

SeiT++: Masked Token Modeling Improves Storage-efficient Training

Minhyun Lee^{1,†}  Song Park^{2,†}  Byeongho Heo²  Dongyoon Han² 
Hyunjung Shim^{3,*} 

¹ Yonsei University ² NAVER AI Lab ³ KAIST AI

[†] Equal contribution. Work done at NAVER AI Lab.

Abstract. Recent advancements in Deep Neural Network (DNN) models have significantly improved performance across computer vision tasks. However, achieving highly generalizable and high-performing vision models requires expansive datasets, resulting in significant storage requirements. This storage challenge is a critical bottleneck for scaling up models. A recent breakthrough by SeiT proposed the use of Vector-Quantized (VQ) feature vectors (*i.e.*, tokens) as network inputs for vision classification. This approach achieved 90% of the performance of a model trained on full-pixel images with only 1% of the storage. While SeiT needs labeled data, its potential in scenarios beyond fully supervised learning remains largely untapped. In this paper, we extend SeiT by integrating Masked Token Modeling (MTM) for self-supervised pre-training. Recognizing that self-supervised approaches often demand more data due to the lack of labels, we introduce TokenAdapt and ColorAdapt. These methods facilitate comprehensive token-friendly data augmentation, effectively addressing the increased data requirements of self-supervised learning. We evaluate our approach across various scenarios, including storage-efficient ImageNet-1k classification, fine-grained classification, ADE-20k semantic segmentation, and robustness benchmarks. Experimental results demonstrate consistent performance improvement in diverse experiments, validating the effectiveness of our method. Code is available at <https://github.com/naver-ai/seit>.

1 Introduction

Recent advancements in Deep Neural Network (DNN) [17, 56] have significantly improved their performance across various computer vision tasks. However, achieving highly generalizable and high-performing vision models demands an extensive dataset, often containing billions of data points [28, 41, 53] with large storage requirements. For example, the LAION-5B [50] dataset requires a storage size of 240TB. This storage challenge emerges as a critical bottleneck for scaling up vision models. Consequently, optimizing storage efficiency becomes essential for the practical advancement of scaled-up vision models.

* Hyunjung Shim is a corresponding author.

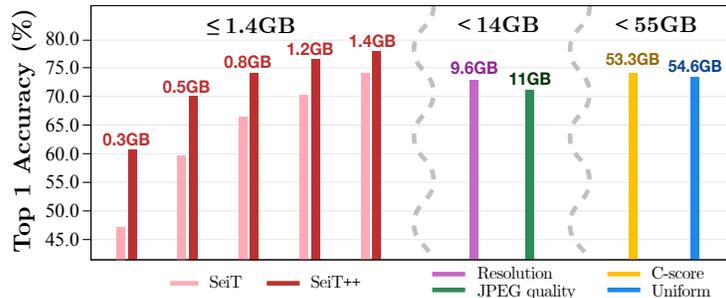


Fig. 1: Over 70% top-1 accuracy is achievable with just 1GB data. On ImageNet-1k, we visualize the trade-off of training data storage vs. top-1 accuracy using the fixed ViT-B/16 for a controlled comparison. Each accuracy metric is individually trained with different data types. While the entire ImageNet-1k dataset requires approximately 140GB for training with images, our approach demonstrates significant storage efficiency over competitors.

Images require more storage than texts due to their intrinsic complexity in converting continuous natural light signals into discrete pixels for computer representation. However, a considerable portion of image data corresponds to redundant or unnecessary details, such as repeated patterns or imperceptible details [47]. These properties lead to large storage requirements and also induce negative impacts on visual recognition tasks as mentioned in [1, 11, 21, 22, 40, 51]. Therefore, several attempts [9, 29, 45, 69, 70] have been made to reduce storage requirements by optimizing the efficiency of image data. A standard approach to reduce storage requirements is to decrease the total number of training data by removing less important samples [9, 29, 45] or condensing the data into smaller synthetic sets [69, 70].

However, as SeiT [44] illustrates, these approaches may incur a significant performance drop or become infeasible for large-scale datasets due to their complexity. Alternatively, some methods aim to reduce the size of each image through compression techniques or resolution adjustments [50]. However, they still face a considerable performance gap with the original dataset. Recently, SeiT [44] has made a significant breakthrough in this area by enabling the direct learning of ViTs [17] from discrete representations (*i.e.*, tokens) pre-extracted by an offline tokenizer [66]. By using vector-quantized features instead of traditional pixel images, SeiT achieves 90% of the performance of a full storage model while requiring only 1% of the storage.

While SeiT [44] has demonstrated the high potential of token-based training, a performance gap with the image-based models still persists, and a prospective solution – exploration beyond fully supervised learning scenarios – remains untapped. To leverage the advantages of large-scale data beyond human-annotated datasets, exploring label-free learning methods is imperative. In this paper, we investigate the compatibility of SeiT with self-supervised learning approaches in token-based frameworks. Among many different self-supervised methods [6, 7, 23],

Masked Image Modeling (MIM) [24, 64] derived from Masked Language Modeling (MLM) [15] has brought significant success and has quickly become a popular approach for learning visual representations.

Inspired by this, we introduce Masked Token Modeling (MTM), a pre-training method that performs token masking and reconstruction in token-based frameworks, and validate its effectiveness for storage efficiency on various benchmarks. To the best of our knowledge, this is the first approach for learning an MLM-based model directly from offline tokens, offering the potential to optimize storage efficiency for large-scale training tasks. As previous studies [12, 13, 56] have highlighted, data augmentation plays a crucial role in large-scale model training and is a fundamental element in various training setups [6, 7, 23, 24, 37, 56, 62, 64]. However, applying existing data augmentation techniques [13, 67, 71] designed for pixel images to tokens poses a significant challenge due to the input domain shift. Therefore, addressing this challenge is key to overcoming major performance bottlenecks and improving the storage efficiency of token-based frameworks.

To this end, we introduce TokenAdapt and ColorAdapt, two novel augmentation strategies designed for token-based training. TokenAdapt facilitates the integration of traditional pixel-based augmentation policies by transforming token embeddings into a feature space compatible with pixel-based augmentations. After applying augmentation to the transformed token embeddings, the augmented token embeddings are then reverted to their original token space. ColorAdapt, inspired by Adaptive Instance Normalization (AdaIN [27]), alters the color attributes of token embeddings while preserving object structure by adjusting their statistics. By integrating TokenAdapt and ColorAdapt into our token augmentation process, we enhance the effectiveness of token-based training across both fully supervised and unsupervised learning scenarios, demonstrating significant improvements in model performance.

Our SeiT++ framework combines Masked Token Modeling (MTM) with the proposed token augmentation strategies, TokenAdapt and ColorAdapt. We validate the effectiveness of SeiT++ on various scenarios: (1) storage-efficient ImageNet-1k classification, (2) transfer learning on fine-grained classification, (3) ADE-20k semantic segmentation, and (4) evaluation on robustness benchmarks. As shown in Figure 1, SeiT++ consistently outperforms comparison methods in storage-efficient ImageNet-1k classification with limited storage sizes, achieving a top-1 accuracy of 77.8% with only 1.4GB of data. These results confirm that token-based training can take advantage of self-supervised learning to improve storage efficiency. We also demonstrate that our proposed token augmentations significantly contribute to storage efficiency by augmenting data effectively. We reveal that our token-based augmentation methods basically yield significant boosts in supervised learning and further amplify improvements in MTM learning. Furthermore, we present how our strategies enhance the generalizability of trained models, as evidenced by our experimental results on fine-grained classification and robustness evaluation. Lastly, the applicability of our method to an alternative input format (*e.g.*, DCT coefficients) demonstrates its extensibility.

2 Related Work

Learning with Tokenization. The use of discrete tokens for training vision models has recently gained attention, particularly in self-supervised representation learning. BeiT [2] employs the Masked Image Modeling (MIM) framework to recover discrete tokens from masked images and CIM [19] utilizes tokens in Corrupted Image Modeling. MAGE [34] demonstrates the incorporation of tokens from VQGAN [18] for both generative training and representation learning within a unified training framework. However, existing studies mainly focus on self-supervised representation learning and rely on online tokenization, converting images to tokens during each iteration of model training. This approach incurs notable memory consumption and computational costs, limiting the full exploitation of the storage efficiency of discrete tokens.

Storage-efficient Vision Training. Efficient model learning is a crucial research area, with studies focusing on optimizing target models [35, 36] or minimizing training data and labels [39, 42]. Storage-efficient training aims to reduce the dataset storage size. To this end, dataset distillation (DD) [60] creates a condensed dataset by leveraging the knowledge of the original dataset. Recent research [33, 48, 49, 69, 70] highlights the effectiveness of dataset distillation in improving training efficiency. However, its high computational complexity poses challenges, particularly when applied to large-scale datasets like ImageNet-1k [14]. Sampling-based methods [9, 29, 45] address this by selecting a subset of the most representative samples from the full dataset. Paul *et al.* [45] identify important samples very early in training. Jiang *et al.* [29] utilize a consistency score based on statistics collected during training. However, the diverse sample selection is not guaranteed, particularly in low data regimes, leading to suboptimal performance compared to dataset distillation (DD) methods [70]. To tackle these issues, SeiT [44] proposes a storage-efficient vision training framework by leveraging the storage efficiency of tokens. They utilize the ViT-VQGAN [66] tokenizer to extract tokens from images and store them for vision model training. As a result, they outperform other storage-efficient training methods, significantly reducing the dataset storage size.

Data Augmentation. Data augmentation serves as a fundamental building block to improve the model performance across diverse computer vision tasks. Various data augmentation techniques have been employed for improved performance [16, 71] or model robustness [20, 38]. Mixup [68] introduces convex combinations of pairs of images and their labels to create a smoother decision boundary. CutMix [67] replaces certain regions of an image with a patch from another image, enhancing the model’s localization ability. AutoAugment [12] leverages reinforcement learning to automate optimal augmentation policy discovery. RandAugment [13] introduces a potent policy based on random combinations of different data augmentation methods. DeiT [56] emphasizes the crucial role of data augmentation in vision transformers for improved model performance.

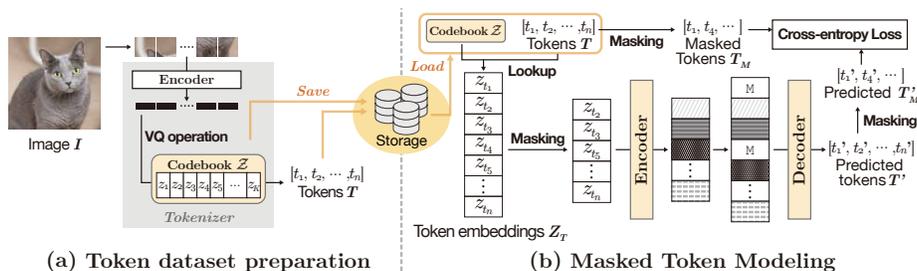


Fig. 2: Masked Token Modeling (MTM) pipeline. MTM is a self-supervised learning approach in token-based frameworks. (a) The tokenized dataset is saved on storage before model training. Then, (b) using only the pre-stored tokens, a storage-efficient vision model (MTM) is trained without relying on labeled datasets.

Previous studies extensively explore the effective data augmentation policies for pixel-based vision model training. However, we observe that pixel-based data augmentations are not well-suited for tokens, posing a critical performance bottleneck. Motivated by this, we tackle the data augmentation challenge, which has not been explored in previous studies for vision model training with tokens.

3 Method

3.1 Preliminary: Storage-efficient Vision Training (SeiT)

SeiT [44] aims to develop a scalable and high-performance vision classifier with minimal storage requirements. It introduces a storage-efficient vision training framework that enables the direct learning of ViT [17] from discrete representations (*i.e.*, tokens) pre-extracted by an offline tokenizer. In addition, SeiT further improves training by incorporating two token-specific augmentations: Token-EDA, drawing inspiration from easy data augmentation for language [61], and Emb-Noise, which introduces Gaussian noise into discrete representations. By using vector-quantized features instead of pixel images, SeiT retains 90% of the precision of a model trained on full-pixel images while using only 1% of the storage space. SeiT consists of two parts: (1) preparing the token dataset and (2) training a model using the token dataset.

Preparing the Token Dataset. SeiT extracts tokens using the ImageNet-1k-trained ViT-VQGAN tokenizer [66] to prepare the token dataset for training. As shown in Figure 2(a), each pixel image I is transformed into a sequence of tokens $T = [t_1, t_2, \dots, t_n]$. Since the tokens are stored in quantized form, each token means an embedding vector in a codebook $\mathcal{Z} = \{z_k \in \mathbb{R}^d | k = 1, \dots, K\}$. *i.e.*, $t_i \in \mathbb{Q}$ where $\mathbb{Q} = \{1, 2, \dots, K\}$ for a codebook with size K . Here, given the tokens T and the codebook \mathcal{Z} , the corresponding token embeddings $Z_T = [z_{t_1}, z_{t_2}, \dots, z_{t_n}]$ can be computed by lookup process. This conversion compresses each image

$I \in \mathbb{R}^{H \times W \times 3}$ into a more compact token sequence $T \in \mathbb{R}^n$, where n is much smaller than $H \times W \times 3$, thus significantly reducing storage requirements. For example, storing the full-pixel ImageNet-1k requires 140GB, but its tokenized form needs only 1.4GB. Notably, the tokens are stored on storage before model training with the codebook \mathcal{Z} . Thus, during the training phase, the pre-stored tokens are used directly as input instead of the original pixel images.

Training Pipeline. To train a vision classifier using tokens, SeiT initially loads the token dataset and applies Token-EDA to the tokens. The tokens are then converted to a one-hot format and randomly resized and cropped. Subsequently, the processed one-hot tokens are transformed into token embeddings using a pre-trained codebook. Techniques such as CutMix and Embedding-noise are applied to these token embeddings, and the resulting augmented tensor is fed into the model for training. Through the proposed framework, SeiT demonstrates the high potential of token-based training, but it lacks exploration beyond fully supervised learning scenarios. In addition, while token-specific augmentation is proposed, the types of token augmentation are limited, and the analysis of data augmentation compared to pixel-based training is insufficient.

3.2 Masked Token Modeling

In this section, we introduce Masked Token Modeling (MTM), a pre-training method in token-based frameworks without relying on labeled datasets. MTM shares a similar concept to Masked Language Modeling (MLM) [15], which works by reconstructing original tokens from their masked versions, employing an autoencoding strategy. Figure 2(b) presents an overview of how our MTM operates.

Masking. We employ a variable masking ratio that follows a truncated normal distribution, as investigated in previous studies [4,5,34]. The tokens are randomly dropped according to the sampled masking ratio, and then the *visible, unmasked* token embeddings are fed into the model. By processing only a subset of token embeddings in the encoder, we significantly reduce the overall pre-training time and memory consumption, which aligns with the findings of MAE [24].

MTM Encoder & Decoder. We use ViT [17] as our encoder because it is able to process both the *visible* token embeddings and the original token embeddings, seamlessly accommodating their different lengths. During training, the encoder takes the *visible* token embeddings and outputs latent features that become the decoder input. Before decoding, the latent features are padded to the length of the original token embeddings using mask tokens M . Each mask token is a shared, learned vector indicating the presence of a missing token to be predicted. The decoder then estimates the original tokens T from the padded features.

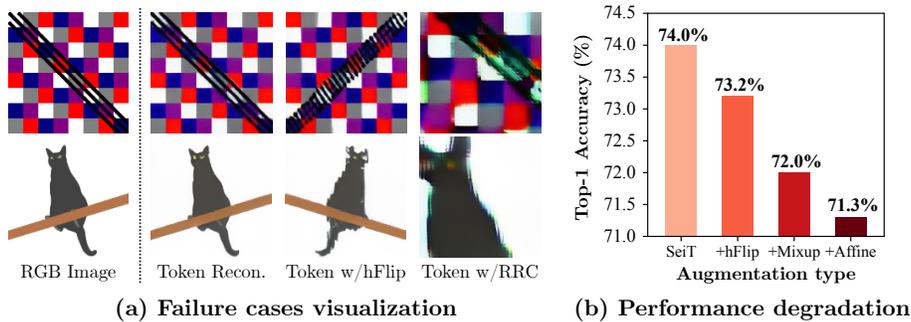


Fig. 3: Data augmentations with tokens. Each ViT-VQGAN decoded image is reconstructed from a given RGB image after undergoing specific data augmentations. We observe that naively adopting these methods results in incorrect tokenization: 1) Token w/ hFlip demonstrates spatial information collapse during tokenization; 2) Token w/ RRC shows interdependence between neighboring token embeddings. We note the reconstructed images fail to preserve the images’ details, suggesting that this incurs ineffectiveness of tokenization. Furthermore, we report top-1 accuracies (ViT-B) on ImageNet-1k according to the augmentations applied during training. All the augmentations are de-facto default training setups for vision transformers [37, 54, 56].

Training Objective. MTM learns to accurately predict the original tokens of each masked token from the visible tokens. In particular, we use a cross-entropy loss between the decoder output and the ground truth tokens as the training objective, which is denoted by \mathcal{L}_{recon} :

$$\mathcal{L}_{recon} = \text{CE}(T'_M, T_M), \quad (1)$$

where T'_M denotes the predictions of the model for the masked tokens, and T_M corresponds to the ground truth tokens for T'_M . Following MAE [24], this loss is optimized only for the masked tokens.

3.3 Data Augmentation for Tokens

While Masked Token Modeling (MTM) improves the storage efficiency of the token-based framework, there is still a performance bottleneck in improving token-based training. Since self-supervised approaches often require more data due to the lack of labels, exploring data augmentation is important to take full advantage of MTM. In this section, we analyze why existing data augmentations designed for pixel images are ineffective in the token domain. We then present TokenAdapt and ColorAdapt as solutions for token-based data augmentation.

Limited Data Augmentations for Tokens. We explore the challenges of pixel-based augmentation in the token domain and identify two main reasons: (1) spatial information collapse during tokenization and (2) inter-token dependencies. The discrete tokenization process condenses a $n \times n$ 2D image patch

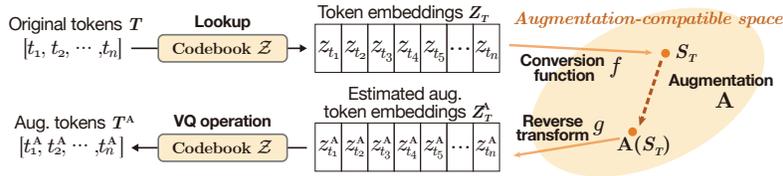


Fig. 4: TokenAdapt processing pipeline. TokenAdapt aims to enhance the compatibility of token embeddings with pixel-based data augmentations by converting them into augmentation-compatible space, applying augmentations, and reverting them back to the original token embedding space.

into a single d -dimensional 1D vector (*i.e.*, token embedding), causing a collapse of spatial information. This collapse hinders augmentations that rely on spatial details (*e.g.*, horizontal flip). Furthermore, unlike independent image pixels, token embeddings are inter-correlated. This complicates augmentations involving interpolation (*e.g.*, Resize or Mixup [68]), by introducing undesired artifacts.

Figure 3 shows the limitations of pixel-based augmentations in the token domain. A horizontal Flip (hFlip) distorts both the diagonal line (top row) and the cat’s silhouette (bottom row) due to spatial information collapse. This occurs because only the inter-token spatial relationship is flipped, while the inner-token spatial arrangement remains unchanged. Augmentations involving interpolation can lead to undesired artifacts, such as color leakage in the case of Random Resized Crop (RRC), due to interdependence between neighboring embeddings. Therefore, we suspect that the direct application of pixel-based augmentations to token embedding may negatively affect token-based training.

TokenAdapt. The concept of TokenAdapt is to establish a feature space that allows for conventional pixel-based augmentations while facilitating the conversion and inversion of pre-extracted tokens with minimal computational overhead. This approach offers several advantages for our TokenAdapt. Firstly, it enables the direct reuse of well-designed pixel-based data augmentation policies. Additionally, the augmented token can be easily computed through the conversion-augmentation-reverse process. Finally, it ensures that the feature spaces for both the original and augmented tokens are preserved identically. This allows for the seamless use of existing model architectures or pre-trained weights.

The goal of TokenAdapt is to design an augmentation function that converts the input tokens T to augmented tokens $T^A = [t_1^A, t_2^A, \dots, t_n^A]$. Note that the augmented tokens should also be quantized $t_i^A \in \mathbb{Q}$. Figure 4 depicts the overall pipeline of TokenAdapt. The initial step of TokenAdapt is transforming token embeddings Z_T , computed from T and codebook \mathcal{Z} , into a representation $S_T \in \mathbb{R}^{h \times w \times d}$ in augmentation-compatible space S . This transformation is achieved using a conversion function f , *i.e.*, $S_T = f(Z_T)$. We apply augmentation A at augmentation-compatible space S and build an augmented representation $S_T^A = A(S_T)$. Then, the augmented representation is reverted to

token embedding space by a reverse transformation g , $Z_T^{\mathbf{A}} = g(S_T^{\mathbf{A}})$. At last, the augmented tokens in embedding space $Z_T^{\mathbf{A}}$ are quantized to codebook index as $T^{\mathbf{A}} = \mathbf{q}_{\mathcal{Z}}(Z_T^{\mathbf{A}})$, where $\mathbf{q}_{\mathcal{Z}}$ indicates the vector quantization process. In summary, the augmentation process of TokenAdapt is as follows:

$$Z_T^{\mathbf{A}} = g(\mathbf{A}(f(Z_T))), \quad T^{\mathbf{A}} = \mathbf{q}_{\mathcal{Z}}(Z_T^{\mathbf{A}}). \quad (2)$$

The conversion function f and the reverse transformation g are learned from paired data of tokens $(T_x, T_{\mathbf{A}(x)})$, where T_x and $T_{\mathbf{A}(x)}$ denote tokens from an image x and its augmented image $\mathbf{A}(x)$, respectively. We use the cross-entropy loss between $Z_{T_x}^{\mathbf{A}}$ and $Z_{T_{\mathbf{A}(x)}}$ as our objective function \mathcal{L}_{TA} for learning our TokenAdapt module:

$$\mathcal{L}_{\text{TA}} = \text{CE}(Z_{T_x}^{\mathbf{A}}, Z_{T_{\mathbf{A}(x)}}). \quad (3)$$

TokenAdapt deals with geometric data augmentations (*e.g.*, flip, resize, crop, affine) and Mixup [68]. These augmentations are commonly used in various training methods, but they can cause a performance drop in token-based training, as shown in Figure 3(b).

Once trained, the TokenAdapt module is used to augment a sequence of tokens T into an augmented sequence $T^{\mathbf{A}}$ across different datasets and downstream tasks. Importantly, the augmentation-compatible feature space, which does not rely on specific semantic knowledge or alignment with downstream tasks, makes it dataset and task-agnostic. The inherent flexibility of TokenAdapt makes it highly generalizable, allowing direct application to other datasets without the need for additional training procedures.

ColorAdapt. As witnessed in Cubuk *et al.* [12], color-based augmentations are often more effective than geometric ones in certain datasets and are an important component of data augmentation strategies. For example, color-based augmentations are useful for dealing with various environmental changes caused by different lighting conditions (*e.g.*, day and night) or weather conditions (*e.g.*, snow). However, integrating them into TokenAdapt is challenging because most color-based augmentations are designed to work only in the pixel domain. To address this, we propose ColorAdapt, a color-based token augmentation inspired by the method [27]. Formally, the ColorAdapt function \mathcal{C} is defined as:

$$\mathcal{C}(Z_{T_1}, Z_{T_2}) = \sigma(Z_{T_2}) \left(\frac{Z_{T_1} - \mu(Z_{T_1})}{\sigma(Z_{T_1})} \right) + \mu(Z_{T_2}), \quad (4)$$

where $\mu(Z_T)$ and $\sigma(Z_T)$ are the mean and standard deviation of each token embedding Z_T across its spatial dimension, computed for each channel.

ColorAdapt mimics color-changing augmentations in the pixel domain while preserving the encoded object structure by transforming the global statistics within each channel of the token embedding. Figure 5 shows the images decoded from various color-augmented tokens. This illustrates how ColorAdapt affects the visual properties of the augmented tokens. By incorporating ColorAdapt

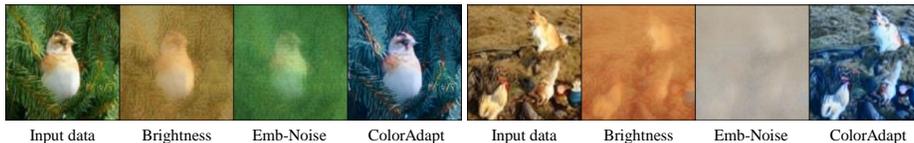


Fig. 5: ColorAdapt provides more reasonable color changes. We present ViT-VQGAN decoded images to verify the quality of tokenizations after color changes. We use the brightness function with a factor of 0.2 following the implementation [46]. Emb-Noise [44] is the color-based token augmentation. Notably, our ColorAdapt effectively preserves object structure in contrast to the failure of the counterparts.

Table 1: Storage-efficient ImageNet-1k classification. We report top-1 accuracies (ViT-B) and compression ratio on ImageNet-1k using various data storage reduction methods. Our SeiT++ shows significant storage efficiency compared to competitors.

Method	Input #	# of images	Top 1 Acc.	Dataset storage size
Full pixels	Image	1,281 k	81.8	140.0 GB
Uniform random sampling	Image	512 k	74.0	54.6 GB (39%)
C-score [29] based sampling		512 k	73.3	53.3 GB (38%)
Adjusting JPEG quality factor to 5		1,281 k	74.6	11.0 GB (8%)
Adjusting image resolution ($\times 0.2$)		1,281 k	75.2	9.6 GB (7%)
SRe2L [65]	Image	200 k	65.9	2.6 GB (2%)
G-VBSM [52]		100 k	63.7	1.6 GB (1%)
SeiT [44]	Token	1,281 k	74.0	1.4 GB (1%)
SeiT++		1,281 k	77.8	1.4 GB (1%)

into our token augmentation strategy, we improve the data diversity, exploring new possibilities for training more robust and generalizable vision models.

4 Experiment

In this section, we validate the effectiveness of SeiT++ on various scenarios. First, we evaluate the performance of our method in the context of storage-efficient ImageNet-1k classification and transfer learning on fine-grained datasets. Then, we demonstrate the adaptability of token-based vision training to ADE-20k semantic segmentation, showing additional improvements facilitated by SeiT++. In addition, we present experimental results illustrating how our strategies contribute to enhanced model generalizability in robustness evaluation. We also explore the applicability of our augmentation strategies to alternative input formats, demonstrating its extensibility. Unless otherwise stated, SeiT++ refers to the combination of MTM and our token augmentation strategies.

Table 2: Token-based ImageNet-1k Classification. We report top-1 accuracies (ViT-B) on ImageNet-1k. Note that SeiT++ w/o MTM means the original SeiT [44] training with our token augmentation strategies.

Storage	SeiT		SeiT++	
	w/o MTM	w/ MTM	w/o MTM	w/ MTM
1.4 GB	74.0	75.1	75.5 (+1.5)	77.8 (+2.7)
1.1 GB	70.2	74.1	73.1 (+2.9)	76.5 (+2.4)
0.8 GB	66.3	70.6	69.1 (+2.8)	74.1 (+3.5)
0.5 GB	59.7	64.5	63.8 (+4.1)	70.0 (+5.5)
0.3 GB	47.2	53.9	51.2 (+4.0)	60.6 (+6.7)

Table 3: Fine-grained classification. We report top-1 accuracies (ViT-B) on various fine-grained datasets. Our method consistently outperforms its counterpart through effective augmentations for tokens.

Method	Flowers	Cars	iNat18	iNat19
SeiT [44]	93.5	79.7	43.1	50.1
SeiT++	95.2 (+1.7)	86.9 (+7.2)	52.4 (+9.3)	58.9 (+8.8)

4.1 Storage-efficient ImageNet-1k

Table 1 summarizes the performance and compression ratio of storage-efficient ImageNet-1k classification. Our method is compared with image-based approaches, including uniform random sampling, C-score [29]-based sampling, adjusting image resolution, and adjusting JPEG quality factor. Despite a 1% compression ratio, when a ViT-B/16 model is pre-trained on tokenized ImageNet-1k using MTM and the ViT model is fine-tuned, our method achieves a 77.8% top-1 accuracy, outperforming various image-based approaches. In Table 2, our method consistently outperforms the token-based learning baseline SeiT by 1.5–6.7%p in top-1 accuracies. Notably, the effectiveness of our method becomes more pronounced as the available storage size decreases, highlighting its efficiency in low-storage scenarios. Furthermore, SeiT++ shows more significant performance gains with MTM. For a storage size of 1.4GB, SeiT++ achieves an improvement of 1.5%p over SeiT without MTM, while with MTM the improvement increases to 2.7%p. These results confirm that SeiT++ effectively takes advantage of MTM.

4.2 Fine-grained Classification

To further investigate the generalizability of SeiT++, we conduct experiments on four different fine-grained datasets: Flowers [43], StanfordCars [31], iNat-18 [57], and iNat-19 [58]. Specifically, we tokenize each dataset and then fine-tune the token-trained model (ViT-B) on ImageNet-1k with the tokenized dataset. Since

Table 4: ADE-20k Semantic Segmentation. We report mIoU on the ADE-20k validation set, highlighting the effectiveness of our approach in pixel-level classification.

Method	Backbone	SeiT	SeiT++
UPerNet [63]	ViT-B [17]	39.0	43.2 (+4.2)

Table 5: Robustness evaluation. We report top-1 robust accuracies (ViT-B) against corruptions and domain shifts for each model trained on ImageNet-1k. Our method consistently outperforms its counterpart on various robustness benchmarks.

Method	Clean	Gauss. Noise	Gauss. Blur	ImageNet-R	Sketch	ObjectNet
SeiT [44]	74.0	50.7	62.6	25.5	22.6	15.7
SeiT++	75.5 (+1.5)	58.6 (+7.9)	66.8 (+4.2)	30.2 (+4.7)	27.7 (+5.1)	18.3 (+2.6)

fine-grained datasets require the classification of objects with intricate characteristics, they are particularly sensitive to undesirable noise (*i.e.*, RRC in SeiT) or excessive perturbation (*i.e.*, Emb-Noise) resulting from data augmentation in the token embedding space. Table 3 shows the quantitative performance (top-1 accuracies) on fine-grained datasets. SeiT++ consistently outperforms the performance of SeiT across various datasets. These results confirm the effectiveness of our proposed method in fine-grained classification.

4.3 ADE-20k Semantic Segmentation

A limitation of token-based learning is often substantiated by learning dense prediction tasks. We conjecture that our method could facilitate training for those tasks, such as semantic segmentation. We extract tokens from ADE-20k [72] and use the pre-trained ViT-B/16 model on tokenized ImageNet-1k to evaluate the effectiveness of our method in semantic segmentation. We then fine-tune the pre-trained model using the tokenized ADE-20k dataset. Table 4 reports the mean intersection-over-union (mIoU) as an evaluation metric. The results show a significant performance improvement (+4.2%p mIoU) over SeiT. These results indicate that SeiT++ significantly improves model performance in the token-based dense prediction task.

4.4 Evaluation on Robustness Datasets

To validate the effectiveness of our augmentation strategies in terms of robustness, we further evaluate the robustness of our method on a set of benchmarks such as adding Gaussian noise [26] or Gaussian blur [26], ImageNet-R [25], Sketch [59], and ObjectNet [3]. Notably, we adopt SeiT++ w/o MTM to validate the impact of our augmentation strategies on model robustness. The results in Table 5 show that SeiT++ consistently outperforms SeiT in terms of robustness

Table 6: Impact of the proposed augmentation strategies. We report the top-1 accuracies (ViT-S) on the ImageNet-100 validation set for various combinations of our token augmentation strategies. Notably, our TokenAdapt and ColorAdapt consistently improve model performance, exhibiting synergy when used together.

Framework	ColorAdapt	TokenAdapt	Top 1 Acc.
	✗	✗	77.3
SeiT [44]	✓	✗	78.3
	✗	✓	80.4
	✓	✓	81.4

Table 7: Adaptability of our method to different input. Input denotes tokenizers for our model. We report the top-1 accuracies (ViT-S) on ImageNet-100. For VQGAN, we used a publicly available VQGAN [18] trained on OpenImages [32] as a tokenizer.

Input	SeiT	SeiT++
Token (ViT-VQGAN [66])	77.3	81.4 (+4.1)
Token (VQGAN [18])	81.8	83.9 (+2.1)
DCT coefficients	69.5	71.0 (+1.5)

accuracy. Previous studies [8, 55] observe that the key to the input pixel robustness depends on the pixel-level augmentations. Since our method enables pixel-based data augmentation to tokens, it achieves higher robustness than the baseline, which does not effectively exploit the pixel-based data augmentation.

4.5 Ablation Study

Impact of Each Augmentation Strategy. In Table 6, we report the impact of proposed token augmentations. In particular, incorporating ColorAdapt leads to a 1.0%p performance improvement, highlighting the positive effect of preserving object structure during model training. Additionally, incorporating TokenAdapt yields an extra 3.1%p improvement. These results confirm that exploring token augmentations is a key factor in fully exploiting pre-extracted tokens.

Extensibility of Our Augmentation Strategy. In this section, we conjecture the broad applicability of our augmentation strategies across various scenarios. To explore its extensibility, we delve into two distinct situations. Firstly, we examine the integration of our method with an alternative tokenizer—the OpenImages-trained VQGAN tokenizer [18], alongside the ImageNet-1k-trained ViT-VQGAN tokenizer [66]. The results in Table 7 consistently demonstrate improved performance on ImageNet-100 across different tokenizers. Additionally, we investigate the adaptability of our approach to a different input format by focusing on Discrete Cosine Transform (DCT) coefficients. Inspired by JPEG compression, we apply DCT to each 8×8 image patch, yielding quantized DCT

coefficients utilized as input data during model training. As shown in Table 7, our approach enhances top-1 accuracy, similar to the tokenizer scenario. This consistent improvement across diverse input formats underscores the universal applicability of our method, highlighting its potential as a widely applicable solution for diverse input representations.

4.6 Implementation Details

We use the pre-trained ViT-VQGAN tokenizer [66] for token extraction. For our TokenAdapt module, we employ a single transformer block for conversion and reverse function. The TokenAdapt module is trained for one epoch on ImageNet-1k with a batch size of 128, a learning rate of 0.001 with cosine scheduling, and a weight decay of 0.05. We follow the training recipes from MAGE [34], SeiT [44], and mmsegmentation [10] for masked token modeling (MTM), token-based classification, and semantic segmentation, respectively. Detailed hyperparameters and settings are available in the Appendix.

5 Conclusion

In this paper, we investigate the compatibility of SeiT with self-supervised learning approaches in token-based frameworks. Among various self-supervised algorithms, Masked Image Modeling (MIM) [24, 64] has recently gained popularity for learning visual representations. Inspired by this, we extend SeiT by integrating Masked Token Modeling (MTM) for self-supervised pre-training. Additionally, recognizing that self-supervised approaches often require more data due to the lack of labels, we propose simple yet effective token-based augmentation strategies: TokenAdapt and ColorAdapt. TokenAdapt aligns the token embedding space with another embedding space compatible with pixel-based data augmentations. In addition, we introduce ColorAdapt, a color-changing augmentation for tokens. Our experimental results across various scenarios, including storage-efficient ImageNet-1k classification, fine-grained classification, ADE-20k semantic segmentation, and robustness benchmarks, consistently validate the effectiveness of our proposed method. Furthermore, the consistent performance improvement across diverse input formats highlights the potential of our approach as a widely applicable solution for various input representations.

Limitations. While SeiT++ is capable of handling large-scale data without labels with optimized storage requirements, experiments on such large datasets (*e.g.*, ImageNet-21k [14]) and huge models are not contained in this paper due to resource constraints. We believe the advantages of MTM and our data augmentation strategies can be seamlessly integrated with scalability.

Acknowledgements. Experiments are based on the NAVER Smart Machine Learning NSML [30] platform. This research was supported by IITP grant funded by the Korea government (MSIT) and KEIT grant funded by the Korea government (MOTIE) (No. 2022-0-00680, 2022-0-01045).

References

1. Bahng, H., Chun, S., Yun, S., Choo, J., Oh, S.J.: Learning de-biased representations with biased representations. In: International Conference on Machine Learning. pp. 528–539. PMLR (2020)
2. Bao, H., Dong, L., Piao, S., Wei, F.: Beit: Bert pre-training of image transformers. arXiv preprint arXiv:2106.08254 (2021)
3. Barbu, A., Mayo, D., Alverio, J., Luo, W., Wang, C., Gutfreund, D., Tenenbaum, J., Katz, B.: Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *Advances in neural information processing systems* **32** (2019)
4. Chang, H., Zhang, H., Barber, J., Maschinot, A., Lezama, J., Jiang, L., Yang, M.H., Murphy, K., Freeman, W.T., Rubinstein, M., et al.: Muse: Text-to-image generation via masked generative transformers. arXiv preprint arXiv:2301.00704 (2023)
5. Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: Maskgit: Masked generative image transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11315–11325 (2022)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
7. Chen, X., Xie, S., He, K.: An empirical study of training self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9640–9649 (2021)
8. Chun, S., Oh, S.J., Yun, S., Han, D., Choe, J., Yoo, Y.: An empirical evaluation on robustness and uncertainty of regularization methods. arXiv preprint arXiv:2003.03879 (2020)
9. Coleman, C., Yeh, C., Mussmann, S., Mirzsoleiman, B., Bailis, P., Liang, P., Leskovec, J., Zaharia, M.: Selection via proxy: Efficient data selection for deep learning. arXiv preprint arXiv:1906.11829 (2019)
10. Contributors, M.: MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation> (2020)
11. Croce, F., Hein, M.: Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In: International conference on machine learning. pp. 2206–2216. PMLR (2020)
12. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation strategies from data. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 113–123 (2019)
13. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 702–703 (2020)
14. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
15. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
16. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. arXiv preprint arXiv:1708.04552 (2017)

17. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
18. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12873–12883 (2021)
19. Fang, Y., Dong, L., Bao, H., Wang, X., Wei, F.: Corrupted image modeling for self-supervised visual pre-training. arXiv preprint arXiv:2202.03382 (2022)
20. Ford, N., Gilmer, J., Carlini, N., Cubuk, D.: Adversarial examples are a natural consequence of test error in noise. arXiv preprint arXiv:1901.10513 (2019)
21. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. arXiv preprint arXiv:1811.12231 (2018)
22. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 (2014)
23. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Dohersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems* **33**, 21271–21284 (2020)
24. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)
25. Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., Song, D., Steinhardt, J., Gilmer, J.: The many faces of robustness: A critical analysis of out-of-distribution generalization. *ICCV* (2021)
26. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. arXiv preprint arXiv:1903.12261 (2019)
27. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE international conference on computer vision. pp. 1501–1510 (2017)
28. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: International conference on machine learning. pp. 4904–4916. PMLR (2021)
29. Jiang, Z., Zhang, C., Talwar, K., Mozer, M.C.: Characterizing structural regularities of labeled data in overparameterized models. arXiv preprint arXiv:2002.03206 (2020)
30. Kim, H., Kim, M., Seo, D., Kim, J., Park, H., Park, S., Jo, H., Kim, K., Yang, Y., Kim, Y., et al.: Nsm1: Meet the mlaas platform with a real-world case study. arXiv preprint arXiv:1810.09957 (2018)
31. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: Proceedings of the IEEE international conference on computer vision workshops. pp. 554–561 (2013)
32. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., et al.: The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision* **128**(7), 1956–1981 (2020)

33. Lee, S., Chun, S., Jung, S., Yun, S., Yoon, S.: Dataset condensation with contrastive signals. In: International Conference on Machine Learning. pp. 12352–12364. PMLR (2022)
34. Li, T., Chang, H., Mishra, S., Zhang, H., Katabi, D., Krishnan, D.: Mage: Masked generative encoder to unify representation learning and image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2142–2152 (2023)
35. Lin, S., Ji, R., Yan, C., Zhang, B., Cao, L., Ye, Q., Huang, F., Doermann, D.: Towards optimal structured cnn pruning via generative adversarial learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2790–2799 (2019)
36. Liu, J., Guo, J., Xu, D.: Apsnet: Toward adaptive point sampling for efficient 3d action recognition. *IEEE Transactions on Image Processing* **31**, 5287–5302 (2022)
37. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
38. Lopes, R.G., Yin, D., Poole, B., Gilmer, J., Cubuk, E.D.: Improving robustness without sacrificing accuracy with patch gaussian augmentation. arXiv preprint arXiv:1906.02611 (2019)
39. Ma, Y., Li, H., Zhang, Z., Guo, J., Zhang, S., Gong, R., Liu, X.: Annealing-based label-transfer learning for open world object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11454–11463 (2023)
40. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. arXiv preprint arXiv:1706.06083 (2017)
41. Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., Bharambe, A., Van Der Maaten, L.: Exploring the limits of weakly supervised pretraining. In: Proceedings of the European conference on computer vision (ECCV). pp. 181–196 (2018)
42. Mei, Z., Ye, P., Ye, H., Li, B., Guo, J., Chen, T., Ouyang, W.: Automatic loss function search for adversarial unsupervised domain adaptation. *IEEE Transactions on Circuits and Systems for Video Technology* **33**(10), 5868–5881 (2023)
43. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: 2008 Sixth Indian conference on computer vision, graphics & image processing. pp. 722–729. IEEE (2008)
44. Park, S., Chun, S., Heo, B., Kim, W., Yun, S.: Seit: Storage-efficient vision training with tokens using 1% of pixel storage. arXiv preprint arXiv:2303.11114 (2023)
45. Paul, M., Ganguli, S., Dziugaite, G.K.: Deep learning on a data diet: Finding important examples early in training. *Advances in Neural Information Processing Systems* **34**, 20596–20607 (2021)
46. Riba, E., Mishkin, D., Ponsa, D., Rublee, E., Bradski, G.: Kornia: an open source differentiable computer vision library for pytorch. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3674–3683 (2020)
47. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10684–10695 (2022)
48. Rosasco, A., Carta, A., Cossu, A., Lomonaco, V., Bacciu, D.: Distilled replay: Overcoming forgetting through synthetic samples. In: International Workshop on Continual Semi-Supervised Learning. pp. 104–117. Springer (2021)

49. Sangermano, M., Carta, A., Cossu, A., Bacciu, D.: Sample condensation in online continual learning. In: 2022 International Joint Conference on Neural Networks (IJCNN). pp. 01–08. IEEE (2022)
50. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems* **35**, 25278–25294 (2022)
51. Scimeca, L., Oh, S.J., Chun, S., Poli, M., Yun, S.: Which shortcut cues will dnns choose? a study from the parameter-space perspective. *arXiv preprint arXiv:2110.03095* (2021)
52. Shao, S., Yin, Z., Zhou, M., Zhang, X., Shen, Z.: Generalized large-scale data condensation via various backbone and statistical matching. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16709–16718 (2024)
53. Singh, M., Gustafson, L., Adcock, A., de Freitas Reis, V., Gedik, B., Kosaraju, R.P., Mahajan, D., Girshick, R., Dollár, P., Van Der Maaten, L.: Revisiting weakly supervised pre-training of visual perception models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 804–814 (2022)
54. Steiner, A., Kolesnikov, A., Zhai, X., Wightman, R., Uszkoreit, J., Beyer, L.: How to train your vit? data, augmentation, and regularization in vision transformers. *arXiv preprint arXiv:2106.10270* (2021)
55. Taori, R., Dave, A., Shankar, V., Carlini, N., Recht, B., Schmidt, L.: Measuring robustness to natural distribution shifts in image classification. *Advances in Neural Information Processing Systems* **33**, 18583–18599 (2020)
56. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: *International conference on machine learning*. pp. 10347–10357. PMLR (2021)
57. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist challenge 2018 dataset. *arXiv preprint arXiv:1707.06642* (2018)
58. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist challenge 2019 dataset. *arXiv preprint arXiv:1707.06642* (2019)
59. Wang, H., Ge, S., Lipton, Z., Xing, E.P.: Learning robust global representations by penalizing local predictive power. In: *Advances in Neural Information Processing Systems*. pp. 10506–10518 (2019)
60. Wang, T., Zhu, J.Y., Torralba, A., Efros, A.A.: Dataset distillation. *arXiv preprint arXiv:1811.10959* (2018)
61. Wei, J., Zou, K.: Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196* (2019)
62. Wightman, R., Touvron, H., Jégou, H.: Resnet strikes back: An improved training procedure in timm. *arXiv preprint arXiv:2110.00476* (2021)
63. Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 418–434 (2018)
64. Xie, Z., Zhang, Z., Cao, Y., Lin, Y., Bao, J., Yao, Z., Dai, Q., Hu, H.: Simmim: A simple framework for masked image modeling. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9653–9663 (2022)
65. Yin, Z., Xing, E., Shen, Z.: Squeeze, recover and relabel: Dataset condensation at imagenet scale from a new perspective. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023)

66. Yu, J., Li, X., Koh, J.Y., Zhang, H., Pang, R., Qin, J., Ku, A., Xu, Y., Baldrige, J., Wu, Y.: Vector-quantized image modeling with improved vqgan. arXiv preprint arXiv:2110.04627 (2021)
67. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6023–6032 (2019)
68. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017)
69. Zhao, B., Bilen, H.: Dataset condensation with differentiable siamese augmentation. In: International Conference on Machine Learning. pp. 12674–12685. PMLR (2021)
70. Zhao, B., Mopuri, K.R., Bilen, H.: Dataset condensation with gradient matching. arXiv preprint arXiv:2006.05929 (2020)
71. Zhong, Z., Zheng, L., Kang, G., Li, S., Yang, Y.: Random erasing data augmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI) (2020)
72. Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision* **127**, 302–321 (2019)