

Accurate Polygonal Mapping of Buildings in Satellite Imagery

Bowen Xu^{a,b,1}, Jiakun Xu^{b,1}, Nan Xue^{b,*} and Gui-Song Xia^{a,b,*}

^aState Key Lab. of LIESMARS, Wuhan University, Wuhan, 430079, China

^bSchool of Computer Science, Wuhan University, Wuhan, 430072, China

ARTICLE INFO

Keywords:

building extraction
building vectorization
high-resolution satellite imagery

ABSTRACT

This paper studies the problem of polygonal mapping of buildings by tackling the issue of mask reversibility that leads to a notable performance gap between the predicted masks and polygons from the learning-based methods. We addressed such an issue by exploiting the hierarchical supervision (of bottom-level vertices, mid-level line segments and the high-level regional masks) and proposed a novel interaction mechanism of feature embedding sourced from different levels of supervision signals to obtain reversible building masks for polygonal mapping of buildings. As a result, we show that the learned reversible building masks take all the merits of the advances of deep convolutional neural networks for high-performing polygonal mapping of buildings. In the experiments, we evaluated our method on the two public benchmarks of AICrowd and Inria. On the AICrowd dataset, our proposed method obtains unanimous improvements on the metrics of AP, AP^{boundary} and PoLiS. For the Inria dataset, our proposed method also obtains very competitive results on the metrics of IoU and Accuracy. The models and source code are available at <https://github.com/SarahwXU/HiSup>.

1. Introduction

Polygonal mapping of buildings, aiming at precisely extracting the footprint of buildings from high-resolution satellite imagery in the form of polygons, is a core and dynamic problem in photogrammetric computer vision and remote sensing (Sohn and Dowman, 2001; Zorzi et al., 2022), and plays important roles in the Geographic Information Systems (GIS) for making up the basic feature classes (Mayer, 1999).

As in orthorectified satellite imagery, the footprints and roof outlines of buildings coincide under most circumstances, the mapping of buildings could be roughly equivalent to extracting building polygons (Wang et al., 2022). Thus, the problem of polygonal mapping of buildings in satellite imagery has usually been formulated in a multi-step paradigm (Girard et al., 2020; Li et al., 2021), by (1) computing the rasterized binary masks of buildings from the input images, (2) vectorizing the binary masks into polygons with certain heuristic post-processing schemes, and (3) optionally simplifying the initial polygons to reduce the redundancy of vertices. In this paradigm, building segmentation, as a key sub-problem, has recently been significantly advanced by deep neural networks (DNNs) (Mnih, 2013; Yuan, 2018). Thereafter, the quality of polygonal mapping of buildings can be improved (Wei et al., 2020) by leveraging the off-the-shelf ad-hoc algorithms such as Marching Squares (Lorensen and Cline, 1987) for polygon generation and the Douglas-Peucker (Douglas and Peucker, 1973) for the shape simplification.

Although the performance of polygonal mapping of buildings has been significantly improved by achieving higher quality segmentation masks, it is worth noticing that the task of polygonal mapping of buildings largely depends on the quality of learned masks. Furthermore, when we get a

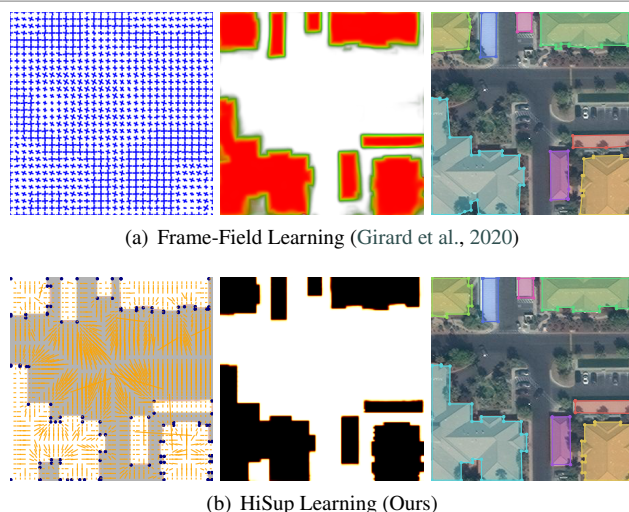


Figure 1: An illustrative comparison between the prior art, Frame-Field Learning (Girard et al., 2020) and our proposed *HiSup* Learning for polygonal mapping of buildings. In contrast to the Frame-Field that is learned as a complementary to handle the uncertainty of the predicted building masks, our proposed method learns the shape-aware feature embedding by end-to-end learning with *Hierarchical Supervisions*, thus minimized the gap between the predictions of masks (in the middle column) and polygons (in the right column).

closer look at the polygonal annotation of buildings, we will find that the learned masks remained an issue of reversibility to the polygonal representation, saying *mask reversibility*. More precisely, with a polygonal annotation of a building, the mask representation is accessible by drawing pixels into an image domain and can be reversed to the polygons as long as the resolution of the image domain is in a reasonable range. However, when we use the mask representation of the polygonal annotations as the supervision signals for the optimization of a convolutional neural network, it would be extremely challenging to obtain such reversible masks.

*Corresponding authors.

¹Equal contributions.

To handle this *mask-reversibility* challenge, the most recent studies either exploited the shape attributes of buildings by Frame-Field (Girard et al., 2020) or improved the polygonal mapping of buildings schemes (Li et al., 2020) with the off-the-shelf mask predictions, pushed the accuracy of predictions to a higher stage. Despite that, it is also worth noticing that there is a notably large performance gap on the AICrowd benchmark (Mohanty et al., 2020) between the polygon and mask predictions. The winning solution (Czakov et al., 2018) achieves the mask average precision of 79.1%, whereas the best-performing model of Frame-Field (Girard et al., 2020) who was applied to the masks achieves the average precision of 67.0%. As the mask predicted in Fig. 1(a) is ambiguous in the shape boundaries, the finally predicted polygons of Frame-Field (Girard et al., 2020) is still far away from that of human annotations. One possible reason leading to such a large performance gap is that the used supervision signals pay more attention to the high-level regional information rather than the detailed geometric shape of buildings.

In this paper, we make efforts to close the gap between the polygonal representations and the mask ones of buildings with deep convolutional neural networks, inspired by an observation:

*The shape composition of polygonal buildings is made up of **points** (or **vertices**) at the bottom, uses **line segments** to associate the points as the middle-level information, and finally forms the **regional instances** of buildings at high-level semantics.*

Therefore, we are motivated to build the full-level supervision to train the neural networks for the sake of shape correctness of the predicted masks. As long as the masks predicted by convolutional neural networks attain the best information about the polygonal shapes, only a simple post-processing scheme is required to achieve the goal of accurate polygonal mapping of buildings.

Towards this goal, we are inspired by the dual representation of line segments, named *regional attraction fields* (Xue et al., 2019, 2021) at the middle-level description as the bridge to get the connection between the unary information of points and complex geometry of regional shapes. Different from the representation of edge pixels, the dual regional representation of line segments demonstrates advantages in two aspects:

- It intrinsically connects to the vertices of polygons as they can be viewed as the endpoint of line segments.
- The non-local support region of each line segment is corresponding to the region-level information of a part of the building mask.

Although the properties of regional attraction fields (Xue et al., 2019, 2021) are appealing and promising for polygonal mapping of buildings, several challenging open problems remain in computing the higher-ordered geometric shapes

from the mid-level line segments (or its dual representation). One may come up with solutions by directly leveraging the detected line segments (or wireframes) as the basic computing clues. However, this will lead to an NP-Hard problem of finding all the closed loops from an undirected graph. Meanwhile, the detection accuracy of line segments and wireframes will also affect the final results. In the most recent, PolyWorld (Zorzi et al., 2022) attempted to solve the problem of polygonal mapping of buildings along this line by using Sinkhorn (Cuturi, 2013) to approximate the step of finding circulars in a graph. However, it still remained the issue of flexibility when handling the complicated buildings that have holes.

In the study of this paper, we present a hierarchical supervision (*HiSup*) learning scheme to address the *mask reversibility* issue and show that the geometric information of polygons should be used inductively in the embedding space of the multi-head predictions by end-to-end learning. Specifically, we design a novel information aggregation mechanism to enhance the feature embedding for both the bottom-level vertex prediction and the high-level mask prediction. By explicitly encoding and aggregating the boundary information into the high-level mask head, we successfully reduced the gap between the mask predictions and the expected polygonal ones. During the inference, we design a simple greedy scheme to polygonize and simplify the masks by using the predicted vertices. Our proposed method, *i.e.*, *HiSup*, achieves the state-of-the-art of polygonal mapping of buildings.

In the experiments, we evaluate our method on the challenging benchmarks of AICrowd (Mohanty et al., 2020) and Inria (Maggiori et al., 2017a). Compared to the current state-of-the-art approaches Frame-Field (Girard et al., 2020) and the method of (Li et al., 2021), the building polygons computed by our *HiSup* obtain a significant improvement of AP (Lin et al., 2014) by 12.4 points and 5.6 points, respectively. In terms of the more challenging metric, AP^{boundary} (Cheng et al., 2021), we push the state-of-the-art performance of 50.0% obtained by PolyWorld (Zorzi et al., 2022) to 66.5%. On the Inria dataset (Maggiori et al., 2017a), our proposed *HiSup* also obtains a considerable improvement over the Frame-Field (Girard et al., 2020). A systematic ablation study is performed to further justify the proposed method.

The main contributions of this paper are summarized as:

1. We address the *mask reversibility* issue of polygonal mapping of buildings by proposing the *HiSup*, which takes the hierarchical supervision signals of vertices, boundaries, and masks to guide the convolutional neural network to learn in a shape-aware fashion.
2. We present a key component of aggregating the embedding of attraction fields into the learning of high-level masks and the bottom-level vertices.
3. We set several new state-of-the-art performances on the challenging AICrowd benchmark in terms of AP and AP^{boundary} by learning semantically-accurate and geometrically-precise masks with our *HiSup*.

The rest of the paper is organized as follows. Section 2 reviews related works. Section 3 describes the hierarchical representation. Section 4 introduces our polygonal mapping of buildings method. Section 5 reports experimental results and some discussions. Finally, conclusions are drawn in section 6.

2. Related work

The problem of (polygonal) mapping of buildings is a longstanding problem, which was extensively studied by the means of using hand-crafted features such as textures, geometry, spectrum, and shadows with a vast body of literature (Pesaresi and Gerhardinger (2011); Xia et al. (2019); Ngo et al. (2016); Senaras et al. (2013); Manno-Kovacs and Sziranyi (2015), to cite a few). These kinds of methods could get good results on certain buildings that satisfy the predefined rules, but they suffer from the degraded performance when buildings have inconsistent appearances in complex environments. Therefore, the research focusing on polygonal mapping of buildings has been gradually moved to the deep learning solutions to handle the more challenging real-world data. In this section, we mainly review the recent advancements for learning-based polygonal mapping of buildings.

2.1. Learning Building Masks with ConvNets

The mask representation is widely used for building extraction in two mainstream pipelines: the building map segmentation (Yang et al., 2018) and the building instance segmentation (Cheng et al., 2019). The former one formulates the problem of building extraction as a pixel-wised binary classification task and the latter one leverages the advances of detected bounding boxes of buildings to directly extract building mask for each instance. Those two pipelines correspond to the semantic segmentation (Long et al., 2015) and instance segmentation (He et al., 2017) respectively in natural images, whereas, they have the same goal for the task of building extraction in satellite images. As the orthorectified satellite imagery eliminates the occlusions between building instances, it is straightforward to delineate the foreground pixels from the instance-agnostic segmentation map into multiple instances by the connected components labeling schemes (Iglovikov et al., 2018; Czakon et al., 2018).

In general, learning a ConvNet is promising to obtain high-quality building masks from satellite images (Saito et al., 2016; Alshehhi et al., 2017; Maggiori et al., 2017b; Yuan, 2018). However, focusing mainly on the high-level semantics makes these approaches perform inaccurate in terms of the building shape correctness. There are some works that focus on alleviating the boundary uncertainties of building maps including multi-scale feature fusion (Maggiori et al., 2017c; Ji et al., 2019), and geometrically-aware learning (Li et al., 2022) with the assistance of attraction fields (Xue et al., 2019, 2021) or edge constraint (Liu et al., 2021a). However, their results are not satisfactory for the accurate polygonal mapping of buildings.

In our method, we concentrate on the geometric correctness of building masks learned by ConvNets. Compared to those methods, our method learns a unified embedding space by exploiting the different level of building representations, which maintains the accuracy of the segmentation masks in both aspects of high-level semantics and the geometric correctness.

2.2. Polygonal Mapping of Buildings

Polygonal mapping of buildings refers to getting vectorized building instances, which is the most compact and meaningful representation format beyond others. Only ordered building vertices are kept to describe individual building instances. Traditionally, building vectorization tends to be a separate process after getting the raster map of building instance segmentation. One of the basic algorithms is the Douglas-Peucker simplification serial techniques (Douglas and Peucker, 1973). Those simplification algorithms are usually rough and have difficulty in identifying right building vertices. Extra refine strategies based on the empirical building shape are proposed in (Wei et al., 2020). Lately, based on both the image and its semantic probability map, a polygonal partition approach shows promising results for building vectorization (Li et al., 2020). However, these multi-staged methods mostly demand separate optimization processes. An overall optimized result is hard to get.

Currently, automatic polygonal mapping of buildings from remote sensing imagery has received considerable attention. Some methods tend to combine the vectorization process with segmentation through the multi-task learning. By learning building segmentation map aligned with proposed frame field vectors, Frame-Field (Girard et al., 2020) gets building polygons through corner-aware contour simplification. The invented frame field map is a 4D vector calculated according to the direction of building edges. In comparison with the frame field (Girard et al., 2020), the attraction fields (Xue et al., 2019, 2021) adopted in our method have more contextual and structural information. Li et al. (2021) generate building masks, corner locations, and edge orientations simultaneously. Then the corner location map and edge orientation map are combined together to help with identifying vertices from pixels belonging to building boundaries. Beyond that, an additional refinement process is adopted to further adjust building vertices on a finer scale. The refinement of building vertices can also be embedded to the whole model and be trained jointly (Chen et al., 2020).

The above methods all get building polygons through simplifying the predicted building contours. Instead, we use detected vertices to rebuild the whole building polygons.

There are other methods that extract building polygons on the basis of object detection formula. Buildings are first marked with bounding boxes. Then, for every building, corners are predicted one by one through recurrent neural networks (RNNs) with convolutional long-short term memory (ConvLSTM) (Li et al., 2019). This framework is further improved by (Zhao et al., 2021) with modified modules including global context block (GCB), boundary refinement




Figure 2: An illustration for the hierarchical representation of vectorized building maps. For the expected building polygons in (a), we learn the high-level segmentation masks (b), the mid-level attraction fields of line segments (c) and the bottom-level vertices (d) with a single convolutional neural network, which guides the neural network learns a unified embedding for accurate polygonal mapping of buildings.

block (BRB) and gated recurrent units (GRU). Following the same detection formula, Liu et al. (2021b) proposed AFM-RNN which introduced the attraction fields (Xue et al., 2019) as embedded features to enhance the building corner detection. In our method, we take advantage of the attraction fields as the bridge to connect both the building masks and vertices implicitly.

Lately, Zorzi et al. (2022) came up with a new proposal termed PolyWorld by considering all building polygons within one image as a whole undirected graph. They predict building polygons by detecting vertices and solving their adjacent connections. The obtained building polygons are very neat in visual. Failed cases appear where buildings have complex structures like inner yards.

3. A Hierarchical Representation of Buildings

Given a satellite image I defined on the image lattice $\Lambda = \mathbb{Z}_W \times \mathbb{Z}_H$ with the resolution of $H \times W$, the n building polygons presented in I are the regions $\Omega_1, \dots, \Omega_n \subset \mathcal{D} \subset \mathbb{R}^2$, where \mathcal{D} is a continuation of the set Λ in the 2D real space. As the buildings are man-made objects, the boundary $\partial\Omega_i$ can be described by a polygon that consists of an array of m_i vertices, denoted by $\partial\Omega_i^{poly} = [\mathbf{x}_i^1, \mathbf{x}_i^2, \dots, \mathbf{x}_i^{m_i}, \mathbf{x}_i^{m_i+1}]$, where $\mathbf{x}_i^{m_i+1} = \mathbf{x}_i^1$. The order of the vertices in $\partial\Omega_i^{poly}$ indicates the adjacent relationship among the vertices.

3.1. High-Level Regional Masks of Buildings

The mask representation is extensively used for the extraction of buildings as a high-level semantic segmentation task. In this sense, the n regions of building polygons are represented by a segmentation map S defined on the image lattice Λ by

$$S(\mathbf{x}) = \begin{cases} 1, & \text{if } \mathbf{x} \in \Omega_i, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Benefiting from the bird-eye view observation of satellites, the n regions can be observed without occlusion between instances, therefore a high-performing convolutional neural network for image semantic segmentation can obtain the instance-level segmentation results in most cases. However, due to the concentration of high-level semantics,

the predicted segmentation maps usually pay more attention to the interior regions of Ω_i rather than their boundaries $\partial\Omega_i$. As a result, it is very challenging to compute accurate polygons from the predicted masks.

3.2. Bottom-Level and Mid-Level Geometries

As one main goal of building segmentation is for the vectorized representation, the polygonal representation about the boundary information $\partial\Omega$ is of great importance for the end task. Here, we briefly discuss both the bottom-level vertices and the mid-level line segments about the boundary information.

When we focus on the polygonal boundary of the buildings, the polygon $\partial\Omega_i^{poly}$ for the i -th building mask contains two kinds of information: the bottom-level vertices $\{\mathbf{x}_i^1, \dots, \mathbf{x}_i^n\}$ and their adjacent relationship represented by mid-level line segments $\{\mathbf{l}_i^1, \dots, \mathbf{l}_i^n\}$. With the bottom-level and mid-level geometries of buildings, there would be another challenging problem of grouping primitives into instances. It was recently approached under the view of graph computation with Sinkhorn approximation (Cuturi, 2013) to avoid the NP-Hard problem of finding circulars in an undirected graph (Zorzi et al., 2022).

In this paper, we believe that although the mask learning is a high-level recognition task, as the polygonal buildings can be lossily converted into mask representation, there should be a unified embedding space to learn and decode consistent multi-level outputs. To approach this, we use the bottom-level vertices and mid-level line segments as the complementary information with the learning of masks, which guides the convolutional neural networks to focus on the bottom-level and mid-level information for the shape correctness of the masks.

Convex and Concave Building Vertices. Inspired by Li et al. (2021), we divide the building vertices (i.e., junctions) according to their convexity. In detail, for the vertices in $\partial\Omega_i^{poly}$, vertices belonging to the minimal convex hull of $\partial\Omega_i^{poly}$ are regarded as the convex vertices, denoted in V_{convex}^i . For the vertices that belong to the set $\partial\Omega_i^{poly} / V_{\text{convex}}^i$, they are denoted by V_{concave}^i . By stacking all the vertices

across instances, we have two sets V_{convex} and V_{concave} as the bottom-level representation of buildings.

Regional Encoding of Line Segments. As discussed before that the line segments are actually the bridge between the bottom-level vertices and the high-level masks, we are interested in the representation of line segments that are well-suited for polygonal mapping of buildings. The most natural choice of the line segments might be the edge map (or line heatmap). However, we found that the learning of edge maps can only improve the mask reversibility marginally. Therefore, we are motivated to use the dual representation, the Attraction Field Representation Map (AFM) (Xue et al., 2019, 2021) for our endtask. Denoted by $L = \{\mathbf{l}_i = (\mathbf{x}_i^s, \mathbf{x}_i^e)\}_{i=1}^m$ (\mathbf{x}_i^s and \mathbf{x}_i^e are the two endpoints), the line segments of the polygonal instances $\{\partial\Omega_i\}$ partition the image lattice Λ into m regions $\mathcal{R} = \{R_i\}_{i=1}^m$ by,

$$R_i = \{\mathbf{x} \in \Lambda | d(\mathbf{x}, \mathbf{l}_i) < d(\mathbf{x}, \mathbf{l}_j), \forall j \neq i, \mathbf{l}_j \in L\}, \quad (2)$$

$$d(\mathbf{x}, \mathbf{l}_i) = \min_{t \in [0,1]} \|\mathbf{x}_i^s + t \cdot (\mathbf{x}_i^e - \mathbf{x}_i^s) - \mathbf{x}\|_2^2. \quad (3)$$

With the region partition \mathcal{R} , the attraction field map $\mathcal{A} \in \mathbb{R}^{2 \times H \times W}$ between any pixel to its close-in line segment \mathbf{l}_j is computed by

$$\mathcal{A}(\mathbf{x}) = \mathbf{x}' - \mathbf{x}, \quad (4)$$

where the point \mathbf{x}' is the projection of \mathbf{x} on \mathbf{l}_j . As the attraction field map encodes the line segments according to \mathcal{R} , the high-level masks are implicitly encoded.

4. Learning Reversible Building Masks

In this section, we present a novel approach that learns building masks in a unified embedding space under the supervision of masks, line segments and the vertices. The unified embedding space takes the merits of different-level of information for the buildings, thereafter enables the convolutional neural networks to learn both the instance-level information and the geometric details with the best reversibility for polygonal mapping of buildings. Then, we present a mask attraction scheme that polygonizes the reversible masks into the expected polygons by attracting the sparse vertices to the boundary of masks.

4.1. Network Architecture Overview

As illustrated in Fig. 3, given an image $I \in \mathbb{R}^{3 \times H \times W}$, we use a backbone network (e.g., HRNets (Wang et al., 2021) or UResNets (Unet (Ronneberger et al., 2015) with ResNet (He et al., 2016) encoder)) to extract the shared feature map $\mathbf{F}_b \in \mathbb{R}^{3 \times H_s \times W_s}$, where $H_s = \lfloor H/s \rfloor$ and $W_s = \lfloor W/s \rfloor$ are the size of \mathbf{F}_b with the down-sampling factor s . Because our method is learned from the bottom-level to high-level information of buildings, three convolution layers (including the BatchNorm (Ioffe and Szegedy, 2015) and ReLU (Nair and Hinton, 2010)) with 256 output channels are used to transform the backbone feature \mathbf{F}_b into the embedding features \mathbf{F}_{ver} of vertices, \mathbf{F}_{afm} of the attraction

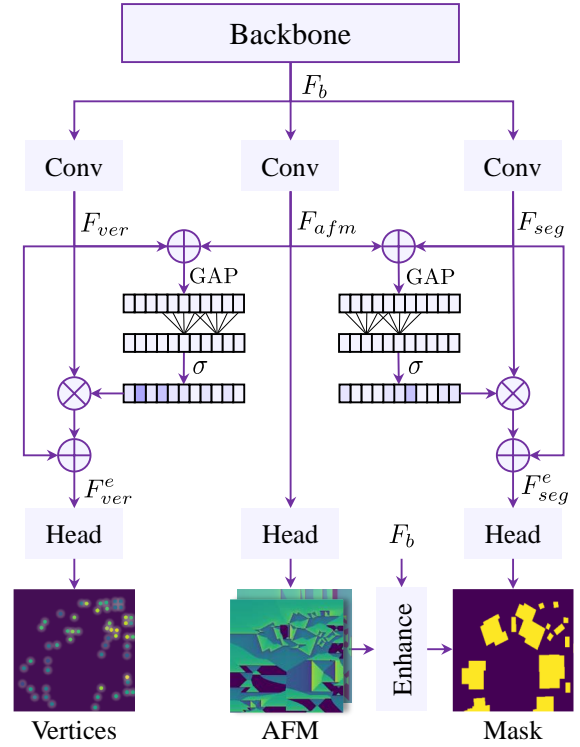


Figure 3: The overall training process. The “Conv” module denotes three continuous sequences including 3×3 convolution, batch normalization and ReLU layers. The “Head” module denotes two 3×3 and 1×1 convolution layers connected with an activation layer. \otimes denotes element-wise product and \oplus denotes element-wise addition. The “Enhance” module denotes the concatenated operations mentioned in Geometric-Aware Mask Learning.

field maps, and \mathbf{F}_{seg} of segmentation maps. The different embedding features are then transformed by the cross-level interaction module (introduced in Sec. 4.2) to predict (1) the heatmap \mathbf{H} and the short-range offset field \mathbf{O} for the vertices, (2) the attraction field map \mathbf{A} of the line segments, and (3) the segmentation mask \mathbf{S} for the segmentation map of the building polygons.

4.2. Learning from Cross-Level Interactions

As the attraction field representation (Xue et al., 2019, 2021) has strong correlations with the vertices and segmentation masks, we are going to regularize the feature maps \mathbf{F}_{seg} and \mathbf{F}_{ver} to emphasize their shape correctness for the mask predictions and the semantic consistency in the learning of vertices.

Efficient Channel Attention with AFM Embedding. Inspired by the success of ECA-Net (Wang et al., 2020), we extend the efficient channel attention mechanism to enhance the feature representation of \mathbf{F}_{seg} and \mathbf{F}_{ver} by

$$\begin{aligned} \mathbf{F}_{\text{seg}}^e &= \sigma(\text{CID}(\text{GAP}(\mathbf{F}_{\text{afm}} + \mathbf{F}_{\text{seg}}))) \cdot \mathbf{F}_{\text{seg}} + \mathbf{F}_{\text{seg}}, \\ \mathbf{F}_{\text{ver}}^e &= \sigma(\text{CID}(\text{GAP}(\mathbf{F}_{\text{afm}} + \mathbf{F}_{\text{ver}}))) \cdot \mathbf{F}_{\text{ver}} + \mathbf{F}_{\text{ver}}, \end{aligned} \quad (5)$$

where $\text{C1D}(\cdot)$, $\text{GAP}(\cdot)$, $\sigma(\cdot)$ are the 1-D convolution, global average pooling and Sigmoid function.

Geometric-Aware Mask Learning. We lift the learned attraction field maps by a linear layer from the predicted attraction field $\mathbf{A} \in \mathbb{R}^{2 \times H_s \times W_s}$ to a 128-channel tensor $\mathbf{F}_{\text{afm}}^*$ and concatenate it to the backbone feature \mathbf{F}_b . The concatenated feature map is subsequently used to predict the building masks \mathbf{S}_{afm} , which guides the backbone feature \mathbf{F}_b to learn the shape correctness from the target of attraction field maps via backpropagation. As the backbone feature map is regularized by the predicted attraction field \mathbf{A} , we are able to use the enhanced feature map $\mathbf{F}_{\text{seg}}^e$ for the final mask prediction, denoted by \mathbf{S}_f . During training, the predicted intermediate mask \mathbf{S}_{afm} and the final mask \mathbf{S}_f are supervised by the binary cross entropy (BCE) loss with the groundtruth masks \mathbf{S}^* by

$$\mathcal{L}_{\text{mask}} = \text{BCE}(\mathbf{S}_{\text{afm}}, \mathbf{S}^*) + \text{BCE}(\mathbf{S}_f, \mathbf{S}^*), \quad (6)$$

while the attraction field \mathbf{A} is learned from the ℓ_1 loss function by

$$\mathcal{L}_{\text{afm}} = \ell_1(\mathbf{A}, \mathbf{A}^*), \quad (7)$$

where \mathbf{A}^* is the ground truth of attraction fields generated by the edge of polygons.

Semantic-Consistent Vertices Prediction. Instead of using the salient edge pixels (from predicted masks) as building vertices, we design a vertex branch to learn the heatmap and short-range offsets of vertices. Taking the enhanced feature $\mathbf{F}_{\text{ver}}^e$, a convolution layer is used to learn the heatmap \mathbf{H} and the offset map \mathbf{O} by the loss functions of cross entropy and ℓ_1 loss by

$$\begin{aligned} \mathcal{L}_{\text{vcls}} &= \text{BCE}(\mathbf{H}, \mathbf{H}^*), \\ \mathcal{L}_{\text{voff}} &= \ell_1(\mathbf{O} \cdot \mathbf{H}^*, \mathbf{O}^* \cdot \mathbf{H}^*), \end{aligned} \quad (8)$$

where \mathbf{H}^* is the ground truth of vertices with the values 0 for background pixels and 1 for the vertices targets, \mathbf{O}^* is the short-range offsets that are in the range of $[-0.5, 0.5]$ for the target pixels.

Total Loss. To train the neural network, we use the predefined loss factors to balance the different types of predictions in the total loss function by

$$\mathcal{L}_{\text{total}} = \lambda_m \cdot \mathcal{L}_{\text{mask}} + \lambda_a \cdot \mathcal{L}_{\text{afm}} + \lambda_c \cdot \mathcal{L}_{\text{vcls}} + \lambda_o \cdot \mathcal{L}_{\text{voff}}, \quad (9)$$

where λ_m , λ_a , λ_c and λ_o are the loss factors of the predicted masks, attraction field maps, the heatmaps and offsets of the vertices. In our experiment, we empirically set $\lambda_m = 1.0$, $\lambda_a = 0.1$, $\lambda_c = 8.0$ and $\lambda_o = 0.25$ according to their loss magnitudes.

4.3. Extracting Polygons by Mask Attraction

Benefiting from our design of cross-level interaction module, the learned masks have a desired property of shape correctness in the boundary geometry. Therefore, the polygonization of the building masks would be straightforward

by tracing the boundary pixels from the predicted masks. For the sake of computing the most simplified polygons without redundant vertices, we found that the learned vertices from \mathbf{H} and \mathbf{O} cooperate very well by attracting the traced boundary pixels of masks to the instance-agnostic vertices. We term our scheme of polygonization as Mask-and-Vertices Attraction, short in *MaV-Attr*. The whole pipeline of our polygon extraction is shown in Fig. 4.

Polygon Initialization. Given a predicted mask \mathbf{S} , the corresponding heatmap \mathbf{H} and offset \mathbf{S} of vertices, we extract the building instances by a given threshold of $\tau \in (0, 1)$ to get the binary mask prediction $\mathbf{S}_\tau = \mathbf{S} > \tau$, which are subsequently processed into n instances by computing the connected components, denoted by $\Omega_1, \dots, \Omega_n$. The redundant polygonal representation of Ω_i is computed by tracing all the boundary pixels, denoted in $\partial\Omega_i^{\text{init}} = [\mathbf{x}_1^i, \dots, \mathbf{x}_r^i, \mathbf{x}_{r+1}^i = \mathbf{x}_1^i]$, where r is the number of boundary pixels of Ω_i . With the initial redundant polygons $\partial\Omega_i^{\text{init}}$, we only need to focus on its simplification by sparse vertices that are computed by leveraging the local non-maximum suppression (NMS) in 3×3 neighborhoods. The local NMS is implemented by a MaxPooling layer with the kernel size 3×3 . After the local NMS, the top- K ($K = 300$ in our experiments) vertices that have a higher classification score than $\tau_v = 0.008$ are retained and refined by the offset vectors for sub-pixel accuracy, denoted by $V = \{\mathbf{v}_1, \dots, \mathbf{v}_m\}$.

Polygon Simplification. For the i -th initial polygon ∂_i^{init} , we match each vertex \mathbf{x}_k^i to the closest vertex $\mathbf{v}_{\pi(k,i)} \in V$ via computing the Euclidean distance by

$$\pi(k, i) = \arg_j \min_{\mathbf{v}_j \in V} \|\mathbf{x}_k^i - \mathbf{v}_j\|_2, \quad (10)$$

where $\pi(k, i)$ is an index mapping from the k -th boundary pixel of the i -th initial polygon to the closest vertex in V . Then, we go through all the boundary pixels \mathbf{x}_k^i and remove the non-minimal pixels that have the same indices. That is to say, for any boundary pixel \mathbf{x}_k^u and \mathbf{x}_k^v , if their indices satisfying $\pi(u, i) = \pi(v, i) = z$ and $\|\mathbf{x}_k^u - \mathbf{v}_z\| < \|\mathbf{x}_k^v - \mathbf{v}_z\|$, the pixel \mathbf{x}_k^v is removed from the initial polygon ∂_i^{init} . In addition to this, for the pixels that are far away from any predicted vertex (under a given threshold $\tau_d = 5$) are also removed. As we already know the topological relationship between any boundary pixels in ∂_i^{init} , a linear scanning is required to efficiently remove the redundant boundary pixels. In the final step, we check the adjacent edges of polygons and merge them if they are parallel up to an angle tolerance of 10° .

The simplified polygon of ∂_i^{init} is denoted by $\partial_i^{\text{final}} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_M, \hat{\mathbf{x}}_{M+1} = \hat{\mathbf{x}}_1]$. For any vertex $\hat{\mathbf{x}}_i \in \partial_i^{\text{final}}$, it must be in the set V . The number of vertices $M = |\partial_i^{\text{final}}|$ is far less than $m = |\partial_i^{\text{init}}|$.

4.4. Implementation details

In our implementation, we mainly use the HRNetV2-W48 (Wang et al., 2021) as the backbone for our method.




Figure 4: The inference process of our method. Building polygons are generated through matching masks with their nearest junctions. The abundant junctions that do not belong to building vertices are removed afterwards.

Meanwhile, we also use the smaller backbones, HRNetV2-W18, HRNetV2-W32 (Wang et al., 2021), and the URNet101 (Ronneberger et al., 2015; He et al., 2016) to show that the possible usage of our method in more practical configurations. The input images are resized to 512×512 during both training and testing. As all the used backbone networks will yield the down-sampled feature maps with the down-sampling factor of $s = 4$, the output maps of our method are with the resolution of 128×128 . In the inference phase, we will resize the polygonal outputs into original image size (e.g., 300×300 in AICrowd dataset (Mohanty et al., 2020)) for evaluation.

During training, we used the data augmentation strategies (as used in (Girard et al., 2020)) including random flip, random rotation, and color jittering. The ADAM optimizer (Kingma and Ba, 2015) is used to train our network on 4 Nvidia RTX 3090 GPUs with the initial learning rate of 10^{-4} and the weight decay 10^{-4} on PyTorch 1.8. The number of training epochs is set to 100 and the learning rate is decayed by 10 after 25 epochs of training. For the ablation study, we reduce the training epochs from 100 to 30 due to the limited computation resources and time.

5. Experiment

5.1. Datasets and evaluation metrics

We compared our method to the state-of-the-art methods on two publicly available datasets.

5.1.1. AICrowd dataset

AICrowd Mapping Challenge dataset (Mohanty et al., 2020) (AICrowd dataset) contains 300×300 pixels RGB images and corresponding annotations in MS-COCO (Lin et al., 2014) format. The training set includes 280741 tiles. For this challenge is already closed, we use the validation subset as testing, whose amount of tiles is 60317. AICrowd dataset also provided users with a small version that contains

smaller amount of samples. We use the small version of AICrowd dataset in the ablation studies.

For evaluating the performance of building instance segmentation, we follow the standard MS-COCO (Lin et al., 2014) evaluation metrics. They are also the official metrics for AICrowd dataset. Particularly, we report the average precision (AP) and average recall (AR) metrics under a range of intersection over union (IoU) thresholds from 0.50 to 0.95 with a step of 0.05. The precision and recall scores under IoU thresholds of 50 and 75 are also listed in Table 1 denoted as AP_{50} , AP_{75} , AR_{50} and AR_{75} . Moreover, to put extra emphasis on the quality of building contours, we introduce the metrics of Boundary IoU (Cheng et al., 2021). Compared with MS-COCO metrics, the Boundary IoU metrics are more reasonable to reveal the accuracy of the detected instance boundaries. For two masks A and B , the Boundary IoU only computes the IoU between sets of pixels that are within distance d from two mask contours. Denoted the contours of masks A and B as A_d and B_d , the Boundary IoU is defined as

$$\text{Boundary IoU}(A, B) = \frac{|(A_d \cap A) \cap (B_d \cap B)|}{|(A_d \cap A) \cup (B_d \cap B)|}. \quad (11)$$

Here the d equals 0.02 which is consistent with the setting of COCO instance segmentation. The average boundary precision denoted by AP^{boundary} under a range of IoU to object boundary is reported in Table 2.

All the above evaluation metrics are based on results at the segmentation level. To directly compare the quality of predicted building polygons, we report the evaluation of PoLiS metric (Avbelj et al., 2015). For two given polygons A and B , the PoLiS is defined as the average distances between each vertex $a_j \in A, j = 1, \dots, q$ of A and its closest point b within the boundary ∂B and vice versa. Suppose the polygon B has vertices $b_k \in B, k = 1, \dots, r$, the PoLiS metric can be

expressed as

$$\begin{aligned} \text{PoLiS}(A, B) = & \frac{1}{2q} \sum_{a_j \in A} \min_{b \in \partial B} ||a_j - b|| \\ & + \frac{1}{2r} \sum_{b_k \in B} \min_{a \in \partial A} ||b_k - a||, \end{aligned} \quad (12)$$

where $\frac{1}{2q}$ and $\frac{1}{2r}$ are the normalization factors. The IoU threshold for filtering predicted building polygons is set to 0.5 as referred to Zhao et al. (2021).

The IoU and complexity aware IoU (C-IoU) (Zorzi et al., 2022) are also calculated for evaluation. For the two polygons A and B , the C-IoU metric is computed as:

$$\text{C-IoU}(A, B) = \text{IoU}(A_m, B_m) \cdot (1 - \text{RD}(N_A, N_B)), \quad (13)$$

where the first term $\text{IoU}(\cdot)$ indicates the normal IoU between two compared polygon masks A_m and B_m . The second term $\text{RD}(N_A, N_B) = |N_A - N_B| / (N_A + N_B)$ is the relative difference between the total number of vertices N_A from polygon A and the number of vertices N_B from polygon B . The C-IoU metric considers both segmentation accuracy and polygonization complexity.

5.1.2. Inria dataset

Inria Aerial Image Labeling dataset (Maggiori et al., 2017a) (Inria dataset) contains aerial orthorectified color imagery of 5000×5000 pixels with a spatial resolution of 0.3 m. Unlike AICrowd dataset, the training and testing images of Inria dataset come from different cities where the building distribution varies. The original Inria dataset provides annotations of pixel-wise semantic segmentation, which is not suitable for getting polygonized results by our method. During the training process, we use the traditional method to convert the raster ground truth labels to vector format. As suggested in (Maggiori et al., 2017a), the first five images of each location from the Inria training set are used for validation.

As the Inria dataset (Maggiori et al., 2017a) did not public the testing samples, we follow their official evaluation protocol by submitting the polygonized results in the format of segmentation masks to their evaluation server. The recommended metrics of IoU and accuracy (Acc) are adopted for comparison. The Acc for Inria dataset is defined as the percentage of correctly classified pixels, while the IoU is computed only for pixels belonging to buildings.

5.2. Results and Analysis

5.2.1. Results on AICrowd dataset

Table 1 and Table 2 demonstrate the quantitative evaluation results on the AICrowd dataset. We compare our method with other competing approaches including Mask R-CNN (He et al., 2017) based on the implementation of (Mohanty et al., 2020), PolyMapper (Li et al., 2019), ASIP (Li et al., 2020), Frame-Field (Girard et al., 2020), PolyWorld (Zorzi et al., 2022) and the recent work proposed by (Li et al., 2021). The Mask R-CNN implemented

by Mohanty et al. (2020) serves as the baseline for AICrowd dataset, which can only obtain building masks instead of polygons. To be able to compare with others, a post-processing with Douglas-Peucker (DP) simplification (Douglas and Peucker, 1973) is adopted to generate vectorized results from the pixel-wise segmentation. The rest of methods all have outputs of polygonal building instances.

It can be seen in Table 1 that our method gets either the highest or comparable scores under the MS-COCO metrics. The average precision termed AP reaches 79.4% by our method. Only the recall score under IoU equals 0.5 is slightly lower than the ASIP (Li et al., 2020) method. It is probably due to the omission of some tiny building pieces that are near the edge of images.

Table 2 reports the additional evaluation results of methods provided with code. The Boundary IoU is a more restricted measurement of the building boundary area. The performance of our method ranks first place on the average boundary precision. It indicates that the building polygons obtained by our method have more precise shapes. Apart from the evaluation of instance segmentation, our method also gets 0.726 on the PoLiS metric, which outperforms the other three vectorization methods. It demonstrates that our method has less overall dissimilarity between vertices of predicted building polygons and the ground truth. For the C-IoU metric, our method surpasses all the other listed methods, which means that the buildings extracted by our method are represented in both concise and accurate polygons. Meanwhile, the learned reversible building masks give our method more flexibility for representing buildings in subtle structure details.

Fig. 5 shows the visual comparison of qualitative results. Each building instance is represented by a different colored polygon. Compared with the mask-related approach Frame-Field (Girard et al., 2020), our method gets building polygons with fewer but more precise vertices. The PolyMapper (Li et al., 2019) and PolyWorld (Zorzi et al., 2022) approaches are based on points inference. They both tend to predict relatively regular simple polygons. However, they are not good at dealing with buildings of complex structures. For instance, they failed in extracting the building marked in green hollow polygon in the image from the last column. While our method could both handle buildings with holes and get compact polygon representation as well.

¹The implementation code from the author's github repository (<https://github.com/lizuoyue/ETH-Thesis/tree/master/building>) is used for training by ourselves. And we perform the testing on the defined testing dataset.

²Here the results are reported by applying the polygonization method mentioned in Frame-Field (Girard et al., 2020) to the the segmentation probability maps from the UNet-Variant (Czikon et al., 2018).

³ASM refers to the Active Skeleton Model polygonization algorithm with the marching squares as initialization in Frame-Field (Girard et al., 2020). Here the parameter of simplification tolerance equals 1. The whole pretrained UResNet101 model is downloaded from <https://github.com/Lydorn/Polygonization-by-Frame-Field-Learning>.

Table 1

Comparison of evaluation results of MS-COCO metrics on the AICrowd dataset. The highest scores are in bold.

Method	Backbone	AP \uparrow	AP $_{50}$ \uparrow	AP $_{75}$ \uparrow	AR \uparrow	AR $_{50}$ \uparrow	AR $_{75}$ \uparrow
Mask R-CNN (He et al., 2017; Mohanty, 2018)	ResNet101	41.9	67.5	48.8	47.6	70.8	55.5
PolyMapper ¹ (Li et al., 2019)	VGG16	50.8	81.7	58.6	58.5	85.4	67.0
ASIP (Li et al., 2020)	-	65.8	87.6	73.4	78.7	94.3	86.1
Frame-Field ² (Girard et al., 2020)	UResNet101	67.0	92.1	75.6	73.2	93.5	81.1
Li et al. (2021)	UResNet101	73.8	92.0	81.9	72.6	90.5	80.7
PolyWorld (Zorzi et al., 2022)	R2U-Net	63.3	88.6	70.5	75.4	93.5	83.1
HiSup (ours)	HRNetV2-W48	79.4	92.7	85.3	81.5	93.1	86.7

Table 2

Comparison of additional results on the AICrowd dataset. The best scores are in bold.

Method	Backbone	AP ^{boundary} \uparrow	PoLiS \downarrow	C-IoU \uparrow	IoU \uparrow
Mask R-CNN (Mohanty, 2018)	ResNet101	15.4	3.454	50.1	61.3
PolyMapper ¹ (Li et al., 2019)	VGG16	22.6	2.215	67.5	77.6
Frame-Field (ASM) ³ (Girard et al., 2020)	UResNet101	34.4	1.945	73.8	84.3
PolyWorld (Zorzi et al., 2022)	R2U-Net	50.0	0.962	88.3	91.2
HiSup (ours)	HRNetV2-W48	66.5	0.726	89.6	94.3

5.2.2. Results on Inria dataset

For the Inria dataset, our method is compared with the other five methods. The top two methods (“Advanced Institute” and “Eugene Khvedchenya”) in Table 3 take the first and second place on the public leaderboard of Inria dataset website. The ICT-Net (Chatterjee and Poullis, 2019) that combines improved UNet with Dense blocks (Huang et al., 2017) and SE blocks (Hu et al., 2020) also got IoU score higher than 80. These three mentioned methods are trained directly on semantic segmentation annotations from the original Inria dataset. They all get segmentation results only and need extra post-processing for vectorization. In addition, we list results of Frame-Field (Girard et al., 2020) and Zorzi et al. (2021), who get vectorized results as ours.

Compared with AICrowd dataset, Inria dataset contains lots of buildings with more complex shapes and structures. And the patterns of building distribution in different cities vary greatly. Table 3 shows the quantitative results of IoU and Acc on the testing images. Compared to the state-of-the-art polygonal mapping of buildings approaches (Girard et al., 2020; Zorzi et al., 2021), our proposed method achieves better performance on both the metrics of IoU and Accuracy.

Fig. 6 shows the qualitative results conducted on an image from the test set of Inria dataset. We produce vectorized polygons that fit the buildings with various shapes.

5.3. Discussion

We perform additional studies to further validate the effect of details of our method.

5.3.1. Effectiveness of mask and vertex attraction

On the basis of the same segmentation masks, we compare our polygon generation algorithm with other methods

Table 3

Comparison of evaluation results on the Inria dataset.

Method	IoU \uparrow	Acc. \uparrow
Advanced Institute ⁴	81.91	97.41
Eugene Khvedchenya ⁴	81.06	97.25
ICT-Net (Chatterjee and Poullis, 2019)	80.32	97.14
Zorzi et al. (2021)	74.40	96.10
Frame-field (Girard et al., 2020)	74.80	95.96 ⁴
HiSup (ours)	75.53	96.27

to validate its effectiveness. We generate the set of segmentation masks by applying our method to images of the AICrowd small validation dataset. The traditional DP simplification (Douglas and Peucker, 1973) algorithm, ASIP (Li et al., 2020), and Frame-Field (Girard et al., 2020) polygonization are chosen as comparison. The tolerance parameter of DP simplification (Douglas and Peucker, 1973) is set to 1. For the ASIP (Li et al., 2020) polygonization, we kept the parameter settings the same as in their released code. For the Frame-Field (Girard et al., 2020) polygonization, we adopt the Active skeleton Model of tolerance 1 with UResNet101 as backbone and marching square as initialization.

As demonstrated in Table 4, in terms of the segmentation evaluation metrics, the segmentation mask obtained by our method gets a fairly high score of 79.3% for the AP. The DP simplification (Douglas and Peucker, 1973) algorithm gets the AP of 69.5% which is comparable to most existing methods. Based on the accurate segmentation masks, the ASIP (Li et al., 2020) and Frame-Field (Girard et al., 2020) polygonization methods get AP of 70.5% and 73.7% respectively. The scores are both higher than the originally reported results as shown in Table 1. Such differences validate our view that the reversibility of building masks is determinant for the final polygonization performance.

⁴Reported on <https://project.inria.fr/aerialimagelabeling/leaderboard/>.



Figure 5: Example of qualitative polygonal mapping of buildings results of AICrowd dataset. Building instances are marked with colors. From the top row to the bottom row: PolyMapper (Li et al., 2019), Frame-Field (Girard et al., 2020) with ASM as polygonization and UNetResNet101 as backbone, PolyWorld (Zorzi et al., 2022), our results and ground truth.

Despite the comparable performance, the polygonization processes of DP and ASIP come with a great loss compared to the original masks. The segmentation AP of the generated polygons drops by almost nine points compared with the masks. The AP^{boundary} also drops more than 10 points. The

introducing of frame field vectors by Frame-Field learning method (Girard et al., 2020) narrows the gap between masks and polygons a little. In contrast, our method has the minimum loss when vectorizing the original masks. The gap of mask AP score is less than one point. By matching with



Figure 6: Crop of polygonal mapping of buildings results by our method on a test image from the Inria dataset. Building instances are marked in different colored polygons. Four zoomed in views of polygonal mapping of buildings are provided as typical.

Table 4

Comparison of our polygon generation with other approaches on the AICrowd small dataset. “Segmentation mask” refers to the predicted mask results by our proposed network model. “DP poly” refers to the traditional Douglas-Peucker simplification algorithm. “ASIP poly” refers to the ASIP polygonization method. “Frame-field poly” refers to the ASM polygonization method from Frame Field with UResNet101. “Our poly” refers to our polygon generation algorithm described.

Method	AP \uparrow	AP ^{boundary} \uparrow	PoLiS \downarrow
Segmentation mask	79.3	65.8	-
DP poly (Douglas and Peucker, 1973)	69.5	51.3	1.178
ASIP poly (Li et al., 2020)	70.5	54.5	1.160
Frame-field poly (Girard et al., 2020)	73.7	57.9	0.984
Our poly	78.8	66.3	0.737

vertices, the polygons obtained by our method have more accurate shapes, which leads to an even higher AP^{boundary} score than masks. Besides, our generated polygons get the minimum PoLiS score.

5.3.2. Effectiveness of hierarchical supervision

We validate the effect of hierarchical supervision by replacing the AFM with edge and frame-field vectors while the remained architectures of the network model stay the same. We also conduct experiments to evaluate the designed Cross-Level Interactions by just predicting AFM as

a separate output to the backbone without the proposed interactions of AFM embedding. The baseline model is set as predicting only masks and junctions from backbone features. All mentioned models are trained and tested on the small version of AICrowd dataset.

The resulting masks and polygons are evaluated by the AP from MS-COCO metrics. The generated junctions are measured by F1-score with the controlled distance threshold set to 5 pixels.

As shown in Table 5, simply predicting AFM representation with a new head helps to improve all the mask, junction and polygon results compared to baseline. When introducing the proposed Cross-Level Interactions with AFM embedding, the performance has been significantly improved. The result also shows that the AFM supervision improves more on mask localization performance and junction detection performance compared to edge supervision and frame-field vectors supervision.

5.3.3. Different backbones

We test our method with different backbone network architectures. The classical UResNet101 model is chosen as a comparison for utilized by most of the previous methods. For the HRNet (Wang et al., 2021) series, we adopt the lightweight models of HRNetV2-W18 and the HRNetV2-W32. Table 6 shows the quantitative evaluation results of different backbone settings on the AICrowd dataset. It can be

Table 5

Ablation study of the proposed hierarchical supervision on the AICrowd small dataset. The comparisons are made by replacing the supervision of AFM with edge and frame-field map respectively. The "AFM (head only)" stands for the simplified version of our model that omits the proposed cross-level interactions.

Baseline	Edge	Frame-field	AFM (head only)	AFM	Mask (AP↑)	Polygon (AP↑)	Junction (F↑)
✓	-	-	-	-	57.1	55.9	59.3
✓	✓	-	-	-	57.7	56.4	60.9
✓	-	✓	-	-	57.7	56.4	61.6
✓	-	-	✓	-	57.6	56.5	61.7
✓	-	-	-	✓	58.1	56.7	62.6

Table 6

Evaluation of our method with different backbones on the AICrowd dataset.

Backbone	AP↑	AP_{50} ↑	AR↑	AR_{50} ↑	AP^{boundary} ↑	PoLiS↓	C-IoU↑
UResNet101	75.8	90.8	77.7	91.7	61.6	0.925	88.2
HRNetV2-W18	71.4	89.6	73.7	90.4	52.6	1.116	85.2
HRNetV2-W32	76.2	91.7	78.5	92.1	61.2	0.872	88.1

seen that our method with UResNet101 gets 75.8% of the AP, which outperforms currently available methods. The larger model with the HRNetV2-W32 gets better performance than the HRNetV2-W18 and UResNet101.

6. Conclusion

This paper studies the problem of polygonal mapping of buildings in a novel viewpoint of *mask reversibility*. By taking the hierarchical-level supervision signals from the bottom-level vertices to the high-level regional masks, we present a novel method, HiSup, to use the mid-level attraction fields of line segments as the most important linkage. The key component of Cross-Level Feature Interaction in our proposed HiSup method learns a unified feature embedding in both aspects of high-level semantics and the shape correctness of buildings, thereafter closes the gap between the segmentation masks and polygons of buildings in satellite imagery. In the experiments, we show that the proposed HiSup outperforms existing polygonal mapping of buildings methods on several challenging metrics on two public benchmarks of AICrowd (Mohanty et al., 2020) and Inria (Maggiori et al., 2017a). The systematic ablation studies further justified our design choices in our HiSup.

References

Alshehhi, R., Marpu, P.R., Woon, W.L., Mura, M.D., 2017. Simultaneous extraction of roads and buildings in remote sensing imagery with convolutional neural networks. *ISPRS Journal of Photogrammetry and Remote Sensing* 130, 139–149. doi:10.1016/j.isprsjprs.2017.05.002.

Avbelj, J., Müller, R., Bamler, R., 2015. A metric for polygon comparison and building extraction evaluation. *IEEE Geoscience and Remote Sensing Letters* 12, 170–174.

Chatterjee, B., Poullis, C., 2019. On building classification from remote sensor imagery using deep neural networks and the relation between classification and reconstruction accuracy using border localization as proxy, in: *Conference on Computer and Robot Vision (CRV)*, pp. 41–48. doi:10.1109/CRV.2019.00014.

Chen, Q., Wang, L., Steven, W.L., Xiuguo, L., 2020. An end-to-end shape modeling framework for vectorized building outline generation from aerial images. *ISPRS Journal of Photogrammetry and Remote Sensing* 170, 114–126.

Cheng, B., Girshick, R., Dollár, P., Berg, A.C., Kirillov, A., 2021. Boundary IoU: Improving object-centric image segmentation evaluation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15334–15342.

Cheng, D., Liao, R., Fidler, S., Urtasun, R., 2019. DARNet: Deep active ray network for building segmentation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7431–7439.

Cuturi, M., 2013. Sinkhorn distances: Lightspeed computation of optimal transport, in: *Advances in Neural Information Processing Systems*, pp. 2292–2300.

Czakon, J., Kaczmarek, A.K., Pyskir, A., Tarasiewicz, P., 2018. Best practices for elegant experimentation in data science projects. <https://github.com/neptune-ai/open-solution-mapping-challenge>.

Douglas, D., Peucker, T., 1973. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. In *Cartographica: The International Journal for Geographic Information and Geovisualization* 10, 112–122.

Girard, N., Smirnov, D., Solomon, J., Tarabalka, Y., 2020. Polygonal building segmentation by frame field learning, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5891–5900.

He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask r-cnn, in: *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, pp. 2980–2988.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778.

Hu, J., Shen, L., Albanie, S., Sun, G., Wu, E., 2020. Squeeze-and-excitation networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42, 2011–2023. doi:10.1109/TPAMI.2019.2913372.

Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2261–2269. doi:10.1109/CVPR.2017.243.

Iglovikov, V., Seferbekov, S.S., Buslaev, A., Shvets, A., 2018. Teraus-netv2: Fully convolutional network for instance segmentation, in: *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops, Computer Vision Foundation / IEEE Computer Society*, pp. 233–237.

- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: International conference on machine learning, PMLR. pp. 448–456.
- Ji, S., Wei, S., Lu, M., 2019. Fully convolutional networks for multisource building extraction from an open aerial and satellite imagery data set. *IEEE Transaction on Geoscience and Remote Sensing* 57, 574–586.
- Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization, in: International Conference on Learning Representations (ICLR).
- Li, M., Lafarge, F., Marlet, R., 2020. Approximating shapes in images with low-complexity polygons, in: Proc. of the IEEE conference on Computer Vision and Pattern Recognition (CVPR), pp. 8633–8641.
- Li, Q., Mou, L., Hua, Y., Shi, Y., Zhu, X.X., 2022. Building footprint generation through convolutional neural networks with attraction field representation. *IEEE Transactions on Geoscience and Remote Sensing* 60, 1–17. doi:10.1109/TGRS.2021.3109844.
- Li, W., Zhao, W., Zhong, H., He, C., Lin, D., 2021. Joint semantic-geometric learning for polygonal building segmentation, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 1958–1965.
- Li, Z., Wegner, J.D., Lucchi, A., 2019. Topological map extraction from overhead images, in: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 1715–1724.
- Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P., 2014. Microsoft coco: Common objects in context, in: Proceedings of European conference on computer vision (ECCV), p. 740–755.
- Liu, Y., Chen, D., Ma, A., Zhong, Y., Fang, F., Xu, K., 2021a. Multi-scale u-shaped CNN building instance extraction framework with edge constraint for high-spatial-resolution remote sensing imagery. *IEEE Transaction on Geoscience and Remote Sensing* 59, 6106–6120.
- Liu, Z., Tang, H., Huang, W., 2021b. AFM-RNN: A sequent prediction model for delineating building rooftops from remote sensing images by integrating RNN with attraction field map, in: Chinese Conference on Pattern Recognition (PRCV), Springer. pp. 474–485. doi:10.1007/978-3-030-88007-1_39.
- Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3431–3440.
- Lorensen, W.E., Cline, H.E., 1987. Marching cubes: A high resolution 3d surface construction algorithm, in: Stone, M.C. (Ed.), Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1987, Anaheim, California, USA, July 27–31, 1987, ACM. pp. 163–169. URL: <https://doi.org/10.1145/37401.37422>, doi:10.1145/37401.37422.
- Maggiori, E., Tarabalka, Y., Charpiat, G., Alliez, P., 2017a. Can semantic labeling methods generalize to any city? the inria aerial image labeling benchmark, in: IEEE International Geoscience and Remote Sensing Symposium (IGARSS), pp. 3226–3229.
- Maggiori, E., Tarabalka, Y., Charpiat, G., Alliez, P., 2017b. Convolutional neural networks for large-scale remote-sensing image classification. *IEEE Transaction on Geoscience and Remote Sensing* 55, 645–657.
- Maggiori, E., Tarabalka, Y., Charpiat, G., Alliez, P., 2017c. High-resolution aerial image labeling with convolutional neural networks. *IEEE Transaction on Geoscience and Remote Sensing* 55, 7092–7103.
- Manno-Kovacs, A., Sziranyi, T., 2015. Orientation-selective building detection in aerial images. *ISPRS Journal of Photogrammetry and Remote Sensing* 108, 94–112.
- Mayer, H., 1999. Automatic object extraction from aerial imagery—a survey focusing on buildings. *Computer Vision and Image Understanding* 74, 138–149. URL: <https://www.sciencedirect.com/science/article/pii/S1077314299907506>, doi:https://doi.org/10.1006/cviu.1999.0750.
- Mnih, V., 2013. Machine Learning for Aerial Image Labeling. Ph.D. thesis. University of Toronto. CAN.
- Mohanty, S.P., 2018. Crowdai mapping challenge 2018: Baseline with mask rcnn. <https://github.com/crowdai/crowdai-mapping-challenge-mask-rcnn>.
- Mohanty, S.P., Czakon, J., Kaczmarek, K.A., Pyskir, A., Tarasiewicz, P., Kunwar, S., Rohrbach, J., Luo, D., Prasad, M., Fleer, S., et al., 2020. Deep learning for understanding satellite imagery: An experimental survey. *Frontiers in Artificial Intelligence* 3, 534696.
- Nair, V., Hinton, G.E., 2010. Rectified linear units improve restricted boltzmann machines, in: ICML, pp. 807–814.
- Ngo, T.T., Mazet, V., Collet, C., Fraipont, P.D., 2016. Shape-based building detection in visible band images using shadow information. *IEEE Journal of Selected Topics in Applied Earth Observations & Remote Sensing*, 1–13.
- Pesaresi, M., Gerhardinger, A., 2011. Improved textural built-up presence index for automatic recognition of human settlements in arid regions with scattered vegetation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 4, 16–26.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, Springer. pp. 234–241.
- Saito, S., Yamashita, Y., Aoki, Y., 2016. Multiple object extraction from aerial imagery with convolutional neural networks. *Journal of Imaging Science and Technology* 60, 10402–1/10402. doi:10.2352/J.ImagingSci.Technol.2016.60.1.010402.
- Senaras, C., Ozay, M., Vural, F.Y., 2013. Building detection with decision fusion. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 6, 1295–1304.
- Sohn, G., Dowman, I., 2001. Extraction of buildings from high resolution satellite data. *Automated Extraction of Man-Made Objects from Aerial and Space Images (III)*. Balkema Publishers, Lisse, 345–355.
- Wang, J., Meng, L., Li, W., Yang, W., Yu, L., Xia, G.S., 2022. Learning to extract building footprints from off-nadir aerial images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* doi:10.1109/TPAMI.2022.3162583.
- Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., Liu, W., Xiao, B., 2021. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43, 3349–3364.
- Wang, Q., Wu, B., Zhu, P., Li, P., Zuo, W., Hu, Q., 2020. ECA-Net: Efficient channel attention for deep convolutional neural networks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 11531–11539.
- Wei, S., Ji, S., Lu, M., 2020. Toward automatic building footprint delineation from aerial images using cnn and regularization. *IEEE Transactions on Geoscience and Remote Sensing* 58, 2178–2189. doi:10.1109/TGRS.2019.2954461.
- Xia, G., Huang, J., Xue, N., Lu, Q., Zhu, X., 2019. Geosay: A geometric saliency for extracting buildings in remote sensing images. *Computer Vision Image Understanding* 186, 37–47.
- Xue, N., Bai, S., Wang, F., Xia, G., Wu, T., Zhang, L., 2019. Learning attraction field representation for robust line segment detection, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Computer Vision Foundation / IEEE. pp. 1595–1603. doi:10.1109/CVPR.2019.00169.
- Xue, N., Bai, S., Wang, F., Xia, G., Wu, T., Zhang, L., Torr, P.H.S., 2021. Learning regional attraction for line segment detection. *IEEE Transaction on Pattern Analysis and Machine Intelligence* 43, 1998–2013.
- Yang, H.L., Yuan, J., Lunga, D., Laverdiere, M., Rose, A., Bhaduri, B., 2018. Building extraction at scale using convolutional neural network: Mapping of the united states. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 11, 2600–2614. doi:10.1109/JSTARS.2018.2835377.
- Yuan, J., 2018. Learning building extraction in aerial scenes with convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40, 2793–2798.
- Zhao, W., Persello, C., Stein, A., 2021. Building outline delineation: From aerial images to polygons with an improved end-to-end learning framework. *ISPRS Journal of Photogrammetry and Remote Sensing* 175, 119–131.
- Zorzi, S., Bazrafkan, S., Habenschuss, S., Fraundorfer, F., 2022. PolyWorld: Polygonal building extraction with graph neural networks in satellite images, in: Proceedings of the IEEE/CVF Conference on Computer

Vision and Pattern Recognition (CVPR).

Zorzi, S., Bittner, K., Fraundorfer, F., 2021. Machine-learned regularization and polygonization of building segmentation masks, in: International Conference on Pattern Recognition (ICPR), pp. 3098–3105. doi:10.1109/ICPR48806.2021.9412866.