

Learning with Unmasked Tokens Drives Stronger Vision Learners

Taekyung Kim*, Sanghyuk Chun, Byeongho Heo, and Dongyoon Han*

NAVER AI Lab

Abstract. Masked image modeling (MIM) has become a leading self-supervised learning strategy. MIMs such as Masked Autoencoder (MAE) learn strong representations by randomly masking input tokens for the encoder to process, with the decoder reconstructing the masked tokens to the input. However, MIM pre-trained encoders often exhibit a limited attention span, attributed to MIM’s sole focus on regressing masked tokens only, which may impede the encoder’s broader context learning. To tackle the limitation, we improve MIM by explicitly incorporating unmasked tokens into the training process. Specifically, our method enables the encoder to learn from broader context supervision, allowing unmasked tokens to experience broader contexts while the decoder reconstructs masked tokens. Thus, the encoded unmasked tokens are equipped with extensive contextual information, empowering masked tokens to leverage the enhanced unmasked tokens for MIM. As a result, our simple remedy trains more discriminative representations revealed by achieving 84.2% top-1 accuracy with ViT-B on ImageNet-1K with 0.5%p gain. We attribute the success to the enhanced pre-training method, as evidenced by the singular value spectrum and attention analyses. Finally, our models achieve significant performance gains at the downstream semantic segmentation and fine-grained visual classification tasks; and on diverse robust evaluation metrics. Code is available at <https://github.com/naver-ai/lut>.

1 Introduction

Triggered by successful transitions of Transformer [52] into vision domains [6, 16], a plethora of effective training strategies for Transformer have emerged [8, 10, 22, 48, 49]. Recent advances in masked image modeling (MIM) [3, 22, 32, 59, 63] noticeably show great success in self-supervised learning (SSL) of Vision Transformers (ViT) by transferring the knowledge of masked language modeling [13]. Conceptually, MIM tasks consist of two parts: randomly masking out a part of inputs (*e.g.*, 75% of input pixels) and predicting the masked inputs by the decoder. This simple strategy enables a model to learn strong representations through the challenging task.

However, MIM strategies often encounter challenges, such as local dependency on attention to understand the entire context of an image. For example,

* Equal contribution

liu *et al.* [37] revealed that Masked Autoencoder (MAE) [22], a state-of-the-art MIM method, exhibits shorter average attention distances. Furthermore, we observe that attention map patterns by MAE substantiate extremely local behavior (See Fig. 1) indeed. In other words, the MAE-trained attention mechanism integrates less information across all image pixels and tends to focus on specific input regions. This is presumably attributed to MIM-pretraining, primarily dedicated to predicting low-level pixel details (*e.g.*, color or texture) without a comprehensive understanding of less-regional information (*e.g.*, the input structure or shape).

We aim to understand the chronic shortage in a limited range of dependencies and how it affects MIM. We illustrate that vanilla MIM methods appear to lack broader-range dependency. Drawing from this, we identify a deficiency inside the vanilla MIM formulation and present a simple solution to its local dependency issue, demonstrating how it enhances MIM pre-training. Our proposed method Learning with Unmasked Tokens (LUT) for MIM complements the sub-optimal representation learning by offering broadly contextualized supervision to unmasked tokens through extracting general context from the entire pixels to strengthen unmasked tokens (for masked tokens to attend), which aims to learn more context-generalized representations for the encoder.

During training, LUT minimizes the discrepancy between the encoded general context representations and the sparse representation processed by the on-line learnable encoder from different views while performing MIM with a decoder. This ensures reinforcing more contextualized unmasked tokens for mask tokens to attend to. A general representation derived from all pixels can effectively utilize a highly augmented view, minimizing reliance on regional changes like color distortion to improve generalization. The learnable network encodes a sparse and unmasked view and matches it to the generalized representation, and the decoder reconstructs the masked pixels using the encoded features, similar to [22]. We presume that our strategy promotes the learning of the encoder by involving unmasked tokens in training, leveraging the target network’s capability to encode a wide context for all pixels.

We verify the effectiveness of LUT by pre-training the ViT architectures [16] on the ImageNet-1K benchmark [45]. We do not solely evaluate our method by linear probing but by fine-tuning results. Given our method’s weight on improving the baseline MIM, LUT-trained ViT-B/16 successfully improves linear evaluation (+2%p) and fine-tuning (+0.5%p) performance gains on ImageNet-1K over MAE. Our fine-tuning result also achieves comparable or outperformed ImageNet-1K validation accuracy (84.2%) compared with other state-of-the-art methods. LUT can be transferred to the multiple fine-grained classifications and show distinguished transferability. LUT further shows superior transferability and tuning robustness on INaturalist datasets. We further transfer our pre-trained model to the semantic segmentation task on ADE20K [62] and show 49.5% mIoU, a solid result. As another benefit, LUT successfully realizes robust training, which results in enhanced robustness results on two in-distribution benchmarks, five out-of-distribution benchmarks, and SI-Score [14].

2 Preliminary

Despite MIM’s strong performance, we claim it still lacks strong attention capability after pre-training, particularly for comprehensive region-wide dependency. The upcoming spatial attention map visualizations motivate our method.

2.1 MIM and Beyond

We begin with a generalized formulation of MIM, addressing the limitation shown in the formulation.

General formulation. Given an image from an augmented view u , we patchify the image into N non-overlapping patches $U = \{\mathbf{u}_i\}_{i=1}^N$. We randomly pick masked patches \mathcal{M} with a high masking ratio $r \in (0, 1)$, where $\mathcal{M} \subset \{1, 2, \dots, N\}$ and $|\mathcal{M}| = \lfloor rN \rfloor$. We denote the masked image patches as $U_m = \{\mathbf{u}_i : i \in \mathcal{M}\}$ and the unmasked patches as $U_r = \{\mathbf{u}_i : i \in \{1, 2, \dots, N\}, i \notin \mathcal{M}\}$. The unmasked patches are fed into the encoder f_θ and become encoded tokens $T_e = f_\theta(U_r)$. The encoded tokens are concatenated with mask tokens \mathbf{m}_i corresponding to the positions of i -th masked patches (entire patches $U_m \cup U_r$ can be fed into the encoder [59] or U_m only is utilized [22]). The only mask tokens predict the image patches through the decoder d_ϕ . We denote i -th decoded mask token and input mask token as \mathbf{m}_i^d and \mathbf{m}_i , where $T_d \cup \{\mathbf{m}_i^d\}_{i \in \mathcal{M}} = d_\phi(T_e \cup \{\mathbf{m}_i\}_{i \in \mathcal{M}})$, informally. Here, T_d is a set of decoded unmasked tokens. Now, the MIM pre-training objective is defined as:

$$\mathcal{L}_{\text{MIM}} = \sum_{i \in \mathcal{M}} \|\mathbf{m}_i^d - \mathbf{u}_i\|_2^2, \quad (1)$$

where m_i are shared for all the positions.

MIM formulation itself falls short in learning broader contexts. We argue that the MIM loss in Eq. (1) may not fully exploit the pre-training capability while it is effectively designed with simplicity. Specifically, owing to the scarcity of unmasked tokens U_r in the encoder f_θ , the formulation does not leverage complete image information; they are encoded through self-attention, being only attended by a few visual tokens. Therefore, the encoding may struggle with low contextualization due to sparse visual tokens from the query.

In the decoder d_ϕ , we argue the restricted local options T_e are only available for the mask tokens to attend to. The loss is primitively to provide limited supervision for masked tokens to reconstruct image patches from the constrained encoded information T_e using the decoder. The mask tokens may eventually attend near visual tokens to reduce the loss, as we will observe in Fig. 1b.

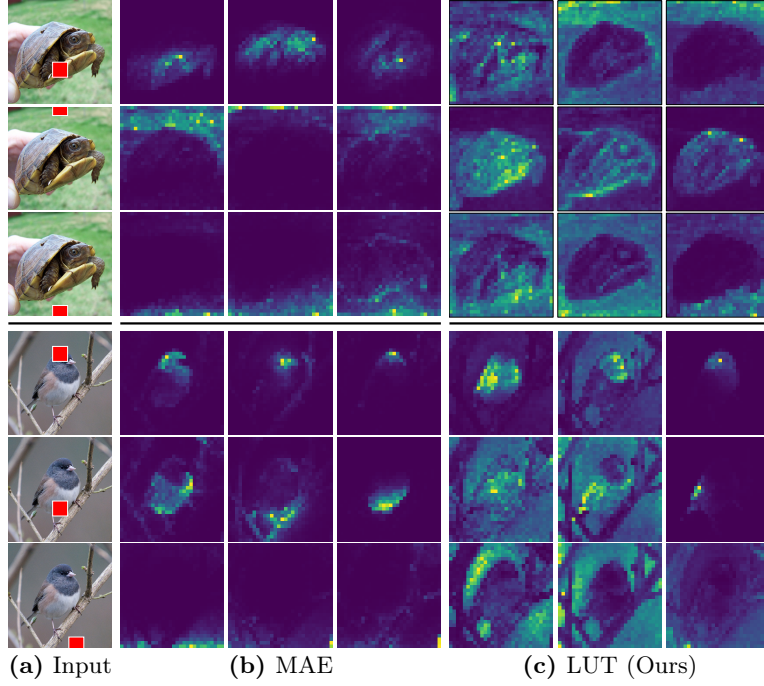


Fig. 1: Motivation - MAE may lack comprehensive region-wide attention.

We observed how attention appears differently in MAE corresponding to given queries. (a) The first column denotes the example images with (**different queries**) randomly picked patch indices. (b), (c) Every set of three columns represents the maps that are the most attended by different heads. The **turtle** images have a foreground, and two upper and lower background queries; the **bird** images have two foreground queries (upper two rows) and one background query. MAE shows localized attention maps but fails to provide comprehensive coverage of either foreground or background.

2.2 Motivation - attention map visualizations

Attention map visualizations qualitatively reveal how a model reacts to queries and reflects the deficiency of representations. We believe the range of attention may suggest the diversity of token dependency that closely links to the encoder’s capability. Fig. 1b shows the attention maps concerning the given query in Fig. 1a by MAE [22]. We exploit self-attention in the last block for visualization in the official ViT-B/16 MAE and visualize maps with 480×480 images from ImageNet-1K. We observe MAE shows narrower highlighted regions for the given queries. Specifically, when a query is selected in the foreground (the 1st, 4th, and 5th rows), MAE only highlights the near patches of the given query; when a query is selected even in the background (the 2nd, 3rd, and 6th rows), we observe the same phenomenon, namely, MAE only focuses on the near patches of the given query. Based on this, we argue that MAE’s attention lacks a broader range of dependency. This may incur a lack of a broad understanding of the entire foreground or background.

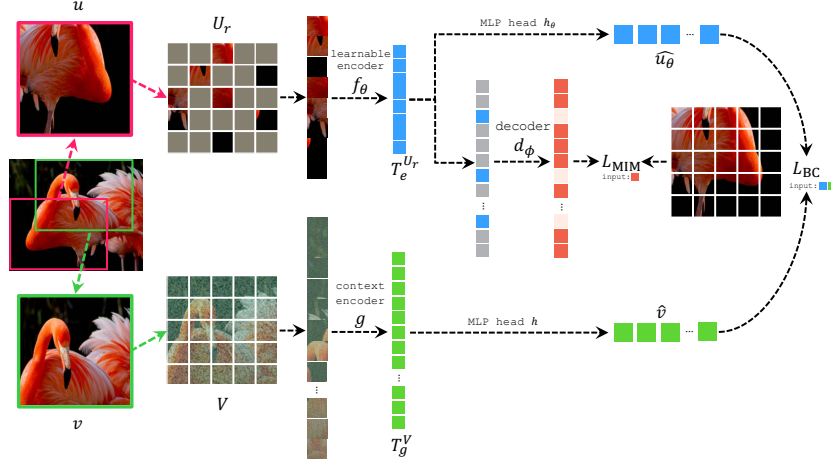


Fig. 2: Framework overview. Our method performs a masked image modeling with masked tokens, complemented by a context encoder that directs the learnable side, giving an augmented complete view for sparse unmasked tokens. We employ distinct, simple MLP heads on top of each encoder to match representations and avoid optimization collapse. Under our macro concept, we opt for a simple choice where the additional context momentum encoder mirrors the online encoder, while we may alternatively use various options. We borrow a flamingo image from n02007558 class in ImageNet-1K.

3 Method

In this section, we introduce our proposed method - Learning with Unmasked Tokens (LUT) for Masked Autoencoder that addresses the local dependency issue. We opt for MAE [22] as the baseline for our formulation.

Our simple solution. Based on the arguments presented, we contend that Eq. (1) computes loss exclusively using reconstructed masked tokens while leaving unmasked tokens being trained implicitly, which may impede learning broad contexts effectively. As a remedy, we employ another loss expected to aid \mathcal{L}_{MIM} by giving expansive supervision to unmasked tokens from the entire tokens:

$$\mathcal{L}_{\text{ours}} = \sum_{i \in \mathcal{M}} \|\mathbf{m}_i^d - \mathbf{u}_i\|_2^2 + \alpha \mathcal{D}(T_e, g(U_m \cup U_r)), \quad (2)$$

where $\mathcal{D}(\cdot, \cdot)$ and $g(\cdot)$ denote a distance function and a context encoder. We straightforwardly give encoded comprehensive supervision from entire tokens to unmasked tokens T_e . During training, the expansively supervised T_e contains extended token information so that mask tokens can leverage it. This potentially gives additional localization capability beyond what the baseline possesses. The options for choosing \mathcal{D} and g are indeed diverse, but we take the simplest way in the next section. Eq. (2) can involve both unmasked and mask tokens, but we focus on unmasked tokens to prevent learning collapse in mask tokens.

Contextualized supervision. The crux of our solution lies in *learning unmasked tokens actively* from a more comprehensive contextualization of entire

visual tokens. We opt for the elements in the newly involved loss (dubbed broader contextualization loss \mathcal{L}_{BC}) in Eq. (2). First, for the context encoder g , we implement this by simply reusing the encoder f_θ to give the supervision back to f_θ . It performs like a token-level regression between the encoders. We employ an efficient yet strong option, momentum networks [7, 10, 20, 23]. The architecture consists of a momentum encoder and MLP head, which is nearly identical to the learnable side.

Additionally, we augment the entire image patches from U to V to enhance the generalization of the encoder and avoid collapse. For the context latent features, the view \mathbf{v} is patchified into $V = \{\mathbf{v}_i\}_{i=1}^N$, respectively. Unlike the general MIM process, all the patches V are encoded by g ; we denote the encoded tokens as $T_g^V = g(V)$, where $T_g^V = \{t_i^V\}_{i=1}^N$. Finally, the MLP head h yields context representations $\hat{\mathbf{v}} = h(\tilde{\mathbf{v}})$, where $\tilde{\mathbf{v}}$ can be $\tilde{\mathbf{v}} = \frac{1}{N} \sum_{i=1}^N t_i^V$ (pooled) or each set of representations $\tilde{\mathbf{v}} = T_g^V$. Alternatively, using aligned tokens [12] for $\tilde{\mathbf{v}}$ could benefit performance, but we simply use a pooled token.

We refer to the process as delivering *contextualized supervision* to unmasked tokens, which encodes information from entire tokens to facilitate training through the broader supervision for unmasked tokens. Note that this supervision branch adds a single forward pass, but ours exhibits a slight time increase of approximately 18% - 0.46s over 0.39s (baseline) at each iteration.

Sparse unmasked tokens that learn broad contexts. Our encoding process obtains regional representations $T_e^{U_r} = f_\theta(U_r)$ from sparsified tokens U_r . Similar to computing general context representations, we aggregate the latent embeddings $\mathbf{u}_\theta = \frac{1}{|T_e^{U_r}|} \sum_{t \in T_e^{U_r}} t$ through averaging. We follow the previous studies preventing training collapse by applying an MLP head h_θ to obtain $\hat{\mathbf{u}}_\theta = h_\theta(\mathbf{u}_\theta)$, forming architectural asymmetry to avoid collapse [9, 20]. LUT can be interpreted as utilizing masked tokens for MIM interacting with sparse visual tokens employed to condense expanded context information.

On contextual discrepancies across views. We aim to provide broader contextualized supervision to unmasked tokens that correspond to the original view of the masked tokens. However, MIMs generally use random resized crop (RRC) [46] for giving geometric variation; we argue that using RRC may not align with our intention and could hinder learning due to divergent views often providing narrower and limited shared information [9, 49]. Thus, we adopt simple resized crop (SRC) [49] instead of RRC. We conjecture the latent features from unmasked tokens can be more reliably guided by the semantics from a broader context. We will show that SRC harms MAE but improves LUT in §6.

Objective function. We finalize our objective by choosing the distance function \mathcal{D} in Eq. (2). We apply the normalized ℓ_2 -distance for the feature distance (*i.e.*, Cosine distance). We have the aggregated context representation $\hat{\mathbf{v}}$ and sparse one $\hat{\mathbf{u}}_\theta$, and their ℓ_2 -normalized version $\bar{\mathbf{u}}_\theta$ and $\bar{\mathbf{v}}$, respectively. Our broader contextualization loss \mathcal{L}_{BC} computes the feature distance between normalized representations $\bar{\mathbf{u}}_\theta$ and $\bar{\mathbf{v}}$, formulated as $\mathcal{L}_{\text{BC}} = \|\bar{\mathbf{u}}_\theta - \bar{\mathbf{v}}\|_2^2$. LUT is

agnostic to the choice of distance function since the fundamental principle of it works regardless of the distance functions - the InfoNCE or Smoothed ℓ_1 losses also show compatibility with LUT. The final objective function is:

$$\min_{\theta} \mathcal{L}_{\text{MIM}} + \alpha \mathcal{L}_{\text{BC}}, \quad (3)$$

where α controls the balance of the broader contextualization loss and the masked image modeling loss. The study on α gives the best fine-tuning performance with 0.25; however, it is insensitive to the choice of α ; for example, $\alpha = 0.25$ and 0.5 shows only 0.1% difference of fine-tuning performance using the ViT-B/16 backbone. We support all our design choices in the ablation studies in Table. 4. Our method is also applicable to SimMIM [59]-like methods with performance improvements (see Appendix for details).

4 Related Work

Prior to transitioning to our experiments, we outline the distinctions between our work and closely related works. Several studies have been conducted recently employing multiple encoders, such as our online and target encoders. For example, a line of research excludes using additional data and employs an additional tokenizer module for reconstruction supervision. Drawing inspiration from the success of masked language modeling [13], milestone MIMs [2, 3, 22, 59, 63] have also gained attention for their ability to train discriminative representations by masking input tokens to reconstruct them during pre-training.

First, iBOT [63] jointly trains the online encoder and the online tokenizer. The main motivation is to align the full representations of multi-view instances among 12 different views while additionally performing masked feature reconstruction. Thus, iBOT needs multi-crops varying in diverse scales and augmentations. iBOT involves only mask tokens to learn target information, which may incur learning partial information, but we nevertheless speculate that leveraging multi-crops diminishes this issue. In contrast, our aim is to employ unmasked tokens to reinforce the understanding of longer-range context by offering complete information from a single view.

data2vec [2] performs patch-wise feature prediction via masked tokens. Except for the tokenizer, the framework is largely based on the BeiT framework [3]. Despite the target momentum features being generated from an entire image, we argue that data2vec still lacks long-range supervision to learn. Specifically, we presume only masked tokens contribute to MIM, suggesting attention between masked and unmasked tokens alone is insufficient for learning contextualized information. Therefore, the patch-wise regression to the token representations may not adequately establish strong neighboring dependencies. We conjecture this was eventually evidenced by inferior localization performance. MSN [1] is a denoising-based method employing masks for denoise without reconstruction loss. Its focus on low-shot learning and prototype-based matching before supervision are major distinctions to our method. CMAE [28] and ConMIM [61] both

focus on contrastive learning along with MIM. However, a major difference of ours stems from the primitive goal. The methods directly combine contrastive learning with all the visual patches recovered by particular reconstruction decoders. On the other hand, our method does not perform contrastive learning but targets to forge the unmasked tokens excluded from MIM training by directing them with comprehensive contextualized information.

5 Experiment

This section demonstrates our method by pre-training on ImageNet-1K, compared with competing SSL methods. We further fine-tune our models to confirm transferability.

5.1 ImageNet-1K Classification

Architecture. We use the standard Vision Transformer (ViT) [16] with a patch size of 16 for all experiments (*i.e.*, ViT-B/16) to fairly compare with prior arts. We use the 8-layer transformer decoder following the MAE [22]’s setup for masked image modeling. We further adopt online and context MLP heads on the top of the encoders to aggregate general context from representations; each consists of fully-connected layers with the embedding dimension of 4096, batch normalization layers [29], and ReLUs [35] following the previous methods [8, 10, 20]. Note that LUT works even with symmetric heads. All the decoder and MLP heads are only used during training.

Pre-training setup. We follow the identical ImageNet-1K [45] pre-training protocol¹ [22]. Our model is pre-trained for 1600 epochs with 40 warmup epochs, batch size of 4096, and input resolution of 224×224 . We use AdamW [39] with momentum (0.9, 0.999). The learning rate is set to 1.5×10^{-4} with cosine learning rate decay [38]. We adopt a layer-wise learning rate decay of 0.65. We set a mask ratio for Eq. (1) to 0.75, a momentum decay rate τ in the target network to 0.996, and the weight of our broader contextualization loss (α in Eq. (3)) to 1.0 and 0.25 for the ViT-S/16 and ViT-B/16 architectures, respectively. We employ the simple resized crop [49] for geometric augmentation, color jittering, and the three augment [48] consists of Gaussian blur, grayscale, and solarization. All models are pre-trained using 8 V100-32GB GPUs.

Results. We compare our method with previous SSL methods [1, 2, 8, 10, 11, 15, 22, 36, 41, 47, 54, 56, 57, 59, 61, 63]. Table 1 shows the evaluation results on the ViT-S/B/L backbones. Our LUT achieves an 82.0%, 84.2%, and 86.0% top-1 accuracy on ViT-S/16, ViT-B/16, and ViT-L/16, which improves 0.6%p, 0.5%p, and 0.4%p over the baseline, respectively. Moreover, LUT outperforms other self-supervised learning methods by a large margin except for some masked feature models. This comes to a head with a smaller ViT-S/16, where most of the

¹ We use the public codebase in <https://github.com/facebookresearch/mae>

Table 1: Comparisons with previous models on ImageNet-1K. We compare LUT with the previous results that used vanilla Vision Transformer architectures. All models were pre-trained and fine-tuned on ImageNet-1K. We use the ViT-S/16, ViT-B/16, and ViT-L/16 architectures and a resolution of 224×224. [†] denotes the models pre-trained using multi-crop augmentation. [‡] denotes our reproduction results. We highlight the best numbers (in boldface) and the second-best numbers (in underlined). For a fair comparison, we do not compare methods using modules trained on extra data, such as CLIP [43] or VQGAN [18].

Method		Supervision	ViT-S	ViT-B	ViT-L	ADE20K
<i>Sup models</i>						
DeiT [48]	ICML 2021	Label	79.9	81.8	-	-
DeiT-III [49]	ECCV 2022	Label	81.4	83.8	84.2	49.3
Cosub [50]	CVPR 2023	Label	81.5	84.2	85.3	49.3
<i>Self-sup models</i>						
MoCo v3 [10]	ICCV 2021	Pixel	81.4	83.2	84.1	47.3
DINO [8]	ICCV 2021	Pixel	81.5 [†]	82.8 [†]	-	46.8
SplistMask [17]	arXiv 2021	Pixel+Feature	81.5	83.6	-	46.8
iBOT [63]	ICLR 2022	Feature	82.0	84.0 [†]	84.8 [†]	50.0[†]
MAE [22]	CVPR 2022	Pixel	81.4 [‡]	83.7 [‡]	85.6 [‡]	48.1
SimMIM [59]	CVPR 2022	Pixel	<u>81.9[‡]</u>	83.8	-	-
MaskFeat [56]	CVPR 2022	Feature	-	84.0	85.7	-
ExtreMa [57]	TMLR 2022	Feature	81.8	83.7	-	47.9
data2vec [2]	ICML 2022	Feature	81.8 [‡]	<u>84.1[‡]</u>	86.6	48.3 [‡]
SemMAE [36]	NeurIPS 2022	Pixel	-	83.3	-	46.3
SdAE [11]	ECCV 2022	Pixel	-	84.1 [†]	-	48.6 [†]
MSN [1]	ECCV 2022	Feature	-	83.4	-	-
BootMAE [15]	ECCV 2022	Pixel+Feature	-	84.2	85.9	49.1
CAN [41]	arXiv 2022	Pixel	-	83.6	84.7	-
ConMIM [61]	ICLR 2023	Dictionary	82.0	83.7	85.5	46.0
SIM [47]	CVPR 2023	Feature	-	83.8	-	-
HPM [54]	CVPR 2023	Pixel	-	84.2	85.8	48.5
MIRL [27]	NeurIPS 2023	Pixel	-	84.1	85.4	-
CrossMAE [19]	arXiv 2024	Pixel	79.3	83.7	85.4	-
Ours	-	Pixel	82.0	84.2	<u>86.0</u>	<u>49.5</u>

results are saturated, but this is presumably due to the low capability of the backbone and the high flexibility of masked feature models. LUT would take advantage of further improvements using masked feature models as the baseline. The results highlight the efficacy of our proposed broader contextualized supervision in enhancing MIM, which showcases its significant potential for further improvements.

Computational costs. Our method includes extra computation from forward inference with images, so there is a slight increase in computational demands, as mentioned above. However, our method achieves a top-1 accuracy of 83.6% at 400 epochs, which matches MAE’s accuracy at 1600 epochs, despite our significantly shorter GPU wall time. Specifically, our method takes 119 hours to complete

400 epochs of training, which is roughly half the training time of MAE’s 1600 epochs, which requires 223 hours.

5.2 ADE20K Semantic Segmentation

To validate the transferability to dense prediction tasks, we evaluate semantic segmentation performances on ADE20K [62]. We follow the standard training protocol [22]; the models are fine-tuned for 160K iterations using UperNet [58] with a batch size of 16 and a resolution of 512×512 . Other detailed hyper-parameters for training are listed in Appendix. The rightmost column in Table 1 shows the mIoU performance comparison. LUT also outperforms the competing methods, including SSL and supervised learning methods. This outcome can be attributed to the improved dense prediction capability.

5.3 Transfer Learning

iNaturalist datasets. To further compare the transferability of learned representations, we measure image classification accuracies by fine-tuning on iNaturalist 2018, iNaturalist 2019, and mini iNaturalist 2021 [51], which are highly imbalanced with different number of images per class. We compare LUT with MoCo v3 [10], BYOL [20], DINO [8], iBOT [63], and MAE [22]. All the models are ImageNet-1k pre-trained ViT-B/16 with a resolution of 224×224 . We report the maximum accuracy with the mean and standard deviation of the accuracies obtained by grid searches of learning rates and weight decay, following the protocol [33]. Table 2 shows LUT outperforms the competitors across all datasets, which reveals superior transferability and tuning robustness.

Fine-Grained Visual Classification (FGVC) datasets. We further validate fine-tuning classification accuracies on CIFAR-10 [34], CIFAR-100 [34], CUB-200 [53], Aircraft [40], Birds [26], Flowers [42], and Dogs [31] following the same evaluation protocol as above. Table 3 showcases LUT achieves the best number on average and outstanding numbers overall, which shows improved transferability and tuning robustness across datasets again.

Table 2: Transfer learning results on iNaturalists. We present the end-to-end fine-tuning accuracies on iNaturalist 2018, iNaturalist 2019, and mini iNaturalist 2021 [51]. We report the best results along with the mean \pm std of the set of accuracies obtained from grid searches for each method. † denotes the models pre-trained using multi-crop augmentations. Our method consistently outperforms the competitors in terms of the best accuracies, further showcasing remarkable tuning robustness.

Method	iNat 2018	iNat 2019	iNat 2021-mini
BYOL	69.8 (68.6 \pm 0.9)	77.4 (76.7 \pm 0.8)	70.5 (69.1 \pm 1.1)
MoCo v3	70.1 (69.4 \pm 0.5)	77.6 (77.2 \pm 0.4)	70.9 (70.5 \pm 0.5)
DINO †	72.1 (71.9 \pm 0.2)	79.4 (79.0 \pm 0.4)	73.0 (72.8 \pm 0.1)
iBOT †	73.8 (73.5 \pm 0.2)	79.9 (79.5 \pm 0.4)	74.5 (74.4 \pm 0.1)
data2vec	75.2 (74.5 \pm 0.7)	80.6 (80.0 \pm 0.5)	76.2 (75.5 \pm 0.9)
MAE	74.6 (74.5 \pm 0.1)	80.2 (80.0 \pm 0.1)	75.7 (75.5 \pm 0.2)
LUT	75.8 (75.3\pm0.3)	81.0 (80.5\pm0.4)	76.7 (76.3\pm0.3)

Table 3: Transfer learning results on FGVC. We present the end-to-end fine-tuning accuracies on multiple datasets, reporting the best results along with the mean \pm std of the accuracies from grid searches. Our method mostly outperforms the competitors at the best accuracies, further showcasing the robustness among different training hyper-parameters. [†] denotes the models pre-trained using multi-crop augmentation.

Method	Aircraft	Birds	CUB-200	CIFAR-10	CIFAR-100	Dogs	Flowers	Avg
DINO [†]	87.0 (86.0 \pm 0.6)	83.9 (83.4 \pm 0.5)	85.1 (84.9 \pm 0.3)	99.0 (98.9 \pm 0.1)	91.3 (90.7 \pm 0.5)	84.8 (84.6 \pm 0.3)	98.8 (98.7 \pm 0.1)	90.0
iBOT [†]	87.3 (86.7 \pm 0.6)	85.5 (85.1 \pm 0.5)	85.9 (85.5 \pm 0.3)	99.2 (98.8 \pm 0.6)	92.0 (91.1 \pm 0.9)	86.0 (85.7 \pm 0.3)	99.0 (99.0 \pm 0.1)	90.7
MAE	88.1 (87.3 \pm 0.9)	84.2 (84.0 \pm 0.3)	84.6 (84.3 \pm 0.2)	98.8 (98.7 \pm 0.1)	90.0 (89.7 \pm 0.3)	86.8 (86.4 \pm 0.3)	98.1 (97.8 \pm 0.3)	90.1
data2vec	87.3 (86.6 \pm 0.7)	84.1 (83.5 \pm 0.5)	84.4 (83.9 \pm 0.4)	98.8 (98.7 \pm 0.1)	91.2 (91.0 \pm 0.2)	85.7 (85.3 \pm 0.3)	96.7 (94.4 \pm 3.3)	89.7
LUT	89.2 (88.3 \pm 0.9)	86.0 (85.3 \pm 0.6)	86.5 (85.7 \pm 0.6)	99.1 (98.9 \pm 0.1)	91.0 (90.7 \pm 0.4)	87.4 (86.7 \pm 0.5)	98.4 (98.2 \pm 0.2)	91.1

6 Analysis and Discussion

6.1 Ablation Study

We conduct ablation studies of LUT pre-training under various available configurations. We use ViT-B/16 and train it for 400 epochs on ImageNet-1K as the fixed pre-training setup. Each model is then individually pre-trained. We report the top-1 fine-tuning and linear probing accuracies for each study.

Steered tokens. We showed that steering unmasked tokens via comprehensive supervision leads to an improved encoder. Additionally, we investigate whether masked tokens also benefit from steering. Table 4a illustrates that training solely with unmasked tokens shows superiority, aligning with our primitive conjecture.

Contextualization for token steering. We mainly used the visual tokens for contextualizing methods, but we studied whether other tokens, such as cls-token, can be an alternative in Table 4b. We observe using pooled visual tokens is preferred for LUT. Considering latent features undergo masked auto-encoding, these results imply that explicitly using general context is more effective than using implicit information via cls-token.

Type of supervision. We study the effectiveness of various supervisions for unmasked tokens in Table 4c. We mainly compare token-wise versus broadly aggregated supervision. While all the types yield performance gains, sole contextualization works the best. It outperforms the combination of token-wise, implying that the additional token-wise supervision may conflict with the aggregated one, which is presumably due to the alignment between the set of tokens.

Loss function. We explore various losses for the broader contextualization loss in Table 4d. While all objectives yield considerable performance as expected above, the cosine distance of latent representations of broader and partial information works best when pre-training by LUT.

Masking ratio for target. We study whether the target encoder needs masked images. Table 4e shows that the model without target masking outperforms all the counterparts. Moreover, the fine-tuning accuracy with masking even underperforms the baseline, implying that transferring coarse information harms the capability of learning representation.

Image crop type. This study highlights how performance is affected by the disparity between the two views. There would be many comparing options; we choose random resized crop (RRC) [46] and simple resized crop (SRC) [49] for

Table 4: Ablation studies. We report fine-tuning (ft) and linear probing (lin) accuracies for each configuration which are pre-trained with ViT-B/16. All models are pre-trained for 400 epochs on ImageNet-1k. We mark our default settings in gray.

case	ft	lin	case	ft	lin
none	82.8	61.5	-	82.8	61.5
unmasked/masked tokens	83.0	67.2	cls-token	83.2	71.0
unmasked tokens only	83.5	67.9	pooled visual tokens	83.5	67.9

(a) **Steered tokens.** Guiding un-masked tokens only performs better. (b) **Context tokens.** Using visual tokens works better.

case	ft	lin
none	82.8	61.5
token-wise	83.2	64.1
contextualization	83.5	67.9
both supervisions	83.1	66.6

(c) **Supervision types.** Contextualization beats others.

case	ft	lin
none	82.8	61.5
infoNCE	83.0	66.7
smoothed- ℓ_1	83.2	60.2
cosine distance	83.5	67.9

(d) **Loss function.** “Cos” works the best.

case	ratio	ft	lin
none (w/o target)	-	82.8	61.5
w/o target masking	0	83.5	67.9
w/ target masking	0.5	82.6	63.2
w/ target masking	0.75	82.5	63.8

(e) **Target’s masking ratio.** Target encoder needs noncorrupted views.

Method	Epochs	Views		ft	lin
		RRC	SRC		
MAE	400	✓		82.8	61.5
MAE	400		✓	82.5	64.2
LUT	400	✓		83.4	67.1
LUT	400		✓	83.5	67.9

(f) **View discrepancy.** Ours benefits reduced differences among views.

comparison. Table 4f shows the model pre-trained with SRC exceeds the fine-tuning accuracy of the case of RRC. Since RRC is more compatible with MAE than SRC, performance improvements are not observed in MAE. Our method benefits from SRC, which indicates that the information that will be encoded needs to align closely with the view of the other side, thereby facilitating training.

6.2 Further Analyses

Grad-CAM visualizations. We analyze the Grad-CAM results for LUT vs. MAE to assess the impacts of our method in terms of context range. Fig. 3 shows LUT better grasps broader foreground object contexts, while MAE has difficulty in fully emphasizing these regions.



(a) Input (b) MAE (c) Ours
Fig. 3: GradCAM visualization.

Analysis of attention distance. We measure the average attention distance (AAD) [16] to explore the dependency range in MAE quantitatively. Formally, AAD is defined by a given self-attention map S as follows:

$$d_{\text{AAD}} = \sum_{j \in \mathcal{M}} \sum_{i \in \mathcal{M}} S(i, j) \cdot \|p_i - p_j\|_2, \quad (4)$$

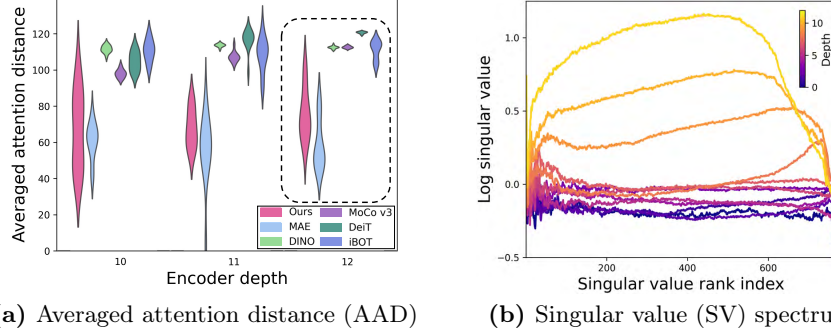


Fig. 4: Visualization Results. (a) We plot AAD values of our LUT, MAE [22], DINO [8], MoCo v3 [10], DeiT [48], and iBoT [63] for the last three layers, which handle high-level semantics. LUT ranges more diverse and broader dependencies overall. (b) We plot the difference of singular values between the baseline and LUT at each layer, showing large gaps (≥ 0), particularly for the later layers. This suggests our method learns more discriminative representations.

where \mathcal{M} denotes a set of patch indices and $S(i, j)$ denotes an attention value between i -th and j -th patches. p_i is a normalized 2D coordinate of the center of the i -th patch. We compute AAD using the entire images from the ImageNet-1K validation set. ViT-B/16 pre-trained by MAE [22], DINO [8], MoCo v3 [10], DeiT [48], and iBoT [63] on ImageNet-1K are included in our study. Considering the crucial role of the final layers in the semantic encoding process, our analysis focuses on the layer-wise AAD for the last three layers. Figure 4a shows that our LUT exhibits diverse ranges, spanning from shorter to longer-range dependencies. Thus, our model generally interacts wider than MAE and also surpasses others in the distance scopes. Notably, the phenomenon is more significant at the final layer, which determines the capacity of contributes most to high-level semantics.

Spectral analysis. We provide additional analysis on the learned layer-wise representations LUT and MAE. Inspired by the previous study [30], we measure the singular values (SVs) of the covariance of features, *i.e.*, how the features are spread in the embedding space. More specifically, we compute a feature covariance matrix on ImageNet-1K validation set (*i.e.*, the covariance matrix has a shape of $50k \times 50k$), and compute the SVs of the covariance matrix. Fig. 4 shows a spectrum of log of singular value gaps between MAE and LUT across the layers. The singular values of LUT surpass the values of MAE across the rank indices in the last layers, while both methods have similar singular values on earlier layers. The results reveal that LUT have larger singular values at the output-side layers, indicating a higher rank of the feature space [21, 60]. In other words, LUT utilizes the output feature space better than MAE, owing to the longer-range understanding prompted by broader contextualization.

Table 5: Robustness evaluation. We evaluate the robustness of the representative ImageNet-1K pre-training methods: DINO, iBOT, MAE, and data2vec with our LUT on in-distribution generalization (IN-V2/Real) and out-of-distribution (IN-A/IN-O/Sketch/R/Cocc/Obj) benchmarks. We also evaluate the capability to detect spurious correlations with background on SI-Score metrics [14]. We highlight the best numbers (in boldface) and the second-best numbers (in underlined). [†] denotes the models pre-trained using multi-crop augmentation. Ours surpasses others than expected by the ImageNet-1K numbers, particularly more on localization-related metrics.

	IN-1k [†]	IN-V2 [†]	IN-Real [†]	IN-A [†]	IN-O [†]	Sketch [†]	IN-R [†]	Cocc [†]	ObjNet [†]	SI-size [†]	SI-loc [†]	SI-rot [†]
DINO [†]	83.1	72.8	87.6	36.3	60.7	35.7	48.2	77.8	36.4	57.8	37.0	43.8
iBOT [†]	83.5	73.5	87.9	39.4	62.0	37.8	50.2	78.6	37.1	<u>58.2</u>	37.6	<u>43.9</u>
MAE	83.7	72.9	88.2	36.7	65.4	35.9	48.9	78.4	37.6	58.0	<u>38.7</u>	42.7
data2vec	<u>84.1</u>	<u>74.2</u>	<u>88.5</u>	<u>41.6</u>	62.2	38.7	53.0	<u>79.1</u>	40.3	57.9	38.6	43.8
LUT	84.2	74.2	88.6	42.5	<u>64.1</u>	<u>38.2</u>	<u>52.1</u>	79.2	<u>38.9</u>	59.8	40.7	44.9

6.3 Robustness Evaluation

We evaluate the robustness of various methods, including DINO [8], iBOT [63], MAE [22], and data2vec [2] with LUT on various robustness benchmarks. We can interpret how our method impacts model robustness. We employ two in-distribution benchmarks including ImageNet-V2 [44] and -Real [5]) and four out-of-distribution benchmarks ImageNet-A [25], -O [25], -R [24], -Sketch [55], and ObjectNet [4]. We further use SI-Score [14] to test spurious correlations with the background. Lastly, we evaluate the center occlusion benchmark that zeroes the center patch in the ImageNet-1K evaluation images. As shown in Table 5, LUT achieves outstanding performance on all the benchmarks.

7 Conclusion

We have introduced a novel framework to address the limited broader understanding of images inherent in MIM. We argued that MIM, such as MAE, learns a narrower range of dependency due to lacking a comprehensive understanding of entire pixels. By visualizing attention maps, we have shown that MAE exhibits incomplete coverage of foreground or background regions. We conjectured this is caused by the potential absence of general context in learned unmasked tokens when interacting with mask tokens in self-attentions. Based on the observation, we have proposed LUT pre-training method, minimizing the discrepancy between the context features and sparse visual tokens through our broader contextualization loss. Our simple remedy enhanced MAE by a large margin on ImageNet-1K and ADE20K, and enables LUT to be very competitive to state-of-the-art methods. LUT further offers significant improvements in transfer learnings, including the iNatrealist and FGVC datasets. Finally, measuring averaged attention distance and spectral analysis demonstrated that LUT can be a simple yet effective supplement for masked image modeling.

Limitations. We verified our method’s applicability only up to ViT-L. Resource constraints prevented us from performing more experiments on larger models. We will further confirm the scalability with ViT-G or ViT-H.

References

1. Assran, M., Caron, M., Misra, I., Bojanowski, P., Bordes, F., Vincent, P., Joulin, A., Rabbat, M., Ballas, N.: Masked siamese networks for label-efficient learning. arXiv preprint arXiv:2204.07141 (2022) [7](#), [8](#), [9](#)
2. Baevski, A., Hsu, W.N., Xu, Q., Babu, A., Gu, J., Auli, M.: data2vec: A general framework for self-supervised learning in speech, vision and language. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S. (eds.) International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 162, pp. 1298–1312. PMLR (2022) [7](#), [8](#), [9](#), [14](#)
3. Bao, H., Dong, L., Piao, S., Wei, F.: Beit: Bert pre-training of image transformers. In: International Conference on Learning Representations (2021) [1](#), [7](#)
4. 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 (2019) [14](#)
5. Beyer, L., Hénaff, O.J., Kolesnikov, A., Zhai, X., Oord, A.v.d.: Are we done with imagenet? arXiv preprint arXiv:2006.07159 (2020) [14](#)
6. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European Conference on Computer Vision. pp. 213–229. Springer (2020) [1](#)
7. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: Proceedings of Advances in Neural Information Processing Systems (2020) [6](#)
8. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the International Conference on Computer Vision (2021) [1](#), [8](#), [9](#), [10](#), [13](#), [14](#)
9. Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15750–15758 (2021) [6](#)
10. Chen, X., Xie, S., He, K.: An empirical study of training self-supervised vision transformers. arXiv preprint arXiv:2104.02057 (2021) [1](#), [6](#), [8](#), [9](#), [10](#), [13](#)
11. Chen, Y., Liu, Y., Jiang, D., Zhang, X., Dai, W., Xiong, H., Tian, Q.: Sdae: Self-distillated masked autoencoder. In: European Conference on Computer Vision. pp. 108–124. Springer (2022) [8](#), [9](#)
12. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. In: Advances in Neural Information Processing Systems (2013) [6](#)
13. 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) [1](#), [7](#)
14. Djolonga, J., Yung, J., Tschannen, M., Romijnders, R., Beyer, L., Kolesnikov, A., Puigcerver, J., Minderer, M., D’Amour, A., Moldovan, D., et al.: On robustness and transferability of convolutional neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16458–16468 (2021) [2](#), [14](#)
15. Dong, X., Bao, J., Zhang, T., Chen, D., Zhang, W., Yuan, L., Chen, D., Wen, F., Yu, N.: Bootstrapped masked autoencoders for vision bert pretraining. arXiv preprint arXiv:2207.07116 (2022) [8](#), [9](#)
16. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.:

- An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021) [1](#), [2](#), [8](#), [12](#)
17. El-Nouby, A., Izacard, G., Touvron, H., Laptev, I., Jegou, H., Grave, E.: Are large-scale datasets necessary for self-supervised pre-training? arXiv preprint arXiv:2112.10740 (2021) [9](#)
 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) [9](#)
 19. Fu, L., Lian, L., Wang, R., Shi, B., Wang, X., Yala, A., Darrell, T., Efros, A.A., Goldberg, K.: Rethinking patch dependence for masked autoencoders. arXiv preprint arXiv:2401.14391 (2024) [9](#)
 20. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., Piot, B., koray Kavukcuoglu, Munos, R., Valko, M.: Bootstrap Your Own Latent - A New Approach to Self-Supervised Learning. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) Advances in Neural Information Processing Systems. vol. 33, pp. 21271–21284. Curran Associates, Inc. (2020) [6](#), [8](#), [10](#)
 21. Han, D., Yun, S., Heo, B., Yoo, Y.: Rethinking channel dimensions for efficient model design. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 732–741 (2021) [13](#)
 22. 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) [1](#), [2](#), [3](#), [4](#), [5](#), [7](#), [8](#), [9](#), [10](#), [13](#), [14](#)
 23. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. arXiv preprint arXiv:1911.05722 (2019) [6](#)
 24. Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.: The many faces of robustness: A critical analysis of out-of-distribution generalization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8340–8349 (2021) [14](#)
 25. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15262–15271 (2021) [14](#)
 26. Horn, G.V., Branson, S., Farrell, R., Haber, S., Barry, J., Ipeirotis, P.G., Perona, P., Belongie, S.J.: Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 595–604 (2015) [10](#)
 27. Huang, G., Fu, H., Bors, A.G.: Masked image residual learning for scaling deeper vision transformers. Advances in Neural Information Processing Systems **36** (2024) [9](#)
 28. Huang, Z., Jin, X., Lu, C., Hou, Q., Cheng, M.M., Fu, D., Shen, X., Feng, J.: Contrastive masked autoencoders are stronger vision learners. arXiv preprint arXiv:2207.13532 (2022) [7](#)
 29. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International Conference on Machine Learning. pp. 448–456. pmlr (2015) [8](#)
 30. Jing, L., Vincent, P., LeCun, Y., Tian, Y.: Understanding dimensional collapse in contrastive self-supervised learning. arXiv preprint arXiv:2110.09348 (2021) [13](#)

31. Khosla, A., Jayadevaprakash, N., Yao, B., Fei-Fei, L.: Novel dataset for fine-grained image categorization. In: First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition. Colorado Springs, CO (2011) 10
32. Kim, T., Han, D., Heo, B.: Morphing tokens draw strong masked image models. arXiv preprint arXiv:2401.00254 (2023) 1
33. Kornblith, S., Shlens, J., Le, Q.V.: Do better imagenet models transfer better? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2661–2671 (2019) 10
34. Krizhevsky, A.: Learning multiple layers of features from tiny images. In: Tech Report (2009) 10
35. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Communications of the ACM* **60**(6), 84–90 (2017) 8
36. Li, G., Zheng, H., Liu, D., Wang, C., Su, B., Zheng, C.: Semmae: Semantic-guided masking for learning masked autoencoders. arXiv preprint arXiv:2206.10207 (2022) 8, 9
37. Liu, X., Zhou, J., Kong, T., Lin, X., Ji, R.: Exploring target representations for masked autoencoders. arXiv preprint arXiv:2209.03917 (2022) 2
38. Loshchilov, I., Hutter, F.: SGDR: Stochastic gradient descent with warm restarts. In: International Conference on Learning Representations (2017) 8
39. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (2019) 8
40. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. arXiv preprint arXiv:1306.5151 (2013) 10
41. Mishra, S., Robinson, J., Chang, H., Jacobs, D., Sarna, A., Maschinot, A., Krishnan, D.: A simple, efficient and scalable contrastive masked autoencoder for learning visual representations. arXiv preprint arXiv:2210.16870 (2022) 8, 9
42. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing (2008) 10
43. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: ICML (2021) 9
44. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do imagenet classifiers generalize to imagenet? In: International Conference on Machine Learning (2019) 14
45. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* **115**(3), 211–252 (2015) 2, 8
46. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1–9 (2015) 6, 11
47. Tao, C., Zhu, X., Su, W., Huang, G., Li, B., Zhou, J., Qiao, Y., Wang, X., Dai, J.: Siamese image modeling for self-supervised vision representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2132–2141 (2023) 8, 9
48. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jegou, H.: Training data-efficient image transformers & distillation through attention. In: International Conference on Machine Learning. vol. 139, pp. 10347–10357 (2021) 1, 8, 9, 13

49. Touvron, H., Cord, M., Jégou, H.: Deit iii: Revenge of the vit. In: Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV. pp. 516–533. Springer (2022) [1](#), [6](#), [8](#), [9](#), [11](#)
50. Touvron, H., Cord, M., Oquab, M., Bojanowski, P., Verbeek, J., Jégou, H.: Co-training 2l submodels for visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11701–11710 (2023) [9](#)
51. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8769–8778 (2018) [10](#)
52. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is All you Need. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 30. Curran Associates, Inc. (2017) [1](#)
53. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset. Tech. Rep. CNS-TR-2011-001, California Institute of Technology (2011) [10](#)
54. Wang, H., Song, K., Fan, J., Wang, Y., Xie, J., Zhang, Z.: Hard patches mining for masked image modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023) [8](#), [9](#)
55. 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 (2019) [14](#)
56. Wei, C., Fan, H., Xie, S., Wu, C.Y., Yuille, A., Feichtenhofer, C.: Masked feature prediction for self-supervised visual pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022) [8](#), [9](#)
57. Wu, Z., Lai, Z., Sun, X., Lin, S.: Extreme masking for learning instance and distributed visual representations. arXiv preprint arXiv:2206.04667 (2022) [8](#), [9](#)
58. Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: European Conference on Computer Vision. Springer (2018) [10](#)
59. 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: International Conference on Computer Vision (2022) [1](#), [3](#), [7](#), [8](#), [9](#)
60. Yang, Z., Dai, Z., Salakhutdinov, R., Cohen, W.W.: Breaking the softmax bottleneck: A high-rank rnn language model. In: International Conference on Learning Representations (2018) [13](#)
61. Yi, K., Ge, Y., Li, X., Yang, S., Li, D., Wu, J., Shan, Y., Qie, X.: Masked image modeling with denoising contrast. International Conference on Learning Representations (2023) [7](#), [8](#), [9](#)
62. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2017) [2](#), [10](#)
63. Zhou, J., Wei, C., Wang, H., Shen, W., Xie, C., Yuille, A., Kong, T.: ibot: Image bert pre-training with online tokenizer. In: International Conference on Learning Representations (2022) [1](#), [7](#), [8](#), [9](#), [10](#), [13](#), [14](#)