

# Contextual Pyramid Attention Network for Building Segmentation in Aerial Imagery

**Citation for published version (APA):**

Sebastian, C., Imbriaco, R., Bondarau, E., & de With, P. H. N. (2021). Contextual Pyramid Attention Network for Building Segmentation in Aerial Imagery. In B. Skoric, & R. van Sloun (Eds.), *Proceedings of the 2021 Symposium on Information Theory and Signal Processing in the Benelux May 20-21, TU Eindhoven* (pp. 8-16). Eindhoven University of Technology.

**Document status and date:**

Published: 01/01/2021

**Document Version:**

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

**Please check the document version of this publication:**

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

**General rights**

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

[www.tue.nl/taverne](http://www.tue.nl/taverne)

**Take down policy**

If you believe that this document breaches copyright please contact us at:

[openaccess@tue.nl](mailto:openaccess@tue.nl)

providing details and we will investigate your claim.

# Contextual Pyramid Attention Network for Building Segmentation in Aerial Imagery

Clint Sebastian Raffaele Imbriaco Egor Bondarev Peter H.N de With  
Eindhoven University of Technology,  
Dept. of Electrical Engineering , SPS-VCA  
Eindhoven, The Netherlands  
{c.sebastian, r.imbriaco, e.bondarev, p.h.n.de.with}@tue.nl

## Abstract

Building extraction from aerial images has several applications in problems such as urban planning, change detection, and disaster management. With the increasing availability of data, Convolutional Neural Networks (CNNs) for semantic segmentation of remote sensing imagery has improved significantly in recent years. However, convolutions operate in local neighborhoods and fail to capture non-local features that are essential in semantic understanding of aerial images. In this work, we propose to improve building segmentation of different sizes by capturing long-range dependencies using contextual pyramid attention (CPA). The pathways process the input at multiple scales efficiently and combine them in a weighted manner, similar to an ensemble model. The proposed method obtains state-of-the-art performance on the Inria Aerial Image Labelling Dataset with minimal computation costs. Our method improves 1.8 points over current state-of-the-art methods and 12.6 points higher than existing baselines on the Intersection over Union (IoU) metric without any post-processing.

## 1 Introduction

The developments in the systematic collection and organization of remote sensing imagery have resulted in several high-resolution aerial imagery datasets. Information from aerial imagery plays a key role in urban planning, disaster aversion, and change detection. Building detection is a crucial aspect for the aforementioned applications. Depending on the geographical region and conditions, building structures have different shapes and sizes. This challenge is particularly addressed by Maggiori *et al.* [1]. They created a dataset of labeled aerial imagery from different locations for this problem, such that a model trained from a variety of sources generalizes to the task of segmentation. Semantic segmentation in aerial imagery is challenging due to variable lighting conditions, shapes/sizes, and large intraclass variations. In this research, we address the problem of improving the building segmentation by utilizing attentive multi-scale pathways. Each of the paths exploits non-local neighborhoods that account for buildings of varying sizes. This allows our network to learn long-range dependencies at various scales with minimal computation costs. In addition to attentive multi-scale pathways, we also incorporate a channel-wise attention module to model interdependencies across channels. In summary, our contributions are as follows.

- We introduce a self-attention based contextual pyramid attention (CPA) module that accounts for various building sizes to segment buildings in aerial images. The proposed module outperforms current state-of-the-art methods by about 2% on the IoU metric and 12.6% over FCN baselines.
- Through experiments, we also show that our base model offers competitive performance to current state-of-the-art methods, while having much lower inference

costs. We also provide ablation studies on the impact of our proposed module and other comparisons.

## 2 Related Work

Convolutional Neural Networks (CNNs) have improved the performance of semantic segmentation significantly in recent years. Encoder-Decoder architectures are a popular choice for segmentation. Using fully convolutional architectures for segmentation is proposed in [2]. Architectures such as FCN, U-Net, DenseNet, etc., are often applied to obtain higher quality segmentations [2, 3, 4]. Nowadays, most segmentation architecture designs are based on using a pre-trained backbone such as VGGNet, ResNet, and a complex decoder such as Pyramid Scene Parsing (PSP) or Atrous Spatial Pyramid Pooling (ASPP)[5, 6, 7, 8, 9, 10].

Due to the unique nature of remote sensing imagery, several custom architectures and loss functions have been developed for aerial image building segmentation [11, 12]. In [11], a multi-task learning approach is introduced, where a distance-transform loss function is applied in conjunction with the cross-entropy loss. Similar to our work, [13] proposes to combine local and global features to improve semantic segmentation of buildings in aerial imagery. However, the combination of two VGGNets makes the overall inference computationally expensive. Our proposed approach circumvents this by splitting only the last convolutional block to operate at multiple scales. In [14], a joint multi-stage multi-task approach is used, where the first stage trains a segmentation network, and the second stage trains for geo-localization using a multi-task loss function. Apart from these, post-processing techniques such as Conditional Random Fields (CRFs) and test-time augmentations are applied to improve segmentation performance. In [15], a recurrent network is applied in a fully convolutional network that exploits a decoder network fusing features from the encoder layer in a similar fashion to feature pyramid networks (FPN). An FPN utilizes a pyramidal hierarchy of features extracted from an encoder that is later combined with the decoder via lateral connections [16]. However, as FPNs utilize only a few of the earlier layers of the encoder for lateral connections, the features might not be rich as extracting from deeper layers.

A common practice to retain feature resolution is using dilated convolutions. Due to the larger feature resolution and the repeated object patterns in aerial imagery, capturing long-range relations is beneficial. However, a single-sized feature map may not be able to capture all objects of variable sizes. For example, spatial relationships across large buildings are easier to capture when feature maps have a low resolution. When buildings are smaller in size, larger feature resolution is suited to capture more fine-grained details. Using a decoder such as FPN is beneficial for high-resolution features, however, as they are not obtained from deeper layers, they lack rich semantics. The proposed approach in this research is suited for these properties that are typically exhibited in aerial imagery. We particularly address the problem of improving semantic segmentation of buildings at various sizes.

## 3 Method

The architecture is composed of an encoder and decoder, similar to other segmentation networks. An overview of the proposed method is shown in Fig. 1. Each component of our model is described in the following subsections.

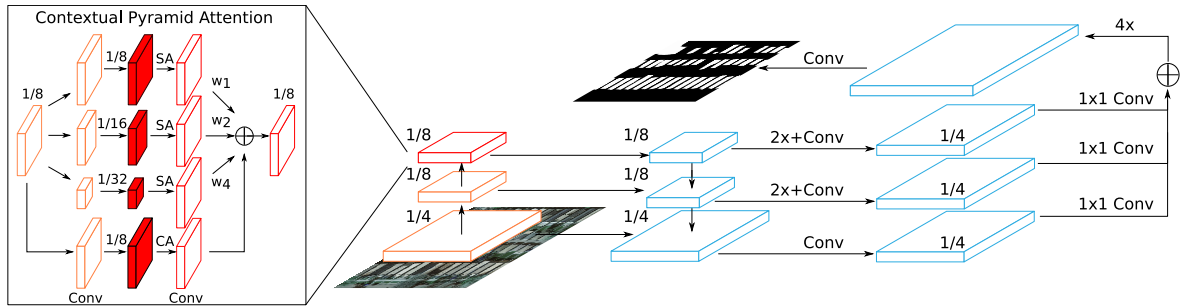


Figure 1: Overview of the proposed network. The network uses a ResNet backbone for the encoder and Feature Pyramid Network decoder. The proposed context pyramid attention (CPA) is applied after Block 4 of ResNet. CPA consists of contextual and channel-wise attention modules, where the contextual attention operates at multiple scales to produce a weighted output that is fed to the FPN decoder (CPA module at the left).

### 3.1 Architecture

**Encoder:** We consider three backbones for training, ResNet18, ResNet101 and Squeeze and Excitation ResNeXt101 (SEResNeXt101). ResNet18 serves as the light architecture that may be suited for fast inference, whereas SEResNeXt101 offers higher performance with additional computation cost. SEResNeXt101 incorporates an additional Squeeze and Excitation module along with aggregated residual connections.

**Decoder:** The decoder is a Feature Pyramid Network (FPN), similar to Semantic FPN in [17]. FPN combines features at different spatial resolutions through lateral connections to a top-down decoder. The lateral connections are established by combining the bottom-up outputs at each corresponding level of the top-down decoder. Each of the combined top-down decoder features is upsampled to  $1/4^{th}$  of the input resolution and is combined via summation to produce rich features that are transformed into a segmentation mask.

### 3.2 Contextual Pyramid Attention block

The proposed Contextual Pyramid Attention block is composed of two parts, a contextual attention module that captures long-range spatial dependencies and a channel-wise attention module. Both proposed attention modules are based on self-attention [18, 19].

**Contextual attention:** Buildings in aerial imagery appear in various sizes, and structural features are often redundant in a given region (e.g. repeated houses structure in a neighborhood). This non-local information is not immediately processed by convolutions as they operate in a limited region defined by the kernel size. Self-attention, on the other hand, can bring in the understanding of long-range dependencies capturing the relation between repeated structures. It is a standard practice in semantic segmentation to use dilated convolutions in deeper layers, to increase the feature resolution. This is beneficial for smaller buildings and furthermore, self-attention blocks can capture fine-grained relations across the image. However, the contextual information of buildings of different sizes becomes challenging for two reasons. First, the dilated convolutions retain the feature resolution and self-attention may then retain spatial relations that are too fine-grained for larger buildings, which may not provide sufficient spatial context. Second, if the dilated convolutions are not applied to retain feature resolution, this may fail to capture relations between smaller buildings. Our

motivation using multi-scale pathways stems from these two reasons.

**Method.** Given a feature  $\mathbf{F} \in \mathbb{R}^{C \times H \times W}$ , we apply convolutions to generate key, query and value features  $\mathbf{K}, \mathbf{Q}, \mathbf{V} \in \mathbb{R}^{C \times H \times W}$ . The tensors  $\mathbf{K}, \mathbf{Q}$  and  $\mathbf{V}$  are reshaped into  $N \times C$ , where  $N = H \times W$ . The self-attention (SA) operation is defined as

$$\mathbf{A}_s = \text{SA}(\mathbf{F}) = \gamma_s \cdot (\text{Softmax}(\mathbf{K}\mathbf{Q}^T))\mathbf{V}, \quad (1)$$

where  $\mathbf{A}_s \in \mathbb{R}^{C \times N}$  is then reshaped into  $C \times H \times W$  and  $\gamma_s$  is a learned parameter. The output of the self-attention is the sum of  $\mathbf{A}_s$  and  $\mathbf{F}$ . The resolution of  $\mathbf{F}$  is retained by dilated convolutions ( $1/8^{\text{th}}$  of the input image resolution). To capture contextual long-range dependencies, we operate the deep features  $\mathbf{F}$  at various scales  $s$ . Contextual attention is

$$\mathbf{C}_s = \sum_{s \in (1,2,4)} w_s \cdot \text{Conv}((\mathbf{F}_{1/s} + \text{SA}(\text{Conv}(\mathbf{F}_{1/s})))_s), \quad (2)$$

where Conv is the convolution operation, followed by ReLU activation and Batch Normalization [20, 21]. The features are downsampled by a factor of  $1/s$ , followed by Conv and SA operations. The resultant features are reshaped and upsampled to the original size, followed by another Conv operation. Each of these paths is weighted by a parameter  $w_s$ . The sum of the pathways results in the contextual attention output  $\mathbf{C}_s$ .

**Channel-wise attention:** Contextual attention employs both spatial axes to establish interdependencies to model spatial information. However, higher-level class or object information is prevalent across channels, and hence, we apply channel-wise attention to model relations across them. The last feature block of ResNet is large with 2,048 channels and is computationally expensive to perform self-attention based operations. Therefore, the features are compressed to  $1/4$  of the total number of channels by  $1 \times 1$  convolutions to obtain the features  $\mathbf{F}_{1/c}$ , where  $c$  is the compression factor across channels (note the difference with scale factor  $s$ ). The key, query and value are all the same in this case ( $\mathbf{F}_{1/c}$ ). However, to apply self-attention, the features are reshaped into  $C \times N$ , where  $N = H \times W$ . The channel-wise attention is now

$$\mathbf{A}_{cw} = \gamma_c \cdot (\text{Softmax}(\mathbf{F}_{1/c}\mathbf{F}_{1/c}^T))\mathbf{F}_{1/c}. \quad (3)$$

The final output is the sum of  $\mathbf{F}_{1/c}$  and  $\mathbf{A}_{cw}$ . Note that the affinity matrices (inputs to the softmax) generated through both attention mechanisms have different shapes. The contextual attention affinity matrix has a size of  $N \times N$ , whereas the channel-wise affinity matrix has a size of  $C \times C$ . The final contextual pyramid attention block is the sum of both contextual and channel-wise attention.

## 4 Experiments

### 4.1 Dataset and Evaluation

**Dataset:** The Inria Aerial Image Labeling Dataset consists of 360 RGB ortho-rectified aerial images at a spatial resolution of 0.3 meters. Each of the ortho-rectified images has a resolution of  $5000 \times 5000$  pixels, covering a region of 1500 square meters per image. The dataset comprises of 10 different cities, out of which 5 cities are available for training. The images are taken at different urban conditions from various cities in Europe and America. The groundtruth is available only for the training set, where information is available for two classes: buildings and non-buildings. For fair comparisons, we follow the same evaluation protocol for testing as in [11, 22, 23]. Images 1–5 of each location are used for validation and Images 6–36 for training.

Method		Austin	Chicago	Kitsap Co.	West Tyrol	Vienna	Overall
FCN + MLP [1]	IoU	61.20	61.30	51.50	57.95	72.13	64.67
	Acc.	94.20	90.43	98.92	96.66	91.87	94.42
SegNet MT-Loss [11]	IoU	76.76	67.06	<b>73.30</b>	66.91	76.68	73.00
	Acc.	93.21	<b>99.25</b>	97.84	91.71	96.61	95.73
MSMT-Stage-1 [22]	IoU	75.39	67.93	66.35	74.07	77.12	73.31
	Acc.	95.99	92.02	99.24	97.78	92.49	96.06
2-levels U-Net + aug. [23]	IoU	77.29	68.52	72.84	75.38	78.72	74.55
	Acc.	96.69	92.40	99.25	98.11	93.79	96.05
ICT-Net [24]	IoU	-	-	-	-	-	75.50
	Acc.	-	-	-	-	-	96.05
ResNet18-FPN -CPA (ours)	IoU	77.89	69.73	65.65	76.95	79.49	75.51
	Acc.	96.77	92.70	99.25	98.13	94.09	96.19
ResNet101-FPN -CPA (ours)	IoU	78.93	<b>71.57</b>	68.06	78.29	81.03	77.03
	Acc.	96.98	93.20	99.27	98.29	94.61	96.47
SEResNeXt101-FPN -CPA (ours)	IoU	<b>80.15</b>	69.54	70.36	<b>80.83</b>	<b>81.43</b>	<b>77.29</b>
	Acc.	<b>97.18</b>	92.78	<b>99.32</b>	<b>98.46</b>	<b>94.67</b>	<b>96.48</b>

Table 1: (Left) Performance comparison of CPA with other state-of-the-art methods.

**Evaluation:** For evaluation, we use both Intersection over Union (IoU) and accuracy. IoU is measured across the building class. IoU is the ratio of true positive pixels to the pixels that are labeled as positive in the ground-truth or predictions. We also utilize accuracy, which is the ratio of correctly classified pixels to the total. Each dataset is evaluated separately and jointly, for both of these metrics.

**Implementation details:** Each of the networks are initialized with weights pre-trained on ImageNet [25]. The models are trained using the Adam optimizer with a learning rate of 0.00001 [26]. The objective function is cross-entropy and is trained for 35 epochs. The input images have a size of  $500 \times 500$  pixels and are randomly rotated by  $90^\circ$  or flipped horizontally and vertically while training. Images are also augmented with minor changes in brightness and contrast. The image augmentations are applied using Albumentations [27], and the models are implemented in PyTorch. Each of the contextual and self-attention weights are initialized to 1 and 0.05, respectively.

## 4.2 Comparison with state-of-the-art

With the proposed module, we observe an improvement of 3-5% IoU across all the datasets except Kitsap Co.. Our method with a simple backbone such as ResNet18 offers competitive performance to current state-of-the-art results. Using a deeper backbone such as ResNet101 or SE-ResNeXt101 further improves performance. Compared to the previous state-of-the-art ICT-Net that uses U-Net with Dense Blocks and SE Blocks [8], we observe a gain of 1.8 IoU points. The improvement is higher than previous increases in recent years. Our method does not rely on post-processing, or multi-task learning methods, however, applying this may further improve performance. A visual comparison of our results is in Fig. 3. We note that on Kitsap Co., which contains the fewest buildings, our performance decreases slightly. We observe very high accuracy (99.32%), whereas IoU is lower, indicating skewness to classifying background better. On visual inspection, we have noticed errors in the ground truth where the background is annotated as buildings. The incorrect ground truth, when combined with the sparsely covered building regions, penalizes the IoU metric, thereby lowering performance.




Figure 2: Comparison of building segmentation on Inria Aerial Image Labeling Dataset. RGB, GT and outputs (bottom row) ResNet101-FPN without (left) and with our module (right).

### 4.3 Ablation studies

To study the impact of attention, an ablation study is conducted without any attention, with self-attention and the proposed CPA. This is shown in Table 1 (top-right). The IoU improves by 0.87 and 1.54 points with the addition of self-attention and CPA over the ResNet-FPN. Furthermore, we conduct experiments to study the impact of different models on computation costs (Table 1 mid-right). For a tile of  $5000 \times 5000$  pixels, ResNet18-FPN-CPA takes only 4.67 seconds on a GTX 1080Ti GPU, whereas ResNet101 and SE-ResNeXt101 take 12.16 and 14.07 seconds, respectively. In [23], a tile takes 160 seconds for processing, however, they use a K80 GPU, which is typically two times slower than our setup. Even so, compared to [23], our largest model (SE-ResNeXt101) is five times faster for inference, while offering higher performance.

## 5 Conclusions

We have presented a novel and effective method to improve building segmentation in aerial imagery. It consists of the contextual pyramid and channel-wise attention blocks to model long-range dependencies across spatial contexts to account for buildings of different sizes. The contextual pyramid attention combines contextual infor-




Figure 3: Prediction, high- and low-resolution attention map. The high-resolution learns semantics related to small buildings and fine-grained details of large buildings whereas low resolution learns semantics related to large buildings.

Model	Accuracy	IoU	Model	Time/tile	IoU
ResNet18-FPN	95.9	73.97	ResNet18	4.67 s	75.51
-Self-attention	96.0	74.64	ResNet101	12.16 s	77.03
-CPA (ours)	96.2	75.51	SE-ResNeXt101	14.07 s	77.29

Table 2: (Right) Comparison of our method with/without the CPA module. (Left) Inference time (in sec.) per tile of  $5000 \times 5000$  pixels with different backbones with CPA module.

mation at multiple scales in an efficient manner with minimal computation overhead, whereas channel-wise attention captures interdependencies across channels. The proposed method achieves state-of-the-art performance on the Inria Aerial Image Labelling dataset by 1.8 and 12.6 IoU points over previous state of the art and baseline. Furthermore, we perform ablation experiments to study the impact of the CPA module and provide comparisons for computation costs. Our model offers high performance without any post-processing with low inference times.

## References

- [1] E. Maggiori, Y. Tarabalka, G. Charpiat, and P. Alliez, “Can semantic labeling methods generalize to any city? the inria aerial image labeling benchmark,” in *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. IEEE, 2017, pp. 3226–3229.
- [2] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.
- [3] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [4] S. Jégou, M. Drozdal, D. Vazquez, A. Romero, and Y. Bengio, “The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation,” in *Proceedings of*



- the *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 11–19.
- [5] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [6] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [7] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, “Aggregated residual transformations for deep neural networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1492–1500.
- [8] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [9] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.
- [10] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, “Encoder-decoder with atrous separable convolution for semantic image segmentation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 801–818.
- [11] B. Bischke, P. Helber, J. Folz, D. Borth, and A. Dengel, “Multi-task learning for segmentation of building footprints with deep neural networks,” in *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2019, pp. 1480–1484.
- [12] C. Sebastian, R. Imbriaco, E. Bondarev, and P. H. de With, “Adversarial loss for semantic segmentation of aerial imagery,” *arXiv preprint arXiv:2001.04269*, 2020.
- [13] A. Marcu and M. Leordeanu, “Dual local-global contextual pathways for recognition in aerial imagery,” *arXiv preprint arXiv:1605.05462*, 2016.
- [14] A. Marcu, D. Costea, E. Slusanschi, and M. Leordeanu, “A multi-stage multi-task neural network for aerial scene interpretation and geolocalization,” *ArXiv*, vol. abs/1804.01322, 2018.
- [15] L. Mou and X. X. Zhu, “Rifcn: Recurrent network in fully convolutional network for semantic segmentation of high resolution remote sensing images,” *arXiv preprint arXiv:1805.02091*, 2018.
- [16] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, “Feature pyramid networks for object detection,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [17] A. Kirillov, R. Girshick, K. He, and P. Dollár, “Panoptic feature pyramid networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6399–6408.
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.

- [19] H. Nam, J.-W. Ha, and J. Kim, “Dual attention networks for multimodal reasoning and matching,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 299–307.
- [20] V. Nair and G. E. Hinton, “Rectified linear units improve restricted boltzmann machines,” in *Proceedings of the 27th international conference on machine learning (ICML-10)*, 2010, pp. 807–814.
- [21] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” *arXiv preprint arXiv:1502.03167*, 2015.
- [22] A. Marcu, D. Costea, E. Slusanschi, and M. Leordeanu, “A multi-stage multi-task neural network for aerial scene interpretation and geolocalization,” *arXiv preprint arXiv:1804.01322*, 2018.
- [23] A. Khalel and M. El-Saban, “Automatic pixelwise object labeling for aerial imagery using stacked u-nets,” *arXiv preprint arXiv:1803.04953*, 2018.
- [24] B. Chatterjee and C. Poullis, “Semantic segmentation from remote sensor data and the exploitation of latent learning for classification of auxiliary tasks,” *arXiv preprint arXiv:1912.09216*, 2019.
- [25] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. IEEE, 2009, pp. 248–255.
- [26] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [27] A. Buslaev, A. Parinov, E. Khvedchenya, V. Iglovikov, and A. Kalinin, “Albumentations: fast and flexible image augmentations,” *arXiv preprint arXiv:1809.06839*, 2018.