

Attention-Guided Masked Autoencoders For Learning Image Representations

Leon Sick
Ulm University

Dominik Engel
Ulm University

Pedro Hermosilla
TU Vienna

Timo Ropinski
Ulm University

Abstract

Masked autoencoders (MAEs) have established themselves as a powerful pre-training method for computer vision tasks. While vanilla MAEs put equal emphasis on reconstructing the individual parts of the image, we propose to inform the reconstruction process through an attention-guided loss function. By leveraging advances in unsupervised object discovery, we obtain an attention map of the scene which we employ in the loss function to put increased emphasis on reconstructing relevant objects. Thus, we incentivize the model to learn improved representations of the scene for a variety of tasks. Our evaluations show that our pre-trained models produce off-the-shelf representations more effective than the vanilla MAE for such tasks, demonstrated by improved linear probing and k -NN classification results on several benchmarks while at the same time making ViTs more robust against varying backgrounds and changes in texture.

1. Introduction

Despite their emergence as a promising alternative to CNNs in computer vision, transformers bring their own training challenges to the table: unstable training, high computational demand, as well as the need for large datasets to achieve competitive results [25]. Recent work has shown that self-supervised Vision Transformer (ViT) pre-training [2, 5, 11, 20, 61] can help mitigate some of these issues and improve their performance for various vision tasks. While there have been different pre-training strategies proposed, one idea that has gained much traction is masking portions of an input image and training a ViT to reconstruct the missing parts of the image. This task has been shown to be most effective in the setting of an MAE [20], which reconstructs the image with equal emphasis on each individual part. While MAEs achieve state-of-the-art representations for end-to-end finetuning tasks, it has also been shown that for linear evaluations, state-of-the-art results cannot be achieved [1, 20]. With our work, we attempt to make progress towards closing this performance gap for MAEs on linear evaluations and k -NN classification. Along this

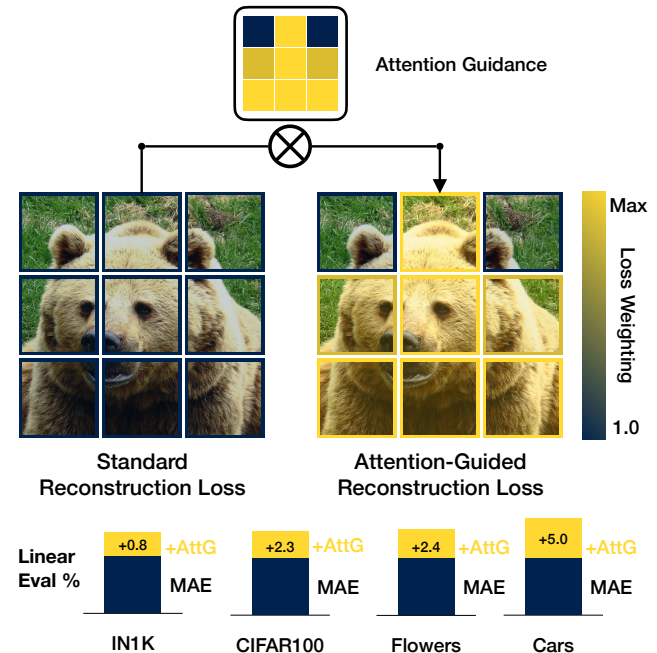


Figure 1. **Attention-Guided Reconstruction Loss.** We leverage an attention map, obtained from a self-supervised object discovery network, to inform our guided reconstruction loss *AttG* towards reconstructing the relevant objects in the scene. Through this process, we guide the MAE to learn more effective off-the-shelf representations.

line, recent works explore linear probing and k -NN evaluations over end-to-end fine-tuning [1, 10, 19, 44, 45], and off-the-shelf feature quality is also shown to be of interest to the community [16, 49, 50, 56, 59]. Methods from the field of explainability [7, 42, 46, 57] suggest that successful image classification models should mainly focus on the relevant objects in the scene to make their predictions. However, recent research has shown [31] that many failure cases in state-of-the-art models are due to the focus of these models on background elements. Therefore, we hypothesize that inducing high-level semantic knowledge, i.e. information about relevant foreground objects, into the MAE pre-training enables the model to learn more effective representations for tasks that require robust image-level understanding. Hence,

with our method, *AttG*, we propose to induce information on the relevant objects in the scene through an attention-guided reconstruction loss. While the vanilla MAE does not exhibit a notable deficiency in its capability to reconstruct objects, we show that our semantic guidance enables the model to output representations which contain more semantically relevant information. By manipulating the reconstruction loss, our model can be guided by an attention map to put increased emphasis on reconstructing relevant objects (see Fig. 1), which according to our ablations produces pre-trained features with improved image-level understanding as compared to changing the established input masking strategy [23, 28]. To generate the employed attention maps, we explore several object discovery methods, both ViT- and CNN-based. With our approach, we aim to shift the training process of the developing a semantic understanding across datasets to the pre-training phase, as to reduce the need for more expensive adaptation mechanisms like end-to-end finetuning. We therefore focus on evaluations that only train a linear layer or directly harness off-the-shelf features. We show the effectiveness of the model’s learned representations across a wide range of tasks and benchmarks that require robust image-level understanding, like linear and k-NN classification, as well as for low-shot finetuning settings or image retrieval tasks. We further demonstrate that these improved properties of the model are also preserved for datasets outside of the pre-training data domain. Even for texture-heavy datasets, where the shape of the object is mostly consistent and only the texture is changed, our model outperforms the off-the-shelf features of the vanilla MAE for linear classification.

2. Related Work

Our work builds upon recent advances in masked modeling, as well as self-supervised pre-training.

Masked modeling. Exploring a reconstruction task of masked regions in the input modality has first seen its successes in language processing with transformer architectures [15, 51] in BERT [14] or GPT [3, 39, 40]. More recently, these methods have been successfully transferred to the computer vision domain. In their work on masked autoencoders, He *et al.* [20] mask a large number of image patches and feed the remaining tokens to the ViT [15] encoder. Afterwards, the masked tokens are added and fed through the ViT decoder with the goal of reconstructing the masked patches. Feichtenhofer *et al.* [17] and Tong *et al.* [48] extend this idea into the spatio-temporal domain for video data, while Girdhar *et al.* [18] propose a unified MAE architecture for both image and video data. Recently, masking strategies have also been applied for generative image transformers in work done by Chang *et al.* [6]. Their model learns to predict the initially masked tokens of the training image. To generate high-fidelity images at inference time,

they sample the model to iteratively predict the tokens to be reconstructed and thereby refine the image. Li *et al.* [28] propose to adapt a semantic-guided masking strategy for pre-training an MAE. They extract segmentation parts by refining a ViT encoder and use these to mask a portion in the input image. While being effective, their approach requires additional finetuning on their segmentation model before extracting the segmentation parts. In contrast, our method does not require this preparation step and we do not alter the original masking strategy, as it has proven to be beneficial by He *et al.* [20]. *AttnMask*, a novel masking strategy proposed by Kakogeorgiou *et al.* [23], is integrated into a student-teacher knowledge distillation training process where the teacher network outputs an attention map that is used to mask the input image for the student network. They also extend their masking approach to keep some patches of the attention map visible, giving the model additional “hints”. In their work on SimMIM, Xie *et al.* [55] successfully design an approach for masked image modelling (MIM), for which they formulate a raw pixel regression task to predict the reconstruction of all masked pixels using a prediction head on top of the latent representation output by their encoder. Bao *et al.* [2] propose BEiT, which in contrast predicts masked patches in the token space, which is created from the image patches by a frozen discrete VAE. Zhuge *et al.* [62] view MIM through the lense of sample selection and introduce a masking strategy which determines a weighted selection score to identify which patches to mask for each tasking iteration.

Self-supervised pre-training. Recent work [4, 5, 8, 9, 30, 47] has shown significant advances in self-supervised pre-training for vision models with various backbone architectures like ResNets [22] and ViTs [15]. Chen *et al.* [8] explore pre-training for ResNets and highlight that effective data augmentation is a critical factor for pre-training on images. Their approach is formulated so that its performance is dependent on the batch size. Caron *et al.* [4] built upon these findings and introduced a multi-crop augmentation strategy as well as a novel learning approach to eliminate the batch size dependence. Pre-training of ResNets is further pushed by works from Mitrovic *et al.* [30] and Tomasev *et al.* [47]. Their works propose to formulate a probabilistic learning objective through an invariance regularization mechanism, built upon the Kullback-Leibler divergence. Research by Caron *et al.* [5] puts special emphasis on exploring self-supervised pre-training with ViTs. They investigate the use of knowledge distillation and find that their ViTs can provide a segmentation of the scene and also perform well on k-NN classification. With MoCo v3 [11], Chen *et al.* also explore the properties of ViTs for pre-training in an evolution of their initial methods [9, 21].

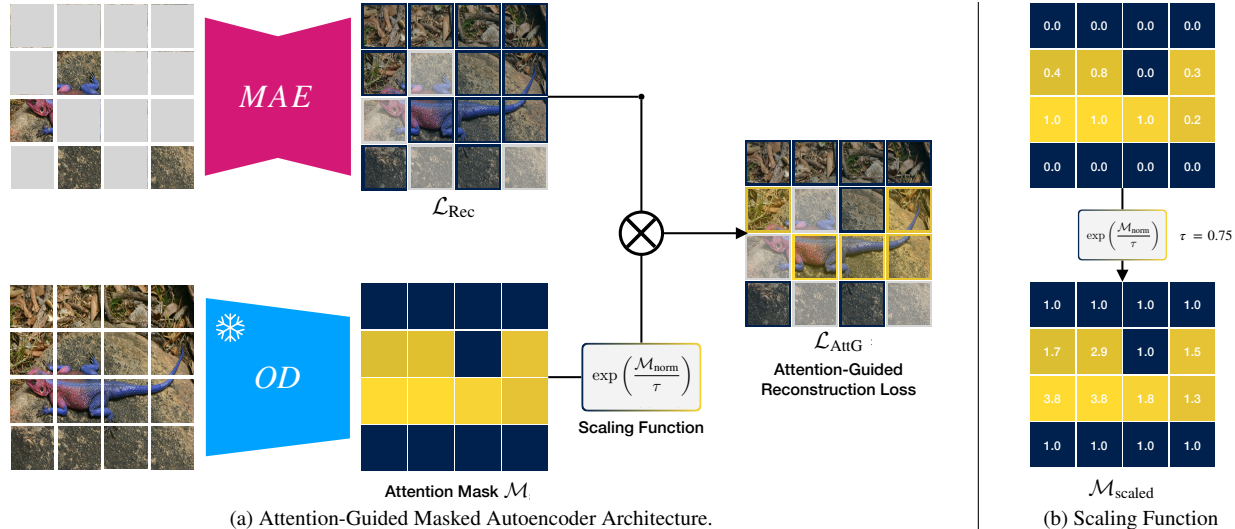


Figure 2. **AttG Architecture Overview And Scaling Function.** Our architecture, displayed in Figure 2a, employs two streams. Besides the main *MAE* backbone, we facilitate an unsupervised object discovery (*OD*) stream with fixed weights to obtain a patch-based attention map. This attention map is then scaled and finally used to inform our attention guided reconstruction loss *AttG* towards relevant objects. Gray overlays depict image patches which are not part of the loss since they have not been masked in the input image. To make the attention map useful for reconstruction guidance, we scale its values, as visualized in Figure 2b. This results in the loss for background patches not being masked out and patches from the relevant object in the scene being further emphasized.

3. Method

In our work, we build upon vanilla MAEs [20]. Our aim is to enable the model to develop a better high-level understanding of the semantics of the scene, i.e. distinguishing foreground from background, and therefore learning representations more effective for a multitude of image-level tasks. As illustrated in Fig. 2, we facilitate unsupervised object discovery (OD) and use the obtained results within a temperature-controlled guidance loss to steer the model towards the relevant objects. In the following subsections, we will detail our learning architecture (Section 3.1), introduce our attention-guided loss (Section 3.2), and discuss unsupervised OD (Section 3.3).

3.1. Learning Architecture

Our learning architecture exploits two input streams as illustrated in Fig. 2. As with conventional MAEs, the backbone of our learning architectures is a ViT-based autoencoder, which is trained to reconstruct an input image based on a masked version of the same image. To achieve this, the autoencoder first encodes the input to a latent representation, which is then fed into a decoder to perform the reconstruction process. To feed the input into the transformer, we first divide it into non-overlapping patches and apply patch-based masking, before the remaining visible patches are fed into the encoder. In the latent space, the masked patches are considered again and the original amount of tokens is fed into the decoder for reconstruction. The reconstruction

loss is computed on the same per-patch basis, as the masking is obtained. To guide the reconstruction loss towards the objects of interest, our architecture facilitates an additional stream for unsupervised OD. Within this stream, a patch-based attention map, focusing on relevant objects, is obtained from the input image. As we want to use this attention map to steer the reconstruction loss, it is important to obtain patches with the same sizes, as the MAE backbone processes.

3.2. Attention-Guided Loss

In order to exploit the attention map obtained in our OD stream, we need to define a guiding mechanism for the model. While directly applying the attention map \mathcal{M} to generate an object-based masking on the input image might seem intuitive, He *et al.* [20] show that a high masking ratio in combination with random sampling significantly outperforms a block-wise masking strategy as proposed by Bao *et al.* [2] with a lower masking ratio. Since the latter is similar to masking objects of interest in the input image, simple object-based patch masking cannot be assumed to be a promising strategy. Therefore, we opt to not interfere with the input masking strategy, and rather steer the learning process by using the attention maps to guide the reconstruction loss. To employ our guided loss instead, we first compute the MAE reconstruction loss [20], i.e. per-patch mean squared error (MSE) between reconstructed patches and their originals $\mathcal{L}_{\text{Rec}} = (\hat{\mathcal{X}} - \mathcal{X}_{\text{norm}})^2$.

This allows us to guide the model by individually weighting different parts of the image. We min-max normalize our attention maps using $\mathcal{M}_{\text{norm}} = (\mathcal{M} - \mathcal{M}_{\text{min}})/(\mathcal{M}_{\text{max}} - \mathcal{M}_{\text{min}})$ with value range $\mathcal{M}_{\text{norm}} \in [0, 1]$, i.e. all patches that are part of the object are close to 1. Using the map in this form would not be useful, since background patches with values of 0 would lead to our attention maps zeroing-out all loss values for background patches, effectively rendering the model unable to learn to reconstruct the entire image. Our goal is instead, to keep loss values of background patches untouched. To achieve this, we scale our attention maps with the exponential function and a temperature parameter $\tau > 0$:

$$\mathcal{M}_{\text{scaled}} = \exp\left(\frac{\mathcal{M}_{\text{norm}}}{\tau}\right) \quad (1)$$

leading to background patches with values of 0 being converted to 1, so that, when weighting the reconstruction, loss values for background patches are not additionally emphasized. At the same time, the attention map values for object patches are scaled, giving them a stronger impact when weighting the loss. By adding a temperature parameter τ , we obtain an additional hyperparameter which gives us fine-grained control over the strength of weighting.

With this, we arrive at our guided loss

$$\mathcal{L}_{\text{AttG}} = \gamma \cdot \mathcal{L}_{\text{Rec}} \cdot \mathcal{M}_{\text{scaled}} \quad (2)$$

with γ being the binary mask of all masked tokens. We schedule τ to increase according a half-cycle cosine temperature schedule starting at $\tau = 0.75$, which we find most effective as detailed in Section 5.4. Our full training process is detailed in Algorithm 1.

Algorithm 1 Pseudocode Attention Map Guidance

Require: image \mathcal{X} , attention map \mathcal{M} ,
masking function **mask**, model **MAE**
for all \mathcal{X}, \mathcal{M} in Dataset **do**
 $\mathcal{X}_{\text{masked}}, \gamma \leftarrow \mathbf{mask}(\mathcal{X}, \text{ratio} = 0.75)$
 $\hat{\mathcal{X}} \leftarrow \mathbf{MAE}(\mathcal{X}_{\text{masked}})$
 $\mathcal{M}_{\text{norm}} \leftarrow (\mathcal{M} - \mathcal{M}_{\text{min}})/(\mathcal{M}_{\text{max}} - \mathcal{M}_{\text{min}})$
 $\mathcal{M}_{\text{scaled}} \leftarrow \exp\left(\frac{\mathcal{M}_{\text{norm}}}{\tau}\right)$ (Eq. 1)
 $\mathcal{X}_{\text{norm}} \leftarrow (\hat{\mathcal{X}} - \mu_{\hat{\mathcal{X}}})/(\sigma_{\hat{\mathcal{X}}} + \epsilon)^{0.5}$
 $\mathcal{L}_{\text{Rec}} \leftarrow (\hat{\mathcal{X}} - \mathcal{X}_{\text{norm}})^2$
 $\mathcal{L}_{\text{Guided}} \leftarrow \gamma \cdot \mathcal{L}_{\text{Rec}} \cdot \mathcal{M}_{\text{scaled}}$ (Eq. 2)
end for

3.3. Object Discovery

The effectiveness of our guidance loss introduced above, stands and falls with the quality of the exploited attention maps. We ensure wide applicability of our method by exploiting state-of-the-art unsupervised OD approaches.

Based on their reported results as well as their widespread use, we have investigated three unsupervised OD approaches to inform our guided loss: a DINO-pretrained ViT [5], a TokenCut ViT [53], and a Grad-CAM-based approach [42]. While the first two are transformer-based approaches, which directly deliver attention maps, Grad-CAM is often praised for its capabilities to extract structures highly relevant for a learner, and has therefore also been considered. With these three approaches, we generate an attention map \mathcal{M} of the object in the scene that serves as a foreground-background segmentation.

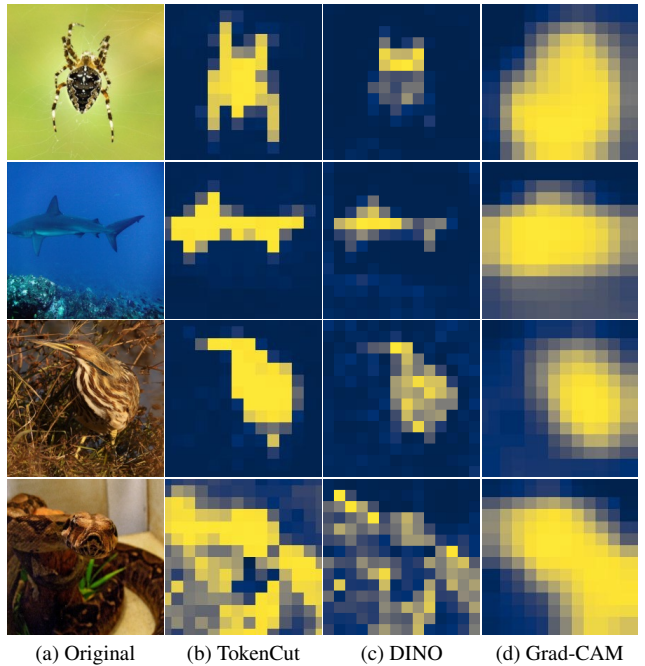


Figure 3. **Visual Comparison Of Our Attention Maps.** We illustrate the differences between our attention maps. Each has been normalized and scaled in the manner as we use them for pre-training our models. From these examples, TokenCut can be perceived to output the best foreground-background separation maps of all presented methods. We study the effect of the different maps on our guidance in Section 5.1.

For the pre-trained DINO [5] ViT, the attention map is obtained, by feeding the image through the network and querying the self-attention module of the last block. This self-attention module with H heads and consequently $j = 1, \dots, H$ sequences takes as input the normalized output sequence from the previous block $Y \in \mathbb{R}^{(n+1) \times d}$ and feeds it through three linear layers to map it to the sequences for query Q_j , key K_j and value V_j , all $\in \mathbb{R}^{(n+1) \times d'}$. The attention matrix is then calculated as follows:

$$\mathcal{A}_j = \text{softmax}(Q_j K_j^T / \sqrt{d'}) \quad (3)$$

with $d' = d/H$, allowing for the maps of all attention heads

to be retrieved. For DINO, we use the attention maps without post processing.

TokenCut ViT [53] post-processes the segmentation-like attention maps \mathcal{A}_j of the pre-trained DINO transformer to obtain a more refined segmentation. Briefly described, TokenCut constructs an undirected graph of patches with each having assigned a feature vector. Each patch is linked to other patches through edges with the cosine similarity as edge capacity. The normalized cut algorithm [43] is then used to obtain the final segmentation output \mathcal{M} .

The Grad-CAM-based approach [42] visualizes the gradients for a certain prediction. The result is a segmentation-like attention map. Differently thought, this approach originally requires labels to inform the process. Since we do not have any labels available, we select the highest scoring category as our target. Since the Grad-CAM model will output a pixel-wise heatmap, we apply average pooling to obtain patches of the same patch-wise format as our MAE and the other two ViTs. In our approach we use a supervised ImageNet-1K-pretrained ResNet-50 as the encoder. Figure 3 shows a visual comparison of the segmentation outputs. From qualitative observations, TokenCut provides the most semantically accurate attention maps with evenly valued patches and a clear distinction between object and background. DINO’s maps also provide a useful semantic layout of the scene but are lacking the same distinctiveness. Grad-CAM on the other hand outputs heatmap-like masks that provides a coarser map. In our ablations (Section 5.1), we explore the effectiveness of these different maps for our guidance mechanism.

4. Experiments

We pre-train our model on ImageNet [13] (IN1K). Our attention maps are generated by TokenCut, as they perform the best according to our ablations in Section 5.1. As emphasized in Section 3, with our attention guidance, we aim to enable the model to learn an improved understanding of higher-level semantics important for image-level tasks. Therefore, we evaluate the learned representations of our model through finetuning a only linear layer on top of them, i.e. linear probing, performing k-NN classification studies, both on IN1K and other relevant benchmarks, as well as low-shot finetuning and image retrieval studies.

4.1. ImageNet Evaluation

We collect several comparable methods in the field and report k-NN classification and linear probing results on ImageNet-1K in Table 1. For a leveled comparison, Table 1 evaluates the vanilla MAE and our attention guidance method, AttG, against other pure MIM-based approaches without a contrastive objective like SimMIM [55], BEiT [2] and SemMAE [28]. We show that our method is able to outperform other purely MIM-based techniques. With our

Method	Size	Epochs	k-NN	Linear
SimMIM	ViT-B/16	800	9.4	56.7
SemMAE	ViT-B/16	800	45.1	65.0
BEiT	ViT-L/16	800	11.4	73.5
MAE	ViT-L/16	800	52.8	73.5
+ AttG (Ours)	ViT-L/16	800	<u>56.2</u>	<u>74.4</u>
MAE	ViT-L/16	1600	50.9	75.1
+ AttG (Ours)	ViT-L/16	1600	59.0	75.9

Table 1. *k*-NN classification and linear probing on IN1K for purely MIM-based techniques. We show results for linear probing and *k*-NN classification, for IN1K we use *k*=20. For both training durations, our method improves upon the results of the MAE for both *k*-NN and linear probing.

Method	Size	Epochs	1%	10%
SimMIM	ViT-B/16	800	23.2	66.3
SemMAE	ViT-B/16	800	52.8	70.6
BEiT	ViT-L/16	800	38.3	71.6
MAE	ViT-L/16	1600	67.1	78.5
+ AttG (Ours)	ViT-L/16	1600	67.5	78.1

Table 2. **Low-shot finetuning on IN1K.** We finetune the pre-trained models with 1% and 10% of IN1K training data and evaluate on the full validation set. For a fair comparison, we list the best performing published model of all other methods. Our model performs especially well in the 1% regime.

attention guidance, our model improves upon the vanilla MAE on top-1 accuracy for linear probing by **+0.9%** after 800 pre-training epochs and **+0.8%** after 1600 epochs, showcasing the effectiveness of our learned off-the-shelf representations. We also perform *k*-NN classification with *k*=20 on ImageNet-1K. At 1600 epochs, we are able to outperform the MAE baseline by **+8.1%**. We also evaluate our model on low-shot finetuning with 1% and 10% of training data and report the results in Table 2. We see that our model performs best in for the 1% subset.

Additionally, we experiment with k-NN classification in a few-shot setting with the 1600-epoch models. We evaluate our model with 1, 5, 10 and 20 training images on the full validation dataset. Our results in Table 3 show our model outperforming for all setups, while the delta increases from **+4.4%** to **+13.5%** with an increasing number of training images, which underlines the quality of our representations. We also evaluate on IN1K in a few-shot linear probing setting. As presented in Table 8, our attention guidance improves the results of the MAE across all subsets.

Method	#	1	5	10	20
MAE		6.5	12.1	15.6	20.2
+ <i>AttG (Ours)</i>		10.9	19.8	26.8	33.7

Table 3. *k*-NN classification with few embedded training examples. We report *k*-NN classification top-1 accuracy for *k*=5 with an IN1K subset of $\# \in \{1, 5, 10, 20\}$ images. With an increasing number of training images, the delta by which our model outperforms the vanilla MAE increases. Both models have been pretrained for 1600 epochs.

4.2. Comparison to Attention-Masking Add-On Methods

We also compare our method against other approaches that use attention maps to build upon a pre-existing baseline, such as AttnMask [24], an extension of iBOT [61], and SemMAE [28], which also builds upon the vanilla MAE. Since these methods build upon different baselines, and therefore also start at varying levels of performance, we attempt to isolate their effects by reporting the delta in performance these methods exhibit versus their respective baselines. Therefore, in Table 4, we report the gains of AttnMask compared to iBOT, whereas for SemMAE and our method, we report additional performance versus the MAE baseline. We show that our method leads to the overall greatest improvement on k-NN performance. While the performance gains for linear probing are below those of AttnMask, our method for guidance with high-level semantic information is the most effective for MAEs.

4.3. Dense Downstream Evaluations

We also evaluate our model’s ability to perform more complex vision tasks. For this, we perform evaluation for semantic segmentation on ADE20k [60] and the NYUv2 [32] dataset. For this, we train a ConvNeXt-like [29] segmentation head on top of our representations. Our results in Table 5 show that our model consistently outperforms the vanilla MAE on pixel accuracy and mIoU. Further, we train a DPT-like [41] adapter on top of our model to perform experiments on the Taskonomy [58] benchmark. Specifically, we train the model to predict Depth, Edges, 2D Keypoints, 3D Keypoints and Occlusion, then report the average loss and rank. Table 5 shows our model outperforming the vanilla MAE again and we present the detailed results for each specific task in the Supplementary.

4.4. Robustness to Background Changes

ImageNet-9 [54] (IN-9) provides a suitable dataset for evaluating robustness against background variations in image classification. It includes validation datasets for nine classes from ImageNet-1K [13] where either the background or the object in the foreground have been altered,

Baseline	Add-On	Epochs	k-NN	Linear
iBOT	+ <i>AttnMask</i>	100	+1.3	+1.7
	+ <i>SemMAE</i>	800	-	+0.7
MAE	+ <i>AttG (Ours)</i>	800	+3.4	+0.9
	+ <i>AttG (Ours)</i>	1600	+8.1	+0.7

Table 4. Comparison of performance deltas from attention-informed approaches that are built on top of different baselines. We compare our method to SemMAE, another method that builds upon MAE, and AttnMask, a technique on top of iBOT. We report gains for linear probing and k-NN on IN1K.

	ADE20k		NYUv2		Taskonomy
	pAcc	mIoU	pAcc	mIoU	Avg Rank
MAE	75.8	35.9	66.9	35.6	1.8
+ <i>AttG (Ours)</i>	75.9	36.2	67.2	35.7	1.2

Table 5. Dense Downstream Evaluations. We conduct semantic segmentation experiments on ADE20K [60] and NYUv2 [32] for both of which our model outperforms the vanilla MAE. Further, we evaluate our model on a variety of tasks in the Taskonomy [58] dataset. On average, our model improves upon the MAE.

resulting in eight different benchmarks. We perform linear probing on the larger training set, ImageNet-9L (IN-9L). IN-9 contains four validation subsets which are suitable for evaluating a model’s robustness to background changes: Only-Foreground (OF), Mixed-Next (MN), Mixed-Rand (MR), and Mixed-Same (MS). These subsets contain images with modified backgrounds, of which MN and MR contain out-of-distribution (OOD) backgrounds. MS is edited with a different in-distribution backgrounds, while OF only has the background removed. We report top-1 accuracy for linear probing in Table 6. Our model is able to outperform the MAE on all four tasks.

4.5. Transfer to Other Datasets

To evaluate the capability of our model to extract relevant features from not only ImageNet, but also other benchmarks, we perform k-NN classification and linear probing on the CIFAR-100 [27], Oxford Flowers [33], and Stanford Cars [26] datasets. We present our results in Table 7. Across all tasks, our model outperforms the vanilla MAE. This yet again shows our model’s capability to extract valuable features, even for datasets other our pre-training dataset. We especially see strong results on CIFAR-100, where the vanilla MAE is outperformed by **+11.5%** on k-NN classification accuracy and **+2.3%** for linear probing.

4.6. Few-Shot Transfer

We also explore effectiveness of our learned representations in different few-shot linear probing settings. For



Method	OF	MS	MR	MN
MAE	93.6	92.8	81.9	78.9
+AttG(Ours)	93.9	93.9	83.9	81.1

Table 6. **Background influence evaluation on ImageNet-9 [54].** Our model has an edge on top-1 accuracy over the vanilla MAE on three benchmarks with modified backgrounds (OF, MR, and MN), showcasing the effectiveness of our method for pre-training robust ViTs. We perform linear probing on IN-9L.

Dataset	Method	Epochs	k-NN	Linear
CIFAR100	MAE	1600	48.9	74.1
	+ AttG (Ours)	1600	60.4	76.4
Flowers	MAE	1600	55.0	72.8
	+ AttG (Ours)	1600	66.2	75.2
Cars	MAE	1600	10.3	24.3
	+ AttG (Ours)	1600	13.2	29.3

Table 7. **k-NN classification and linear probing on different datasets.** We evaluate our method on a wide variety of transfer datasets by performing linear probing and k-NN evaluation and observe consistent performance improvements with our method. Our evaluation settings are equivalent to those in Table 1.

Dataset	Method	#	10	25	50
IN1K	MAE		39.3	56.0	62.3
	+ AttG (Ours)		43.3	57.7	63.1
CIFAR-100	MAE		33.0	39.6	52.0
	+ AttG (Ours)		35.4	42.8	55.7
Cars	MAE		11.0	19.3	23.2
	+ AttG (Ours)		12.9	23.6	29.4
DTD	MAE		46.4	53.6	-*
	+ AttG (Ours)		46.4	53.8	-*
Pets	MAE		59.1	69.9	71.3
	+ AttG (Ours)		69.9	78.9	80.0

Table 8. **Few-shot classification using linear probing on transfer datasets.** We report top-1 accuracy results for few-shot classification when linear probing both models on $\# \in \{10, 25, 50\}$ images per class. Both models have been pretrained for 1600 epochs. *DTD train split contains only 40 images per class.

both CIFAR-100 and Stanford Cars, we limit the amount of training images to 10, 25, and 50. To test our methods robustness against changes in texture, we also evaluate on the Describable Textures Dataset (DTD) [12] and the Oxford-IIIT Pets dataset [34]. We perform linear probing on the reduced datasets and evaluate the models’ performance on the full validation splits. As detailed in Table 8, our ap-

proach is able to outperform the vanilla MAE by substantial margins throughout the entire range of subset configurations and across multiple datasets. These results demonstrate once more that our attention-guided loss is highly effective in enabling MAEs to learn better representations.

4.7. Image Retrieval

We evaluate our learned representations with regard to their effectiveness on the revisited [35] Oxford and Paris image retrieval datasets [37]. Our evaluation protocol follows Caron *et al.* [5]. We freeze our model and retrieve the image pairs using a k -NN based on the 1024-d output features. We report the mAP for the Medium (M) and Hard (H) splits in Table 9. Our model produces representations that improve upon the vanilla MAE across the board.

5. Ablations

For ablations, we pre-train for 400 epochs on ImageNet-100, an IN1K subset, for computational efficiency.

5.1. Attention Map Quality

We experiment with different approaches for generating our attention maps. An overview of qualitative differences can be observed in Figure 3. The TokenCut attention map covers many details, such as the spider legs or the shark fins. DINO on the other hand outputs an attention map mostly covering the object but not with equal emphasis. In contrast, the Grad-CAM output generously overlays the entire animal but also intersects many background patches. In Table 10, we report the influence of using different OD approaches for deriving the attention maps informing our reconstruction loss. All models are pre-trained with an input masking ratio of 0.75 and a half-cycle cosine temperature schedule starting at $\tau = 0.75$. As expected based on the qualitative results, TokenCut maps perform best with a clear advantage over the other two OD approaches. This confirms our intuition that high-quality attention maps should yield the best results, and proves the effectiveness of our approach, since the model seems to be able to extract most valuable semantic knowledge from highest quality maps. In addition, we experiment with attention maps retrieved from Grad-CAM++ [7], MaskCut [52], DINO trained with SimPool [36] and from the CLIP [38] ResNet with Attention. From our experiments, we find that maps from TokenCut are most effective for our method.

5.2. Computational Overhead

Applying our attention maps to the reconstruction of the MAE comes with only marginally increased computational costs. We observed an increase in pre-training runtime of below 1% over the vanilla MAE. Our attention maps can be generated during training, but to ensure reproducibility of our results, we decide to generate them once ahead of

Dataset	\mathcal{ROx}		\mathcal{RPar}	
	M	H	M	H
Method				
MAE	13.1	2.3	24.4	5.2
+ <i>AttG (Ours)</i>	19.1	3.0	33.0	8.4

Table 9. **Image retrieval.** We evaluate on the medium (M) and hard (H) image retrieval tasks of the Oxford and Paris datasets. We show that our off-the-shelf features enable improved performance.

Model	Linear	k-NN
TokenCut [53]	77.1	57.1
DINO [5]	75.7	56.0
Grad-CAM [42]	76.3	57.0
Grad-CAM++ [7]	76.2	56.6
MaskCut [52]	73.8	56.7
SimPool [36]	73.7	53.4
CLIP [38]	75.7	48.0

Table 10. **Comparison of different OD approaches.** The TokenCut maps are most effective for our approach, consistent with Figure 3, where TokenCut provides formidable attention maps.

Attention Guidance	k-NN	Linear
Input Masking	39.2	57.1
Foreground-Only	56.7	74.8
Background-Only	43.5	69.2
Inverted Attention Map	48.7	72.6
<i>AttG</i> Reconstruction Guidance (Ours)	57.0	77.1

Table 11. **Ablation of different uses of the attention map.** Guiding the MAE through the reconstruction loss yields better results, both for k-NN classification with $k=5$ and linear probing.

pre-training. Obtaining an attention map requires a single forward pass of the object discovery network, making it a negligible overhead in relation to the total training time.

5.3. Guidance Mechanism

One alternative to our proposed attention-guided reconstruction loss would be to inform the initial *Input Masking* of the input image through the attention maps [24, 28]. Therefore, we investigate a modified masking strategy with a masking ratio of 75% and sampling the masked patches from the attention map in a descending fashion, removing patches with the highest attention values. We assume this can mask the relevant object as well as some additional patches. Using this approach together with the original MAE loss results in inferior results across our evaluations. We experiment with further alternatives to guiding the reconstruction loss with our attention maps. We try reconstructing *Foreground-Only* by zeroing all patches with attention values below the 10%-quantile threshold. Furthermore, to emphasize background reconstruction, we use the *Inverted Attention Maps* to guide the MAE and also mask

τ	With Schedule	Without Schedule
0.7	76.1	76.3
0.75	77.1	76.6
0.8	76.5	77.0

Table 12. **Experiments with different temperatures and scheduling.** We observe the best linear probing results when $\tau = 0.75$ and is increased with a half-cycle cosine schedule.

out the 10%-quantile threshold of this inverted map to reconstruct *Background-Only*. In Table 11, we report the top-1 accuracy for k-NN classification and linear probing on IN-100, where our guidance outperforms the alternatives.

5.4. Temperature Parameter Influence

In Table 12, we evaluate the influence of different temperature parameters and our half-cycle temperature schedule. For fixed values, we observe that $\tau = 0.8$ leads to the best linear probing results. But when slowly increasing τ from 0.8 on, we see a drop in performance, fueling the belief that further increasing the temperature is ineffective. A decrease of τ leads to worse performance in our experiments. We achieve our best results when increasing the temperature according to a half-cycle temperature schedule, starting at $\tau = 0.75$, outperforming the model with fixed $\tau = 0.75$ by **+0.5%** top-1 accuracy for linear probing.

6. Limitations

As outlined in Section 1, MAEs provide state-of-the-art results for finetuning tasks, while not being able to do so for linear evaluations. As we show that our method narrows the gap for linear evaluations, it is important to also consider its impact on finetuning tasks. Our findings show that while our method does not outperform the vanilla MAE on finetuning, our results, presented in the supplementary, are still competitive. Also, our approach exhibits dependence on the object discovery network used to obtain the attention maps.

7. Conclusion & Discussion

We have proposed an attention-guided reconstruction masked autoencoder, which introduces high-level semantic information into the pre-training through an attention-guided reconstruction loss. This improved image-level properties are reflected in the increased performance when evaluating our model’s off-the-shelf features on a large set of challenging image-level tasks and across multiple data domains. Thus, our method presents a significant step towards enabling MAEs to learn more powerful representations already during pre-training, and with attention maps being the first of potentially many guidance modalities. Further, our approach is transferable to future MAEs since we apply our guidance to the reconstruction loss.

References

- [1] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture, 2023. **1**
- [2] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. In *International Conference on Learning Representations*, 2021. **1, 2, 3, 5**
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. **2**
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020. **2**
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. **1, 2, 4, 7, 8**
- [6] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and Bill Freeman, editors. *MaskGIT: Masked Image Generative Transformers*, 2022. **2**
- [7] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE winter conference on applications of computer vision (WACV)*, pages 839–847. IEEE, 2018. **1, 7, 8**
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. **2**
- [9] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. **2**
- [10] Xinlei Chen, Zhuang Liu, Saining Xie, and Kaiming He. Deconstructing denoising diffusion models for self-supervised learning. *arXiv preprint arXiv:2401.14404*, 2024. **1**
- [11] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9640–9649, 2021. **1, 2**
- [12] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014. **7**
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. **5, 6**
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, 2019. **2**
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. **2**
- [16] Mohamed El Banani, Amit Raj, Kevis-Kokitsi Maninis, Abhishek Kar, Yuanzhen Li, Michael Rubinstein, Deqing Sun, Leonidas Guibas, Justin Johnson, and Varun Jampani. Probing the 3d awareness of visual foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21795–21806, 2024. **1**
- [17] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *arXiv preprint arXiv:2205.09113*, 2022. **2**
- [18] Rohit Girdhar, Alaaeldin El-Nouby, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Omnimae: Single model masked pretraining on images and videos. *arXiv preprint arXiv:2206.08356*, 2022. **2**
- [19] Matthew Gwilliam and Abhinav Shrivastava. Beyond supervised vs. unsupervised: Representative benchmarking and analysis of image representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9642–9652, 2022. **1**
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. **1, 2, 3**
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. **2**
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **2**
- [23] Ioannis Kakogeorgiou, Spyros Gidaris, Bill Psomas, Yannis Avrithis, Andrei Bursuc, Konstantinos Karantzas, and Nikos Komodakis. What to hide from your students: Attention-guided masked image modeling. *arXiv preprint arXiv:2203.12719*, 2022. **2**
- [24] Ioannis Kakogeorgiou, Spyros Gidaris, Bill Psomas, Yannis Avrithis, Andrei Bursuc, Konstantinos Karantzas, and Nikos Komodakis. What to hide from your students: Attention-guided masked image modeling. *CoRR*, abs/2203.12719, 2022. **6, 8**
- [25] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM computing surveys (CSUR)*, 54(10s):1–41, 2022. **1**
- [26] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. **6**

- [27] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. **6**
- [28] Gang Li, Heliang Zheng, Daqing Liu, Chaoyue Wang, Bing Su, and Changwen Zheng. Semmae: Semantic-guided masking for learning masked autoencoders. *arXiv preprint arXiv:2206.10207*, 2022. **2, 5, 6, 8**
- [29] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022. **6**
- [30] Jovana Mitrovic, Brian McWilliams, Jacob C Walker, Lars Holger Buesing, and Charles Blundell. Representation learning via invariant causal mechanisms. In *International Conference on Learning Representations*, 2020. **2**
- [31] Mazda Moayeri, Phillip Pope, Yogesh Balaji, and Soheil Feizi. A comprehensive study of image classification model sensitivity to foregrounds, backgrounds, and visual attributes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19087–19097, June 2022. **1**
- [32] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. **6**
- [33] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, Dec 2008. **6**
- [34] O. M. Parkhi, A. Vedaldi, A. Zisserman, and C. V. Jawahar. Cats and dogs. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012. **7**
- [35] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Lost in quantization: Improving particular object retrieval in large scale image databases. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2008. **7**
- [36] Bill Psomas, Ioannis Kakogeorgiou, Konstantinos Karantzaolos, and Yannis Avrithis. Keep it simple: Who said supervised transformers suffer from attention deficit? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5350–5360, 2023. **7, 8**
- [37] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5706–5715, 2018. **7**
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. **7, 8**
- [39] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. **2**
- [40] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. **2**
- [41] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12179–12188, 2021. **6**
- [42] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. **1, 4, 5, 8**
- [43] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000. **5**
- [44] Yangyang Shu, Anton Van den Hengel, and Lingqiao Liu. Learning common rationale to improve self-supervised representation for fine-grained visual recognition problems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11392–11401, 2023. **1**
- [45] Nina Shvetsova, Felix Petersen, Anna Kukleva, Bernt Schiele, and Hilde Kuehne. Learning by sorting: Self-supervised learning with group ordering constraints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16453–16463, 2023. **1**
- [46] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013. **1**
- [47] Nenad Tomasev, Ioana Bica, Brian McWilliams, Lars Buesing, Razvan Pascanu, Charles Blundell, and Jovana Mitrovic. Pushing the limits of self-supervised resnets: Can we outperform supervised learning without labels on imagenet? *arXiv preprint arXiv:2201.05119*, 2022. **2**
- [48] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022. **2**
- [49] Vadim Tschernezki, Iro Laina, Diane Larlus, and Andrea Vedaldi. Neural feature fusion fields: 3d distillation of self-supervised 2d image representations. In *2022 International Conference on 3D Vision (3DV)*, pages 443–453. IEEE, 2022. **1**
- [50] Grant Van Horn, Elijah Cole, Sara Beery, Kimberly Wilber, Serge Belongie, and Oisín Mac Aodha. Benchmarking representation learning for natural world image collections. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12884–12893, 2021. **1**
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. **2**
- [52] Xudong Wang, Rohit Girdhar, Stella X Yu, and Ishan Misra. Cut and learn for unsupervised object detection and instance segmentation. *arXiv preprint arXiv:2301.11320*, 2023. **7, 8**
- [53] Yangtao Wang, Xi Shen, Yuan Yuan, Yuming Du, Maomao Li, Shell Xu Hu, James L Crowley, and Dominique Vaufreydaz. Tokencut: Segmenting objects in images and videos with self-supervised transformer and normalized cut. *arXiv preprint arXiv:2209.00383*, 2022. **4, 5, 8**

- [54] Kai Yuanqing Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. Noise or signal: The role of image backgrounds in object recognition. In *International Conference on Learning Representations*, 2020. 6, 7
- [55] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022. 2, 5
- [56] Huzheng Yang, James Gee, and Jianbo Shi. Brain decodes deep nets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23030–23040, 2024. 1
- [57] Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson. Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*, 2015. 1
- [58] Amir R. Zamir, Alexander Sax, William Shen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 6
- [59] Renrui Zhang, Lihui Wang, Yu Qiao, Peng Gao, and Hongsheng Li. Learning 3d representations from 2d pre-trained models via image-to-point masked autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21769–21780, 2023. 1
- [60] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017. 6
- [61] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. Image bert pre-training with online tokenizer. In *International Conference on Learning Representations*, 2021. 1, 6
- [62] Zhengyang Zhuge, Jiaying Wang, Yong Li, Yongjun Bao, Peisong Wang, and Jian Cheng. Patch-aware sample selection for efficient masked image modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17245–17253, 2024. 2