerial imagery from the Inria dataset (spatial resolution: 0.3 m/pixel). Pixel-based true positives, false positives, and false negatives are marked in white, green, and red, respectively.

TABLE VII
ABLATION STUDY OF THE IMPOSED CONSISTENCY ON THREE DATASETS. (%)

The type of the imposed consistency	Planet dataset (3 m/pixel) 1600 labeled, 3200 unlabeled		Massachusetts dataset (1 m/pixel) 1100 labeled, 2324 unlabeled		Inria dataset (0.3 m/pixel) 13000 labeled, 26852 unlabeled	
	F1 score	IoU	F1 score	IoU	F1 score	IoU
Feature and output consistency	59.35	42.20	70.26	54.16	85.86	75.22
Output consistency	57.95	40.80	69.15	52.84	85.21	74.23

the results on the Inria dataset. It can be clearly seen that our method is able to avoid more false alarms than its competitors. This suggests that the proposed method has a better capability of using unlabeled data to improve network performance.

VI. DISCUSSION

As shown in the results on three datasets for a semi-supervised setting, our proposed method with a ratio of 2:1 can deliver the best results. Therefore, in this section, we carry out ablation studies of the proposed method under this data split.

A. Ablation Study of the Imposed Consistency

One contribution of our approach worthy of being highlighted is that we introduce a novel objective function by imposing consistency on both features and outputs between the main decoder and the auxiliary decoder.

The statistical results of different types of the imposed consistency are reported in Table VII. The experimental results show that implementing feature and output consistency for this task is helpful to improve the network performance, and we can see nearly 1% gains in IoU on all the datasets when compared with solely output consistency. This may be because more abstract and invariant information are included in feature representations [70], and the network is able to learn more knowledge when feature consistency is additionally imposed.

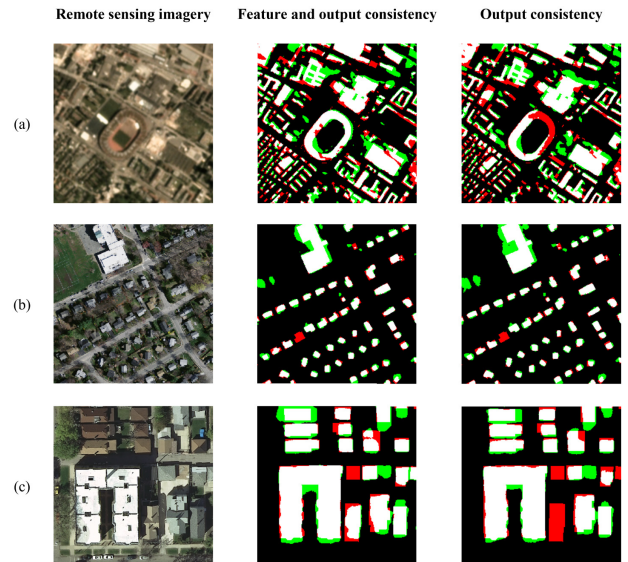


Fig. 12. Results obtained from different methods on (a) Planet dataset (spatial resolution: 3 m/pixel), (b) Massachusetts dataset (spatial resolution: 1 m/pixel), and (c) Inria dataset (spatial resolution: 0.3 m/pixel). Pixel-based true positives, false positives, and false negatives are marked in white, green, and red, respectively.

Fig. 12 illustrates a visual comparison between different types of the imposed consistency. Some buildings are omitted in the results provided by sole output consistency in the example areas of the National Institute for Research in

TABLE VIII
ABLATION STUDY OF THE ASSIGNED PERTURBATION ON THREE DATASETS. (%)

The position of the assigned perturbation	Planet dataset (3 m/pixel) 1600 labeled, 3200 unlabeled		Massachusetts dataset (1 m/pixel) 1100 labeled, 2324 unlabeled		Inria dataset (0.3 m/pixel) 13000 labeled, 26852 unlabeled	
	F1 score	IoU	F1 score	IoU	F1 score	IoU
d=1	57.72	40.47	67.67	51.46	84.65	73.59
d=2	59.35	42.20	68.66	52.27	84.44	73.28
d=3	57.32	40.18	67.02	51.08	84.74	73.73
d=4	56.58	39.45	70.26	54.16	84.67	73.63
d=5	59.63	39.50	67.65	51.11	85.86	75.22

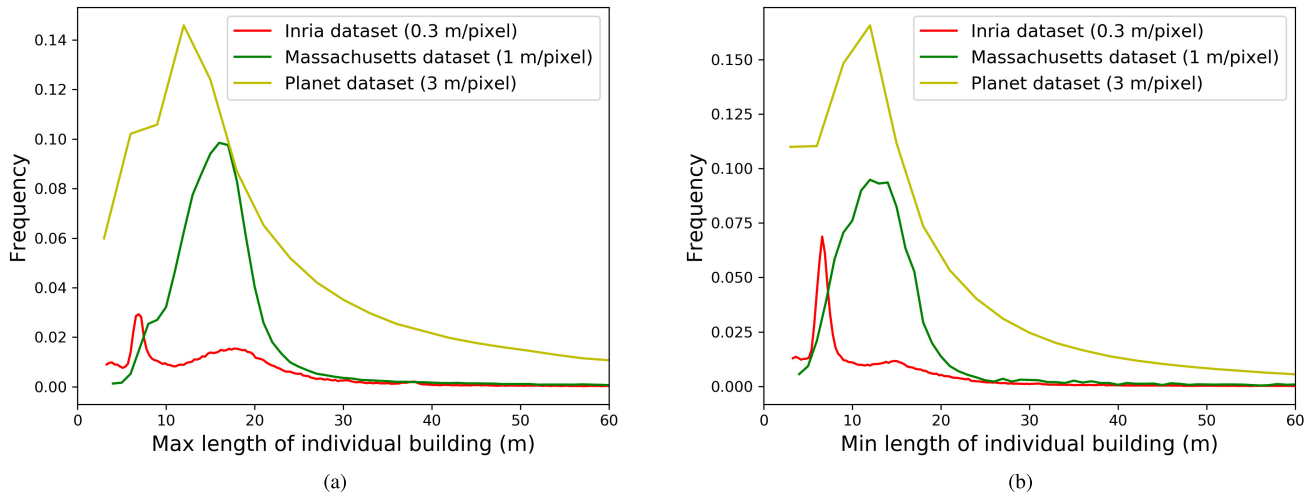


Fig. 13. Summarized statistics of (a) max length and (b) min length of individual buildings on three datasets.

Digital Science and Technology (INRIA) dataset. The reason is that the sole output consistency ignores the rich information in feature representations. On the contrary, building masks obtained by the feature and output consistency are much closer to real building shapes. This suggests that our method can capture information in both feature representations and outputs, enabling the enhancement of semantic information of buildings.

B. Ablation Study of the Assigned Perturbation

For the perturbation being assigned to feature representations within the encoder, we propose an instruction to select the optimal position: the encoder depth d . To verify this instruction, we apply the perturbation to five different positions within the encoder, respectively. Specifically, d is first set as five numbers i.e., 1, 2, 3, 4, and 5, to investigate its impact on final results. The spatial size of their corresponding feature maps is 128×128 , 64×64 , 32×32 , 16×16 , and 8×8 .

The statistical results of the perturbation applied to different depths within the encoder are shown in Table VIII. We can see that the best position to assign the perturbation is varied across different datasets. Moreover, increasing the value of the depth will promote the improvement of the results on higher resolution dataset (Inria dataset). However, we note that a large value of d will lead to a reduction in accuracy metrics on the relatively low-resolution dataset (Planet dataset). The best results are obtained when $d = 2$ for the Planet dataset, $d = 4$ for the Massachusetts dataset, and $d = 5$ for the Inria dataset. This coincides with our proposed instruction to apply the perturbation.

Taking the spatial resolution of remote sensing imagery into consideration, the respective fields of these positions are corresponding to $3 \times 2^2 = 12$ m (Planet dataset), $1 \times 2^4 = 16$ m (Massachusetts dataset), and $0.3 \times 2^5 = 9.6$ m (Inria dataset), which are close to the size of a building that usually has a length within the range from 10 to 20 m. Afterward, we calculate the statistics of individual buildings of all three datasets, i.e., max length and min length (cf. Fig. 13). We found that the mean values of the max length of individual buildings are 19 m for the Planet dataset, 17 m for the Massachusetts dataset, and 16 m for the Inria dataset. The mean values of the min length of individual buildings are 17 m for the Planet dataset, 14 m for the Massachusetts dataset, and 12 m for the Inria dataset. That is to say, the mean values of max length and min length of individual buildings also range from 10 to 20 m among all the datasets. This indicates that the geometrical characteristics of the building are related to the effective receptive field of the network, which may place an emphasis on how to select the optimal position to assign the perturbation in the whole framework. Therefore, we infer that the perturbation should be assigned to different positions within the encoder according to the spatial resolution of remote sensing imagery and the mean size of the individual buildings in the study area.

C. Ablation Study of the Auxiliary Decoder

In our approach, an auxiliary decoder is used to train the unlabeled set, and additional training signals can be extracted by enforcing the consistency of features and predictions

TABLE IX
ABLATION STUDY OF THE AUXILIARY DECODER ON THREE DATASETS. (%)

Method	Planet dataset (3 m/pixel) 1600 labeled, 3200 unlabeled		Massachusetts dataset (1 m/pixel) 1100 labeled, 2324 unlabeled		Inria dataset (0.3 m/pixel) 13000 labeled, 26852 unlabeled	
	F1 score	IoU	F1 score	IoU	F1 score	IoU
With auxiliary decoder	59.35	42.20	70.26	54.16	85.86	75.22
Without auxiliary decoder	58.37	41.20	68.73	52.34	84.84	73.67




Fig. 14. Results obtained from different methods on (a) Planet dataset (spatial resolution: 3 m/pixel), (b) Massachusetts dataset (spatial resolution: 1 m/pixel), and (c) Inria dataset (spatial resolution: 0.3 m/pixel). Pixel-based true positives, false positives, and false negatives are marked in white, green, and red, respectively.

between the main decoder and the auxiliary decoders. To validate the effectiveness of the auxiliary decoder, we perform an ablation study with another competitor, i.e., the proposed method without auxiliary decoder. That is to say, the auxiliary decoder is removed, and the main decoder takes as input both uncorrupted and perturbed versions of the encoder’s output to impose consistency on their features and outputs.

The ablation study is carried out on the Planet, Massachusetts, and Inria datasets. The numerical results are shown in Table IX. As can be seen in the statistical results on all three datasets, an auxiliary decoder brings a nearly 1% improvement in IoU, leading to a positive influence on the performance of our network. Fig. 14 shows a visual comparison of segmentation results, which demonstrates that the performance of our approach can be boosted up by the leverage of an auxiliary decoder. In Fig. 14 (e) and (h), the method without auxiliary decoder wrongly identifies cars as buildings on both the Massachusetts and Inria datasets. This is because the colors of cars are similar to those of buildings, which leads to a misjudgment. The use of an auxiliary decoder is able to avoid such false alarms. The main reason is that supervision from the same decoder might guide the network to better approximate the features and outputs of the perturbed inputs, making the network converge in the wrong direction. In contrast, supervision by the features and predictions from the other decoder is able to avoid overfitting the wrong direction.

VII. CONCLUSION


Considering that the performance of semantic segmentation networks is limited when the annotated training samples are insufficient, a novel semi-supervised building footprint generation method with feature and output consistency training is proposed in this article. The proposed model comprises three modules: a shared encoder, a main decoder, and an auxiliary decoder. More specifically, the shared encoder and the main decoder are designed to learn from labeled data in a fully supervised manner. Afterward, we assign the perturbation at intermediate feature representations within the encoder and aim to encourage the auxiliary decoder to give consistent predictions for unlabeled inputs as the main decoder. The consistency is imposed between the outputs and features of the main decoder and those of the auxiliary decoder. The performance of the proposed end-to-end network is assessed on three datasets with different resolutions: Planet dataset (3 m/pixel), Massachusetts dataset (1 m/pixel), and Inria dataset (0.3 m/pixel). The experimental results suggest that the incorporation of both feature and output consistency in our method can offer more satisfactory building footprints, where omission errors can be alleviated to a large extent. Therefore, we believe that our method is a robust solution for building footprint generation when dealing with scarce training samples. Furthermore, the best position to assign the perturbation has been investigated that the perturbation should be applied to different depths within the encoder according to the spatial resolution of input remote sensing imagery and the mean size of the individual buildings in the study area. This practical strategy is beneficial to other semi-supervised building footprint generation works that use remote sensing imagery. A subsequent study will intend to investigate the potential of the feature and output consistency training in the instance segmentation of buildings.

REFERENCES

- [1] X. X. Zhu *et al.*, “Deep learning in remote sensing: A comprehensive review and list of resources,” *IEEE Geosci. Remote Sens. Mag.*, vol. 5, no. 4, pp. 8–36, Dec. 2017.
- [2] Y. Shi, Q. Li, and X. X. Zhu, “Building footprint generation using improved generative adversarial networks,” *IEEE Geosci. Remote Sens. Lett.*, vol. 16, no. 4, pp. 603–607, Apr. 2018.
- [3] Y. Hua, L. Mou, and X. X. Zhu, “Relation network for multilabel aerial image classification,” *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 7, pp. 4558–4572, Jul. 2020.
- [4] Y. Shi, Q. Li, and X. X. Zhu, “Building segmentation through a gated graph convolutional neural network with deep structured feature embedding,” *ISPRS J. Photogramm. Remote Sens.*, vol. 159, pp. 184–197, Jan. 2020.
- [5] Q. Li, Y. Shi, X. Huang, and X. X. Zhu, “Building footprint generation by integrating convolution neural network with feature pairwise conditional random field (FPCRF),” *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 11, pp. 7502–7519, Nov. 2020.


- [6] Q. Li, L. Mou, Y. Hua, Y. Shi, and X. X. Zhu, "Building footprint generation through convolutional neural networks with attraction field representation," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–17, 2022.
- [7] B. Athiwaratkun, M. Finzi, P. Izmailov, and A. G. Wilson, "There are many consistent explanations of unlabeled data: Why you should average," in *Proc. Int. Conf. Learn. Represent.*, 2018, p. 1.
- [8] G.-J. Qi and J. Luo, "Small data challenges in big data era: A survey of recent progress on unsupervised and semi-supervised methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 4, pp. 2168–2187, Apr. 2022.
- [9] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Cham, Switzerland: Springer, 2016, pp. 694–711.
- [10] G. French, S. Laine, T. Aila, M. Mackiewicz, and G. Finlayson, "Semi-supervised semantic segmentation needs strong, varied perturbations," in *Proc. Brit. Mach. Vis. Conf.*, no. 31, 2020, pp. 1–21.
- [11] Y. Ouali, C. Hudelot, and M. Tami, "Semi-supervised semantic segmentation with cross-consistency training," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 12674–12684.
- [12] X. Qin, S. He, X. Yang, M. Dehghan, Q. Qin, and J. Martin, "Accurate outline extraction of individual building from very high-resolution optical images," *IEEE Geosci. Remote Sens. Lett.*, vol. 15, no. 11, pp. 1775–1779, Nov. 2018.
- [13] X. Huang and L. Zhang, "A multidirectional and multiscale morphological index for automatic building extraction from multispectral GeoEye-1 imagery," *Photogramm. Eng. Remote Sens.*, vol. 77, no. 7, pp. 721–732, 2011.
- [14] A. O. Ok, "Automated detection of buildings from single VHR multispectral images using shadow information and graph cuts," *ISPRS J. Photogramm. Remote Sens.*, vol. 86, pp. 21–40, Dec. 2013.
- [15] M. Turker and D. Koc-San, "Building extraction from high-resolution optical spaceborne images using the integration of support vector machine (SVM) classification, Hough transformation and perceptual grouping," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 34, pp. 58–69, Feb. 2015.
- [16] V. Mnih, *Machine Learning for Aerial Image Labeling*. Princeton, NJ, USA: Citeseer, 2013.
- [17] R. Alshehhi, P. R. Marpu, W. L. Woon, and M. D. Mura, "Simultaneous extraction of roads and buildings in remote sensing imagery with convolutional neural networks," *ISPRS J. Photogramm. Remote Sens.*, vol. 130, pp. 139–149, Aug. 2017.
- [18] E. Maggiori, Y. Tarabalka, G. Charpiat, and P. Alliez, "Can semantic labeling methods generalize to any city? The inria aerial image labeling benchmark," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2017, pp. 3226–3229.
- [19] X. Liu *et al.*, "Building instance extraction method based on improved hybrid task cascade," *IEEE Geosci. Remote Sens. Lett.*, vol. 19, pp. 1–5, 2022.
- [20] H. Huang, Y. Chen, and R. Wang, "A lightweight network for building extraction from remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–12, 2022.
- [21] H. Guo, Q. Shi, B. Du, L. Zhang, and H. Ding, "Scene-driven multitask parallel attention network for building extraction in high-resolution remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 5, pp. 4287–4306, May 2020.
- [22] Q. Zhu, C. Liao, H. Hu, X. Mei, and H. Li, "MAP-Net: Multiple attending path neural network for building footprint extraction from remote sensed imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 7, pp. 6169–6181, Jul. 2021.
- [23] H. Jung, H.-S. Choi, and M. Kang, "Boundary enhancement semantic segmentation for building extraction from remote sensed image," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–12, 2022.
- [24] K. Lee, J. H. Kim, H. Lee, J. Park, J. P. Choi, and J. Y. Hwang, "Boundary-oriented binary building segmentation model with two scheme learning for aerial images," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–17, 2022.
- [25] D. Peng, H. Guan, Y. Zang, and L. Bruzzone, "Full-level domain adaptation for building extraction in very-high-resolution optical remote-sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–17, 2022.
- [26] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 3431–3440.
- [27] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 801–818.
- [28] B. Baheti, S. Innani, S. Gajre, and S. Talbar, "Eff-UNet: A novel architecture for semantic segmentation in unstructured environment," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 358–359.
- [29] S. Jegou, M. Drozdal, D. Vazquez, A. Romero, and Y. Bengio, "The one hundred layers tiramisu: Fully convolutional DenseNets for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jul. 2017, pp. 11–19.
- [30] J. Lin, W. Jing, H. Song, and G. Chen, "ESFNet: Efficient network for building extraction from high-resolution aerial images," *IEEE Access*, vol. 7, pp. 54285–54294, 2019.
- [31] S. Wei, S. Ji, and M. Lu, "Toward automatic building footprint delineation from aerial images using CNN and regularization," *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 3, pp. 2178–2189, Mar. 2019.
- [32] L. Xu *et al.*, "HA U-Net: Improved model for building extraction from high resolution remote sensing imagery," *IEEE Access*, vol. 9, pp. 101972–101984, 2021.
- [33] B. Bischke, P. Helber, J. Folz, D. Borth, and A. Dengel, "Multi-task learning for segmentation of building footprints with deep neural networks," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2019, pp. 1480–1484.
- [34] Y. Liu, D. Chen, A. Ma, Y. Zhong, F. Fang, and K. Xu, "Multiscale U-shaped CNN building instance extraction framework with edge constraint for high-spatial-resolution remote sensing imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 7, pp. 6106–6120, Jul. 2021.
- [35] Q. Li *et al.*, "Instance segmentation of buildings using keypoints," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Sep. 2020, pp. 1452–1455.
- [36] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2961–2969.
- [37] Y. Wei, J. Feng, X. Liang, M.-M. Cheng, Y. Zhao, and S. Yan, "Object region mining with adversarial erasing: A simple classification to semantic segmentation approach," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1568–1576.
- [38] Y. Wei, H. Xiao, H. Shi, Z. Jie, J. Feng, and T. S. Huang, "Revisiting dilated convolution: A simple approach for weakly- and semi-supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7268–7277.
- [39] J. Lee, E. Kim, S. Lee, J. Lee, and S. Yoon, "FickleNet: Weakly and semi-supervised semantic image segmentation using stochastic inference," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 5267–5276.
- [40] Z. Huang, X. Wang, J. Wang, W. Liu, and J. Wang, "Weakly-supervised semantic segmentation network with deep seeded region growing," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7014–7023.
- [41] J. Chen, F. He, Y. Zhang, G. Sun, and M. Deng, "SPMF-Net: Weakly supervised building segmentation by combining superpixel pooling and multi-scale feature fusion," *Remote Sens.*, vol. 12, no. 6, p. 1049, Mar. 2020.
- [42] Z. Li, X. Zhang, P. Xiao, and Z. Zheng, "On the effectiveness of weakly supervised semantic segmentation for building extraction from high-resolution remote sensing imagery," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 14, pp. 3266–3281, 2021.
- [43] M. U. Rafique and N. Jacobs, "Weakly supervised building segmentation from aerial images," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Jul. 2019, pp. 3955–3958.
- [44] J.-H. Lee, C. Kim, and S. Sull, "Weakly supervised segmentation of small buildings with point labels," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 7406–7415.
- [45] W.-C. Hung, Y.-H. Tsai, Y.-T. Liou, Y.-Y. Lin, and M.-H. Yang, "Adversarial learning for semi-supervised semantic segmentation," in *Proc. Brit. Mach. Vis. Conf. (BMVC)*. Newcastle, U.K.: BMVA Press, Sep. 2018, pp. 1–65. [Online]. Available: <http://bmvc2018.org/contents/papers/0200.pdf>
- [46] N. Souly, C. Spampinato, and M. Shah, "Semi supervised semantic segmentation using generative adversarial network," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 5688–5696.
- [47] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, 2014, pp. 1–9.
- [48] X. Yao, Y. Wang, Y. Wu, and Z. Liang, "Weakly-supervised domain adaptation with adversarial entropy for building segmentation in cross-domain aerial imagery," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 14, pp. 8407–8418, 2021.

- [49] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2017, pp. 214–223.
- [50] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "Mixup: Beyond empirical risk minimization," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2018, pp. 1–13.
- [51] J. Wang, C. H. Q. Ding, S. Chen, C. He, and B. Luo, "Semi-supervised remote sensing image semantic segmentation via consistency regularization and average update of pseudo-label," *Remote Sens.*, vol. 12, no. 21, p. 3603, Nov. 2020.
- [52] J. Kang, Z. Wang, R. Zhu, X. Sun, R. Fernandez-Beltran, and A. Plaza, "PiCoCo: Pixelwise contrast and consistency learning for semisupervised building footprint segmentation," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 14, pp. 10548–10559, 2021.
- [53] A. Oliver, A. Odena, C. A. Raffel, E. D. Cubuk, and I. Goodfellow, "Realistic evaluation of deep semi-supervised learning algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 31, 2018, pp. 3239–3250.
- [54] M. Belkin, P. Niyogi, and V. Sindhwani, "Manifold regularization: A geometric framework for learning from labeled and unlabeled examples," *J. Mach. Learn. Res.*, vol. 7, no. 11, pp. 2399–2434, 2006.
- [55] T. Miyato, S.-I. Maeda, M. Koyama, and S. Ishii, "Virtual adversarial training: A regularization method for supervised and semi-supervised learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 8, pp. 1979–1993, Aug. 2019.
- [56] V. Verma, A. Lamb, J. Kannala, Y. Bengio, and D. Lopez-Paz, "Interpolation consistency training for semi-supervised learning," in *Proc. 28th Int. Joint Conf. Artif. Intell.*, Aug. 2019, pp. 1–23.
- [57] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning," 2016, *arXiv:1610.02242*.
- [58] Q. Xie, Z. Dai, E. Hovy, M.-T. Luong, and Q. V. Le, "Unsupervised data augmentation for consistency training," in *Proc. NIPS*, vol. 33, 2020, pp. 6256–6268.
- [59] O. Chapelle and A. Zien, "Semi-supervised classification by low density separation," in *Proc. Int. Workshop Artif. Intell. Statist. (AISTATS)*, Barbados, Caribbean, 2005, pp. 57–64, Jan. 2005.
- [60] Y. Luo, J. Zhu, M. Li, Y. Ren, and B. Zhang, "Smooth neighbors on teacher graphs for semi-supervised learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8896–8905.
- [61] W. Luo, Y. Li, R. Urtasun, and R. Zemel, "Understanding the effective receptive field in deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, 2016, pp. 1–9.
- [62] *PlanetScope*. Accessed: Aug. 6, 2020. [Online]. Available: <https://www.planet.com>
- [63] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 6105–6114.
- [64] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1251–1258.
- [65] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2117–2125.
- [66] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Int. Conf. Learn. Represent. (ICLR)*, 2015, pp. 1–14.
- [67] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [68] V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 12, pp. 2481–2495, Jan. 2017.
- [69] Z. Huang, G. Cheng, H. Wang, H. Li, L. Shi, and C. Pan, "Building extraction from multi-source remote sensing images via deep deconvolution neural networks," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2016, pp. 1835–1838.
- [70] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.



Qingyu Li (Student Member, IEEE) received the bachelor's degree in remote sensing science and technology from Wuhan University, Wuhan, China, in 2015, and the master's degree in Earth Oriented Space Science and Technology (ESPACE) from the Technische Universität München (TUM), Munich, Germany, in 2018. She is currently pursuing the Ph.D. degree with TUM and the German Aerospace Center (DLR), Wessling, Germany.

Her research interests include deep learning, remote sensing mapping, and remote sensing applications.



Yilei Shi (Member, IEEE) received the Dipl.-Ing degree in mechanical engineering and the Dr.-Ing degree in signal processing from the Technische Universität München (TUM), Munich, Germany, in 2010 and 2019, respectively.

In April and May 2019, he was a Guest Scientist with the Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, U.K. He is a Senior Scientist with the Chair of Remote Sensing Technology, TUM. His research interests include fast solver and parallel computing for large-scale problems, high-performance computing and computational intelligence, advanced methods on synthetic-aperture radar (SAR) and InSAR processing, machine learning, and deep learning for a variety of data sources, such as SAR, optical images, and medical images, and partial differential equation (PDE)-related numerical modeling and computing.



Xiao Xiang Zhu (Fellow, IEEE) received the M.Sc., Dr.-Ing., and Habilitation degrees in signal processing from the Technical University of Munich (TUM), Munich, Germany, in 2008, 2011, and 2013, respectively.

She is a Professor of data science in Earth Observation (EO) (former: Signal Processing in EO) at TUM, and the Head of the Department "EO Data Science," Remote Sensing Technology Institute, German Aerospace Center (DLR). Since 2019, she has been a Co-Coordinator of the Munich Data Science Research School (www.mu-ds.de). Since 2019, she also heads the Helmholtz Artificial Intelligence—Research Field "Aeronautics, Space and Transport." Since May 2020, she has been the Director of the International Future AI Laboratory "AI4EO—Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond," Munich. Since October 2020, she has been a Co-Director of the Munich Data Science Institute (MDSI), TUM. She was a Guest Scientist or a Visiting Professor with the Italian National Research Council (CNR)-Institute for Electromagnetic Sensing of Environment (IREA), Naples, Italy; Fudan University, Shanghai, China; the University of Tokyo, Tokyo, Japan; and the University of California at Los Angeles, Los Angeles, CA, USA, in 2009, 2014, 2015, and 2016, respectively. She is a Visiting AI Professor at European Space Agency (ESA's) Phi-Laboratory. Her main research interests are remote sensing and earth observation, signal processing, machine learning, and data science, with their applications in tackling societal grand challenges, e.g., global urbanization, united nation (UN's) sustainable development goals (SDGs), and climate change.

Dr. Zhu is a member of the Young Academy (Junge Akademie/Junges Kolleg) at the Berlin-Brandenburg Academy of Sciences and Humanities and the German National Academy of Sciences Leopoldina and the Bavarian Academy of Sciences and Humanities. She serves in the Scientific Advisory Board in several research organizations, among others the German Research Center for Geosciences (GFZ) and the Potsdam Institute for Climate Impact Research (PIK). She is an Associate Editor of the *IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING* and serves as the Area Editor responsible for special issues of the *IEEE Signal Processing Magazine*.