
Where Should I Spend My FLOPS?

Efficiency Evaluations of Visual Pre-training Methods

Skanda Koppula¹ Yazhe Li¹ Evan Shelhamer¹ Andrew Jaegle¹ Nikhil Parthasarathy^{1,2}
Relja Arandjelovic¹ João Carreira¹ Olivier Hénaff¹
¹DeepMind ²NYU Center for Neural Systems
{skandak,yazhe,shelhamer,drewjaegle,nikparth,relja,joaoluis,henaff}@google.com

Abstract

Self-supervised methods have achieved remarkable success in transfer learning, often achieving the same or better accuracy than supervised pre-training. Most prior work has done so by increasing pre-training computation by adding complex data augmentation, multiple views, or lengthy training schedules. In this work, we investigate a related, but orthogonal question: given a *fixed* FLOP budget, what are the best datasets, models, and (self-)supervised training methods for obtaining high accuracy on representative visual tasks? Given the availability of large datasets, this setting is often more relevant for both academic and industry labs alike. We examine five large-scale datasets (JFT-300M, ALIGN, ImageNet-1K, ImageNet-21K, and COCO) and six pre-training methods (CLIP, DINO, SimCLR, BYOL, Masked Autoencoding, and supervised). In a like-for-like fashion, we characterize their FLOP and CO₂ footprints, relative to their accuracy when transferred to a canonical image segmentation task. Our analysis reveals strong disparities in the computational efficiency of pre-training methods and their dependence on dataset quality. In particular, our results call into question the commonly-held assumption that self-supervised methods inherently scale to large, uncurated data. We therefore advocate for (1) paying closer attention to dataset curation and (2) reporting of accuracies in context of the total computational cost.

1 Introduction

Self-supervised learning (SSL) has emerged as the dominant pre-training approach to attain state-of-art performance on complex visual tasks. The complexity and diversity of SSL methods have exploded in recent years, gradually improving to supersede both supervised pre-training and from-scratch training. Today’s visual transfer learning practitioners have a menagerie of pre-training methods, models, and datasets from which to choose. Contrastive methods (e.g. CLIP [27], SimCLR [7], and ReLIC [23, 32]), self-distillation methods (e.g. BYOL [12], data2vec [2], DINO [6], and ODIN [16]), and masked autoencoding (MAE) [13, 36] are three of the best-known flavours of modern SSL for visual pre-training. Despite its strong accuracy on downstream tasks, modern SSL is computationally expensive. Most methods train for a constant factor more (sometimes an order of magnitude more) [7, 12] than their supervised counterparts, necessitate multiple forward passes per gradient step [6, 7, 12, 16, 32], or use expensive Transformer-based decoders that scale quadratically with image resolution [13].

In this work, we take a measured look at the computational efficiency of various visual pre-training methods across multiple datasets and model sizes. The increasing monetary and carbon cost of training large deep networks forces us to ask the question: where should I spend my FLOPs for greatest effect? Our paper takes the first steps for determining the strongest methods, datasets, and models for pre-training on a fixed budget, as those in academia and industry are both ultimately

limited by total computation. The computational expense of the (pre-)training stage is often ignored, but its performance has material impact: (1) faster pre-training reduces iteration time for research, (2) less computationally-expensive methods are accessible to more research groups, (3) with larger data and models becoming more accessible, compute becomes the bottleneck, and training efficiency gains therefore translate into greater performance gains [17] (4) these improvements compound as many modalities are combined during pre-training [1, 35], and (5) efficiency is simply better for the environment. Lastly, accounting for computation is more empirically rigorous: more computation often yields better accuracy (more steps, views, etc.), and just as scaling laws for models show improvement with more parameters [17], reporting computation-normalized performance enables more like-for-like comparisons.

Only a few prior works have compared pre-training datasets across SSL methods [10, 20, 29, 37]. Notably, SEER [11] presents strong results pre-training on the large-scale uncurated IG dataset, but do not compare pre-training datasets while controlling for others factors. Additionally, while model scaling curves are popular in the SSL literature [11, 12, 32], backbone model size is not equivalent to computational efficiency, as other aspects like the number of steps can dominate overall FLOP cost. We have yet to find prior work comparing (or reporting) effectiveness as a function of FLOPs. We address this gap by profiling four popular self-supervised methods (BYOL [12], SimCLR [7], DINO [6], and MAE [13]) and two supervised methods (CLIP [27] and standard classification). We compute their per gradient step FLOP cost (post-compilation, running on the same accelerator hardware), and use this to cross-compare the three axes of pre-training method, model size (ViT-B vs. ViT-L [10]), and dataset choice (COCO, ImageNet 1K, ImageNet-21K, ALIGN, and JFT-300M).

We find that the computational efficiency of SSL methods differs by up to an order magnitude, and that FLOP-efficient methods can train significantly larger backbones than FLOP-inefficient ones, while being cheaper to run. The best efficiency/accuracy trade-off (that is, the pareto-optimal curve) of CLIP dominates that of classification supervised methods on large-scale datasets, nearly matched or exceeded by MAE, depending on the dataset and downstream task. Below all these methods are SimCLR, BYOL, and DINO, in that order. Dataset quality and curation level also matter significantly: when controlling for FLOPs, most methods perform significantly worse on uncurated data. While SSL offers the promise of learning effectively on in-the-wild data, most methods are tuned explicitly for, or simply perform better on, the highly-curated ImageNet-1K and 21K datasets [30].

We hope this work sparks future research into visual SSL methods that learn more effectively and scalably on uncurated data, following the initial promise of [11]. As we progress towards richer media and larger datasets, the need for efficient SSL algorithms will grow, and we hope to encourage methods that achieve the promise of label-efficient and computationally scalable learning.

2 Methods

2.1 Pre-training Methods and Models

We run four common self-supervised methods (BYOL [12], SimCLR [7], DINO [6], and MAE [13]) and two supervised methods (CLIP [27] and standard softmax classification). These methods span the contrastive, self-distillation, and reconstruction-based method families popular in modern SSL. These methods are pre-trained with four computational budgets (from 62,500 to 750,000 steps), with a batch size of 4096 (except for DINO, which is run with a batch size of 1024 and for $4\times$ longer). For context, 250,000 steps at batch size 4096 is equivalent to seeing the same number of training images as in the commonly used 800-epochs of ImageNet-1K training in SSL. We use two sizes of vision transformer as the backbone vision encoder: ViT-B/16 (87M parameters) and ViT-L/16 (305M parameters). More detail on pre-training and optimization details is in Section A.3.

Because different masking ratios in MAE affect both accuracy and FLOP cost, we sweep across four different masking ratios (55%, 65%, 75%, and 85%) in the neighborhood of the default 75%. Because the decoder in MAE can occupy up to 85% of the total per-step time and FLOP cost (Section A.5), we experiment with both the default 8-layer MAE decoder, and a ‘tiny’ one-layer decoder that enables 72% faster MAE training. Our CLIP model includes two sizes of a BERT-style language model: either a ‘mini’ 4-layer, 256D-hidden encoder (12M parameters), or a ‘medium’ 8-layer, 512D-hidden encoder (44M parameters). This setup mimics the architectures used in [1]. To train on the hierarchical classification-based JFT, we use the softmax-based averaging method described in Section A.3. For CLIP, we use a pseudo-caption passed to the language encoder of the form ‘This is a photo of <label1> and <label2> ... <label_n>.’.

2.2 Datasets

We use five pre-training datasets: COCO (118K images), ImageNet-1K (1.2M images, ‘11K’), ImageNet-21K (15M images, ‘I21K’, a superset of ImageNet-1K), JFT-300M (302M images), and ALIGN (1.6B images). These datasets span both size and amounts of curation. 11K possesses human-guided views typically framed around a centered object, with 1000 hand-picked classes. ALIGN and JFT are two mildly-filtered web-scale datasets.

2.3 Measurement and Evaluation

We evaluate downstream performance by finetuning the pre-trained encoders for semantic segmentation on the popular dataset, ADE20K [38]. In our experiments, we find this task to be reasonably representative of other complex visual tasks, such as object detection (A.6), but also demand more spatial understanding from representations as compared to more simple objectives such as ImageNet-1K classification. We use the same finetuning segmentation head as in [12, 14, 15, 16], which takes as input the reshaped last layer tokens from the ViT encoder. We finetune each model for a fixed 80 epochs similar to a standard ADE20K setup [25]. More details can be found in Section A.4.

We use a profiling tool for JAX to trace and count the floating point operations of a single pre-training step, for each pre-training method and model size. FLOP measurements occur after model code is lowered and compiled to a specific accelerator setup (8 trays of TPUs [18]). This provides us with an accurate FLOP estimate that correlates well with seconds per training step (Section A.2). CO₂ emissions are estimated by translating the total FLOP count to hypothetical training time on a 32 TPUv3 group. From this, power consumption and carbon emissions are estimated based on the accelerator power draw and datacenter location and efficiency using an online tool [34].

3 Results

3.1 Compute Efficiency of Supervised and Self-supervised Learning

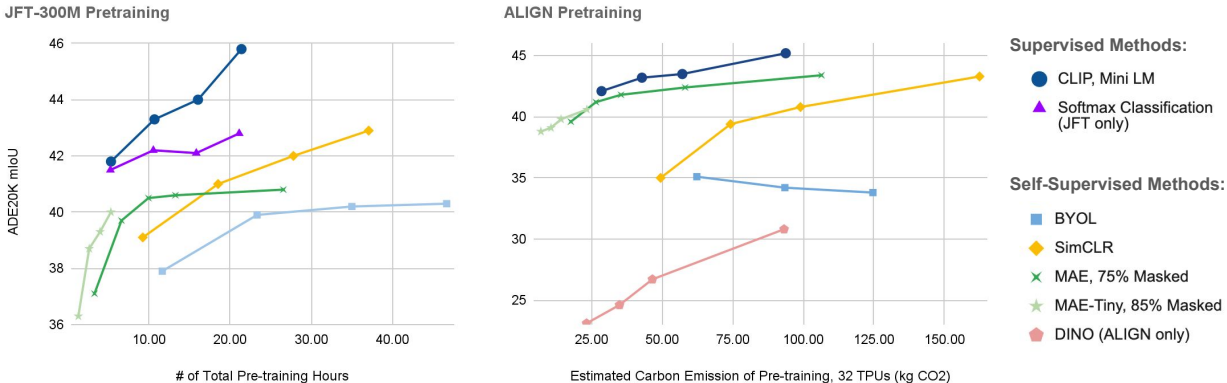


Figure 1: Segmentation accuracy after pre-training ViT-B models on ALIGN or JFT. Supervision consistently outperforms self-supervision in both accuracy (up) and efficiency (left) for this image segmentation evaluation, with MAE closely following supervised pre-training. Best viewed in color.

Our first set of accuracy-efficiency curves are shown in Figure 1. Self-supervised methods are generally less FLOP efficient and supervised representations dominate the efficiency pareto-front. In particular, CLIP pre-training demonstrates strong compute scaling and outperforms the standard classification objective. BYOL performs poorly on these benchmarks and we hypothesize that this is due to the ‘strong’ normalization required by BYOL [28] (which is not provided in a standard ViT backbone). While SimCLR, BYOL, and DINO are substantially less efficient than supervised learning, MAE closely follows the pareto-efficiency front of the supervised methods, and in our object detection results (A.6), MAE even slightly exceeds to supervised curves.

3.2 Scaling Model Size or Training Length?

In our second set of accuracy-efficiency curves, Figure 2, we show the compute scaling of ViT-B and ViT-L which contains roughly 3× more parameters. For a given pre-training budget, it is not always beneficial to scale the model size to obtain better accuracy: for example, with CLIP, a ViT-B trained

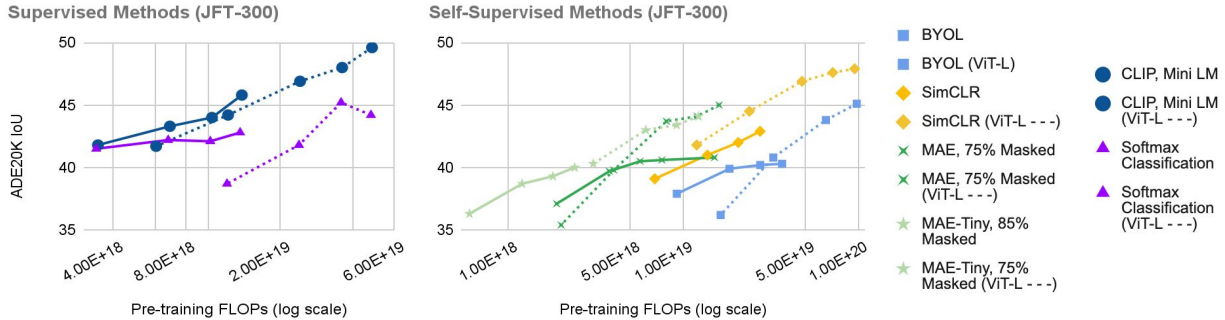


Figure 2: Segmentation accuracy-efficiency curves for JFT pre-training with ViT-B (solid) and ViT-L (dashed). Up is more accurate and left is more efficient. Best viewed in color.

with 3×10^{18} FLOPs outperforms a ViT-L trained with 6×10^{18} (a $3\times$ gain in efficiency). More generally, ViT-B’s FLOP-efficiency exceeds that of ViT-L across the computational regime we tested.

The story is slightly different for self-supervised methods: larger models tend to have an on-par or better compute/accuracy curve for the FLOP regimes tested. This is similar to what is observed in language model scaling [17, 19]. For most methods, the small and large model curves intersect, indicating the point at which it is better to switch to larger model sizes for a given FLOP budget.

3.3 Dataset Choice and Curation

We also evaluate our SSL methods across multiple pre-training datasets, as described in Sections 2.1 and A.3. Downstream results for our ViT-B/16 models are shown in Table 1. Only BYOL (pre-trained on ImageNet-21K) outperforms ImageNet-1K pre-training. We suspect the power of I1K pre-training stems from a combination of (1) I1K’s more even class labels and label curation, (2) implicit human supervision (centering of the photos, framing the object) [31], and (3) tuning methods around the I1K pre-training benchmark (e.g. in choosing augmentations) [8].

Only SimCLR demonstrates consistent improvement as dataset size increases (apart from I1K). For other methods, such as BYOL and DINO, the best results are not with the largest pre-training dataset, ALIGN. Furthermore, ALIGN (which has mild pre-processing, a caption frequency-based filtering step) yields consistently better results than less curated (albeit smaller) JFT. We had difficulty stabilizing long-schedule training for DINO on JFT, longer than the listed 500K steps @ 1K batch size. Our efforts tuning gradient clipping, decay, and learning rate were unsuccessful.

4 Conclusion

We take a first step towards more rigorously characterizing the computational efficiency of modern supervised and self-supervised pre-training methods along the axes of pre-training methods, dataset, and model size. We advocate that for sake of research accessibility, environmental impact, and robust comparisons, more attention be paid to the computational efficiency of visual pre-training methods. Future directions include more fine-grained characterization of our efficiency curves into power scaling laws, design of new SSL objectives and curation methods that overcome dataset sensitivity, and methods that push the learning efficiency of SSL to be competitive with supervised pre-training.

Method	Pre-training Dataset				
	ImageNet-1K	COCO	ImageNet-21K+1K	JFT	ALIGN
BYOL	41.7	37.1	42.4	40.3	35.1
SimCLR	44.0	38.9	42.9	42.9	43.3
MAE	42.3	41.0	40.5	40.6	42.3
DINO	36.7	31.4	34.2	29.2*	32.2

Table 1: ADE20K finetuning results (mIoU) after pre-training on various datasets with four SSL methods (fixed 250K steps). Bold indicates the best pre-training dataset for that method. DINO training with long schedules on JFT (marked with an asterisk) exhibits significant training instability.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a Visual Language Model for Few-shot Learning. *arXiv preprint arXiv:2204.14198*, 2022.
- [2] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. data2vec: A General Framework for Self-supervised Learning in Speech, Vision and Language. *arXiv preprint arXiv:2202.03555*, 2022.
- [3] Andy Brock, Soham De, Samuel L Smith, and Karen Simonyan. High-performance Large-scale Image Recognition Without Normalization. In *International Conference on Machine Learning*, pages 1059–1071. PMLR, 2021.
- [4] BYOL Color Augmentation. <https://github.com/deepmind/deepmind-research/blob/master/byol/utils/augmentations.py#L362>, 2022.
- [5] BYOL Official Implementation. https://github.com/deepmind/deepmind-research/blob/master/byol/byol_experiment.py, 2022.
- [6] 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.
- [7] 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.
- [8] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020.
- [9] DINO Official Implementation. <https://github.com/facebookresearch/dino>, 2022.
- [10] 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. *arXiv preprint arXiv:2010.11929*, 2020.
- [11] Priya Goyal, Mathilde Caron, Benjamin Lefaudeux, Min Xu, Pengchao Wang, Vivek Pai, Man- nat Singh, Vitaliy Liptchinsky, Ishan Misra, Armand Joulin, et al. Self-supervised Pretraining of Visual Features in the Wild. *arXiv preprint arXiv:2103.01988*, 2021.
- [12] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33:21271–21284, 2020.
- [13] 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.
- [14] 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.
- [15] Olivier J Hénaff, Skanda Koppula, Jean-Baptiste Alayrac, Aaron Van den Oord, Oriol Vinyals, and João Carreira. Efficient Visual Pretraining with Contrastive Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10086–10096, 2021.
- [16] Olivier J Hénaff, Skanda Koppula, Evan Shelhamer, Daniel Zoran, Andrew Jaegle, Andrew Zisserman, João Carreira, and Relja Arandjelović. Object Discovery and Representation Networks. *arXiv preprint arXiv:2203.08777*, 2022.

- [17] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- [18] Norman P Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. In-datacenter performance analysis of a Tensor Processing Unit. In *Proceedings of the 44th annual International Symposium on Computer Architecture*, pages 1–12, 2017.
- [19] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [20] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General Visual Representation Learning. In *European Conference on Computer Vision*, pages 491–507. Springer, 2020.
- [21] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. *arXiv preprint arXiv:2203.16527*, 2022.
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014.
- [23] Jovana Mitrovic, Brian McWilliams, Jacob Walker, Lars Buesing, and Charles Blundell. Representation learning via invariant causal mechanisms. *arXiv preprint arXiv:2010.07922*, 2020.
- [24] mmseg FCN Head. github.com/open-mmlab/mmdetection/blob/master/mmdet/models/decode_heads/fcn_head.py, 2022.
- [25] open mmlab. mmdetection. https://github.com/open-mmlab/mmdetection/blob/master/configs/vit/upernet_vit-b16_mln_512x512_80k_ade20k.py, 2022.
- [26] PyTorch ImageNet Pixel Normalization. <https://github.com/pytorch/examples/blob/97304e232807082c2e7b54c597615dc0ad8f6173/imagenet/main.py#L197-L198>, 2022.
- [27] 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.
- [28] Pierre H Richemond, Jean-Bastien Grill, Florent Altché, Corentin Tallec, Florian Strub, Andrew Brock, Samuel Smith, Soham De, Razvan Pascanu, Bilal Piot, et al. Byol works even without batch statistics. *arXiv preprint arXiv:2010.10241*, 2020.
- [29] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21K Pretraining for the Masses. *arXiv preprint arXiv:2104.10972*, 2021.
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [31] Yonglong Tian, Olivier J Henaff, and Aäron van den Oord. Divide and contrast: Self-supervised learning from uncurated data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10063–10074, 2021.
- [32] 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.
- [33] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training Data-efficient Image Transformers & Distillation through Attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021.

- [34] <https://mlco2.github.io/impact/#compute>. Machine learning CO2 impact calculator. *Machine Learning CO2 Impact Calculator*, 2022.
- [35] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022.
- [36] Jiantao Wu and Shentong Mo. Object-wise Masked Autoencoders for Fast Pre-training. *arXiv preprint arXiv:2205.14338*, 2022.
- [37] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Yixuan Wei, Qi Dai, and Han Hu. On data scaling in masked image modeling. *arXiv preprint arXiv:2206.04664*, 2022.
- [38] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene Parsing through ADE20K Dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 633–641, 2017.

A Appendix

A.1 CO₂ estimation

To estimate CO₂ emissions of pre-training, we first must estimate total training time on some fixed hardware (32 TPUv3s, in our case). Unfortunately, total training time is not something that typically can be measured in an apples-to-apples fashion, because each of our pre-training jobs runs on different hardware (various flavours and counts of TPUs and GPUs), based on availability. So, we follow a three step process to obtain more accurate estimates on standardized hardware:

1. We measure the total per step FLOP cost of each pre-training job, through a hardware-targeted JAX profiler.
2. We convert this FLOP cost to the steps per second on 32 TPUv3s. This is done by collecting steps/second measurements from short-schedule pre-training jobs across all pre-training methods. From this, we’re able to model the power-law correlation between per step FLOP cost with steps per second on TPUv3 hardware ($R^2 = 0.97$), and translate more easily between FLOP cost and estimated steps per second. This also has the added benefit of smoothing the small performance measurement variability between runs.
3. With the steps per second of a training job, and the training length (in steps), we calculate the total estimated training on the target hardware topology.

Given the total training time, we use the online calculator from [34], to compute an estimated equivalent power and carbon cost for a single TPUv3 chip in the europe-west1 region. We scale this by 32 (using 32 TPUs in our estimate) to obtain the power and carbon cost numbers we report in the paper.

We note, however, that the datacenters in which our experiments are run are completely carbon offset.

A.2 FLOP estimation

Post-compiled FLOPs can slightly differ from estimated FLOPs, as the compiler can excise or expand parts of the computation graphs, and/or leverage hardware-specific specific SIMD compute instruction ops, based on characteristics of the target hardware. Hence, rather than numerically adding up the theoretical FLOP count based on the architecture implementation, we use AX’s ability to trace through the lowered, post-compilation code to measure the FLOP count.

Figure A.2 reports our measured FLOP cost per training step for the evaluated pre-training methods.

A.3 Pre-training configuration

To gather a wide sample of points along FLOP/accuracy curve, we train all pre-training methods for a varying number of steps. Each pre-training length needs to be trained separately to fit with

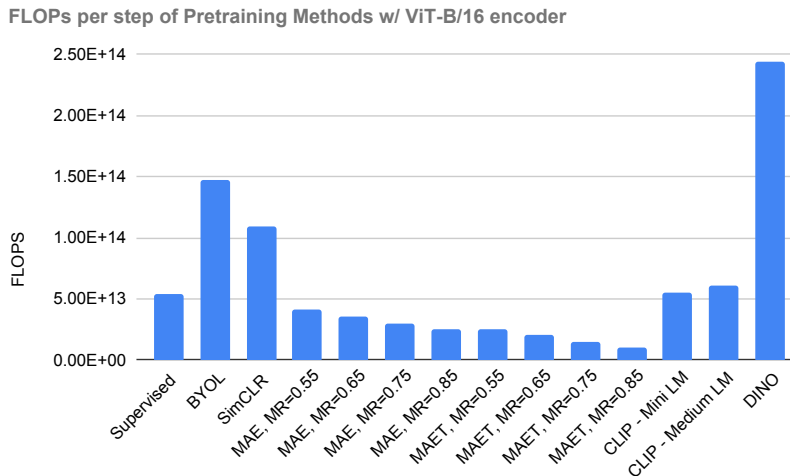


Figure 3: Per step FLOP count of various supervised and self-supervised pre-training methods.

the optimal cosine learning rate schedule that warms up and decays back to zero at the end of the pre-training length. On JFT, we pre-train the all methods except DINO for 62500, 125000, 187500, and 250000 steps at batch size 4096. To match the official implementation, we train DINO with a 1024 batch size, for 125000, 250000, and 500000 steps. To obtain an additional data point further along the curve, in a higher FLOP regime, we train an MAE for 500K steps as well. On ALIGN, we pre-train all methods except DINO for 125000, 187500, 250000, and 410000 steps at batch size 4096. To match the official implementation, we train DINO with a 1024 batch size, for 125000, 250000, 500000, and 750000 steps. To obtain an additional data point along the curve, we train an MAE for 750K steps as well. On I21K, I1K, and COCO, we train all methods for 250K at batch size 4096.

For all methods, we apply standard ImageNet image normalization for both pre-training and finetuning [26], with image size of (224, 224). We use the MAE ViT [13] as the backbone in SimCLR, BYOL, and CLIP (using e.g., the cosine position encodings). Supervised classification pre-training uses the backbone from [10], and DINO uses the standard [6] ViT implementation.

For training softmax classification on the hierarchical JFT, we use the multi-label one-hot averaging scheme from [3]. This scheme performs similarly to a sigmoid-based classification. In pseudo-code:

```

y = one_hot(jft_labels, num_jft_classes)
y = np.sum(y, -2)
y = y / np.sum(y, -1, keepdims=True)
y = smooth_labels(y, 0.1)
classification_loss = softmax_cross_entropy(logits, y, reduction='mean')

```

For CLIP and softmax classification methods, we set the maximum number of allowable labels in JFT to 8. For images with more labels than this, the label list will be truncated; for images with fewer labels, it will be padded.

For CLIP, we re-use the same training crop and rescaling augmentations as [27] and the contrastive pre-training for [1]. For all datasets, we use AdamW with weight decay of $2e-2$, $b1=0.9$, $b2=0.95$, and a gradient clipping thresholding the global norm to 1.0. We use warm-up cosine decay learning rate schedule, with base learning rate of $1e-3$, starting learning rate of $1e-7$, and 10K warmup steps. The joint embedding space has dimension 768 and the language encoder is fed a maximum caption length of 55 characters.

For BYOL, we re-use the same training crop and rescaling augmentations as [12] and the official I1K implementation [5]. For JFT and ALIGN, we use a target EMA of 0.99, base learning rate of $3e-4$, and weight decay of $1e-1$. For I1K and I21K, we use a target EMA of .996 and weight decay of $1e-1$. For COCO, we use a target EMA of .97, weight decay of $1e-6$, and learning rate of $3e-4$. For all datasets, we use the Adam optimizer with $b1=0.9$, $b2=0.999$. We use a warmup cosine schedule, with warmup steps equal to 5% of the total steps in the entire schedule. Our SimCLR implementation uses these same optimization settings (but no target EMA).

For DINO, we use the same augmentations as in the official implementation [9]. We use the Adam optimizer with $b1=0.9$, $b2=0.999$. For most datasets, we use gradient clipping of 1.0 (although this is lowered for larger more unstable datasets like JFT to down to 0.1), 10 local views, and a cosine LR schedule. Our weight decay starts at 0.04 and ends at 0.4. We use an LR of 0.00075 for most datasets, except for JFT long schedules where most of our best results and most stable training is with an LR of 0.0001. In the LR schedule, we use a warmup of 12K steps and minimum LR of $2e-6$.

For classification-based supervised ViT training, we use the standard crop-resize augmentation as is standard in I1K training and drop path rate of 0.1. We use a cosine schedule for the learning rate, with a base value LR of $1e-4$, and 6K warmup steps. We use a weight decay of 0.3. We use the Adam optimizer with $b1=0.9$, $b2=0.95$. These settings are the same across datasets. We use gradient clipping thresholding the global norm at 1.0. Importantly, we use a common initialization trick of initializing values to keep the initial range within 0.02, and initializing the head bias to -10 [10, 33].

For all datasets, we use AdamW with weight decay of $2e-2$, $b1=0.9$, $b2=0.95$, and a gradient clipping thresholding the global norm to 1.0. We use warm-up cosine decay learning rate schedule, with base learning rate of $1e-3$, starting learning rate of $1e-7$, and 10K warmup steps. The joint embedding space has dimension 768 and the language encoder is fed a maximum caption length of 55 characters.

A.4 Fine-tuning configuration

We finetune all models for a fixed 80 epochs on ADE20K, using 16 TPUv1s and a per device batch size of 2. For both training and evaluation, we use an image size of [512, 512]. We add photometric augmentations and randomly rescaled image cropping (minimum scale factor of 0.5, maximum of 2.0, sampling the scaling factor in steps of 0.25) during training. For photometric augmentations, we use the color augmentation function in [4], with settings $brightness=0.1$, $contrast=0.1$, $saturation=0.1$, $hue=0.1$, $apply_prob=0.5$. For evaluation, we take a center crop of the desired size. We use the train split for finetuning and report numbers on the val split.

We use nearly the same optimization settings as [25]. We use a warmup+polynomial LR schedule, with 1500 warmup steps, and the same configuration (`lr_config`) as [25]. We use ADAM optimizer with $b1=0.9$, $b2=0.999$, $eps=1e-8$ and constant weight decay. The segmentation head on top of the ViT encoder takes as input the last layer ViT tokens, reshaped into a square grid, and has the same architecture as implemented in [24] (two 3×3 convolutional layers, and channel dimension of 256). We apply gradient clipping with global norm cutoff of 1.

Similar to [13], we ran an exploratory learning rate (LR) search across $\{1e-5, 5e-5, 1e-4, 5e-4\}$ and weight decay (WD) values $\{0.1, 0.05, 0.01\}$ on the 250K-step pre-trained models. We used the best performing values as the fine-tuning settings for that method and dataset. Generally, we found $LR=1e-4$, $WD=0.05$ as the best setting for MAE pre-trained networks and $LR=5e-5$, $WD=0.1$ as optimal for BYOL, CLIP, and SimCLR pre-trained networks. For DINO, the optimal settings were generally $LR=5e-05$, $WD=0.05$. To obtain final numbers, we ran five fine-tuning jobs with different seeds using the optimal LR/WD settings and reported the median.

From our experiments, this same ViT+segmentation head architecture trained from scratch on this benchmark yields 20.8 mIoU, worse than all supervised or self-supervised based initialization.

A.5 Profiling of MAE

The 75% masking ratio that is commonly used for image tasks is not always optimal across datasets and MAE variants: e.g. on COCO, find that 65% masking to reach highest absolute downstream accuracy. For all other datasets, we find that 75% is optimal. On ALIGN, 85% masking also results in similar performance. With the one-layer decoder, we find that slightly higher masking rates also sometimes have better absolute downstream accuracy and efficiency-normalized accuracy. In Figures 1 and 2, we report masking ratios that result in the highest efficiency/accuracy curve (that envelopes the other ratio's curves, by being further to the left and higher). For Table 1, we report results for the masking ratios that result in highest absolute downstream accuracy.

By masking a significant portion of input tokens, masked autoencoding has a particularly lightweight encoder. From our profiling, we find that the wall-clock performance improvements a over a baseline ViT encoder is often significantly lower than the masking ratio, because of the computational overhead of the decoder. For example, with an 85% masking ratio, wall-clock time savings are slightly under 30%, because computation becomes dominated by the decoder. For example, with an 75% masking

Method	Pre-training Dataset				
	ImageNet-1K	COCO	ImageNet-21K+1K	JFT	ALIGN
BYOL	40.0	39.2	40.3	39.1	27.9
SimCLR	42.0	41.7	41.2	40.3	35.6
MAE	44.8	44.6	43.1	42.6	44.2
DINO	32.8	29.8	33.2	30.9*	32.3

Table 2: COCO finetuning results (mAP) after pre-training on various datasets with four SSL methods. Bold indicates the best pre-training dataset for that method. DINO pre-training with long schedules on JFT (marked with an asterisk) exhibited some training instability.

ratio, roughly 65% of the time to compute a forward and backward pass is occupied by operations for the decoder. With a 55% masking ratio, this decoder percent falls to roughly 45%. This motivated our use of the ‘tiny’ MAE variant, with a simple one-layer decoder. For a 75% masking ratio, the wall-clock speed of tiny-MAE is 95% faster.

A.6 Transfer to Object Detection

To verify the robustness of our overall conclusions, in addition to semantic segmentation, we also test our pre-trained representations on the canonical COCO-2017 [22] object detection task. For finetuning on COCO, we use a Mask-RCNN architecture with a modified ViT-B/16 backbone (instead of a ResNet-50 backbone), similar to [21]. The Mask-RCNN has the same configuration and optimization settings as the ResNet-50 based Mask-RCNN in [15, 16], finetuning for 22K steps. We slightly modify the [21] architecture for our evaluations, and remove their windowed attention and extra convolutional pooling operations inserted into each transformer block (and instead use the ViT backbone unchanged). This reduces the computation efficiency during finetuning, but enables us to use the same exact pre-trained backbone with little modification to the attention mechanism/masking. We report `Box AP@[IoU=0.50:0.95|area=all|maxDets=100]` in our results. For reference, from scratch training for the same number of steps obtains 22.0 mAP.

Table 2 reports COCO object detection results of SSL pre-training on various datasets. Similar to our segmentation results, we find that ImageNet-1K pre-training often outperforms much larger pre-training datasets, with exception of the “full ImageNet” (I21K+I1K), where the performance improvement over I1K alone is small (≤ 0.5 mAP). More broadly, we find in this downstream task as well, that pre-training dataset size does not correlate with downstream performance.

Figure A.6.4 and A.6.5 shows the FLOP efficiency of each method, measured on the object detection task. We obtain similar results to segmentation, with the exception that MAE overtakes one or two supervised methods, depending on the pre-training dataset. MAE performs strongly both in absolute terms, and with respect to its compute efficiency.

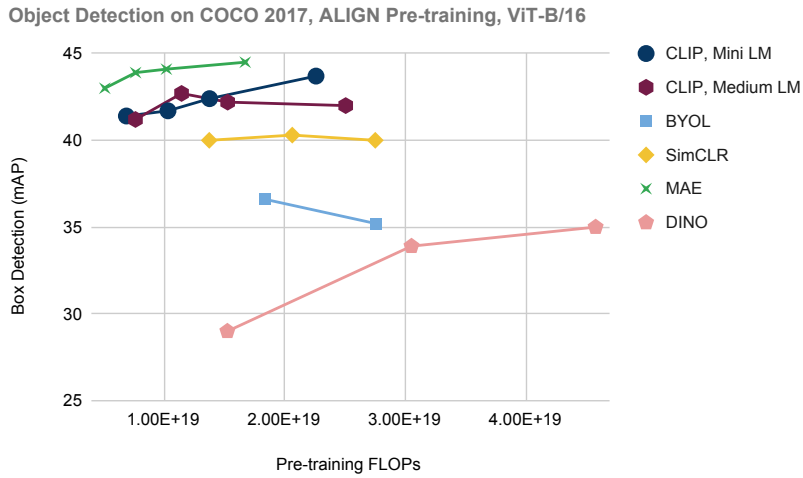


Figure 4: Object detection performance after pre-training ViT-B models on ALIGN. MAE outperforms CLIP-based supervision which in turn outperforms other self-supervision methods in both accuracy (up) and efficiency (left). Best viewed in color.

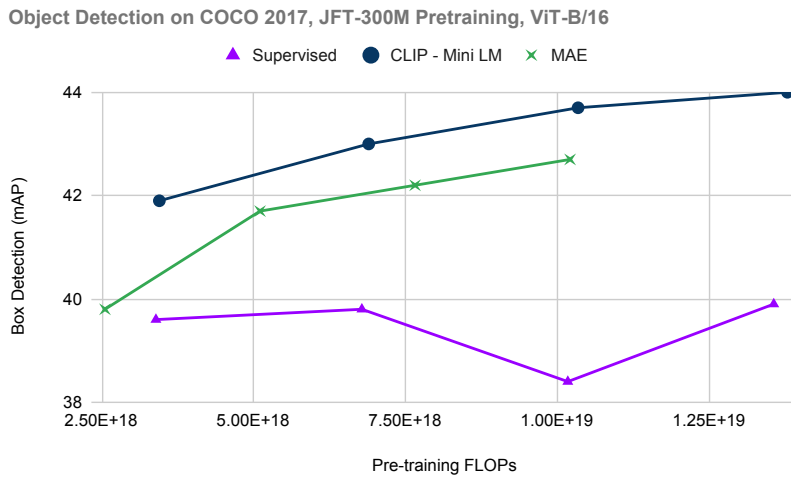


Figure 5: Object detection performance after pre-training ViT-B models on JFT-300M. CLIP outperforms MAE which outperforms standard classification supervision. Best viewed in color.