



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Leveraging Registers in Vision Transformers for Robust Adaptation

S. M. Yellapragada, K. Thopalli, V. Narayanaswamy, W. Sakla, Y. Liu, Y. Mubarka, D. Samras, J. J. Thiagarajan

October 17, 2024

International Conference on Acoustics, Speech and Signal Processing
Hyderabad, India
April 6, 2025 through April 11, 2025

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Leveraging Registers in Vision Transformers for Robust Adaptation

Srikar Yellapragada¹, Kowshik Thopalli², Vivek Narayanaswamy², Wesam Sakla²,
Yang Liu², Yamen Mubarka², Dimitris Samaras¹, Jayaraman J. Thiagarajan²

¹Stony Brook University

²Lawrence Livermore National Laboratory

Abstract—Vision Transformers (ViTs) have shown success across a variety of tasks due to their ability to capture global image representations. Recent studies have identified the existence of high-norm tokens in ViTs, which can interfere with unsupervised object discovery. To address this, the use of “registers” which are additional tokens that isolate high norm patch tokens while capturing global image-level information has been proposed. While registers have been studied extensively for object discovery, their generalization properties particularly in out-of-distribution (OOD) scenarios, remains underexplored. In this paper, we examine the utility of register token embeddings in providing additional features for improving generalization and anomaly rejection. To that end, we propose a simple method that combines the special CLS token embedding commonly employed in ViTs with the average-pooled register embeddings to create feature representations which are subsequently used for training a downstream classifier. We find that this enhances OOD generalization and anomaly rejection, while maintaining in-distribution (ID) performance. Extensive experiments across multiple ViT backbones trained with and without registers reveal consistent improvements of 3-4% in top-1 OOD accuracy and a 2-3% reduction in false positive rates for anomaly detection. Importantly, these gains are achieved without additional computational overhead.

Index Terms—Vision Transformer, Robustness, OOD Generalization, Anomaly Rejection, Registers.

I. INTRODUCTION

Vision Transformers (ViTs) [1], which adapt the transformer architecture from natural language processing [2], have demonstrated exceptional performance across a range of vision tasks including classification and object detection [3]–[5]. At a high level, ViTs process images by dividing them into fixed-size patches, each treated as a token, with a special [CLS] token appended to aggregate information across tokens for downstream tasks.

In a recent study on the behavior of large scale ViTs, Darcet *et al.* [6] observed that patch tokens associated with low-informative regions (e.g., background) tend to have significantly higher ℓ_2 norm values compared to other tokens. While this tendency does not impact their utility in image-level prediction tasks, it increases the risk of compromising the performance in dense prediction tasks (e.g., object discovery). To circumvent this behavior, the authors explored the use of “registers”, which are additional tokens appended to the input sequence during training. Interestingly, these registers were found not only to isolate the behavior of the high norm patch




Fig. 1: **Impact of token embedding choices on linear probing on frozen Dino-V2 ViT-G backbones.** Each color indicates the token embeddings chosen for optimizing a linear classifier on ImageNet (IN)-1K along with the protocol adopted for pre-training the backbone (w/o registers or w registers). Here, [CLS] represents the classification token, μ_P and μ_R represents the mean patch and register token embeddings respectively. While it is common to utilize [CLS] and μ_P to train the classifier, we obtain improved generalization if the backbone is trained with registers (red vs green). Training a classifier naively with register tokens results in drop in generalization. However, our approach (yellow) maintains ID accuracy while providing substantial gains in OOD (ImageNet-Adversarial, ImageNet-Sketch) accuracies.

tokens but also contain global image-level information. This was validated when a linear classifier was trained on register token embeddings extracted from in-distribution (ID) data, demonstrating enhanced classification performance.

While the utility of registers in achieving improving object discovery has been well studied, their efficacy in downstream tasks such as classification remains under explored, which is the primary focus of our paper. To thoroughly assess the role of registers in these tasks, it is critical to evaluate their generalization capabilities under a variety of out-of-distribution (OOD) scenarios. As shown in Figure 1, naively

training a linear classifier with register token embeddings results in a non-trivial drop in generalization performance on OOD datasets implying that the registers by themselves do not capture all the information required to achieve robustness.

Consequently, in this paper, we examine the utility of register token embeddings, which are shown to contain global image-level information, as auxiliary features that supplement the [CLS] token embedding. We believe that such a combination will yield in enriched representations for robust adaptation. This hypothesis aligns with the findings of Zhang *et al.* [7], where seemingly redundant features with respect to in-distribution data provided significant benefits in OOD scenarios when combined together.

To that end, we train a linear classifier by combining the representations from the [CLS] token with the average-pooled register token embeddings μ_R . We compare our approach with the standard practice of combining [CLS] token with average-pooled patch tokens μ_P [8]. Strikingly, our method significantly improves both OOD generalization, while maintaining ID performance, as shown in Fig 1. Moreover, our approach also boosts anomaly rejection performance, compared to baselines. Our comprehensive experiments across multiple ViT backbones reveal a consistent improvement of **3-4%** in top-1 accuracy for OOD generalization and an average reduction of **2-3%** in the false positive rate for anomaly detection. These findings support our hypothesis that the concatenation of registers with [CLS] yields more robust feature representations. Importantly, these gains are obtained without incurring any additional computational overhead.

In summary, our contributions are as follows:

- We develop a novel method that leverages the register token embeddings as auxiliary features for building robust adaptation protocols.
- Through extensive experiments, we demonstrate consistent improvements in OOD generalization and anomaly detection across multiple ViT architectures.
- Our method achieves these gains with negligible computational overhead during training or inference, offering an efficient solution.

II. RELATED WORK

Large activations in transformers: Understanding the behavior of transformers has been an important research topic in the last few years, owing to their success in various NLP and vision tasks. A common approach employed to achieve this is to analyze the activations and the features extracted by the model. Recent studies have uncovered the presence of outlier weights and large-magnitude activations in transformers, affecting both language [9]–[12] and vision models [6], [13]. These phenomena manifest as concentrated attention weights in LLMs [10], [11] and in vision models [6] **Implications and solutions for massive activations:** These large activations are found to pose challenges for model quantization [4], [14] and image segmentation [6]. Proposed remedies to address these challenges include appending “register” tokens to the input patch sequence [6] or explicit

attention biases [15] during training. While these approaches typically discard the additional tokens during inference, our work identifies that these tokens contain valuable auxiliary information to that captured by the [CLS] token. Therefore, we propose leveraging this information to obtain richer feature representations, enabling superior performance across a diverse range of tasks.

III. METHOD

Preliminaries: Vision Transformers (ViTs) process images by dividing them into patches and viewing them as a sequence of tokens, where each token corresponds to a patch in the image. An additional token, the [CLS] token, is appended to the patch sequence to aggregate the global information from the image patches. The [CLS] embedding is often used as the global representation of the image in downstream tasks.

Let $X \in \mathbb{R}^{H \times W \times 3}$ represent a 3 channel RGB input image, where H and W are the height and width of the image, respectively. A ViT backbone, denoted by F , divides X into L patches, each of size $K \times K$, and projects them into D -dimensional embeddings. The [CLS] token $c \in \mathbb{R}^D$ is added to aggregate global information. The backbone produces:

- [CLS] token $c \in \mathbb{R}^D$
- Patch tokens $P = \{p_1 \dots p_L\} \in \mathbb{R}^{L \times D}$
- Register tokens $R = \{r_1 \dots r_M\} \in \mathbb{R}^{M \times D}$

where D is the hidden dimensionality of the transformer, and L and M are the numbers of patch tokens and register tokens, respectively.

In typical ViT models, downstream tasks use either the [CLS] token or the mean of patch tokens $\mu_P = \frac{1}{L} \sum_{j=1}^L p_j$ [16], [17]. In Dino-v2, the [CLS] token is concatenated with μ_P , and fed to a linear classifier h_θ which is optimized with the following objective

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^N \ell(h_\theta(f_i), y_i) \quad (1)$$

where N is the total number of samples, $f_i = [c_i; \mu_P^i]$ is the concatenated [CLS] token c_i and the mean of the patch tokens μ_P^i for the i^{th} sample, $h_\theta(f_i) = f_i^\top \theta$ denotes the linear probe with parameters θ , and $\ell(\cdot, \cdot)$ is the cross-entropy loss computed between the prediction $h_\theta(f_i)$ and the target y_i .

Proposed Approach: We propose to leverage the auxiliary information in register tokens to augment the information found in the [CLS] token, thereby enhancing out-of-distribution (OOD) detection performance. As previously stated, Darcet *et al.* [6] noted that register tokens capture global information similar to the [CLS] token, by training linear probes on register tokens without observing significant drop in ID performance. Furthermore, Zhang *et al.* [7] showed that although similar representations from models trained with different random seeds yield comparable ID performance, they can encode different information under OOD conditions. Crucially, combining these representations leads to richer feature sets that outperform either representation alone in OOD tasks.

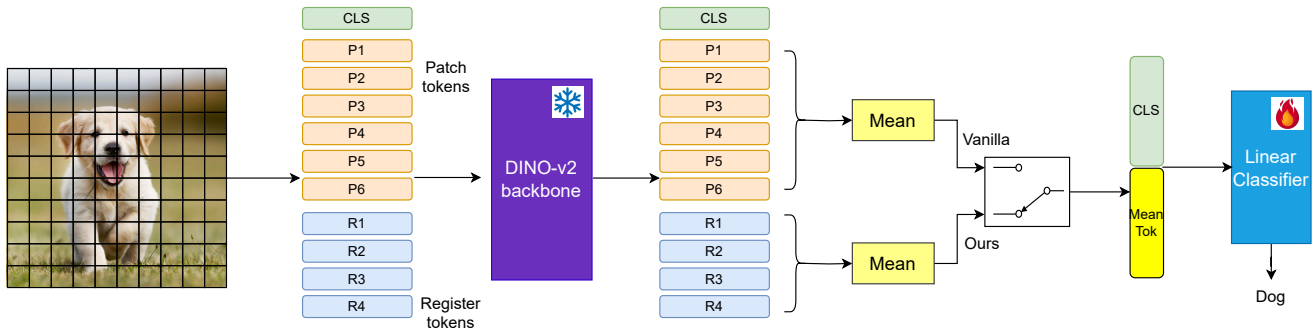


Fig. 2: **Overview of our proposed method:** For large-scale vision transformer backbones (e.g., DINO-v2) pre-trained with “registers”, we find that concatenating registers (Mean (R_1, R_2, R_3, R_4)) along with [CLS] is critical for obtaining rich features that enable robust adaptation. In particular, we train a linear classifier on these concatenated features and observe improved generalization and anomaly rejection capabilities.

Building on this insight, we concatenate the [CLS] token with the mean of register tokens, $\mu_R = \frac{1}{M} \sum_{k=1}^M r_k$, as opposed to discarding the register tokens as described in previous works [6], [18]. This combined representation exploits the auxiliary information captured by the register tokens, allowing us to create a richer feature representation. The concatenated representation is then fed to the linear classifier h_θ which is optimized with the same objective as presented in Eq. 1 except that the f_i is now given by $f_i = [c_i; \mu_R^i]$. Figure. 2 provides an overview of the proposed approach.

Note our method does not introduce additional computational overhead because the ViT backbone is already trained and kept frozen during the experiments. We only need to train a linear layer, however, just like the baseline approach that concatenates the [CLS] token with the patch token mean. Since both methods involve concatenating two vectors ($[CLS; \mu_R]$ vs $[CLS; \mu_P]$), the linear layer remains the same size. Thus, given a backbone already trained with register tokens, the only step required is training this linear layer, making our approach computationally efficient.

IV. EXPERIMENTS

Datasets. We evaluate our proposed approach on the following datasets: (i) Training/In-Distribution (ID) data: We train the Linear Classifiers (LC) on ImageNet-1K [19] for all experiments. ImageNet-1K is a large-scale imaging benchmark comprising 1.3 million training images and 50,000 validation images across 1000 diverse categories. (ii) Out-of-Distribution (OOD) generalization benchmarks: (a) ImageNet-A (In-A) [20] is a dataset comprising 7,500 adversarially filtered images across a 200-classes subset of ImageNet-1K’s 1,000 classes that were found to be challenging to current ImageNet classifiers, (b) ImageNet-R (In-R) [21] contains 30,000 images of different real-world renditions from 200 classes of ImageNet; (c) ImageNet-S (In-S) [22] contains 50 black and white sketch images for each class in ImageNet for a total of 50,000. (iii) Anomaly rejection benchmarks: Following standard practice [23], we consider (a) Texture (DTD), (b) SVHN, (c)

Places 365, (d) LSUN, (e) LSUN (Resized) and (f) iSUN datasets. It must be noted that, every image in these datasets is appropriately resized and passed as input into the networks. **Setup.** We utilize the open-source¹ Dino-v2 ViT backbones trained with registers [6] for the experiments involving register tokens, and Dino-v2 ViT backbone trained without registers for other experiments. We evaluate three variants: ViT Giant, Large, and Base. In all cases, we only train the last linear layer of the network while keeping the backbone frozen. We train the linear layers with the SGD optimizer for 10,000 iterations, using random-resized-crop data augmentation. The training process was conducted on 16 NVIDIA V100 GPUs, with a batch size of 256 per GPU and learning rate of 0.01. **Metrics.** We report the Top@1 accuracy to measure OOD generalization performance. For anomaly rejection, we utilize Maximum Softmax Probability (MSP) [24] and energy [23] scoring functions and report the FPR@TPR95 and AU-ROC [23] for performance evaluation.

Baselines. We compare our proposed approach with the following methods to understand the impact of choosing different tokens for training the LC. (i) Concatenation of the [CLS] token and the patch token mean ($CLS; \mu_P$ W/o reg) when the model is trained without registers, as in [8], and (ii) the same concatenation when the model is trained with registers ($CLS; \mu_P$ with reg). To ensure consistent feature scales during concatenation, we use features extracted after the final LayerNorm from the Dinov2 backbones.

A. OOD Generalization

As shown in Table I, our experiments on OOD generalization demonstrate the effectiveness of combining the [CLS] token with the mean of Register tokens ($CLS; \mu_R$). Across all three ViT architectures (Giant, Large and Base), our approach consistently outperforms both the baselines. For ViT-Giant, our method achieves the highest accuracy on all three ImageNet OOD generalization benchmarks. This represents improvements of 4.44%, 4.2% and 2.87% respectively over

¹<https://github.com/facebookresearch/dinov2>

TABLE I: **Evaluating ID and OOD generalization as well anomaly rejection performance across different ViT architectures.** We highlight the best performing method in every case with green. While the ID accuracies remain comparable to a baseline variant, we find that our proposed approach consistently outperforms the existing baselines on OOD generalization as well as anomaly rejection across a variety of benchmarks and architectures.

ViT-Giant													
Method	Dino-v2 with registers	ID Acc	ImageNet OOD			Score	Anomaly Rejection (FPR ↓ / AUROC ↑)						
			In-A	In-R	In-S		DTD	SVHN	Places-365	LSUN	LSUN-R	ISUN	Mean
CLS ; μ_P	×	86.5	71.24	74.52	59.46	MSP	48.81/85.32	22.68/95.71	59.65/83.0	18.76/96.04	43.32/90.95	41.99/90.32	39.19/90.22
CLS ; μ_P	✓	87.1	73.65	76.31	61.86		46.42/86.16	27.16/94.57	58.58/ 83.09	16.3/96.5	40.41/91.14	38.27/91.25	37.86/90.45
CLS ; μ_R	✓	86.57	78.09	80.51	64.43		45.53/ 86.07	16.71/ 96.71	58.88/82.79	11.78/97.48	37.06/91.59	33.76/92.22	33.95/91.14
CLS ; μ_P	×	-	-	-	-	Energy	31.83/ 92.12	10.11/97.66	43.04/ 89.72	6.34/98.60	30.30/94.92	26.73/94.74	24.73/94.63
CLS ; μ_P	✓	-	-	-	-		30.48/91.45	6.81/98.52	43.17/88.84	5.03/98.85	21.77/ 95.78	20.25/95.70	21.25/ 94.86
CLS ; μ_R	✓	-	-	-	-		32.41/90.64	4.96/98.92	45.5/87.64	3.85/99.06	19.95/95.76	18.48/95.95	20.86/94.66

ViT-Large													
Method	Dino-v2 with registers	ID Acc	ImageNet OOD			Score	Anomaly Rejection (FPR ↓ / AUROC ↑)						
			In-A	In-R	In-S		DTD	SVHN	Places-365	LSUN	LSUN-R	ISUN	Mean
CLS ; μ_P	×	86.3	71.24	74.52	59.46	MSP	51.4/84.59	27.2/94.6	59.54/83.24	20.63/95.49	40.81/91.61	40.46/90.9	40.1/90.07
CLS ; μ_P	✓	86.7	73.65	76.31	61.86		48.12/85.67	22.44/95.42	58.13/83.54	17.36/96.32	37.35/92.39	37.49/91.48	36.82/90.81
CLS ; μ_R	✓	85.82	75.33	79.2	62.82		43.71/86.25	15.93/96.76	55.88/83.35	12.65/97.12	33.49/92.38	31.35/92.77	32.17/91.44
CLS ; μ_P	×	-	-	-	-	Energy	47.27/87.26	48.95/93.03	55.77/86.15	17.8/96.13	49.99/91.78	42.85/91.91	43.77/91.04
CLS ; μ_P	✓	-	-	-	-		35.83/90.51	10.35/97.73	45.21/ 89.05	8.08/98.18	24.83/95.76	24.56/95.21	24.81/94.41
CLS ; μ_R	✓	-	-	-	-		34.06/90.87	7.05/98.55	43.44/88.79	6.86/98.5	19.24/96.21	18.87/96.15	21.59/94.84

ViT-Base													
Method	Dino-v2 with registers	ID Acc	ImageNet OOD			Score	Anomaly Rejection (FPR ↓ / AUROC ↑)						
			In-A	In-R	In-S		DTD	SVHN	Places-365	LSUN	LSUN-R	ISUN	Mean
CLS ; μ_P	×	84.5	55.54	63.9	50.92	MSP	55.66/83.95	18.35/96.44	64.65/81.76	21.94/95.42	39.51/91.57	39.9/91.24	40.01/90.06
CLS ; μ_P	✓	84.21	54.56	65.02	53.61		53.12/84.45	20.31/96.27	62.87/82.34	16.84/96.31	33.83/92.98	35.38/92.24	37.06/90.76
CLS ; μ_R	✓	83.84	57.87	68.75	55.14		48.09/85.42	19.83/95.8	59.6/82.48	13.55/97.09	30.24/93.66	29.56/93.36	33.48/91.3
CLS ; μ_P	×	-	-	-	-	Energy	60.09/85.47	47.68/93.64	74.61/82.53	33.55/94.50	70.7/88.61	64.44/89.28	58.51/89.01
CLS ; μ_P	✓	-	-	-	-		42.32/89.58	18.22/96.34	52.38/87.91	8.73/98.06	24.85/95.82	24.03/95.44	28.59/93.86
CLS ; μ_R	✓	-	-	-	-		37.54/90.71	12.02/97.46	46.38/88.83	7.07/98.57	17.91/96.97	16.87/96.74	22.96/94.88

the best performing baseline of (CLS ; μ_P with reg). Similar trends are observed for ViT-Large and ViT-Base. These results suggest that the combination of by viewing register tokens as auxiliary features and combining it with CLS embeddings yield substantially superior capture more features in terms of generalization and robustness.

B. Anomaly Rejection

Our approach demonstrates even significant improvements for the task of anomaly rejection and consistently outperforms other ablations. For instance, with ViT-Base, FPR@TPR95 drops dramatically on SVHN (from 18.22 % to 12.02%) and LSUN-R (from 24.85 % to 17.91 %). On an average for ViT-Large, our method reduces the mean FPR by 22.18 and 3.22 percentage points compared to the baselines respectively. These improvements hold across all ViT architectures, demonstrating the robustness of our approach. Combining [CLS] and Register tokens enhances the model’s ability to detect anomalies by capturing auxiliary information that helps distinguish ID from OOD samples. These findings indicate potential for effective real-world anomaly rejection.

V. CONCLUSION

In this work, we explored the utility of register token embeddings as auxiliary features that supplement the [CLS]

token in Vision Transformers (ViTs) to improve the feature richness for out-of-distribution (OOD) generalization and anomaly rejection. While previous approaches discard register tokens after training, we demonstrated that these tokens contain valuable information that augments the global features captured by the [CLS] token. By concatenating the [CLS] token with the mean of register tokens, we produced richer representations for the linear classifier to perform better on OOD tasks across multiple ViT architectures pre-trained with or without registers. Our approach consistently outperformed baseline methods in both generalization and anomaly rejection, achieving notable improvements in top-1 accuracy and reductions in false positive rates. Importantly, these gains were achieved without introducing additional computational overhead. These findings suggest that register tokens are not redundant but instead play a crucial role in enhancing the robustness and adaptability of ViTs.

ACKNOWLEDGMENT

This work was performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract No. DE-AC52-07NA27344, Lawrence Livermore National Security, LLC. This work is supported by LDRD projects 22-ERD-006 and 24-FS-002.

REFERENCES

- [1] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [2] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- [3] Hugo Touvron, Matthieu Cord, and Hervé Jégou. Deit iii: Revenge of the vit. In *European conference on computer vision*, pages 516–533. Springer, 2022.
- [4] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. GPT3.int8(): 8-bit matrix multiplication for transformers at scale. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [5] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023.
- [6] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. In *The Twelfth International Conference on Learning Representations*, 2024.
- [7] Jianyu Zhang and Léon Bottou. Learning useful representations for shifting tasks and distributions. In *International Conference on Machine Learning*, pages 40830–40850. PMLR, 2023.
- [8] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- [9] Stefan Heimersheim and Alex Turner. Residual stream norms grow exponentially over the forward pass. In *AI Alignment Forum*, page 23, 2023.
- [10] Olga Kovaleva, Saurabh Kulshreshtha, Anna Rogers, and Anna Rumshisky. Bert busters: Outlier dimensions that disrupt transformers. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3392–3405, 2021.
- [11] Jun Zhao, Zhihao Zhang, Yide Ma, Qi Zhang, Tao Gui, Luhui Gao, and Xuanjing Huang. Unveiling a core linguistic region in large language models. *arXiv preprint arXiv:2310.14928*, 2023.
- [12] William Timkey and Marten Van Schijndel. All bark and no bite: Rogue dimensions in transformer language models obscure representational quality. *arXiv preprint arXiv:2109.04404*, 2021.
- [13] Feng Wang, Jiahao Wang, Sucheng Ren, Guoyizhe Wei, Jieru Mei, Wei Shao, Yuyin Zhou, Alan Yuille, and Cihang Xie. Mamba-r: Vision mamba also needs registers. *arXiv preprint arXiv:2405.14858*, 2024.
- [14] Bobby He, Lorenzo Noci, Daniele Paliotta, Imanol Schlag, and Thomas Hofmann. Understanding and minimising outlier features in neural network training. *arXiv preprint arXiv:2405.19279*, 2024.
- [15] Mingjie Sun, Xinlei Chen, J Zico Kolter, and Zhuang Liu. Massive activations in large language models. *arXiv preprint arXiv:2402.17762*, 2024.
- [16] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12104–12113, June 2022.
- [17] Xianhang Li, Zeyu Wang, and Cihang Xie. An inverse scaling law for clip training. *Advances in Neural Information Processing Systems*, 36, 2024.
- [18] Charlie Saillard, Rodolphe Jenatton, Felipe Llinares-López, Zeldia Mariet, David Cahané, Eric Durand, and Jean-Philippe Vert. H-optimus-0, 2024.
- [19] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [20] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15262–15271, 2021.
- [21] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. *ICCV*, 2021.
- [22] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems*, pages 10506–10518, 2019.
- [23] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems*, 33:21464–21475, 2020.
- [24] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations*, 2017.