This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

Large-Scale Fine-Grained Building Classification and Height Estimation for Semantic Urban Reconstruction: Outcome of the 2023 IEEE GRSS Data Fusion Contest

Guozhang Liu, Baochai Peng, Ting Liu, Pan Zhang, Mengke Yuan, Chaoran Lu, Ningning Cao, Sen Zhang, Simin Huang, Tao Wang, Xiaoqiang Lu, Licheng Jiao, Qiong Liu, Lingling Li, Fang Liu, Xu Liu, Yuting Yang, Kaiqiang Chen, Zhiyuan Yan, Deke Tang, Hai Huang, Michael Schmitt, *Senior Member, IEEE,* Xian Sun, *Senior Member, IEEE,* Gemine Vivone, *Senior Member, IEEE,* Claudio Persello, *Senior Member, IEEE,* Ronny Hänsch, *Senior Member, IEEE*

This article presents the scientific outcomes of the 2023 Data Fusion Contest (DFC23) organized by the Image Analysis and Data Fusion Technical Committee of the IEEE Geoscience and Remote Sensing Society. The contest consists of two tracks investigating the fusion of optical and synthetic aperture radar data for 1) fine-grained roof type classification and 2) height estimation. During the development phase, 1,000 people registered for the contest, while at the end 55 and 35 teams competed during the test phase in the two tracks, respectively. This paper presents the methods and results obtained by the first and second-ranked teams of each track. In Track 1, both winning teams leveraged pre-training, modern network architectures, model ensembles, and measures to cope with the imbalanced class distribution. The solutions to Track 2 are more diverse and are characterized by modern multi-task learning approaches. The data of this contest is openly available to the community for further research, development, and refinement of machine learning methods.

***Index Terms*—Transformers, convolutional neural networks, deep learning, data fusion, fine-grain building classification, and monocular height estimation.**

I. Introduction

Buildings play a vital role in urban areas, yet the focus of research in the extraction and 3D reconstruction from remote sensing data often neglects detailed information about the specific roof types of buildings. This oversight restricts the potential for in-depth analysis, particularly in the context of urban planning applications. Classifying building roof types at a fine-grained level using satellite imagery poses significant challenges due to the presence of ambiguous visual features within the images. Furthermore, the lack of datasets specifically designed for fine-grained building roof type classification exacerbates the difficulty of this task.

The 2023 IEEE GRSS Data Fusion Contest (DFC23), organized by the Image Analysis and Data Fusion Technical Committee (IADF TC) of the IEEE Geoscience and Remote Sensing Society (GRSS), the Aerospace Information Research Institute under the Chinese Academy of Sciences, the Universität der Bundeswehr München, and GEOVIS Earth Technology Co., Ltd. aims to push current research on building extraction, classification, and 3D reconstruction towards urban reconstruction with fine-grained semantic information of roof types and height information.

To this aim, the DFC23 establishes a large-scale, fine-grained, and multi-modal benchmark for the classification of building roof types [1]. It consists of two challenging competition tracks investigating the fusion of optical and Synthetic Aperture Radar (SAR) data focusing on roof type classification and building height estimation, respectively. The data provided by the DFC23 includes several novel properties:

- **Globally Distributed and Large-Scale:** Buildings are distributed across seventeen cities on six continents to cover a wide range of different building styles. This allows capturing the heterogeneity of cities in different continents with various landforms.
- **Fine-Grained Roof Type Categories:** The buildings are labeled according to a detailed (fine-grained) categorization of roof types. The DFC23 provides nearly 300k instances with twelve different types of building roofs, which renders building roof-type classification significantly more challenging.
- **Multimodal Data:** To facilitate multimodal data fusion, optical imagery and SAR images are provided. The information captured by these different modalities can be jointly exploited, potentially resulting in the development of more accurate building extraction and classification models.

The contest is designed as a benchmark competition following previous editions [2]–[7] and consists of two parallel tracks:

- Track 1: Building Detection and Roof Type Classification;
- Track 2: Multi-Task Learning of Joint Building Extraction and Height Estimation.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

Track 1: Building Detection and Roof Type Classification

This track focuses on the detection and classification of building roof types from high-resolution optical satellite imagery and SAR images. The SAR and optical modalities are expected to provide complementary information. The given dataset covers seventeen cities worldwide across six continents. The classification task consists of twelve fine-grained, predefined roof types. Figure 1 shows an example.

Track 2: Multi-Task Learning of Joint Building Extraction and Height Estimation

This track defines the joint task of building extraction and height estimation. Both are fundamental tasks for building reconstruction. As in Track 1, the input data are multimodal optical and SAR satellite imagery. Building extraction and height estimation from single-view satellite imagery depend on semantic features extracted from the imagery. Multi-task learning provides a potentially superior solution by jointly analyzing features from the different data sources and forming implicit constraints between multiple tasks compared to conventional single-mode methods. Satellite images are provided with reference data, i.e., building annotations and normalized Digital Surface Models (nDSMs). The participants are required to reconstruct building heights and extract building footprints. Figure 2 shows an example.

II. The Data set

The images of the DFC23 dataset are acquired by the SuperView-1 (or "GaoJing" in Chinese), Gaofen-2 and Gaofen-3 satellites, with spatial resolutions of 0.5 m, 0.8 m, and 1 m, respectively. Normalized Digital Surface Models (nDSMs) provided for reference in Track 2 are produced from stereo images captured by Gaofen-7 and WorldView 1 and 2 with a ground sampling distance of roughly 2 m. Data was collected from seventeen cities on six continents to provide a large and representative dataset of high diversity regarding landforms, architecture, and building types. Roof type categories are organized according to twelve fine-grained roof type classes based on the geometry of the roof.

The data of this contest remains openly available to the community[1].

III. Contest Organization and Submissions

The contest consisted of two phases:

Phase 1: Participants are provided training data and additional validation images (without corresponding reference data) to train and validate their algorithms. Participants can submit results for the validation set to the Codalab competition websites for Track 1[2] and Track 2[3] to get feedback on their performance. The performance of the best submission from each account will be displayed on the leaderboard. In parallel, participants submit a short description of the approach used to be eligible to enter Phase 2.

Phase 2: Participants receive the test dataset (without the corresponding reference data) and submit their results within seven days from the release of the test data set. After evaluation of the results, four winners for each track are announced.

We received 1,000 registrations at the CodaLab competition website during the development phase (see Table I).

For Track 1, there were 680 unique registrations at the CodaLab competition website during the development phase, and 55 teams entered the test phase after screening the descriptions of their approaches submitted by the end of the development phase. 4,921 submissions were received during the development phase, with active participation from all registered teams. During the test phase, the maximum number of submissions per team was limited to 5 per day, and 420 submissions were received.

For Track 2, there were 320 unique registrations at the CodaLab competition website during the development phase, and 35 teams entered the test phase after screening the descriptions of their approaches submitted by the end of the development phase. In total, 1,692 submissions were received during the development phase, with active participation from all registered teams. During the test phase, the maximum number of submissions per team was limited to 5 per day, and 274 submissions were received.

TABLE I: Registration and submission statistics of the two tracks

	development phase		test phase	
	registrations	submissions	teams	submissions
Track 1	680	4921	55	420
Track 2	320	1692	35	274

Participants of the two tracks come from 32 countries. The specific number and proportion of participants in each country are shown in Figure 3.

The first- to fourth-ranked teams were awarded as winners of the DFC2023 for each track and were invited to present their solutions during the 2023 IEEE International Geoscience and Remote Sensing Symposium (IGARSS 2023). However, the fourth-ranked team of Track 1 decided to withdraw from the contest. In the following, we list the winning teams of the DFC2023 in Track 1

- **1st place**: *PIESAT-AI* team - Guozhang Liu, Baochai Peng, Ting Liu, Pan Zhang, Mengke Yuan, Chanran Lu, Ningning Cao, Sen Zhang, Simin Huang, Tao Wang from PIESAT Information Technology Co, Ltd., Beijing, China [8].
- **2nd place**: *IPIU-XDU* team; Xiaoqiang Lu, Licheng Jiao, Qiong Liu, Lingling Li, Fang Liu, Xu Liu, Yuting Yang from Xidian University, China [9].
- **3rd place**: *carryhjr* team; Jiarui Hu (Wuhan University), Zijun Huang (Guangdong University of Technology), Fei Shen (Nanjing University of Science and Technology), Dian He (Tsinghua University), Qingyu Xian (Tsinghua University) [10].

[1]https://ieee-dataport.org/competitions/2023-ieee-grss-data-fusion-contest-large-scale-fine-grained-building-classification

[2]CodaLab Track 1: https://codalab.lisn.upsaclay.fr/competitions/8987

[3]CodaLab Track 2: https://codalab.lisn.upsaclay.fr/competitions/8988

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

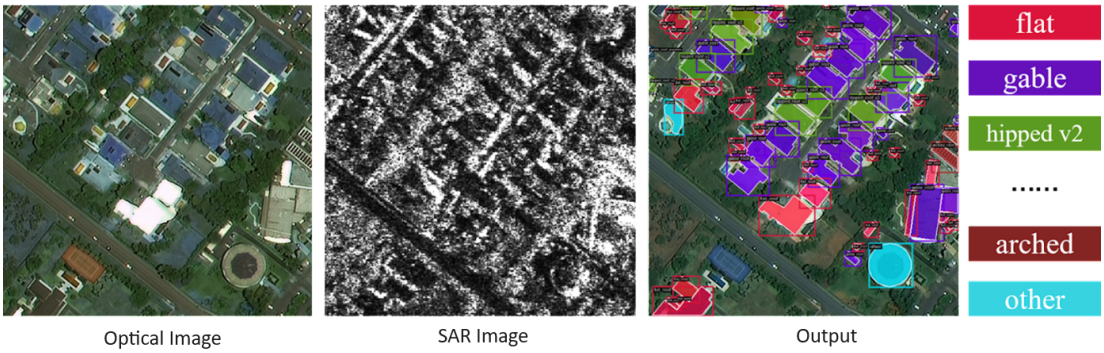

Fig. 1: Example image tile of the provided multi-modal data (optical and SAR) for building detection and roof type classification.

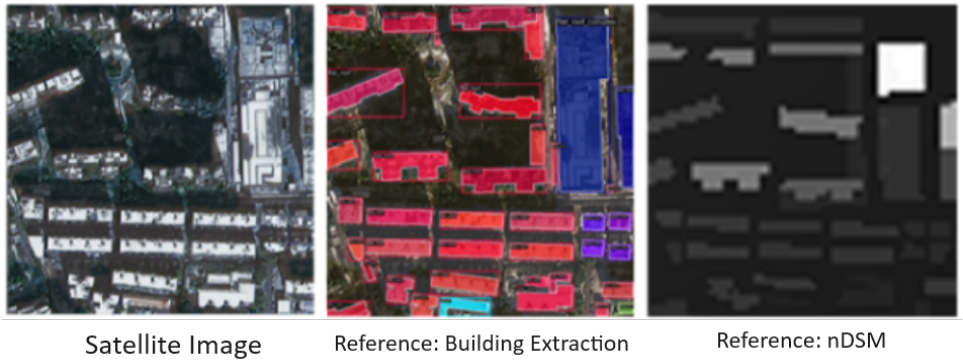

Fig. 2: Example of joint building extraction and height estimation.

and in Track 2

- **1st place**: *PIESAT-AI* team; Chaoran Lu, Ningning Cao, Pan Zhang, Ting Liu, Baochai Peng, Guozhang Liu, Mengke Yuan, Sen Zhang, Simin Huang, Tao Wang, from PIESAT Information Technology Co, Ltd., Beijing, China [11].
- **2nd place**: *IPIU-XDU* team; Xiaoqiang Lu, Licheng Jiao, Qiong Liu, Lingling Li, Fang Liu, Xu Liu, Yuting Yang from Xidian University, China [12].
- **3rd place**: *ZheWang* team; Yuxuan Guo, Zhe Wang from Wuhan University, China [13].
- **4th place**: *carryhjr* team; Jiarui Hu (Whuhan University), Zijun Huang (Guangdong University of Technology), Fei Shen (Nanjing University of Science and Technology), Dian He (Tsinghua University), Qingyu Xian (Tsinghua University) [14].

The two best-ranked teams in both tracks are invited to provide a deeper discussion on their respective approaches to win the DFC23 in this paper.

IV. Track 1 - First Place: Team PIESAT-AI

A. Method

Team PIESAT-AI introduced an effective instance segmentation framework designed to tackle the complex challenges of fine-grained building roof classification. The approach builds upon the Cascade Mask R-CNN [15] architecture and is enhanced with domain-adapted pretraining and a modernized dual-backbone structure. These techniques collectively address a range of issues, such as the significant inter-class imbalance of roof types, the diversity of building styles of global cities, and the distinction of imaging satellites. The overall structure is shown in Figure 4. To improve training stability and to realize effective weight initialization, a domain-adaptive pretraining model is employed to optimize the parameter initialization. A composite dual-branch backbone structure is integrated as a more robust and discriminative feature extractor. The newly designed backbone mitigates problems associated with small-size instance segmentation and classification of minority categories. The dual-branch backbone is applicable for optical as well as SAR inputs, which enhances the multi-modal data fusion performance. A tailored data augmentation pipeline is adopted that incorporates modified copy-paste techniques, Stochastic Weight Average (SWA [16]) training strategy, and a novel instance segmentation model aggregation method during inference. These strategies enhance the precision of the fine-grained roof instance segmentation.

Domain adapted pretraining. The strategy for initializing model weights holds a pivotal role within the optimization framework. The utilization of pre-trained weights from ImageNet-22k has two advantages: convergence acceleration and overall performance improvement. The Fully Convolutional Masked Autoencoder (FCMAE) [17] demonstrates a superior performance by improving the domain adaptation capacity for specific datasets with self-supervised pretraining tailored to convolution-based models. To this end, the pretrain-

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

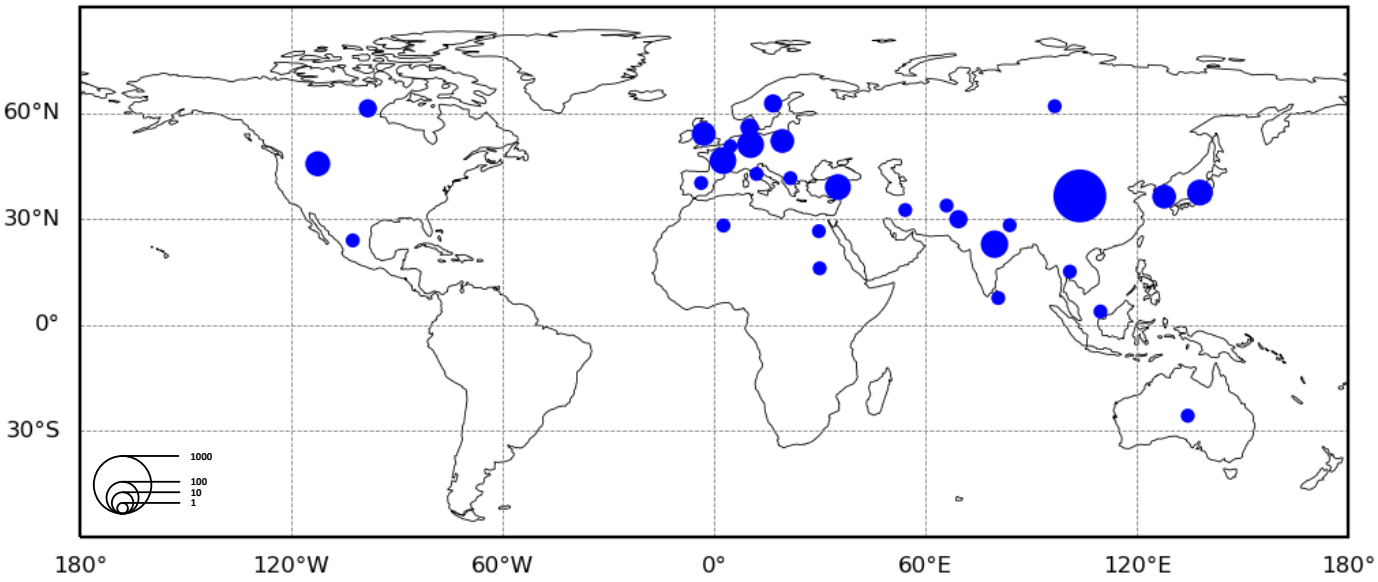

Fig. 3: Nationality statistics of participants of the two tracks.

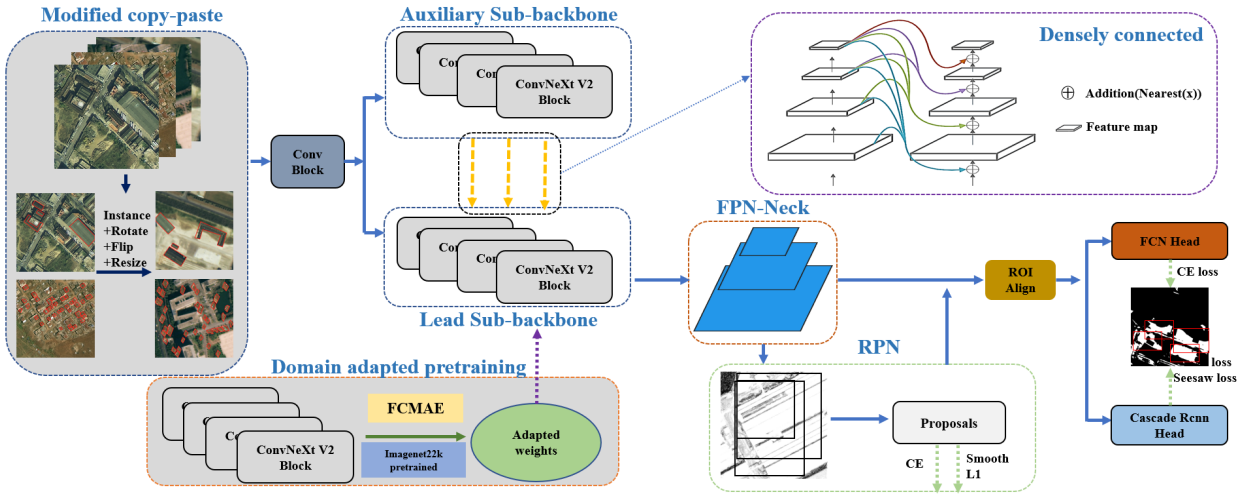

Fig. 4: The overview training workflow and network structure of the Track 1 winning approach.

ing weights of ConvNeXt V2 on the RGB modality dataset are leveraged and form the bedrock for initializing the dual sub-backbones. This elaborated weight initialization, coupled with the FCMAE pretraining, boosts the detection performance and domain-specific adaptability.

Composite dual-backbone feature extractor. To obtain a robust feature extractor, a dual-branch backbone is used encompassing two interconnected sub-backbones. This structure is validated in CBNet [18]. The configuration exhibits remarkable adaptability across both single-modality and multi-modality inputs. The lead sub-backbone and auxiliary sub-backbone are densely connected, where high-level low-resolution feature maps undergo nearest interpolation before the fusion with low-level high-resolution ones. The fusion scheme is defined by

$$F_{lead}^{j} = G_{lead}^{j}(\sum_{i=j}^{L} N(F_{aux}^{i}) + F_{lead}^{j-1}), \quad (1)$$

where L represents the number of downsampling stages of the sub-backbone, F_{aux}^{i}, F_{lead}^{j} represent the feature map of the i-th and j-th stages in the auxiliary sub-backbone and lead sub-backbone, respectively, $G_{lead}^{j}(x)$ represents a model block of the j-th stage in lead sub-backbone, and N represents the nearest neighbour interpolation operator.

The composite dual-backbone design promotes the extraction of global context and local details and achieves notable feature extraction capability.

Modified copy-paste. Copy-paste [19] is a simple yet powerful technique for instancing-level data augmentation. To enhance the diversity of the dataset, PIESAT-AI presents an improved modification of this approach. The design involves an initial extraction of instances from images, followed by the application of randomized resizing, rotation, and flipping to these pixel-level instances. Subsequently, these augmented instances are seamlessly integrated into the selected images. This synthesizing process preserves the authenticity of the

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited, content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

nadir-viewing images. Due to the risk that synthetic images might introduce distributional shifts, PIESAT-AI implements a two-phased training pipeline. In the initial phase, the model is trained using the dataset enriched by the modified copy-paste technique, covering 90% of the total training epochs. The subsequent phase focuses on fine-tuning, during which the model refines its performance over the remaining 10% of epochs without the modified copy-paste data augmentation. In addition to the distinctive augmentation strategies, PIESAT-AI's data augmentation pipeline encompasses various routine techniques, including random rotation, resizing, cropping, etc.. This holistic augmentation approach synergizes diverse methodologies to ensure the model's adaptability to a wide range of real-world scenarios.

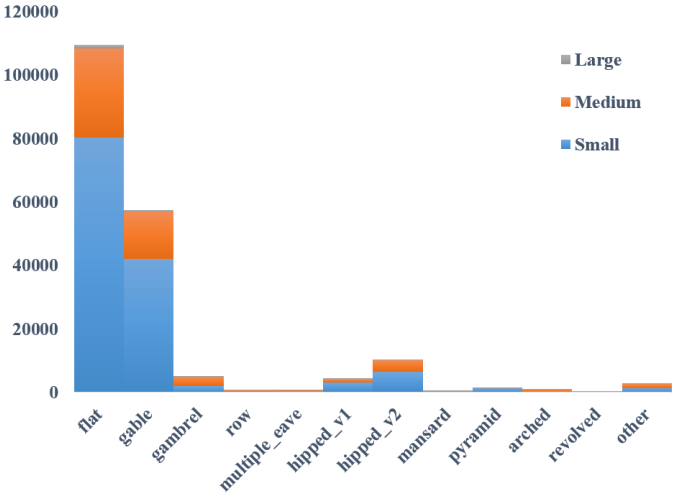

Fig. 5: Distribution of instance pixels among predefined roof categories of DFC 2023. Gray, orange and blue represent large, medium and small instances, respectively.

Additional training strategy. To effectively tackle the challenges arising from the long-tailed distribution of roof categories and the distributional disparities between the training and validation datasets, the approach focuses on the construction of a more generalized model involving the strategic incorporation of diverse training techniques to strengthen the model's performance:

- **Tactics for long-tail distribution:** The statistics of the Track 1 training dataset (see Figure 5) reflect the practical difficulty that there are significant proportion disparities for different classes. Thus, a balanced sampling strategy is utilized with the aim of increasing the probability of selecting minority-class samples. Furthermore, the seesaw loss [20] mechanism is employed to alleviate the impact of gradients stemming from less common negative class samples.
- **Stochastic Weight Average (SWA):** The SWA [16] technique with cyclical learning rates is employed during 12 further epochs after all training epochs are completed. The weights of the models of these 12 epochs are averaged to obtain the final model weights.
- **Model ensemble:** To improve the generalization capability of the model, a model ensemble strategy is employed to fuse several different models. When it comes to object detection, the widely adopted methodology of weighted bounding box fusion (WBF) [21] has emerged as a prevalent strategy for orchestrating model ensembles. PIESAT-AI introduces Weighted Segmentation Fusion (WSF), which is suitable for instance segmentation: First, the WBF strategy is adapted to obtain fused bounding boxes. Then, the masks from different models are fused to obatin the final results.
- **Multi-modality input:** Several experiments have been carried out to evaluate the potential of employing multi-modal inputs, i.e. using both RGB and SAR data. However, the outcomes of the multi-modal approach fail to match the high performance achieved through utilizing the RGB modality only.

B. Results

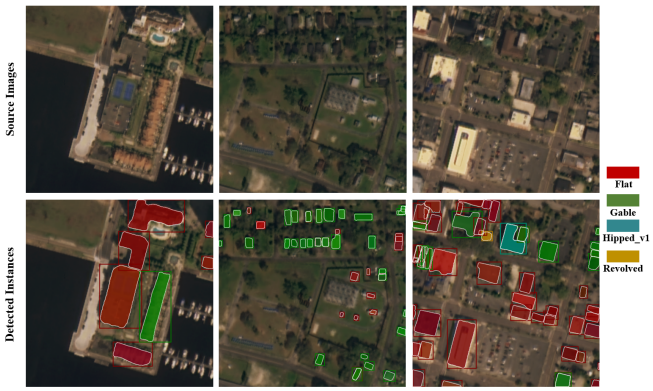

Fig. 6: Qualitative results of the first place of Track 1.

Table II contains the accuracy of different experiments. These ablation studies show that the best result with an mAP_{50} of 50.6% is achieved by fusing several strong detectors trained under different hyperparameters and backbones with WSF. Notably, SAR data does not enhance model performance but accuracy drops significantly compared to the single RGB modality input. The visualization results depicted in Figure 6 demonstrate the model's capability to accurately identify and delineate the boundaries of roofs, even for small-scale and poorly defined targets. Despite these promising results, some instances of misclassification persist and indicate the potential of further refinement and optimization of the model.

C. Discussion

An extensive series of experiments led to a robust and successful model tailored for fine-grained building roof instance segmentation. The approach integrates domain-adapted pretraining and a dual backbone framework. It utilizes optical satellite imagery as the primary input source. The inherent misalignment and heterogeneity between optical and SAR data presents great challenges to enhance performance in a multi-modal data fusion setting.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

TABLE II: The details of Team PIESAT-AI's ablation study, CMR=Cascade Mask Rcnn, SCP=Simple Copy-Paste, MCP=Modified Copy-Paste, db=Dual-Backbone, DAP=Domain Adapted Pretraining, CNV2=ConvNeXt V2, WSF= Weighted Segmentation Fusion.

Method	Development phase (mAP50)	Test phase (mAP50)
CMR + db-swin-base	0.443	/
CMR + db-swin-base **+ SAR-input**	0.371	/
CMR + db-swin-base**+ seesaw**	0.461	/
CMR + db-swin-base + seesaw **+ SCP**	0.469	/
CMR **+ db-CNV2-base-DAP**+ seesaw + SCP	0.476	/
CMR + db-CNV2-base-DAP + seesaw + SCP **+ SWA**	0.485	/
CMR + db-CNV2-base-DAP + seesaw **+ MCP** + SWA	0.496	0.426
CMR + db-CNV2-base-DAP + seesaw **+ MCP+finetune** + SWA	/	0.441
CMR + **db-CNV2-large-DAP** + seesaw + MCP+finetune + SWA	/	0.448
WSF based model ensemble	0.5510	0.5060

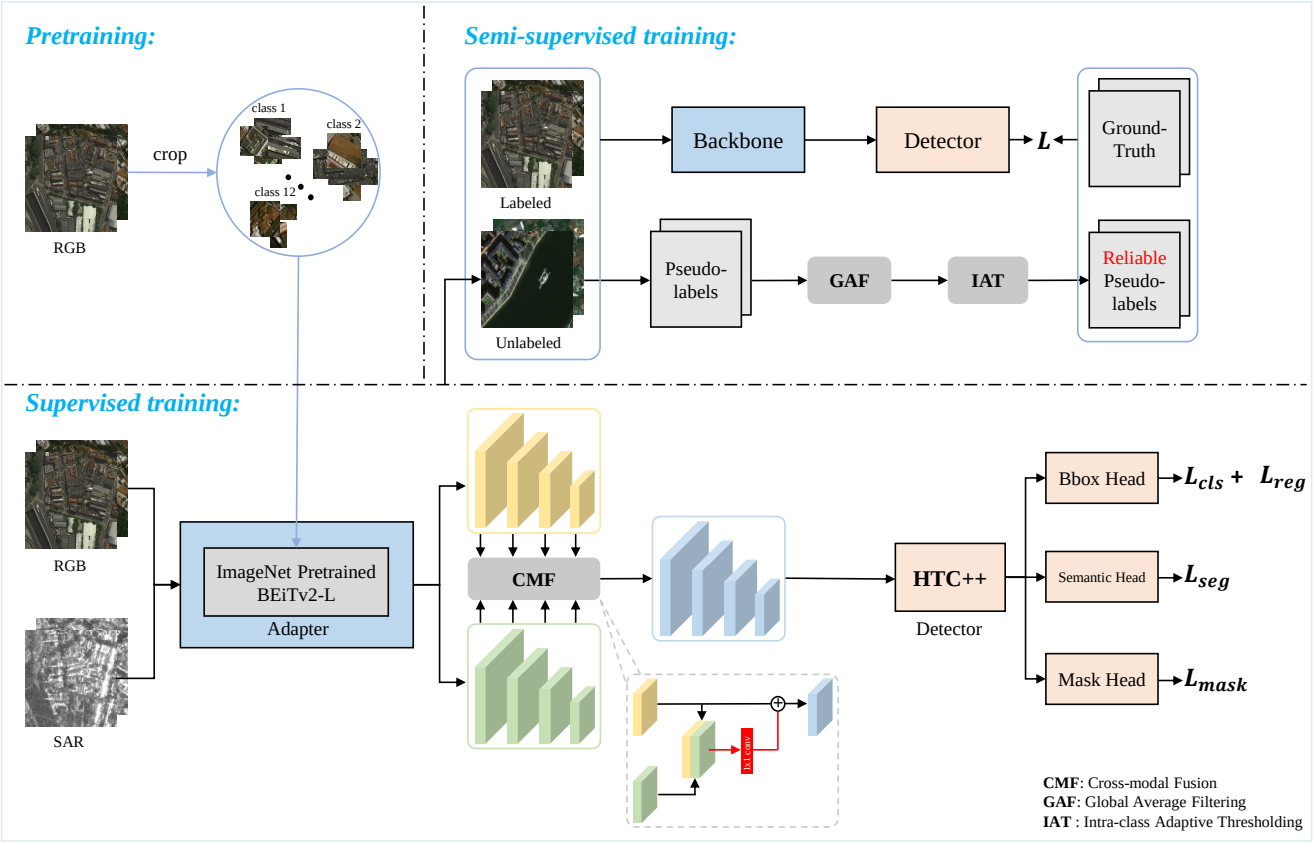

Fig. 7: Illustration of the three-stage learning framework of the second place for Track 1.

V. Track 1 - Second Place: Team IPIU-XDU

A. Method

To address the problem of a long-tailed class distribution, general data insufficiency, inter-class similarity, and intra-class differences in the DFC23 dataset, team IPIU-XDU proposes a three-stage learning framework consisting of pretraining, supervised training, and semi-supervised training shown in Figure 7.

Pre-training. Several computer vision approaches make extensive use of models pre-trained on ImageNet. However, there are considerable differences between natural and remote sensing images, such as less-defined shapes and a rather dispersed distribution of objects in remote sensing imagery. Team IPIU-XDU cropped every building box to create a fine-grained classification dataset used to fine-tune the ImageNet pre-trained model into a building-specific model. BEiTv2-L [22] is used as a feature extractor. Following the process of fine-tuning, a vision transformer adapter [23] is used to enhance the representations of the model.

Supervised training.The following three techniques aim to improve the supervised model to derive a better result for semi-supervised training.

- **Long-tailed data distribution:** Table III shows that the UBC [24] dataset is severely imbalanced. To address this issue, image-level resampling and object-level balanced copy-paste strategies are introduced. Specifically, the

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

number N_{ins} of each class instance is used to compute the ratio r_{ins} of class instances and the ratio r_{img} of class images. Two thresholds $\delta_{\text{ins}} = 0.005$ and $\delta_{\text{img}} = 0.1$ are defined to determine whether to copy-paste and resample, respectively, i.e. depending on r_{ins} or r_{img} being less than δ_{ins} or δ_{img}.

- **Cross-modal fusion:** The key obstacles in building detection and classification include occlusions, shadows, and diverse roof types of urban coverings. In terms of the benefits and drawbacks of various forms of data, RGB images provide excellent spatial resolution and textural information but are vulnerable to weather, whereas SAR images work in all weather conditions and highlight foreground items. As a result, using them both augments the input feature space and might potentially improve the resulting model. To that end, a simple but effective cross-modal fusion (CMF) module is created that independently concatenates hierarchical features extracted from RGB and SAR. Feature fusion and channel dimensionality reduction are achieved by using 1×1 convolutions and residual connections to improve the RGB feature representation.
- **MHEM loss and training optimization:** The majority of currently used fine-grained visual classification techniques overfit challenging examples in the training set but do not learn to generalize to unseen examples in the test set. In order to effectively modulate the hard instances and encourage the model to avoid overfitting them, a moderate hard example modulation (MHEM) loss is incorporated [25]. Except for the bounding box classification head loss L_{cls}, the losses of the bounding box regression head L_{reg}, semantic head L_{seg}, and mask head L_{mask} are set according to [26].

TABLE III: Class-instance and class-image distributions.

Category	id	N_{ins}	r_{ins}	N_{img}	r_{img}
Flat	1	109460	0.5635	3461	0.9304
Gable	2	57426	0.2956	2965	0.7970
Gambrel	3	5087	0.0262	524	0.1409
Row	4	943	0.0049	477	0.1282
Multiple	5	738	0.0038	176	0.0473
Hipped_v1	6	4320	0.0222	1083	0.2911
Hipped_v2	7	10220	0.0526	1422	0.3823
Mansard	8	510	0.0026	339	0.0911
Pyramid	9	1520	0.0078	631	0.1696
Arched	10	833	0.0043	536	0.1441
Revolved	11	154	0.0008	138	0.0371
Other	12	2872	0.0148	1099	0.2954

Semi-supervised training. Semi-supervised learning (SSL) for visual dense prediction in the downstream tasks has been established to be effective [27]–[29]. Self-training, which is a key branch of SSL, can help a model that was only trained on a small amount of labeled data by expanding it to new situations and training samples. However, because of the complex unlabeled data and model bias, it suffers greatly from noise being present in the initial pseudo-labels. To this end, two core denoising techniques are proposed to explore a safe boundary for using unlabeled data: Image-level denoising by global average filtering (GAF) and object-level denoising by intra-class adaptive thresholding (IAT) are performed in a global-to-instance manner.

- **Global average filtering:** Given the i-th unlabeled image x_i and a trained model F, the initial pseudo-labels of x_i can be formulated as $y_i = F(x_i)$. The maximum total number of image objects is set to $N_{dets} = 2000$. The pseudo-label of the j-th object in y_i consists of a corresponding confident score s_j and a 4-dimensional coordinate representing the position. An image-level confidence S_{x_i} of x_i is computed as

$$S_{x_i} = \frac{1}{N_{dets}} \sum_{j=1}^{N_{dets}} s_j. \tag{2}$$

  A threshold $\delta_{\text{GAF}} = 0.01$ is used to determine whether to drop x_i depending on S_{x_i} being less than δ_{GAF}.
- **Intra-class adaptive thresholding:** For class c, a class-level confidence $S^c_{x_i}$ of x_i can be formulated as

$$S^c_{x_i} = \frac{1}{N^c_{dets}} \sum_{j=1}^{N^c_{dets}} s^c_j, \tag{3}$$

  where N^c_{dets} represents the number of objects belonging to class c. A majority class m_i of x_i can be derived by $m_i = \arg\max(\mathbf{S}^c_{x_i})$. For m_i, by comparison with a threshold $\delta_{\text{IAT}} = 0.2$, the ratio r_i is computed as

$$r_i = \frac{1}{N^{m_i}_{dets}} \sum_{j=1}^{N^{m_i}_{dets}} \mathbb{1}(s^{m_i}_j \geq \delta_{\text{IAT}}), \tag{4}$$

  where $\mathbb{1}(\cdot)$ is the indicator function. Finally, for any other class c ($c \neq m_i$), $\mathbf{S}^c_{x_i}$ is sorted in descending order and only the top r_i objects are kept ensuring reliable pseudo-labels with the same confidence-level. The corresponding thresholds can be selected for each class and each unlabeled image in an adaptive manner.

Two-stage results ensemble pipeline. As shown in Figure 8, for n predictions from n trained models, Weighted Boxes Fusion (WBF [21]) is used to ensemble the corresponding bounding boxes predictions, resulting in more accurate proposals. These are then passed to the best model's mask head to perform segmentation, deriving the final result. Fifteen models trained with different hyper-parameters are assigned equal weights for integration in the test phase of the contest.

B. Results

All experiments are conducted on 8 NVIDIA V100. The pretraining setting is based on [22], except for batch size and learning rate using a linear scale. For supervised and semi-supervised training, the model is optimized by AdamW with a base learning rate of 0.0002, and batch size of 8. The training image base size is set to 1024-1536, and the test multi-scale size is increased by $0, 64, 128, 192, 256$ based on the base size. Random flipping contains horizontal and vertical flipping, both with a probability of 0.25. The training assigner is set to 3 times the default setting. The test-time non-maximum

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited, content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

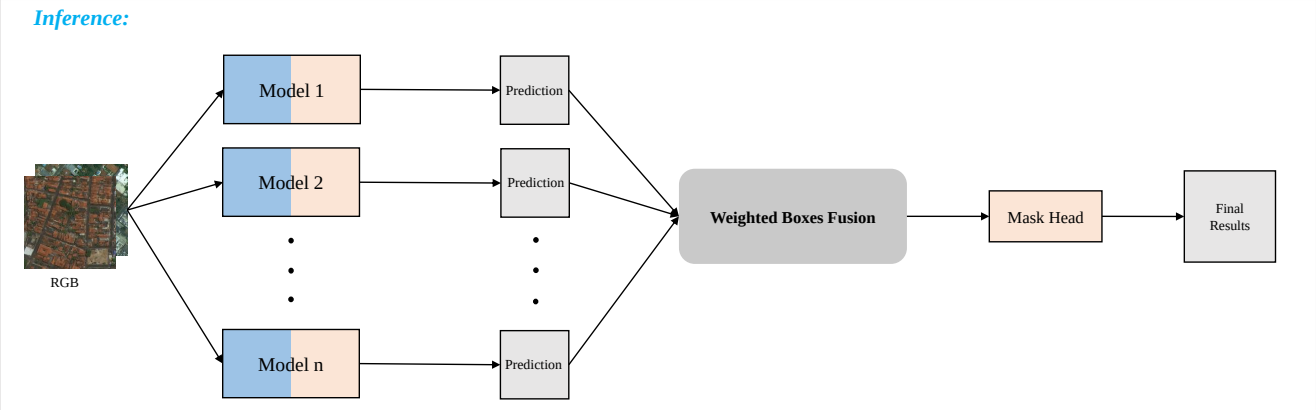

Fig. 8: Two-stage test-time augmentation results ensemble pipeline used by the second place of Track 1.

TABLE IV: Ablation study on various components of Track 1's second place. Pre: pretraining. Sup: supervised training. Semi: semi-supervised training. TTA: test-time augmentation.

Baseline	Pre	Sup	Semi	TTA	AP_{50}(val)	AP_{50}(test)
✓					0.438	-
✓	✓				0.447	-
✓	✓	✓			0.508	-
✓	✓	✓		✓	0.530	-
✓	✓	✓	✓		0.534	-
✓	✓	✓	✓	✓	**0.559**	**0.495**

TABLE V: Ablation study on the three supervised training strategies of Track 1's second place. Pre: pretraining. BCP: balanced copy-paste. CMF: cross-modal fusion. MHEM: moderate hard example modulation loss.

Baseline + Pre	BCP	CMF	MHEM	AP_{50}(val)
✓				0.447
✓	✓			0.489
✓	✓	✓		0.501
✓	✓	✓	✓	**0.508**

suppression uses a score threshold of 0.00001 and a maximum object number of $2,000$.

To investigate each component of the proposed method, extensive experiments are performed on the UBC dataset. Table IV shows the results which led to the first place ($AP_{50} = 0.559$) in the development phase and second place ($AP_{50} = 0.495$) in the test phase. Table V shows the effectiveness of the three supervised training strategies.

Figure 9 visually demonstrates that the method of Track 1's second place can accurately extract and identify various building instances. Moreover, thanks to the cross-modal fusion module, the method is capable of mitigating the negative effects caused by cloud occlusion.

C. Discussion

Team IPIU-XDU addressed several challenges that hinder a good performance of building classification detectors, including long-tailed class distributions, data insufficiency, inter-class similarity, and intra-class differences. Specifically, IPIU-XDU proposed a three-stage training framework consisting of pretraining on the cropped buildings, supervised training with a strong vision transformer adapter with added cross-modal fusion module, and semi-supervised training with two image- and object-level denoising techniques. In addition, a two-stage test-time ensemble pipeline is designed to further boost the performance and generalization of the model. Extensive experiments on the DFC23 data demonstrate the effectiveness of the method.

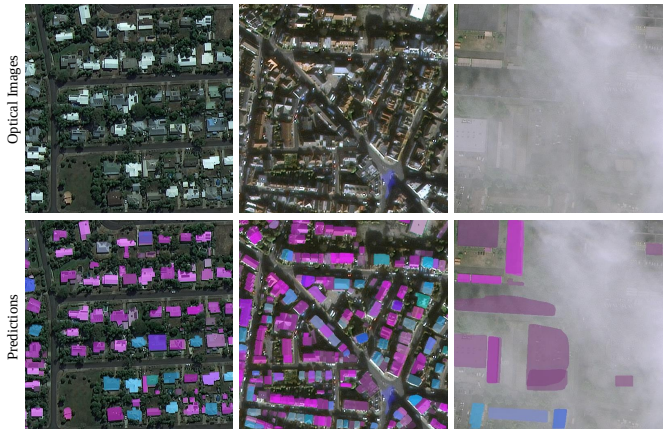

Fig. 9: Qualitative results of the second place of Track 1.

VI. Track 2 - First Place: Team PIESAT-AI

A. Method

As depicted in Figure 10, there are several spatial misalignments between building footprints and stereo-reconstructed normalized Digital Surface Model (nDSM) height labels in the DFC2023 datasets. To address this issue, firstly, a Height-hierarchy Guided Dual-decoder Network (HGDNet) is proposed to estimate building height. In contrast to methods like PopNet [30] and SCENet [31], which rely on high-quality aligned semantic labels for height estimation, HGDNet performs the joint task of height estimation and hierarchical classification by utilizing hierarchical information about

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited, content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

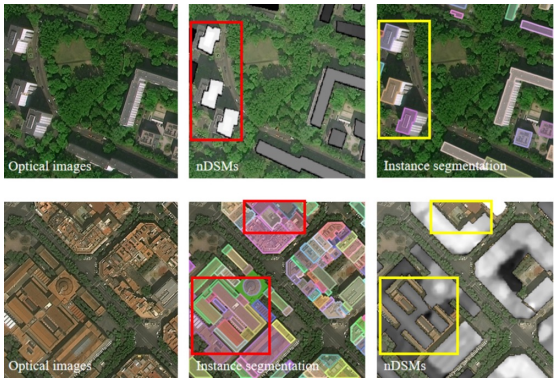

Fig. 10: A Misaligned Example in DFC2023 Dataset

heights. Under the guidance of a synthesized discrete height-hierarchical nDSM, an auxiliary height-hierarchical building extraction branch enhances the height estimation branch with implicit constraints, yielding an accuracy improvement of more than 6% on the DFC2023 Track 2 dataset. Secondly, to achieve a precise building instance segmentation, an additional two-stage cascade architecture is adopted to extract individual building contours. The improved feature extraction and the fusion of outcomes from distinct structural models significantly contribute to the performance increase.

Height Estimation: The architecture of HGDNet is shown in Figure 11. A ConvNeXt V2 base is adopted as the encoder module. The dual decoders consist of the height estimation branch and the height-hierarchical segmentation branch.

- **Height Estimation Branch:** Building heights exhibit an uneven distribution, primarily concentrated within the range of 0-50m and even as low as in the range of 0-10m. To address this issue, a logarithmic transformation is applied to the nDSMs labels to approximate a Normal distribution as

$$DSM_{\text{re-norm}} = \frac{\ln\left(nDSMs\right)}{\max(\ln\left(nDSMs\right))}. \tag{5}$$

  UPerNet [32] is employed as the foundational decoding layer. The up-down pathway and lateral connections work cooperatively to enhance multi-scale features extracted from the backbone. These augmented features are subsequently upsampled to a predetermined scale and then fused. This fusion process generates the final feature for accurate height estimation. Additionally, to accommodate the height regression task, an additional 1-channel convolutional layer is appended to the height branch, followed by a sigmoid layer at the end of the network.
- **Height-hierarchical Segmentation Branch:** An additional height-hierarchical segmentation branch is seamlessly integrated into the network architecture to enhance the precision of height estimation. This branch leverages the same UPerNet decoder employed in the height estimation branch. However, the fused features are fed to a single convolutional layer comprising n channels, where n corresponds to the number of height hierarchies. The proposed methodology avoids the direct employment of instance segmentation annotations featuring only a single class. Instead, the nDSM is partitioned into a discrete hierarchy. This partitioning is achieved through the distribution analysis of nDSMs and the application of clustering algorithms. Subsequently, nDSMs are classified into n classes, thereby establishing height-hierarchy labels. These distinct labels are used to steer the process of building height estimation.
- **The Weighted Loss Function:** The height-hierarchical segmentation branch uses a cross-entropy loss while the height branch uses a smoothed $L1$ loss. Given the disparity in magnitudes between the different loss functions, weighting coefficients α, β are introduced to balance the influence of the two branches. The final form of the loss function is

$$Loss = \alpha \cdot L_{CE} + \beta \cdot L_{smooth_{L1}}. \tag{6}$$

Building Extraction: To extract accurate contours of buildings, various two-stage cascade detection frameworks and backbones are employed as shown in Figure 11.

- **Cascade Framework:** The approach integrates several efficient cascade networks, including Cascade Mask R-CNN [33] and HTC++ [34]. These networks perform target extraction in two stages. The initial stage emphasizes locating the targets to generate proposals and then classifying, segmenting, and refining the proposals.
- **Backbone:** After conducting comparative experiments, ConvNeXt V2-Base [17] and BEiT V2 [22] networks are selected to extract semantic features. These backbone networks extract profound features across distinct scales from the input data, ensuring robust extraction for buildings of various styles.
- **Loss Function:** To enhance the fidelity of bounding box prediction, the model employs the smoothed $L1$ and GIOU loss functions. For building mask learning, the cross-entropy loss function is employed.
- **Model Ensemble:** A modified version of WBF ensembles all results from different models. Firstly, WBF is employed to fuse bounding boxes. Subsequently, the segmentation masks are combined. Finally, the fused masks are cropped using the merged boxes. This strategy facilitates the integration of diverse models and is applicable for multi-scale scenarios.

B. Results

The performance of HGDNet is compared with other methods. Table VI shows the specific quantitative comparison. Notably, while "ConvNextv2+Uper" and the proposed "HGDNet:ConvNeXt V2" differ only in the presence of a height-hierarchical semantic branch, the discernible 6% performance gap underscores the clear effectiveness of the height-hierarchical semantic branch.

The accuracy metrics of each model are shown in Table VII. The optimal performance is attained through a selection of seven models, not necessarily the top seven highest-scoring ones, but rather a set encompassing diverse architectural configurations.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited, content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

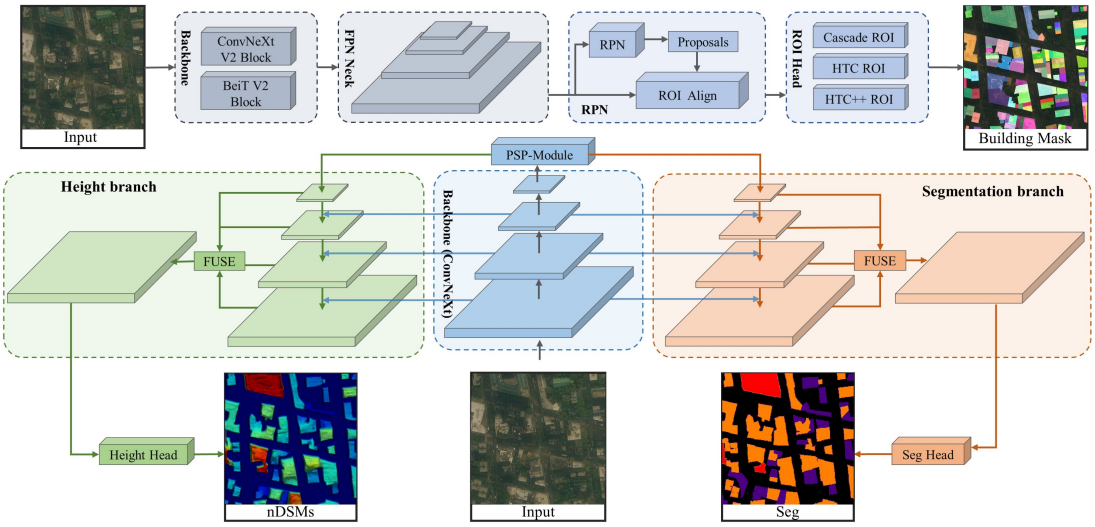

Fig. 11: Overview of the building extraction (top) and height estimation (bottom) framework proposed for the winning solution in Track 2.

TABLE VI: Comparison of Team PIESAT-AI's method and other methods on the validation set.

Method	δ_1	δ_2	δ_3
$Baseline^4$	0.7042	0.7613	0.7922
DSMNet [35]	0.7339	0.7866	0.8110
SCENet-50 [31]	0.7731	0.8214	0.8438
ConvNextv2+Uper	0.7313	0.7822	0.8084
HGDNet:Without seg branch	0.7313	0.7822	0.8084
HGDNet:ResNet-50	0.7824	0.8037	0.8502
HGDNet:ConvNeXt V1 Base	0.7930	0.8371	0.8544
HGDNet:ConvNeXt V2 Base	**0.7966**	**0.8383**	**0.8581**

TABLE VII: Metrics for 7 building extraction models and their fusion result on the test set

Method	id	AP_{50}	mAP
HTC++_BEiT V2	1	0.730	0.402
HTC++_ConvNeXt V2-Base	2	0.710	0.388
HTC_BEiT V2	3	0.722	0.400
Cascade_BEiT V2	4	0.751	0.428
Cascade_ConvNeXt V2-Base	5	0.737	0.422
HTC_ConvNeXt V2-Base	6	0.732	0.420
Cascade_ ConvNeXt V2-Base	7	0.721	0.410
WSF based model ensemble		**0.773**	**0.450**

Figure 12 illustrates the high performance of the proposed method in building extraction and height estimation. For downstream tasks such as 3D reconstruction of buildings, it is essential to further refine the segmentation of building edges and to ensure the smoothness of building heights.

C. Discussion

A straightforward yet highly efficient architecture termed Height-hierarchy Guided Dual decoders Network (HGDNet) is introduced for accurate building height estimation. Using an advanced backbone, the architecture unifies semantic segmentation and height estimation tasks and leverages shared multi-scale features and implicit constraints to enhance the overall performance. Furthermore, the fusion of outcomes from multiple two-stage instance segmentation models boosts the model's capability to discern buildings of various styles. The efficacy of the proposed approach is validated through extensive experimentation on the DFC2023 dataset.

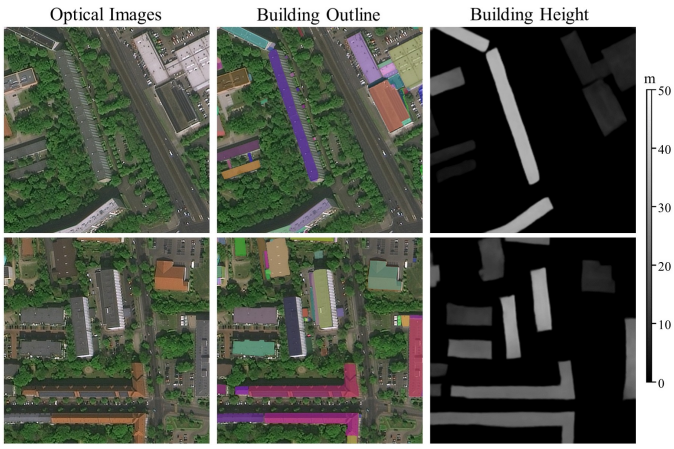

Fig. 12: Qualitative results of the first place of Track 2.

VII. Track 2 - Second Place: Team IPIU-XDU

A. Method

To address the problems of data limitation and foreground-background confusion of the DFC23 datasets, team IPIU-XDU proposes a novel trident cooperation network (TCNet) to perform end-to-end building extraction and height estimation using RGB and SAR data, as shown in Figure 13.

Trident cooperation network. The proposed TCNet contains a cross-model fusion (CMF) module and three task-specific heads [26], [36], [37]. To perform end-to-end building extraction and height estimation, the shared encoded features

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

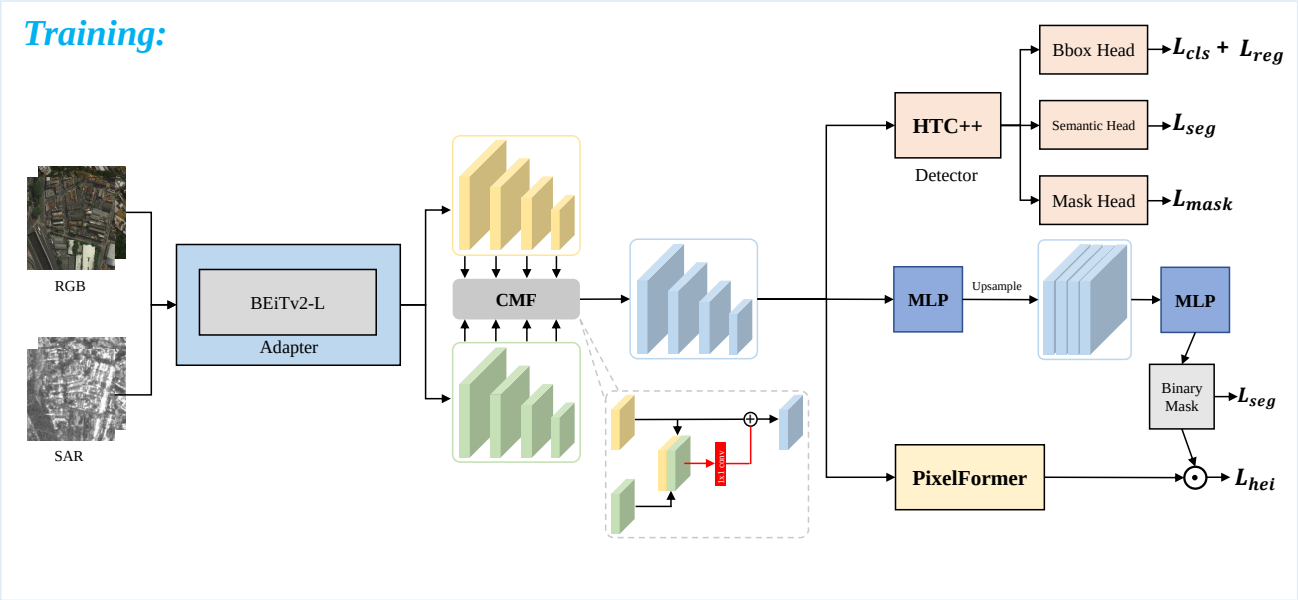

Fig. 13: Illustration of the proposed trident cooperation network (TCNet). The backbone is pre-trained on the DFC23 Track-1 data for initialization. A cross-modal fusion (CMF) module provides more informative shared features for downstream visual tasks. Finally, the trident task-specific heads are leveraged, consisting of a detector [26], a lightweight all-MLP decoder [36], and a pixel query regression head [37]. The binary mask predicted by the all-MLP is additionally used to alleviate fore-background confusion existing in the estimated map.

play an essential role in multi-task learning. Therefore, a simple yet effective CMF module is proposed to derive more representative features. To perform each visual sub-task, three task-specific decoders are introduced:

- **Hybrid task cascade:** HTC [26] is an efficient cascade architecture for instance segmentation, which is mainly used in natural image interpretation. The key components include a cascade structure that detects instances at different scales and refines the results from the previous stages, a semantic segmentation branch that helps to improve the detection accuracy of small objects in crowded scenes, and a feature alignment module that aligns the features of the semantic segmentation and the bounding box regression tasks to generate spatial context to better utilizing the shared features. Based on this strong detection baseline, a robust training pipeline consisting of single-scale training is constructed leveraging random flipping and edge enhancement as used in [28]. The losses of the bounding box classification head L_{cls}, bounding box regression head L_{reg}, semantic head L_{seg}, and mask head L_{mask} are defined according to [26].
- **All-MLP:** Unlike traditional segmentation decoder architectures that use convolutional layers, the All-MLP [36] decoder utilizes a series of pure multi-layer perceptrons to decode the features from the encoder, which avoids the hand-crafted and computationally demanding components. This additional binary segmentation branch is introduced to highlight the foreground predictions and to suppress the wrong background predictions in the height estimation maps. The used loss L_{seg} is the standard cross-entropy loss.
- **PixelFormer:** Monocular height estimation (MHE), similar to monocular depth estimation (MDE), aims to predict pixel-wise height given a single-view image. PixelFormer [37] is used as an efficient MDE network. It proposes a skip attention module to fuse the encoder and decoder features and a bin center predictor module that estimates bin centers adaptively per image using the global information from the coarsest-level feature maps. Since PixelFormer uses a typical encoder-decoder architecture, its original input features are replaced with the shared hierarchical features extracted from RGB and SAR data. Its optimization target is further weighted by the binary mask predicted by the All-MLP segmentation head to force the model to focus on the foreground building objects. The height estimation loss L_{hei} is the scale-invariant loss used in [37].

Finally, the total loss is the uniformly weighted sum of each sub-task loss.

Self-training. Semi-supervised learning (SSL) has been proven to be effective for downstream dense prediction visual tasks [27]–[29]. Self-training, as a main branch of SSL, can achieve better performance and robustness by expanding various scenarios and training samples to assist the model trained only on limited labeled data. Since the dataset of the DFC23 Track 2 is a subset of the dataset of the DFC23 Track 1, the difference set of them can be a suitable unlabeled dataset as it has similar scenarios and data distribution.

Given several trained models, the following two-stage trident cooperation results fusion pipeline is used to derive more reliable initial pseudo-labels. From the pseudo-labels of the detection task, predicted bounding boxes with confidence

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

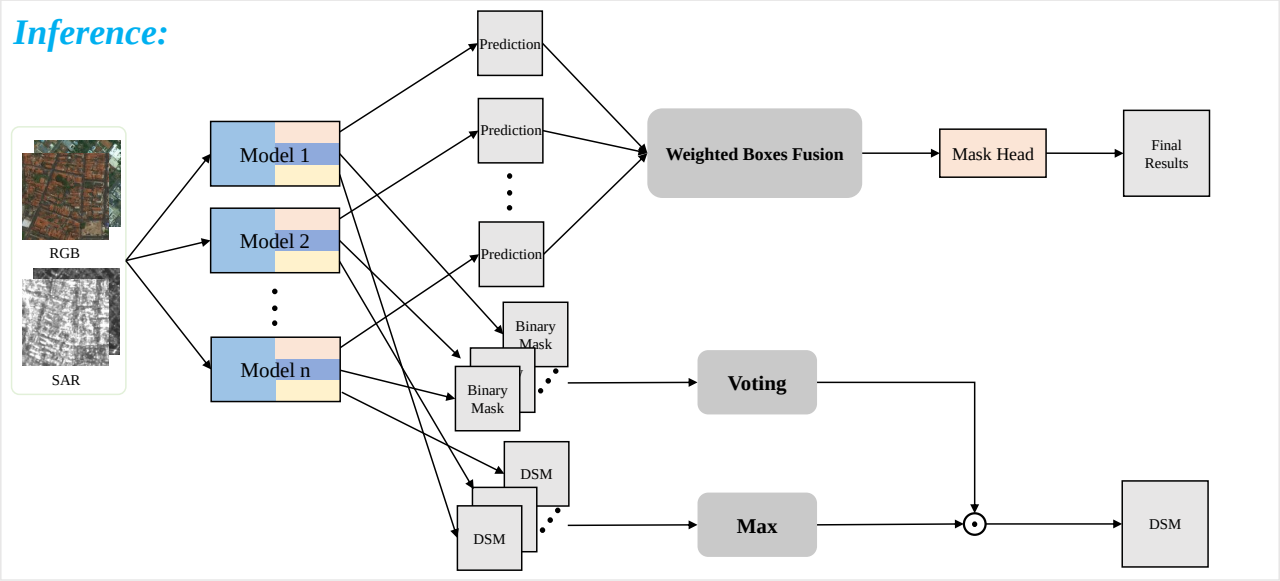

Fig. 14: Illustration of the two-stage trident cooperation results fusion framework proposed by the second place of Track 2 consisting of task-specific ensemble strategies, including weighted boxes fusion [21], hard voting, and pixel-wise maximum for merging the detection results, the binary masks, and the DSM results, respectively. The merged binary masks are employed to element-wise multiply with the height nDSM maps to filter the confusing background.

TABLE VIII: Ablation study on various components. Sup: supervised training. Semi: semi-supervised training. TTA: test-time augmentation.

Baseline	Sup	Semi	TCNet	TTA	AP_{50}(val)	AP_{50}(test)	δ_1(val)	δ_1(test)
✓					0.7230	-	0.7561	-
✓	✓				0.7460	-	0.7708	-
✓	✓	✓			0.7510	-	0.7722	-
✓	✓	✓	✓		0.7570	-	0.7746	-
✓	✓	✓	✓	✓	0.7790	0.7830	0.7766	0.7855

scores lower than $\delta_d = 0.8$ are removed while their initial pseudo-labels are maintained for the binary segmentation and height estimation tasks. Finally, by combining both labeled data and unlabeled data with their corresponding processed pseudo-labels, the proposed TCNet is retrained to boost the performance and generalization capabilities.

Two-stage trident cooperation results fusion pipeline. As shown in Figure 14, for n detection predictions from n trained models, the proposed pipeline uses weighted boxes fusion to only ensemble their bounding boxes predictions, resulting in more accurate proposals, which are then passed to the best model's mask head to perform segmentation, deriving final detection results. For n binary segmentation masks, a hard voting strategy with a mask threshold of 0.5 is used to obtain majority predictions. These are further employed to remove wrong background predictions of the height estimation.

B. Results

Table VIII shows the effectiveness of each component of Team IPIU-XDU's method, indicating that the proposed TCNet can benefit from shared information of multi-task learning.

Figure 15 visually demonstrates that the method of Track 2's second place can accurately extract various building instances

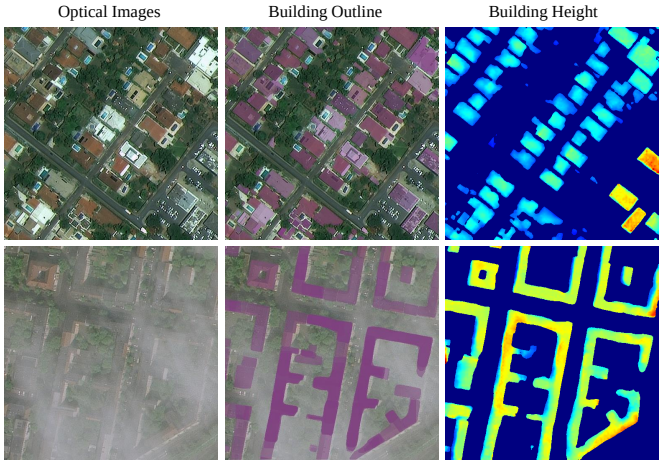

Fig. 15: Qualitative results of the second place of Track 2.

and estimate their corresponding height. Moreover, thanks to the cross-modal fusion module, the method is capable of mitigating the negative effects caused by cloud occlusion.

C. Discussion

Team IPIU-XDU proposed a novel trident cooperation network, TCNet, to perform end-to-end building extraction and height estimation. TCNet has a strong vision transformer adapter backbone that added a cross-modal fusion module, which provides more informative shared hierarchical features for the following trident task-specific decoders. To boost the performance and robustness of TCNet, team IPIU-XDU introduced a simple yet effective self-training. In addition, a two-stage trident cooperation results fusion framework in-

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully e
content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

cluding three task-specific ensemble strategies is designed to further improve the proposed TCNet. Extensive experiments on DFC23 demonstrate the effectiveness of the proposed method.

VIII. Conclusion

The capability to classify fine-grained building types according to different roof typologies promises to offer valuable insights into the urban fabric that are useful for urban planning and management. Furthermore, the capability to estimate building heights provides an additional layer of spatial information, enabling a three-dimensional understanding of urban structures. Elevation data, when integrated with other geospatial information, contributes to the creation of realistic 3D city models and digital twins, fostering a virtual representation that mirrors the physical world.

This article summarized the outcomes of the 2023 IEEE GRSS Data Fusion Contest (DFC), describing the methodological choices made by the teams that ranked first and second in the contest. The strategies adopted by the winning teams provide valuable insights, including advanced deep learning methods as well as best practices to boost the performance of instance segmentation and monocular height estimation techniques.

For Track 1, team PIESAT-AI adopted an instance segmentation framework based upon the Cascade Mask R-CNN [15] architecture enhanced with domain-adapted pretraining and a modernized dual-backbone structure. Team IPIU-XDU resorted to a three-stage learning framework containing pretraining, supervised training, and semi-supervised training. They adopted the BEiTv2-L [22] network for feature extraction, followed by fine-tuning with a vision transformer adapter [23] used to enhance the representations of the model. Both teams paid attention to the imbalance of class distributions adopting resampling and copy-paste data balancing/augmenting strategies. Interestingly, both teams improved the model performance adopting an ensemble strategy to combine multiple models.

For Track 2, team PIESAT-AI proposed a Height-hierarchy Guided Dual-decoder Network (HGDNet) to estimate building heights. HGDNet performs a joint task of height estimation and hierarchical classification by utilizing hierarchical height information. Team IPIU-XDU proposed a novel trident cooperation network, TCNet, to perform end-to-end building extraction and height estimation. In addition, a fusion framework including three task-specific ensemble strategies is designed to further improve the results.

All in all, DFC23 has fostered the development of advanced image analysis and data fusion techniques to tackle complex tasks such as building instance segmentation with fine-grained roof type classification and height estimation. The methods adopted by the two winning teams reveal many relevant contributions in terms of deep learning architectures, learning strategies, class balancing techniques, ensemble modelling and several tricks and tweaks to boost the model performance.

The datasets remain open to the scientific community to further advance the development of image analysis and data fusion methods.

Acknowledgements

The IADF TC chairs would like to thank the Aerospace Information Research Institute under the Chinese Academy of Sciences, the Universität der Bundeswehr München, and GEOVIS Earth Technology Co., Ltd. for providing the data and the IEEE GRSS for continuously supporting the annual Data Fusion Contest through funding and resources.

References

[1] C. Persello, R. Hänsch, G. Vivone, K. Chen, Z. Yan, D. Tang, H. Huang, M. Schmitt, and X. Sun, "2023 ieee grss data fusion contest: Large-scale fine-grained building classification for semantic urban reconstruction [technical committees]," *IEEE Geoscience and Remote Sensing Magazine*, vol. 11, no. 1, pp. 94–97, 2023.

[2] N. Yokoya, P. Ghamisi, J. Xia, S. Sukhanov, R. Heremans, I. Tankoyeu, B. Bechtel, B. Le Saux, G. Moser, and D. Tuia, "Open data for global multimodal land use classification: Outcome of the 2017 ieee grss data fusion contest," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 11, no. 5, pp. 1363–1377, May 2018.

[3] Y. Xu, B. Du, L. Zhang, D. Cerra, M. Pato, E. Carmona, S. Prasad, N. Yokoya, R. Hänsch, and B. Le Saux, "Advanced multi-sensor optical remote sensing for urban land use and land cover classification: Outcome of the 2018 ieee grss data fusion contest," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 6, pp. 1709–1724, June 2019.

[4] S. Kunwar, H. Chen, M. Lin, H. Zhang, P. Dangelo, D. Cerra, S. M. Azimi, M. Brown, G. Hager, N. Yokoya, R. Hansch, and B. Le Saux, "Large-scale semantic 3d reconstruction: Outcome of the 2019 ieee grss data fusion contest - part a," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, pp. 1–1, 2020.

[5] Y. Lian, T. Feng, J. Zhou, M. Jia, A. Li, Z. Wu, L. Jiao, M. Brown, G. Hager, N. Yokoya, R. Hansch, and B. Le Saux, "Large-scale semantic 3d reconstruction: Outcome of the 2019 ieee grss data fusion contest - part b," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, pp. 1–1, 2020.

[6] C. Robinson, K. Malkin, N. Jojic, H. Chen, R. Qin, C. Xiao, M. Schmitt, P. Ghamisi, R. Hänsch, and N. Yokoya, "Global land-cover mapping with weak supervision: Outcome of the 2020 ieee grss data fusion contest," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 3185–3199, 2021.

[7] Y. Ma, Y. Li, K. Feng, Y. Xia, Q. Huang, H. Zhang, C. Prieur, G. Licciardi, H. Malha, J. Chanussot, P. Ghamisi, R. Hänsch, and N. Yokoya, "The outcome of the 2021 ieee grss data fusion contest - track dse: Detection of settlements without electricity," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 12 375–12 385, 2021.

[8] G. Liu, B. Peng, T. Liu, P. Zhang, M. Yuan, C. Lu, N. Cao, S. Zhang, S. Huang, and T. Wang, "Fine-grained building roof instance segmentation based on domain adapted pretraining and composite dual-backbone," in *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023.

[9] X. Lu, L. Jiao, Q. Liu, L. Li, F. Liu, X. Liu, and Y. Yang, "A strong vision transformer adapter with adaptive thresholding for fine-grained building classification," in *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023.

[10] J. Hu, Z. Huang, F. Shen, D. He, and Q. Xian, "A bag of tricks for fine-grained roof extraction," in *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023.

[11] C. Lu, N. Cao, P. Zhang, T. Liu, B. Peng, G. Liu, M. Yuan, S. Zhang, S. Huang, and T. Wang, "A height-hierarchy guided dualdecoder network for single view building extraction and height estimation," in *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023.

[12] X. Lu, L. Jiao, Q. Liu, L. Li, F. Liu, X. Liu, and Y. Yang, "Trident cooperation network for building extraction and height estimation," in *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023.

[13] Y. Guo and Z. Wang, "Height estimation based on semantic segmentation," in *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. This is the author's version which has not been fully e
content may change prior to final publication. Citation information: DOI 10.1109/JSTARS.2024.3403201

[14] J. Hu, Z. Huang, F. Shen, D. He, and Q. Xian, "A robust method for roof extraction and height estimation," in *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023.
[15] Z. Cai and N. Vasconcelos, "Cascade r-cnn: Delving into high quality object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 6154–6162.
[16] H. Zhang, Y. Wang, F. Dayoub, and N. Sünderhauf, "Swa object detection," *arXiv preprint arXiv:2012.12645*, 2020.
[17] S. Woo, S. Debnath, R. Hu, X. Chen, Z. Liu, I. S. Kweon, and S. Xie, "Convnext v2: Co-designing and scaling convnets with masked autoencoders," *arXiv preprint arXiv:2301.00808*, 2023.
[18] T. Liang, X. Chu, Y. Liu, Y. Wang, Z. Tang, W. Chu, J. Chen, and H. Ling, "Cbnet: A composite backbone network architecture for object detection," *IEEE Transactions on Image Processing*, vol. 31, pp. 6893–6906, 2022.
[19] Y. Zhang, K. Chen, X. Wang, H. Li, and D. Lin, "Simple copy-paste is a strong data augmentation method for instance segmentation," *arXiv preprint arXiv:2006.09845*, 2020.
[20] J. Wang, W. Zhang, Y. Zang, Y. C. andJiangmiao Pang, T. Gong, K. Chen, Z. Liu, C. C. Loy, and D. Lin, "Seesaw loss for long-tailed instance segmentation," *arXiv preprint arXiv:2008.10032*, 2021.
[21] R. Solovyev, W. Wang, and T. Gabruseva, "Weighted boxes fusion: Ensembling boxes from different object detection models," *Image and Vision Computing*, pp. 1–6, 2021.
[22] Z. Peng, L. Dong, H. Bao, Q. Ye, and F. Wei, "Beit v2: Masked image modeling with vector-quantized visual tokenizers," *arXiv preprint arXiv:2208.06366*, 2022.
[23] Z. Chen, Y. Duan, W. Wang, J. He, T. Lu, J. Dai, and Y. Qiao, "Vision transformer adapter for dense predictions," *arXiv preprint arXiv:2205.08534*, 2022.
[24] X. Huang, L. Ren, C. Liu, Y. Wang, H. Yu, M. Schmitt, R. Hänsch, X. Sun, H. Huang, and H. Mayer, "Urban building classification (ubc)-a dataset for individual building detection and classification from satellite imagery," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2022, pp. 1413–1421.
[25] Y. Liang, L. Zhu, X. Wang, and Y. Yang, "Penalizing the hard example but not too much: A strong baseline for fine-grained visual classification," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
[26] K. Chen, J. Pang, J. Wang, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Shi, W. Ouyang *et al.*, "Hybrid task cascade for instance segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4974–4983.
[27] X. Lu, G. Cao, and T. Gou, "Semi-supervised landcover classification with adaptive pixel-rebalancing self-training," in *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2022, pp. 4611–4614.
[28] X. Lu, L. Jiao, F. Liu, S. Yang, X. Liu, Z. Feng, L. Li, and P. Chen, "Simple and efficient: A semisupervised learning framework for remote sensing image semantic segmentation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–16, 2022.
[29] X. Lu, L. Jiao, L. Li, F. Liu, X. Liu, S. Yang, Z. Feng, and P. Chen, "Weak-to-strong consistency learning for semisupervised image segmentation," *IEEE Transactions on Geoscience and Remote Sensing*, 2023.
[30] Z. Zheng, Y. Zhong, and J. Wang, "Pop-net: Encoder-dual decoder for semantic segmentation and single-view height estimation," in *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2019, pp. 4963–4966.
[31] S. Xing, Q. Dong, and Z. Hu, "Sce-net: Self-and cross-enhancement network for single-view height estimation and semantic segmentation," in *Remote Sensing*, vol. 14. MDPI, 2022, p. 2252.
[32] T. Xiao, Y. Liu, B. Zhou, and et al., "Unified perceptual parsing for scene understanding," in *ECCV*, 2018, pp. 418–434.
[33] Z. Cai and N. Vasconcelos, "Cascade r-cnn: high quality object detection and instance segmentation," *PAMI*, vol. 43, no. 5, pp. 1483–1498, 2019.
[34] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10 012–10 022.
[35] M. Elhousni, Z. Zhang, and X. Huang, "Height prediction and refinement from aerial images with semantic and geometric guidance," *IEEE Access*, vol. 9, pp. 145 638–145 647, 2021.
[36] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo, "Segformer: Simple and efficient design for semantic segmentation with transformers," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12 077–12 090, 2021.
[37] A. Agarwal and C. Arora, "Attention attention everywhere: Monocular depth prediction with skip attention," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 5861–5870.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/