

Efficient Building Roof Type Classification: A Domain-Specific Self-Supervised Approach

Guneet Mutreja¹, Ksenia Bittner¹

¹Remote Sensing Technology Institute, German Aerospace Center (DLR), Weßling, Germany –
(guneet.mutreja, ksenia.bittner)@dlr.de

Keywords: Building Roof Type Classification, Self-Supervised Learning, EfficientNet, Domain-Specific Pretraining, Aerial Imagery, Remote Sensing.

Abstract

Accurate classification of building roof types from aerial imagery is crucial for various remote sensing applications, including urban planning, disaster management, and infrastructure monitoring. However, this task is often hindered by the limited availability of labeled data for supervised learning approaches. To address this challenge, this paper investigates the effectiveness of self-supervised learning with EfficientNet architectures, known for their computational efficiency, for building roof type classification. We propose a novel framework that incorporates a Convolutional Block Attention Module (CBAM) to enhance the feature extraction capabilities of EfficientNet. Furthermore, we explore the benefits of pretraining on a domain-specific dataset, the Aerial Image Dataset (AID), compared to ImageNet pretraining. Our experimental results demonstrate the superiority of our approach. Employing Simple Framework for Contrastive Learning of Visual Representations (SimCLR) with EfficientNet-B3 and CBAM achieves a 95.5% accuracy on our validation set, matching the performance of state-of-the-art transformer-based models while utilizing significantly fewer parameters. We also provide a comprehensive evaluation on two challenging test sets, demonstrating the generalization capability of our method. Notably, our findings highlight the effectiveness of domain-specific pretraining, consistently leading to higher accuracy compared to models pretrained on the generic ImageNet dataset. Our work establishes EfficientNet-based self-supervised learning as a computationally efficient and highly effective approach for building roof type classification, particularly beneficial in scenarios with limited labeled data.

1. Introduction

1.1 Background

The proliferation of remote sensing technologies has made aerial imagery readily available, opening up a wide range of applications for understanding and managing our built environment. Accurate and up-to-date information about buildings, including their roof structures, is crucial for urban planning, disaster response, infrastructure management, and the creation of detailed digital twins. For instance, precise 3D building models, which rely on accurate roof type classification, are essential for simulating urban environments, analyzing shadows, and generating realistic visualizations. Furthermore, knowing the distribution of different roof types is valuable for damage assessment after natural disasters and for energy modeling to optimize energy consumption and mitigate urban heat islands.

Despite its importance, obtaining sufficient labeled data for training robust roof type classification models remains a persistent challenge in remote sensing. Manually annotating aerial imagery to identify diverse roof structures is time-consuming, requires expert knowledge, and can be cost-prohibitive. This challenge is amplified by the global variability in roof architectures, ranging from simple flat roofs to complex combinations of gables, hips, and dormers. Additionally, dense urban environments often present difficulties in accurately delineating individual roofs due to closely spaced buildings and complex building shapes.

To address the limitations of relying on scarce labeled data, this paper explores the potential of self-supervised learning for building roof type classification. Self-supervised learning allows models to learn from readily available unlabeled data by



Figure 1. Classification results of BEit model pretrained using Aerial Image Dataset on building roof types in the Braunschweig area.

leveraging inherent structures and patterns within the data itself to generate supervisory signals. This approach holds significant promise for remote sensing applications where labeled data is often limited. Specifically, we propose a novel framework that combines the computational efficiency of EfficientNet architectures, enhanced with a Convolutional Block Attention Module (CBAM), and domain-specific pretraining on the Aerial Image Dataset (AID) to achieve accurate and efficient roof type classification, even with limited labeled examples.

1.2 Related work

Self-Supervised Learning. The scarcity of labeled data in remote sensing has spurred the development of self-supervised learning methods, enabling models to learn representations from unlabeled data. These methods have gained significant

traction due to their ability to leverage the abundance of unlabeled aerial imagery and reduce the reliance on costly manual annotations. A prominent category of self-supervised learning is contrastive learning, where models are trained to distinguish between augmented views of the same image (positive pairs) and different images (negative pairs). Chen et al. (2020), He et al. (2020), and Caron et al. (2020) are notable examples of contrastive learning methods that have shown success in learning useful representations for various downstream tasks, including image classification and object detection. These methods encourage models to learn invariant features that are robust to different viewpoints, lighting conditions, and other variations commonly encountered in remote sensing imagery. Other approaches leverage pretext tasks, where models are trained to solve a predefined task using the unlabeled data, as a form of self-supervision. Examples include Gidaris et al. (2018), Larsson et al. (2017), or Noroozi and Favaro (2016). While these pretext tasks might not be directly related to the final downstream task, they encourage models to learn meaningful representations that can be transferred and fine-tuned for specific applications.

More recently, transformer-based models have emerged as powerful tools for self-supervised learning in computer vision. He et al. (2021), inspired by the success of Devlin et al. (2018) in natural language processing, employ a masking strategy where portions of the input image are masked, and the model is trained to reconstruct the missing pixels. Bao et al. (2021) and Peng et al. (2022) leverage masked image modeling for pretraining, achieving state-of-the-art results on various image classification benchmarks. Chen et al. (2021) combines contrastive and generative objectives by incorporating masked image modeling into the momentum contrastive (MoCo) framework.

This paper investigates the effectiveness of several of these self-supervised learning approaches, including Simple Framework for Contrastive Learning of Visual Representations (SimCLR), Swapping Assignments between Views (SwAV), and BEit, for the specific task of building roof type classification from aerial imagery.

Building Roof Type Classification. Classifying building roof types from aerial imagery is a crucial task in remote sensing, with applications in urban planning, 3D city modeling, disaster response, and energy modeling. However, developing accurate and robust methods for this task is challenging, primarily due to the difficulty of obtaining large, high-quality labeled datasets. Traditional approaches to roof type classification often relied on handcrafted features and rule-based systems (Pu and Vosselman, 2009). However, these methods typically struggle to generalize to diverse roof structures and are sensitive to variations in image quality and scene complexity.

With the advent of deep learning, supervised learning approaches using convolutional neural networks (Convolutional Neural Network (CNN)) have become dominant in building roof type classification. Interestingly, Buyukdemircioglu et al. (2021) showed that even a shallow CNN architecture can achieve acceptable accuracy for roof type classification from very high-resolution orthophotos, particularly when dealing with sparse data. They compared their shallow CNN to fine-tuned versions of VGG-16 (Simonyan and Zisserman, 2015), EfficientNetB4 (Tan and Le, 2019), and ResNet-50 (He et al., 2015), finding that while the shallow model had slightly lower

overall accuracy, it still performed adequately for many applications. This suggests that computationally efficient CNN models can be effective for roof type classification. Huang et al. (2023) introduced the Urban Building Classification (UBC) V2-A Benchmark, a large-scale dataset for global building detection and fine-grained classification from satellite imagery, demonstrating the capabilities of CNNs for this task. While promising, these supervised methods require substantial amounts of labeled data for training.

Despite the advancements made with deep learning, the reliance on large labeled datasets remains a significant bottleneck for supervised approaches. Creating such datasets is labor-intensive, time-consuming, and expensive, especially when considering the global diversity of roof structures and the challenges posed by dense urban environments. This limitation motivates the exploration of self-supervised learning methods, which can learn effective representations from unlabeled data, as investigated in this paper.

EfficientNets for Object Recognition. EfficientNets have evolved as a compelling alternative to more resource-intensive architectures like ResNets (He et al., 2015). Introduced by Tan and Le (2019), EfficientNets offer a balanced trade-off between computational complexity and performance. Their unique scaling method optimizes depth, width, and resolution, thereby enabling performance gains with fewer parameters and less computational cost. Given the specific challenges posed by aerial object recognition, wherein the availability of computational resources can be a limiting factor, EfficientNets stand as a viable alternative to traditional architectures. Recent works have started to examine the integration of EfficientNets within self-supervised frameworks (Gao et al., 2021), which has influenced our choice for this study.

Domain-Specific Pretraining. While pretraining models on large, diverse datasets like ImageNet (Deng et al., 2009) has been a common practice, there is a growing realization that domain-specific pretraining can offer distinct advantages (Gu et al., 2020). This is particularly true in specialized fields such as aerial object recognition, where the statistical properties of images can differ considerably from generic datasets. By utilizing a domain-specific dataset for pretraining - like the AID (Xia et al., 2016) in our case - the model is better aligned with the characteristics inherent to aerial imagery, thereby potentially improving performance and generalizability.

While previous research has explored self-supervised learning, EfficientNets, and domain-specific pretraining separately, this paper presents a novel framework that integrates these advancements to address the challenges of accurate and efficient building roof type classification from aerial imagery, especially in scenarios with limited labeled data.

2. Methodology and Experiments

This section presents the methodology employed for classifying building roof types from aerial imagery using self-supervised learning. We begin with an overview of our two-phase approach: pretraining and linear evaluation. We then provide a detailed description of the datasets used for both phases, followed by explanations of the self-supervised learning methods and the architecture choices employed for pretraining. Finally, we outline the linear evaluation protocol used to assess the performance of our approach on the roof type classification task.




Figure 2. Visual representation of representative architecture configurations employed in the study: ResNet50, EfficientNetB3, and EfficientNetB3 enhanced with CBAM. These architectures serve as examples of the diverse backbones explored, ranging from ResNet34 to ResNet50 and EfficientNetB0 to B3 with and without CBAM.

2.1 Pretraining with Self-Supervised Learning

AID for Pretraining. The AID (Xia et al., 2016) is a large-scale dataset explicitly designed for aerial scene classification, making it suitable for pretraining models on a diverse range of aerial imagery. It consists of 10,000 images categorized into 30 different scene classes, including airports, beaches, bridges, and various urban and rural landscapes. The images in AID have a spatial resolution of 600×600 pixels. No additional preprocessing was applied to the images beyond the standard normalization procedures provided by the *PyTorch* library. We selected AID for pretraining because it contains a wide variety of aerial scenes, enabling our models to learn generic features and representations relevant to the domain of aerial imagery. We hypothesize that pretraining on AID, a domain-specific dataset, will be more beneficial for the downstream task of roof type classification compared to pretraining on a more general-purpose dataset like ImageNet.

Backbone Architectures. For this study, we chose EfficientNets (Tan and Le, 2019) as our backbone architectures for their computational efficiency and high performance in image recognition tasks. EfficientNets are a family of convolutional neural networks that achieve state-of-the-art accuracy with significantly fewer parameters and computational resources compared to traditional CNNs like ResNets (He et al., 2015). Their effi-

ciency is particularly important for processing large-scale aerial imagery datasets, which are often computationally demanding. We evaluated EfficientNet models ranging from B0 to B3, offering a trade-off between model complexity and accuracy. For comparison, we also included ResNet architectures (ResNet34 and ResNet50) in our experiments. Figure 2 provides a visual illustration of the ResNet-50 architecture, EfficientNet-B3 architecture, and its variant enhanced with the CBAM, which will be described in the next section.

Attention Mechanism. To further enhance the feature extraction capabilities of our EfficientNet models, we incorporate the CBAM (Woo et al., 2018). CBAM is a lightweight attention mechanism that can be seamlessly integrated into CNN architectures. It refines the feature maps by selectively attending to important spatial and channel-wise information. In our approach, we append CBAM to the end of the EfficientNet backbone before the classification head. By integrating CBAM, we aim to improve the ability of the model to focus on the most relevant features for roof type classification, potentially leading to higher accuracy.

Self-Supervised Learning Methods. We explore a variety of self-supervised learning methods to pretrain our EfficientNet models. These methods can be broadly categorized into two groups: contrastive learning methods and transformer-based

Method	ResNet		EfficientNet				EfficientNet+CBAM			
	34	50	B0	B1	B2	B3	B0	B1	B2	B3
DenseCL	81.8 ± 0.6	72.1 ± 0.4	74.9 ± 1.9	75.4 ± 0.2	72.2 ± 0.7	66.1 ± 0.8	76.9 ± 1.01	71.5 ± 0.3	70.8 ± 3.01	66.79 ± 2.6
MoCoV3	88.8 ± 0.39	93.7 ± 0.2	83.3 ± 3.1	72.6 ± 3.2	77.9 ± 4.8	76.6 ± 0.85	79.9 ± 0.6	77.7 ± 4.6	75.4 ± 3.3	76.7 ± 3.6
SimCLR	83.6 ± 0.2	86.8 ± 0.1	94.9 ± 0.1	94.6 ± 0.2	94.8 ± 0.1	94.7 ± 0.5	95.1 ± 0.26	95.5 ± 0.07	94.9 ± 0.1	94.7 ± 0.12
SwAV	72.5 ± 3.8	72.1 ± 4.2	50.4 ± 3.94	73.4 ± 0.58	72.2 ± 0.43	72.0 ± 0.89	76.8 ± 0.5	78.6 ± 0.2	75.7 ± 0.2	75.8 ± 0.3

Table 1. Classification Accuracies for Building Roof Types Using CNN-Based Self-Supervised Models pretrained on AID: A Comparison Across ResNet, EfficientNet, and CBAM-Enhanced EfficientNet Encoders. The best performing models from each of the methods are highlighted.




Figure 3. Sample images from the dataset illustrating the four roof types: Gable, Hip, Flat, and Complex.

methods.

- Contrastive methods: SimCLR, MoCoV3, SwAV employ contrastive learning, where models learn to distinguish between augmented views of the same image (positive pairs) and different images (negative pairs). For these methods, we employed the LARS optimizer (You et al., 2017) with a learning rate initially determined using a linear scaling rule. The learning rate increased during the first 10 epochs and decayed following a cosine schedule. A weight decay of $1e - 6$ was applied. For DenseCL, we used Stochastic Gradient Descent (SGD) with a learning rate of 0.03, a momentum of 0.9, and a weight decay of $1e - 4$.
- Transformer-based methods: For BEit, BEitV2, Masked Autoencoder (MAE), we used the AdamW optimizer (Loshchilov and Hutter, 2019) with a weight decay of 0.1 and an initial learning rate of $1e - 4$ with a warmup period of 40 epochs.

We pretrain our models using each of these methods for 300 epochs on the AID dataset. The specific hyperparameters used for each method, including optimizer, learning rate, batch size, and data augmentation techniques, will be provided in the following section on Linear Evaluation Protocol. All experiments were conducted using the MMPreTrain library (Contributors, 2023) on an NVIDIA Titan RTX GPU with 24GB of memory.

Method	ViTBase
BEit	95.1 ± 0.2
BEitV2	93.7 ± 0.1
MAE	77.1 ± 0.3
MoCoV3	73.3 ± 4.9

Table 2. Classification Accuracies for Building Roof Types Using Transformer-Based Self-Supervised Models pretrained on AID: A Comparison on ViTBase Encoder.

2.2 Linear Evaluation

Dataset and labeling. To evaluate the effectiveness of the pre-trained EfficientNet models for building roof type classification, we perform a linear evaluation. This involves fine-tuning the pretrained models on a labeled dataset of aerial images and assessing their performance on the roof type classification task. We utilized building roof vectorization dataset (Hensel et al., 2021), originally lacking roof type information. We manually annotated the images to classify them into four roof types: Gable, Hip, Flat, and Complex, as illustrated in Figure 3. The images are RGB orthophotos with a spatial resolution of 0.1 meters per pixel. The labeled dataset is divided into a training set of 7500 images and a validation set of 765 images.

Data Augmentation and Training Details. For data augmentation, random resize crops and horizontal flips were applied. The models were then linearly evaluated for 100 epochs on this dataset using the following optimization strategies:

- For SimCLR, SwAV, and MoCoV3, the LARS optimizer was used with a learning rate of 0.6 and a momentum of 0.9.
- For DenseCL, we used SGD with a learning rate of 0.1, a momentum of 0.9, and a weight decay of $1e - 4$.
- For BEit and MAE, we used SGD with a learning rate of 0.1, a momentum of 0.9, and a weight decay of $1e - 3$.

The learning rates for all models were decayed according to cosine schedule.

3. Results and Discussion

3.1 Performance of Convolutional Models

We evaluated various convolutional models, including ResNet and EfficientNet variants, with and without the CBAM. As shown in Table 1, incorporating CBAM consistently improved the performance of the EfficientNet models. For instance, DenseCL with EfficientNet-B0 and CBAM achieved an accuracy of 76.9%, compared to 74.9% for the same model without CBAM.




Figure 4. Performance visualization: BEit (Row 1) and SimCLR (Row 2) on a manually labeled Roof Graph subsets, pretrained on AID.

Input image	ResNet		EfficientNet				EfficientNet+CBAM			
	34	50	B0	B1	B2	B3	B0	B1	B2	B3

Table 3. Class Activation Maps generated by the SimCLR model pretrained on the Aerial Image Dataset. Each row represents a different roof category: the first row contains Complex, the second Flat, the third Gable, and the fourth Hip.

This trend suggests that CBAM’s selective attention mechanism effectively enhances feature extraction for building roof type classification.

3.2 Comparative Performance: SimCLR and Transformer-Based Models

Notably, the SimCLR method combined with EfficientNet and CBAM consistently outperformed other architectures, even exceeding the accuracy of transformer-based models like BEit and BEitV2 in certain cases, as shown in Table 1 and Table 2. For example, SimCLR with EfficientNet-B1 and CBAM achieved a remarkable accuracy of 95.5%, representing a substantial 12% improvement over its ResNet-pretrained counterpart. This result underscores the effectiveness and efficiency of combining

EfficientNet backbones with CBAM for self-supervised learning tasks in building roof type classification.

The Class Activation Maps generated by SimCLR dramatically validate our observation about its exceptional performance when combined with EfficientNet and CBAM as shown in Table 3. ResNet 34 and 50 exhibit lower activations compared to all variants of EfficientNet. Most notably, the CBAM-enhanced EfficientNets show the highest activations, supporting our claim that SimCLR optimized with EfficientNet and CBAM exceeds its ResNet versions in performance.

Method	ViTBase/ResNet50
BEit	96.4 ± 0.7
BEitV2	96.8 ± 0.3
DenseCL	94.1 ± 0.4
MAE	77.1 ± 2.0
MoCoV3 ResNet	95.8 ± 0.08
MoCoV3 ViT	92.9 ± 0.2
SimCLR	80.78 ± 0.2
SwAV	87.9 ± 0.3

Table 4. Classification Accuracies for Building Roof Types Using Transformer-Based and CNN-Based Self-Supervised Models pretrained on ImageNet-1K dataset: A Comparison on ViTBase and ResNet50 Encoder.

Method	Braunschweig dataset	Roof graph Subset
BEit	90.7	81
SimCLR	87.2	75.2

Table 5. Classification Accuracies of SimCLR (EfficientNet-B1 + CBAM) and BEit model pretrained on AID dataset on test datasets.

3.3 Impact of Pretraining Datasets: ImageNet-1k vs. Domain-Specific

While models pretrained on the ImageNet-1k dataset generally demonstrated good performance, as shown in Table 4, certain methods exhibited significant improvements when pretrained on the domain-specific AID. SimCLR and MAE, in particular, showed substantial gains with domain-specific pretraining. For instance, SimCLR’s accuracy increased from 80.78% with ImageNet-1k pretraining to 83.6% with AID pretraining for ResNet50 backbone. This finding underscores the importance of leveraging domain-specific data for pretraining, especially in specialized applications like building roof type classification.

3.4 Generalization to Test Datasets

To assess the generalizability of our approach, we evaluated the two best-performing models, SimCLR (with EfficientNet-B3 + CBAM) and BEit, both pretrained on the AID dataset, on two independent test sets. The first test set, the Braunschweig dataset, consists of approximately 350 building polygons extracted from the Braunschweig city dataset (acquired from CrowdAI) and manually labeled for roof type information. The images are 3-band RGB orthophotos with a spatial resolution of 0.3 meters from our flight campaign. The second test set, the Roof Graph subset, comprises 300 manually labeled images randomly selected from the Roof Graph dataset (Hensel et al., 2021), which covers the city of Detmold, Germany. The images are 3-band RGB orthophotos with a resolution of 0.1 meters per pixel.

As shown in Table 5, BEit achieved higher accuracies (90.7% and 81%) on both the Braunschweig and Roof Graph subsets, respectively, compared to SimCLR (87.2% and 75.2%). These results demonstrate the robust performance of both methods on unseen data, supporting the generalization capabilities of our self-supervised pretraining approach. Figures 4 and 1 provide visual examples of the classification results on these test sets.

3.5 Computational Efficiency

In practical remote sensing applications, computational efficiency is crucial, especially when dealing with large-scale aer-

ial imagery datasets. Table 6 compares the computational requirements of ResNet50, EfficientNet-B3, and EfficientNet-B3 with CBAM. EfficientNet-B3 demonstrates significant advantages, requiring substantially fewer parameters (13.255M vs. 27.969M) and FLOPS (0.992G vs. 4.109G) compared to ResNet50 while achieving comparable or even better accuracy. The addition of CBAM only marginally increases the computational cost of EfficientNet-B3. These results highlight EfficientNet’s efficiency, making it highly suitable for large-scale aerial imagery processing tasks where resource constraints and processing time are critical considerations.

Arch.	Params (M)	FLOPS (G)	Avg. iter. time (s/iter)
ResNet50	27.969	4.109	0.967
EfficientNetB3	13.255	0.992	0.3588
EfficientNetB3 + CBAM	13.274	0.992	0.3597

Table 6. **Compute requirements.** We present the total number of parameters, FLOPS (floating point operations per second), and the average iteration time required for the three architectures.

4. Conclusion

Our results demonstrate the significant potential of self-supervised learning for building roof type classification from aerial imagery. The combination of EfficientNet architectures and the Convolutional Block Attention Module (CBAM) proved particularly effective. Incorporating CBAM consistently enhanced the performance of EfficientNet models, highlighting the benefits of selective attention for this task. Moreover, EfficientNet consistently outperformed ResNet, achieving higher accuracy while requiring substantially fewer parameters and computational resources as demonstrated in 3.5.

Among the self-supervised methods evaluated, Simple Framework for Contrastive Learning of Visual Representations (SimCLR), when combined with EfficientNet-B3 and CBAM, achieved the highest classification accuracy. This method not only outperformed other Convolutional Neural Network (CNN)-based approaches but also demonstrated comparable or superior performance to transformer-based models in certain cases.

Our findings strongly emphasize the advantage of domain-specific pretraining. Models pretrained on the Aerial Image Dataset (AID) consistently achieved higher accuracy than those pretrained on the general-purpose ImageNet-1k dataset. This highlights the importance of leveraging datasets that reflect the unique characteristics of aerial imagery for specialized tasks like building roof type classification.

The promising results obtained in this study lay the foundation for further exploration of self-supervised learning in remote sensing. Future work could investigate the application of these methods to more complex roof type categories or explore the integration of other advanced architectures and attention mechanisms. Additionally, expanding the domain-specific pretraining dataset with more diverse aerial imagery could further improve performance and generalization capabilities.

References

- Bao, H., Dong, L., Wei, F., 2021. BEiT: BERT Pre-Training of Image Transformers. *CoRR*, abs/2106.08254.
- Buyukdemircioglu, M., Can, R., Kocaman, S., 2021. DEEP LEARNING BASED ROOF TYPE CLASSIFICATION USING VERY HIGH RESOLUTION AERIAL IMAGERY. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLIII-B3-2021, 55–60.
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A., 2020. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. *CoRR*, abs/2006.09882.
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G. E., 2020. A Simple Framework for Contrastive Learning of Visual Representations. *CoRR*, abs/2002.05709.
- Chen, X., Xie, S., He, K., 2021. An Empirical Study of Training Self-Supervised Vision Transformers. *CoRR*, abs/2104.02057.
- Contributors, M., 2023. Openmmlab's pre-training toolbox and benchmark. <https://github.com/open-mmlab/mmpretrain>.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Devlin, J., Chang, M., Lee, K., Toutanova, K., 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR*, abs/1810.04805.
- Gao, Y., Zhuang, J., Li, K., Cheng, H., Guo, X., Huang, F., Ji, R., Sun, X., 2021. DisCo: Remedy Self-supervised Learning on Lightweight Models with Distilled Contrastive Learning. *CoRR*, abs/2104.09124.
- Gidaris, S., Singh, P., Komodakis, N., 2018. Unsupervised Representation Learning by Predicting Image Rotations. *CoRR*, abs/1803.07728.
- Gu, Y., Tinn, R., Cheng, H., Lucas, M., Usuyama, N., Liu, X., Naumann, T., Gao, J., Poon, H., 2020. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. *CoRR*, abs/2007.15779.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R. B., 2021. Masked Autoencoders Are Scalable Vision Learners. *CoRR*, abs/2111.06377.
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R., 2020. Momentum contrast for unsupervised visual representation learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9726–9735.
- He, K., Zhang, X., Ren, S., Sun, J., 2015. Deep Residual Learning for Image Recognition. *CoRR*, abs/1512.03385.
- Hensel, S., Goebbels, S., Kada, M., 2021. Building Roof Vectorization with Ppgnet. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 46W4, 85–90.
- Huang, X., Chen, K., Tang, D., Liu, C., Ren, L., Sun, Z., Hänsch, R., Schmitt, M., Sun, X., Huang, H., Mayer, H., 2023. Urban Building Classification (UBC) V2—A Benchmark for Global Building Detection and Fine-Grained Classification From Satellite Imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 61, 1–16.
- Larsson, G., Maire, M., Shakhnarovich, G., 2017. Colorization as a Proxy Task for Visual Understanding. *CoRR*, abs/1703.04044.
- Loshchilov, I., Hutter, F., 2019. Decoupled weight decay regularization.
- Noroozi, M., Favaro, P., 2016. Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles. *CoRR*, abs/1603.09246.
- Peng, Z., Dong, L., Bao, H., Ye, Q., Wei, F., 2022. Beit v2: Masked image modeling with vector-quantized visual tokenizers.
- Pu, S., Vosselman, G., 2009. Knowledge based reconstruction of building models from terrestrial laser scanning data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 64(6), 575–584.
- Simonyan, K., Zisserman, A., 2015. Very deep convolutional networks for large-scale image recognition.
- Tan, M., Le, Q. V., 2019. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. *CoRR*, abs/1905.11946.
- Woo, S., Park, J., Lee, J., Kweon, I. S., 2018. CBAM: Convolutional Block Attention Module. *CoRR*, abs/1807.06521.
- Xia, G., Hu, J., Hu, F., Shi, B., Bai, X., Zhong, Y., Zhang, L., 2016. AID: A Benchmark Dataset for Performance Evaluation of Aerial Scene Classification. *CoRR*, abs/1608.05167.
- You, Y., Gitman, I., Ginsburg, B., 2017. Scaling SGD Batch Size to 32K for ImageNet Training. *CoRR*, abs/1708.03888.