


# Masked Image Modeling as a Framework for Self-Supervised Learning across Eye Movements

Robin Weiler<sup>1,\*</sup>, Matthias Brucklacher <sup>1,\*</sup>, Cyriel M. A. Pennartz<sup>1</sup>, and Sander M. Bohte<sup>1,2</sup>

<sup>1</sup> Cognitive and Systems Neuroscience Group, University of Amsterdam, Science Park 904, 1098XH Amsterdam, Netherlands

<sup>2</sup> Machine Learning Group, CWI Amsterdam, Science Park 123, 1098 XG Amsterdam, Netherlands  
`matthias.brucklacher@gmail.com`

**Abstract.** To make sense of their surroundings, intelligent systems must transform complex sensory inputs to structured codes that are reduced to task-relevant information such as object category. Biological agents achieve this in a largely autonomous manner, presumably via self-supervised learning. Whereas previous attempts to model the underlying mechanisms were largely discriminative in nature, there is ample evidence that the brain employs a generative model of the world. Here, we propose that eye movements, in combination with the focused nature of primate vision, constitute a generative, self-supervised task of predicting and revealing visual information. We construct a proof-of-principle model starting from the framework of masked image modeling (MIM), a common approach in deep representation learning. To do so, we analyze how core components of MIM such as masking technique and data augmentation influence the formation of category-specific representations. This allows us not only to better understand the principles behind MIM, but to then reassemble a MIM more in line with the focused nature of biological perception. We find that MIM disentangles neurons in latent space without explicit regularization, a property that has been suggested to structure visual representations in primates. Together with previous findings of invariance learning, this highlights an interesting connection of MIM to latent regularization approaches for self-supervised learning. The source code is available under <https://github.com/RobinWeiler/FocusMIM>

**Keywords:** Self-supervised learning · Representation learning · Generative model

## 1 Introduction

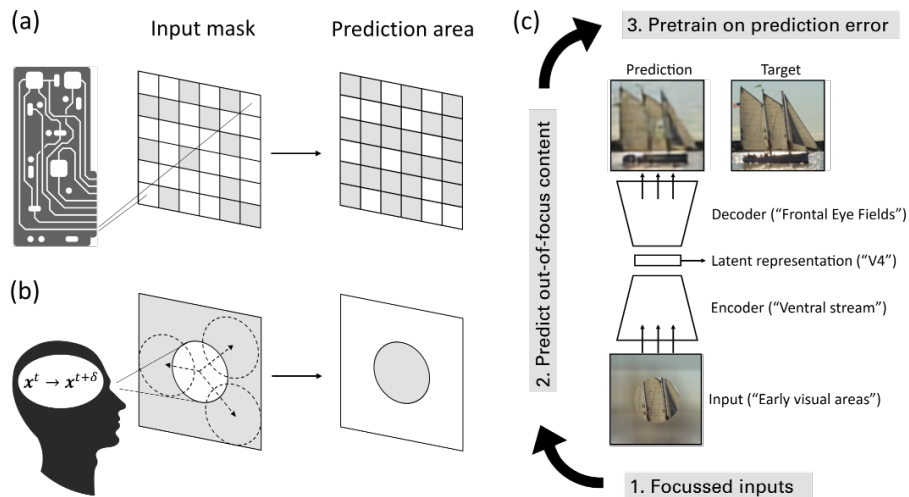
Both biological and artificial intelligent systems must construct useful object representations (in the general case equivalent to classifiable) from large amounts of unlabeled data [1]. In self-supervised learning (SSL), two main approaches have

---

\* Equal contribution

crystalized over the last years: multi-view based learning (via contrasting [2], latent regularization [3] or distillation [4]), and masked image-modeling (MIM) [5, 6] which predicts occluded image content usually at the pixel-level, but to which also the more recently developed latent predictions [7] can be counted. While contrastive methods have been suggested as a model of SSL in the brain, structuring representational geometry via temporal [8] and spatial cues [9], we postulate that the brain may be engaged in MIM via eye movements and attention shifts (Fig. 1): Perception, at a singular point in time, is selective, both through foveal vision and selective attention [10]. Sequentially, saccadic eye movements and attention shifts then flood the sensory stream with new information - in a manner that is to some degree predictable from knowledge about the direction and magnitude of the gaze shift (via a corollary discharge [11, 12]) and image structure.

Studying the principles behind MIM is attractive, as it requires fewer assumptions about temporal sequences and data augmentations than the mentioned multi-view approaches to biological SSL and achieves performance on par or close to contrastive methods [13]. MIM fits well with predictive theories of cortical processing [14–16]. Both conceptualize perception as the process of iden-



**Fig. 1.** Representation learning through eye movements. (a) Random patch masks used in artificial MIM approaches, gray patches are hidden from the network. (b) In our approach, selective masking is achieved through the inhomogeneous nature of foveal vision. Eye movements depicted as vectors reveal previously inaccessible or distorted information that is compared to the prediction  $x^{t+\delta}$ . (c) After pretraining on the prediction task, the quality of latent representations in the latent representation is assessed through classification accuracy in linear probing.

tifying a hierarchical latent variable model that could have caused the sensory inputs [17, 18] (this generative aspect is lacking in discriminative SSL models) and are predictive of future inputs. Experimental evidence for predictions of visual content across saccades comes, e.g., from [19]. Anatomically, the frontal eye fields have been proposed to predict across eye movements to stabilize perception [11]. We postulate a larger role of these predictions: the learning of visual features useful for downstream tasks such as classification.

Here, we investigate the influence of multiple variations to MIM in convolutional neural networks (CNNs) that align well with biological constraints. First, the influence of different masking strategies is tested, resulting in a peripheral masking strategy that learns strong representations and is more in line with the focused nature of biological perception. In doing so, we discover that representations are markedly improved, and that neurons become more decorrelated, when peripheral information is completely suppressed during learning. Second, the influence of data augmentation is tested, showing that these are especially relevant for the peripheral masking strategy. Third, we show that the relevant loss signal originates from the main object: although networks reconstruct the context too, this is not necessary for representation learning. In fact, we find that under some conditions networks learn faster when being only rewarded for predicting the foreground, which is interesting as segmentation may naturally be available from geometric cues or motion [20, 21]. Indeed, availability of segmentation has been shown to aid contrastive SSL if leveraged during pretraining [22], while its effect on MIM has previously not been studied to our knowledge.

## 2 Related Work

The breakthrough for MIM in representation learning arguably was achieved using random patch masks (Fig. 1a) in a vision transformer [6], an architecture that differs from information processing in biological networks. Investigations into the influence of masking strategies found that the optimal masking ratio falls roughly between 40% and 80%, being sufficiently large but not overly aggressive [6, 18]. A more biologically compatible architecture (CNNs) for MIM was used by Pathak et al. [5], as do we for this reason. Pathak et al. found that a central mask (predicting from peripheral context to a small central region) performed worse than randomly distributed patch or region masks, although quantitative reports were not provided. Here, we address this shortcoming and systematically study the choice of the masking paradigm. A recent CNN-based MIM is ConvNeXt-v2 [13], which utilizes random patch masks like the transformer-based methods in combination with a shallow decoder, whereas we strive to make the masking paradigm more plausible. Xie et al. [23] further demonstrated the potential for MIM in CNNs, achieving high accuracy in a ResNet-based model. In contrast to their evaluation paradigm that employs the same (full) dataset both during pretraining and fine-tuning [see also 13], we return to the ethologically relevant

setting with limited amounts of labeled data.

Biological approaches to SSL derive mostly from multi-view approaches. They avoid representation collapse, a catastrophic network state in which different inputs are projected onto the same latent activity pattern and thus cannot be distinguished anymore, by introducing an expanding force that pushes representations apart. Illing et al. [9] used a single negative sample at a time, interpreted as being obtained by larger eye movements. Their approach converts contrastive predictive coding [2] into a fully local learning rule. Interestingly, the model also has a generative component (predicting the next representation in a sequence of views) as a means to representation learning that is not investigated further. Local predictive learning (LPL) [8], on the other hand, relies on pushing representations apart based on their temporal distance, converting variance regularization [3] into a local learning rule. In contrast, MIM, as we employ it, prevents representation collapse in a more minimalist way, i.e., without the necessity of negative examples [9] or accumulated temporal statistics [8], but instead solves a generative problem in which loss minimization is incompatible with a collapsed solution. As the approach presented here is an initial step, we use the spatially non-local backpropagation, however, see [24, 25] and subsequent work for biologically plausible solutions yielding comparable weight updates.

Previous functional models have combined focused vision with eye movements in a variety of ways, from controlling information gain [26] and active inference [27], to achieving representation invariance and a reduced need for detail in generative model learning [28]. The model of Thompson et al. [29] combines spatially structured eye movements with visual information processing in an enactive approach to counting. The approach of Illing et al. [9] uses on- vs. off-object eye movements as a heuristic to shape representational geometry in latent space. In comparison, our approach is more reduced in dynamics and focuses instead on the consequences for neural representation learning when casting eye movements as a self-supervised predictive task. It should also be noted that all mentioned functional interpretations are mutually compatible.

### 3 Method

#### 3.1 Model

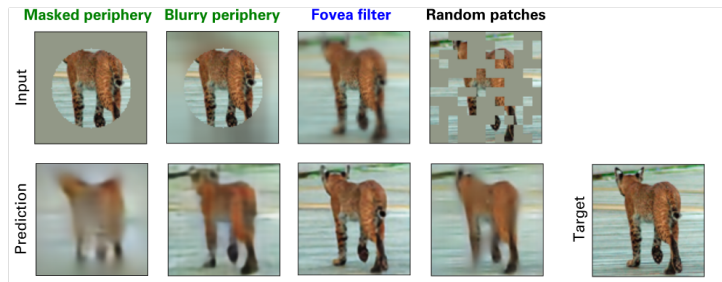
We reduce the task of predicting across eye movements to its essence and avoid learning a remapping of spatial receptive fields, which we argue is orthogonal to the process of information restoration that drives representation learning [18]. Thus, inputs and predictions take place in the same reference frame (Fig. 1c), as in other MIM approaches [6, 13, 18]. The network architecture, shown in Fig. 1c, resembles an autoencoder. Unless explicitly noted otherwise, we use the following setup: The encoder reduces the spatial dimension from  $96 \times 96$  (with three color channels) to  $12 \times 12$  ( $\times 128$  channels). Due to the large receptive fields of the latent space neurons at the output of the encoder, they map roughly onto neurons

in visual area V4 or higher in primates [30]. The inversely structured decoder (with deconvolutions instead of convolutions) maps the latent representations back onto a  $96 \times 96 \times 3$  output, constituting the model’s expectation of image content. Both encoder and decoder comprise nine ResNet18-style blocks [31] of two convolutions passed by a skip connection. Convolutions are followed by a ReLU activation function except the last decoder layer, which is followed by a sigmoid function. The convolution at the entry to each third block has stride = 2, which is where we double and half the number of channels in the encoder and decoder, respectively, starting with 32 channels at the first encoder block.

### 3.2 Self-Supervised Pretraining

As other work in MIM, we define pretraining as the self-supervised training phase in which the model is presented with a partial view of an image and attempts to predict the image pixel by pixel in full resolution. The loss function is given as the mean squared reconstruction error in parts of visual space with incomplete inputs (masked or blurred, depending on the paradigm), normalized by the area over which the loss is calculated. This generative loss is reminiscent of predictive coding [15]. Pretraining was conducted on the unlabeled split of the STL-10 dataset [32], which was specifically constructed for SSL with few labeled examples, and is sufficiently naturalistic and large to exclude the need for another dataset. We investigated a broad spectrum of input masking strategies (Fig. 2) that pertain to the original hypothesis of peripheral masking:

**Random patches:** The most common masking strategy, used, e.g., in [13, 18, 31] (cf. Fig. 1a) serves as an optimal baseline. We split the image into 144 patches of  $8 \times 8$  pixels and replaced the masked image content with the average color of the image.



**Fig. 2.** The masking strategies define the pretraining task. Top row: Exemplary inputs for different pretraining paradigms, with masked areas covered in the image’s gray average color. Masked periphery and random patches are shown with 60% masking ratio. Bottom row: Predictions from the given inputs after pretraining. Right image: Ground truth

**Masked periphery:** A central circle in full resolution, with the periphery masked with the sample’s average color (see also Fig. 1b). Biologically, this strict discarding of out-of-focus information is imaginable through inattentiveness [10].

**Blurry periphery** and **Blurry random patches:** Variations of the above paradigms with local Gaussian blurring instead of a uniformly colored mask to test the influence of incomplete masking.

**Foveal filter:** A centered circle with a small radius in full resolution. Outside the circle, the level of blur increases with distance to the center, mimicking the physiological constraints of human foveal vision, i.e., the distribution of cone photoreceptors across the retina [33, p. 244]. The implementation follows [34, 35].

For both random patches and masked periphery (using a square mask instead of a circle), we experimented with sparse convolutions as in [13] to avoid processing of borders between masks and image content. However, we found no significant difference in linear probing accuracy (as described below), so we continued with standard convolutions.

### 3.3 Choice of Pretraining Hyperparameters

During pretraining, we employed random-resized cropping (scale  $\in \{0.08, 1.0\}$ , aspect ratio  $\in \{0.75, 1.33\}$ ). We used the Adam optimizer [36] with weight decay =  $1e-8$ , batch size = 512, and selected the learning rate from  $\{5e-5, 1e-4, 5e-4\}$  that led to the best classification performance (described below) for each of the major pretraining tasks (masked periphery, random patches, foveal filter) to render the obtained findings independent of the choice. This value was then kept when blurry information was introduced. For each of the major masking paradigms, we additionally varied the masking ratio from 10-90% in steps of 10 to choose a model with optimal classification accuracy. For the random patches, we found a masking ratio of 60% to lead to the best representations. For the masked periphery, a higher ratio of 80% performed slightly better. We thus continued with this fraction for the remaining results, and refer to [18, 31] for more detailed studies on the influence of the masking ratio. After 500 epochs of pretraining, we continued with the model that achieved the lowest reconstruction loss for further analysis.

### 3.4 Linear Probing to Assess Representation Quality

To analyze the quality of the learned representations after pretraining, linear probing was conducted. Here, we followed the standard SSL procedure of discarding the decoder and appending a linear layer to the frozen encoder. The  $128 \times 12 \times 12 \times 10$  (latent dimension \* classes of the STL-10 labeled splits) parameters of the linear layer were then fit without further augmentations on the training split of the STL-10 dataset (5,000 images), of which we used 1,000 for

training validation and early stopping at the minimum validation loss. Irrespective of pretraining paradigm, we used images without mask or blurring in this phase, as these consistently led to higher readout accuracy. As in pretraining, we employed the Adam optimizer [36] (learning rate = 1e-4, weight decay = 1e-8, batch size = 512). Classification accuracy was then evaluated on the 8,000 test images, averaged across five randomly seeded runs. We also experimented with multilayer perceptron readout and fine-tuning (retraining the whole network), but neither of these approaches led to large differences in accuracy.

### 3.5 Generalization to a different network architecture

To test in how far the influence of masking strategy depends on the network architecture, we pretrained a modern CNN, the ConvNeXt-v2-Pico [13]. Key differences to the ResNet described above are increased receptive field size, separate depthwise and pointwise convolutions, the use of GELU activation functions, and a shallow decoder consisting of a single block. Pretraining on the STL-10 dataset was conducted in two configurations: with random patch masks or with masked periphery. Due to the ConvNeXt-v2’s initial stage that splits the input image into 7x7 patches, we implemented peripheral masking as predicting from the central 3x3 patches to the surround. For the random patches, we used a masking ratio of 60%, as suggested in [13]. Models were then pretrained for 370 epochs with hyperparameters close to [13]: batch size = 512, learning rate selected for optimal performance out of {1.6e-4, 5e-4, 1.5e-3}, AdamW optimizer [37] with 20 epoch linear warmup and subsequent cosine decay, weight decay = 0.05 as well as random resized cropping and horizontal flipping augmentations. For linear probing and fine-tuning, random erasing was additionally applied, and warmup was reduced to three epochs, with a base learning rate of 0.0006 (fine-tuning) and 0.008 (linear probing). While fine-tuning led to the best performance when using global average pooling before the classification head, linear readout profited from skipping this, so we proceeded with the optimal settings for each.


### 3.6 Correlation of Neurons in Latent Space

In addition to linear discriminability, representations can also be analyzed based on their statistical properties. One desirable property here is decorrelation [3, 38, 39]. To quantify this, the covariance matrix is required [3, 39]:

$$C(R) = \frac{1}{n-1} \sum_{i=1}^n \mathbf{r}_i \mathbf{r}_i^T \quad (1)$$

with  $R$  defined as a tensor containing the  $n$  latent representation vectors  $\mathbf{r}_i$  that are z-scored across the batch dimension, with  $n$  equal to the number of samples (we compute across one mini-batch with  $n = 512$ ). As in [3], the sum of the squared off-diagonal entries in  $C$  then quantifies the degree to which neurons are correlated:

$$c(R) = \sum_{i \neq j} [C(R)]_{i,j}^2 \quad (2)$$



**Fig. 3.** Representative presegmentation masks used to investigate how discarding loss signals from the background affects network performance. The masks were obtained with the *rembg*-library, with brightness indicating confidence, overlaid onto the respective STL-10 images.

### 3.7 Disentangling Predictions of Object vs. Background

Unlike assumed in section 3.6, eye movements reveal information sequentially, often targeting salient points such as objects [40, 41]. Objects may also be segmented already before recognition, either from motion signals [20, 21] or occlusion geometry [20], making them prioritized prediction targets. We thus evaluated the influence of reduced learning signal availability from the background. First, we obtained a figure-background segmentation mask using the *rembg*-library (<https://github.com/danielgatis/rembg>). We discarded  $\sim 8,000$  images for which no clear segmentation mask could be obtained. In these, less than 100 pixels (out of  $96 \times 96$ ) received a confidence score of 0.8 (on a scale from 0 to 1) or higher. In the remaining images, the masks were generally high-quality and only small object parts were missed (Fig. 3). During pretraining, the reconstruction loss was then multiplied with the foreground confidence score at each point in space and normalized to compensate the reduced prediction area.

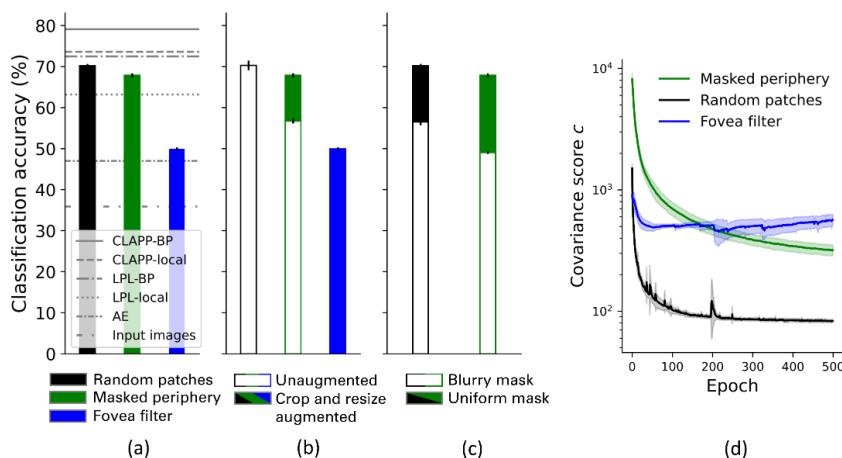
## 4 Results

To identify the minimally necessary components for biological MIM, we applied various combinations of masking strategy and data augmentation, and plot their influence on representation quality in Fig. 4a-c.

### 4.1 Influence of Masking Strategy

Out of the used masking strategies, we identified the masked periphery condition as the most promising candidate for a biological model. This masking technique inspired by foveal visual perception led to the second-highest classification accuracy in linear readout (mean value  $67.9 \pm 0.4\%$ ) when crop-and-resize transformations were used during pretraining, outperformed only by the random patch masks ( $70.2 \pm 0.4\%$ ). Similarly, in the ConvNeXt-v2 architecture, peripheral masking led to a high linear readout accuracy of  $71.6 \pm 0.6\%$  (fine-tuning:





**Fig. 4.** Linear readout accuracy quantifies representation quality after pretraining. a) Influence of masking strategy in comparison to the multi-view approaches CLAPP [9] and LPL [8], that both come in two variants trained with a local learning rule or backpropagation. The remaining baselines are pure autoencoding (AE) from full image to full image in the same network as the masked methods, and directly conducting linear probing on the input images. b) Influence of data augmentation. In the left- and rightmost bar, the filled and outlined bars overlap, i.e., classification accuracy is unaffected by augmentation. c) Incomplete masking by Gaussian blurring, instead of uniform average coloring, drastically decreased representation quality. d) Masked image modeling implicitly decorrelates latent space neurons. Error bars and shaded regions indicate the standard deviation across five randomly seeded runs.

$90.9 \pm 0.3\%$ ), compared to  $68.2 \pm 0.3\%$  (fine-tuning:  $89.4 \pm 0.4\%$ ) attained with random patch masks. Strong performance in these settings supports the assumption that solving the pretraining tasks of predicting hidden image areas leads to discovery of object representations [18]. In contrast, deblurring the peripheral region from the foveal filter resulted in comparably poor representations, with readout accuracy not much higher than pure autoencoding. Presumably, this condition allows the network to find a local solution that does not require global integration [18]. To more closely investigate the effect of input information content on classification accuracy, we replaced uniform masking with Gaussian blurring (Fig. 2, third column). Despite the low information content, the introduced information was sufficient to substantially degrade performance in both the masked periphery and random patch condition (Fig. 4c).

## 4.2 Decorrelation of Latent Space Neurons

To investigate the degree to which neurons in latent space develop independent tuning properties, we computed the correlation term  $c$  using Eq. 2. Interestingly, both masked periphery and random patches decorrelated neural representations

(Fig. 4d), and more strongly than the foveal filter. This decorrelation, reminiscent of sparse coding [42], emerged without being explicitly enforced as in other approaches [3, 39, 42].

### 4.3 Influence of Data Augmentation

Applying crop-and-resize augmentations during pretraining proved necessary for representation learning in the masked periphery condition. When leaving out this augmentation, readout accuracy dropped from  $67.9 \pm 0.4\%$  to  $56.7 \pm 0.5\%$  (Fig. 4b). Biologically, cropping and resizing can be approximately obtained by changes in viewpoint and viewing angle, especially since no assumption about temporal co-occurrence is made: differently augmented versions of an image are presented in different epochs. Interestingly, neither the foveal filter nor the random patches were affected by ablating the crop-and-resize transformation. Maintaining the position of the patch masks across epochs alone did not reduce readout accuracy ( $70.3 \pm 0.9\%$ ), but in combination with ablation of the crop-and-resize augmentation, accuracy reduced to  $62.5 \pm 0.8\%$ . We conclude that the model requires exposure to a variety of prediction tasks from each object, whether through data augmentation or through changes in mask position.

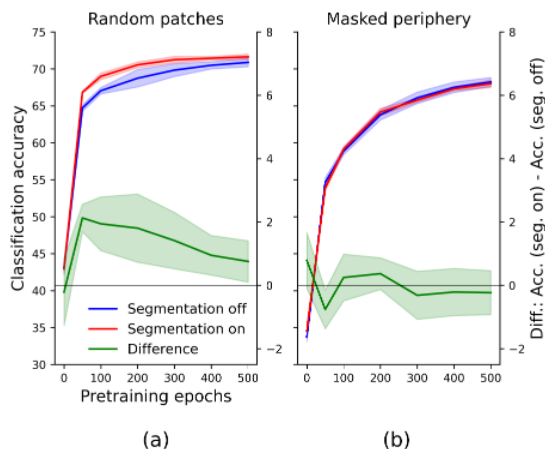
### 4.4 Predicting Object vs. Background

Restricting the pretraining loss to the main object instead of the background (cf. section 3.7) did not affect linear probing accuracy negatively, suggesting that on-object predictions are sufficient for classification. While there was no substantial difference in readout accuracy when disregarding, as opposed to when using, loss signal from the background in the masked periphery condition (Fig. 5b), the network pretrained using random patch masks profited from disregarding the background, especially in early epochs (Fig. 5a).

### 4.5 Reconstruction Quality

Beyond the network’s discriminative function, its generative capacity bears functional relevance, e.g., for counterfactual reasoning, predictions of occluded image content, and imagination [17]. An interesting effect from generative model learning could be observed in the networks pretrained with uniform (average colored) and blurry masks: These networks also reconstructed in locations in which no loss was calculated. While this would be expected to some degree when the masked area varies across image presentations (random patches), it is surprising in the masked periphery condition (Fig. 2, second column). Here, also the central region was reconstructed reasonably well, although it did not contribute to the pretraining loss. This holistic percept fits well with proposals on the role of visual predictions in creating a stable visual representation [11, 17].

Across masking paradigms, reconstruction accuracy was uncorrelated with readout accuracy: while the foveal filter achieved the best reconstructions (Fig. 2),



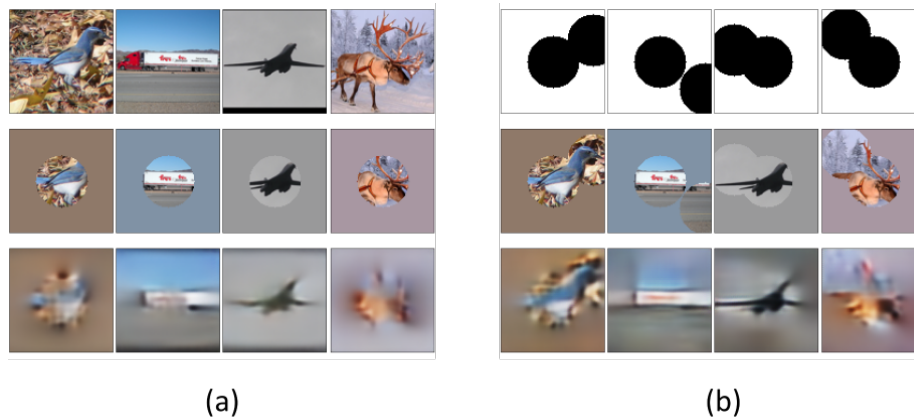
**Fig. 5.** Influence of presegmentation and masking strategy on pretraining efficiency. (a) When using random patches, using presegmentation to weight the reconstruction loss accelerated pretraining. (b) This was not the case when restricting the reconstruction loss to the foreground object. Shaded regions indicate the standard deviation across five randomly seeded runs.

it performed poorly in linear probing (Fig. 4a). Given a pretraining paradigm, however, longer training on the reconstruction task also improved readout (Fig. 5).

A third angle to analyze the generative capacity of the network stems from the fact that eye movements are often preceded by covert attention shifts, which increase information sampling from these areas [43]. The sampled information can then be used to improve predictions about the expected image content. To test whether the model could exploit such a mechanism for improved reconstructions, we provided the model with a second circular input sampling area (covert attention) on top of the peripheral masking while constraining the error to the foreground (section 3.7). The additionally provided information indeed led to more truthful predictions (Fig. 6), i.e., reduced prediction errors (from  $2.6e-2$  to  $2.3e-2$ ), which ties the approach back to earlier theories of predictions as a stabilizing mechanism [11]. Applying this twofold sampling during pretraining, however, did not alter readout accuracy significantly, irrespective of whether the covert attention circle was placed at random or on the main object.

## 5 Conclusion

In this paper, we investigated key components of masked image modeling (MIM) through the lens of biological vision. As a similarly performant alternative to the random patch masks common in artificial representation learning [6, 13], we identified the more biologically plausible peripheral masking (Fig. 4a). Provided with



**Fig. 6.** Sampling via covert attention improves predictions and stabilizes perception. (a) Top row: ground truth. Middle row: input in the masked periphery condition with a masking ratio of 80%. Bottom row: reconstructions after pretraining. (b) Same as (a) with additionally provided information via a second, non-central input circle (covert attention). The top row depicts the input mask consisting of these two circular areas (fovea and attention).

differently augmented views of the input images (Fig. 4b), this approach does not pose restrictions on the temporal sequence in which inputs are encountered, as does [8], or on knowledge about the spatial extent of an object to obtain negative samples, as does [9]. We conclude that peripheral masking and prediction is a candidate mechanism for self-supervised learning in the brain, with saccadic prediction errors computed in the early (but likely not primary) visual cortex [19] where bottom-up signals are compared to feedback from the frontal eye fields [11]. However, further experimental studies are required to definitely link these error signals to plasticity in the ventral stream, and to elucidate whether covert attention shifts can elicit similar prediction errors. Functionally, an additional advantage of the employed approach in contrast to purely discriminative methods is the acquisition of a generative model (Fig. 2, Fig. 6). That multi-view SSL has been implemented with more local learning rules than backpropagation [8, 9, 44, 45] prompts the question of whether the same can be done with MIM. From a theoretical perspective, implicit decorrelation of latent space neurons (Fig. 4d), together with previous observations of implicit invariance learning [46], connect the MIM framework to SSL via latent regularization [3], where these objectives are explicitly enforced. This raises the question of whether and how synergies could be efficiently exploited in models combining both learning paradigms [7, 47]. As a general requirement for high classification accuracy across all investigated masking strategies, we identified opacity of the input masks (Fig. 4c), just as humans sometimes do not perceive objects in their field of view when attending to another point in space [10]. As primates employ eye movements in a

strategic and enactive way [29], an interesting extension would be to learn from multiple peripherally masked glimpses in a sequence through spatiotemporally masked videos [48].

**Acknowledgements** This project has received funding from the European Union’s Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3 to C.P. and S.B.). We thank SURF ([www.surf.nl](http://www.surf.nl)) for the support in using the National Supercomputer Snellius, especially Thomas von Osch for the discussions.

## References

1. LeCun, Y.: A path towards autonomous machine intelligence version 0.9.2, 2022-06-27. Open Review 62(1) (2022)
2. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. Preprint arXiv:1807.03748 (2018)
3. Bardes, A., Ponce, J., LeCun, Y.: VICReg: Variance-invariance-covariance regularization for self-supervised learning. Preprint arXiv:2105.04906 (2021)
4. Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M.: Bootstrap your own latent—a new approach to self-supervised learning. *NeurIPS* 33, 21271–21284 (2020)
5. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: *CVPR*, pp. 2536–2544 (2016)
6. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *CVPR*, pp. 16000–16009 (2022)
7. Assran, M., Duval, Q., Misra, I., Bojanowski, P., Vincent, P., Rabbat, M., LeCun, Y., Ballas, N.: Self-supervised learning from images with a joint-embedding predictive architecture. In: *CVPR*, pp. 15619–15629 (2023)
8. Halvagal, M.S., Zenke, F.: The combination of Hebbian and predictive plasticity learns invariant object representations in deep sensory networks. *Nature Neuroscience* 26(11), 1906–1915 (2023)
9. Illing, B., Ventura, J., Bellec, G., Gerstner, W.: Local plasticity rules can learn deep representations using self-supervised contrastive predictions. *NeurIPS* 34, 30365–30379 (2021)
10. Simons, D.J., Chabris, C.F.: Gorillas in our midst: Sustained inattention blindness for dynamic events. *perception* 28(9), 1059–1074 (1999)
11. Crapse, T.B., Sommer, M.A.: The frontal eye field as a prediction map. *Progress in brain research* 171, 383–390 (2008)
12. Von Holst, E., Mittelstaedt, H.: Das Reafferenzprinzip: Wechselwirkungen zwischen Zentralnervensystem und Peripherie. *Naturwissenschaften* 37(20), 464–476 (1950)
13. Woo, S., Debnath, S., Hu, R., Chen, X., Liu, Z., Kweon, I.S., Xie, S.: ConvNeXt v2: Co-designing and scaling convnets with masked autoencoders. In: *CVPR*, pp. 16133–16142 (2023)
14. Dayan, P., Hinton, G.E., Neal, R.M., Zemel, R.S.: The helmholtz machine. *Neural computation* 7(5), 889–904 (1995)

15. Rao, R.P., Ballard, D.H.: Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience* 2(1), 79–87 (1999)
16. Friston, K.: A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences* 360(1456), 815–836 (2005)
17. Pennartz, C.M.A.: *The brain's representational power: on consciousness and the integration of modalities*. MIT Press (2015)
18. Kong, L., Ma, M.Q., Chen, G., Xing, E.P., Chi, Y., Morency, L.-P., Zhang, K.: Understanding Masked Autoencoders via Hierarchical Latent Variable Models. In: CVPR, pp. 7918–7928 (2023)
19. Ehinger, B.V., König, P., Ossandón, J.P.: Predictions of visual content across eye movements and their modulation by inferred information. *Journal of Neuroscience* 35(19), 7403–7413 (2015)
20. Tsao, T., Tsao, D.Y.: A topological solution to object segmentation and tracking. *PNAS* 119(41), e2204248119 (2022)
21. Brucklacher, M., Pezzulo, G., Mannella, F., Galati, G., Pennartz, C.M.: Learning to segment self-generated from externally caused optic flow through sensorimotor mismatch circuits. *bioRxiv preprint*. doi:10.1101/2023.11.15.567170v2 (2023)
22. Hénaff, O.J., Koppula, S., Alayrac, J.-B., Van den Oord, A., Vinyals, O., Carreira, J.: Efficient visual pretraining with contrastive detection. In: CVPR, pp. 10086–10096 (2021)
23. 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: CVPR, pp. 9653–9663 (2022)
24. Lillicrap, T.P., Santoro, A., Marris, L., Akerman, C.J., Hinton, G.: Backpropagation and the brain. *Nature Reviews Neuroscience* 21(6), 335–346 (2020)
25. Whittington, J.C., Bogacz, R.: An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. *Neural computation* 29(5), 1229–1262 (2017)
26. Renninger, L., Coughlan, J., Verghese, P., Malik, J.: An information maximization model of eye movements. *NeurIPS* 17 (2004)
27. Friston, K., Adams, R.A., Perrinet, L., Breakspear, M.: Perceptions as hypotheses: saccades as experiments. *Frontiers in psychology* 3, 151 (2012)
28. Larochelle, H., Hinton, G.E.: Learning to combine foveal glimpses with a third-order Boltzmann machine. *NeurIPS* 23 (2010)
29. Thompson, J.A., Sheahan, H., Summerfield, C.: Learning to count visual objects by combining “what” and “where” in recurrent memory. In: *NeurIPS 2022 Workshop on Gaze Meets ML*, pp. 199–218 (2022)
30. Smith, A.T., Singh, K.D., Williams, A.L., Greenlee, M.W.: Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cerebral cortex* 11(12), 1182–1190 (2001)
31. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR, pp. 770–778 (2016)
32. Coates, A., Ng, A., Lee, H.: An analysis of single-layer networks in unsupervised feature learning. In: *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 215–223 (2011)
33. Purves, D., Augustine, G., Fitzpatrick, D., Hall, W., LaMantia, A., McNamara, J., White, L.: *Neuroscience, 3rd Edition*. Sinauer Associates (2004)
34. Perry, J.S., Geisler, W.S.: Gaze-contingent real-time simulation of arbitrary visual fields. In: *Human vision and electronic imaging VII*, pp. 57–69 (2002)

35. Jiang, M., Huang, S., Duan, J., Zhao, Q.: SALICON: Saliency in context. In: CVPR, pp. 1072–1080 (2015)
36. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
37. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
38. Barlow, H.B.: Possible principles underlying the transformation of sensory messages. *Sensory communication* 1(01), 217–233 (1961)
39. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. In: ICML proceedings, pp. 12310–12320 (2021)
40. Hayhoe, M., Ballard, D.: Eye movements in natural behavior. *Trends in cognitive sciences* 9(4), 188–194 (2005)
41. Yarbus, A.L.: *Eye movements and vision*. Plenum Press (1967)
42. Olshausen, B.A., Field, D.J.: Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381(6583), 607–609 (1996)
43. Houtkamp, R., Spekreijse, H., Roelfsema, P.: A gradual spread of attention. *Perception & psychophysics* 65, 1136–1144 (2003)
44. Löwe, S., O’Connor, P., Veeling, B.: Putting an end to end-to-end: Gradient-isolated learning of representations. *NeurIPS* 32 (2019)
45. Siddiqui, S.A., Krueger, D., LeCun, Y., Deny, S.: Blockwise self-supervised learning at scale. Preprint arXiv:2302.01647 (2023)
46. Kong, X., Zhang, X.: Understanding masked image modeling via learning occlusion invariant feature. In: CVPR, pp. 6241–6251 (2023)
47. Jing, L., Zhu, J., LeCun, Y.: Masked siamese convnets. arXiv preprint arXiv:2206.07700 (2022)
48. Feichtenhofer, C., Li, Y., He, K.: Masked autoencoders as spatiotemporal learners. *NeurIPS* 35, 35946–35958 (2022)