

CoDo: Contrastive Learning with Downstream Background Invariance for Detection

Bing Zhao Jun Li Hong Zhu

Department of AI and HPC Inspur Electronic Information Industry Co., Ltd
 Beijing, China

{zhaobing01, lijun09, zhuhongbj}@inspur.com

Abstract

The prior self-supervised learning researches mainly select image-level instance discrimination as pretext task. It achieves a fantastic classification performance that is comparable to supervised learning methods. However, with degraded transfer performance on downstream tasks such as object detection. To bridge the performance gap, we propose a novel object-level self-supervised learning method, called **Contrastive learning with Downstream background invariance (CoDo)**. The pretext task is converted to focus on instance location modeling for various backgrounds, especially for downstream datasets. The ability of background invariance is considered vital for object detection. Firstly, a data augmentation strategy is proposed to paste the instances onto background images, and then jitter the bounding box to involve background information. Secondly, we implement architecture alignment between our pretraining network and the mainstream detection pipelines. Thirdly, hierarchical and multi views contrastive learning is designed to improve performance of visual representation learning. Experiments on MSCOCO demonstrate that the proposed CoDo with common backbones, ResNet50-FPN, yields strong transfer learning results for object detection.

1. Introduction

The paradigm of supervised pretraining and finetuning has been dominant in computer vision for a period of time. Typically, the pretraining is optimized on large-scale labeled datasets, and then regarded as initialized weights to finetune for various downstream tasks [12] [14] [32]. Instead, self-supervised learning (SSL) aims to learn the generic pretraining representations, independent of manual labels [18]. Recently, SSL has achieved a performance comparable to supervised pretraining in image classification [10] [2] [1] [9]. However, it suffers a performance degradation when applied to downstream tasks, such as




Figure 1. The performance inconsistency of related SSL methods in image classification and object detection. Specifically, the classification performance of the relevant SSL methods on ImageNet is increasing, while there is no corresponding trend in object detection when finetuning the pre-trained weights on MSCOCO.

object detection. It indicates that the existing SSL approaches mainly focus on image classification, without considering the location modeling ability for object detection. The performance gap of mainstream SSL methods in image classification and object detection is as suggested in Figure 1. It can be observed that the linear classification accuracy of the relevant SSL methods on ImageNet dataset is increasing from 67.5% to 75%. While compared with MOCO [10], which is proposed in 2020, the detection performance (mAP) of recent approaches is reduced when finetuning on MSCOCO dataset.

Existing works that contribute to bridge the gap usually focus on pretext task and architectural alignment. Firstly, instance discrimination, the typical SSL pretext task, usually assumes that different data-augmentations (views) of the same image should be similar but discriminable from other images. Generally speaking, researchers tend to be-

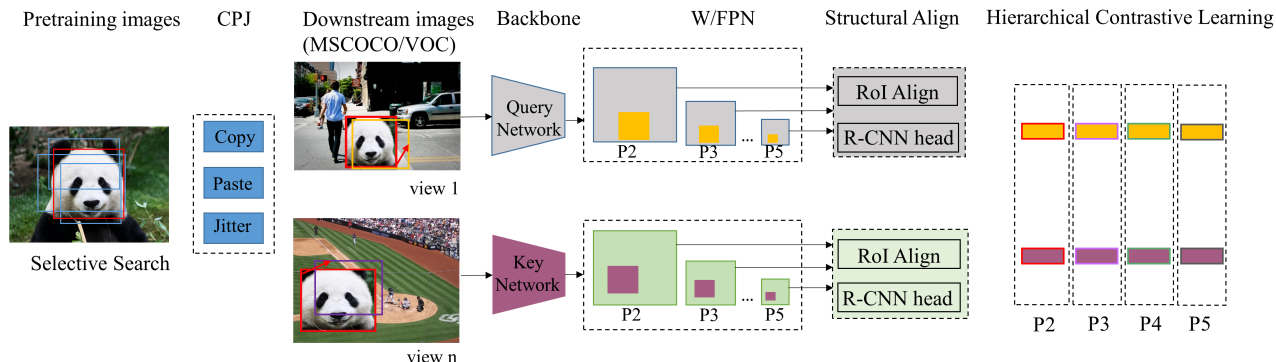


Figure 2. The overview of CoDo. For a pretraining image, we utilize selective search to generate a series of proposals. Only one proposal is randomly selected, then pasted onto various downstream background images.

lieve that image-level pretext task is not suitable for object-level task. Namely, instance discrimination is suit for image classification datasets, such as ImageNet [5], which is usually single-centric-object. However, not for object detection datasets, which mainly consist of non-iconic images, such as MSCOCO [16], there are multiple instances on an image. And secondly, the spatial modeling required for object detection can be optimized during self-supervised pretraining by aligning model structures, such as introducing feature pyramid network [15], RoIAlign [11] and so on. Although SSL makes it possible to involve downstream datasets during pretraining, we identify that prior works still overlook the function of downstream datasets. Specifically, obtaining the location ability for same foreground objects of pretraining datasets, when the various background images are from object detection datasets.

Motivated by this, we present a new object-level contrastive learning method that fuses pretraining and downstream datasets, called **Contrastive learning with Downstream background invariance (CoDo)**. We firstly generate object proposals for pretraining images by selective search [21], and paste them at arbitrary aspect ratios and scales onto various downstream background images. Then, by introducing bounding box jitter, proposals with background information will be regarded as views for object-level contrastive learning.

Involving bounding box in pretraining, we are allowed to refer to the object detectors for structural alignment. Apart from pre-training the backbone, our approach realizes a better initialization for all components in detectors. Experimentally, we study mainstream detection backbone network, ResNet50-FPN on MSCOCO. Our approach shows impressive improvements over the baseline.

2. Related work

Self-supervised learning refers to learn visual features from unlabeled data without human annotations [13]. A mainstream solution is to utilize various pretext tasks to generate pseudo labels in order to obtain the generalizable representations, such as Rotation [8], Colorization [30], Inpainting [19], Jigsaw Puzzle [6] and so on. Recently, contrastive learning becomes the most popular image-level pretext task for self-supervised learning [4]. The optimization goal is to maximize the representation of image instances and their corresponding views, while minimizing the remaining other image instances [25]. The further contrastive learning approaches focus on a better construction of negative sample and a simplification of network structure. For instance, Momentum Contrast (MoCo) [10] builds a moving-averaged encoder and maintain a negative sample queue to minimize a contrastive loss. SimCLR [2] verifies the effectiveness of data augmentation strategy and increases the batch size of training samples for contrastive learning. Bootstrap Your Own Latent (BYOL) [9] minimizes a similarity loss between online and target networks without using negative pairs. Barlow Twins [29] introduces a cross-correlation matrix to avoid trivial constant solutions, the simplified model does not need predictor network, momentum encoder and stop-gradients any more.

The further goal of SSL is to learn general representations which can be transferred for downstream tasks. Some modified approaches have been proposed to bridge the gap between pretraining and downstream tasks. Detco [26] and Insloc [28] separately design a detection-friendly pretext task. Detco contrasts the global image and local image patches to improve object detection. Insloc constructs a new view by randomly pasting foreground objects in different background images for contrastive learning. OLR [27] and Soco [23] propose object-level unsupervised representation learning framework for object detection respectively. Self-

EMD [17] performs the Earth Mover’s Distance as a metric to measure the spatial similarity between two image representation in order to learn spatial visual representations for object detection. DiLo [31] selects saliency estimation to localize the foreground object in a data-driven approach. Self-EMD [17] directly pre-trains on MSCOCO and achieves a higher detection performance, instead of commonly used ImageNet.

3. Method

We propose a new object-level contrastive learning framework as shown in Figure 2, which introduces the unlabeled downstream datasets information during pre-training. The proposals of the pretraining images are pasted to the downstream images, the bonding box of proposals is deformed to generate new views with downstream information. It enables localization of the foreground object during pretraining, and can be regarded as a simulation of object detection during pretraining. Considering that no supervised information of downstream dataset is introduced into the pre-training, so there is no risk of data leakage. Meanwhile, we introduce architectural alignment to pre-train the essential properties of object detector, such as FPN, RoIAlign and R-CNN head. Our proposed method is detailed below.

3.1. Copy, paste and jitter (CPJ)

We design a data augmentation method for proposals to realize location modeling, terms as CPJ. In order to realize object-level contrastive learning, we select selective search to generate proposal in unsupervised way. Considering ImageNet is usually regarded as a single-centric-object dataset, most proposals are similar. So we only randomly select one proposal for each image, those proposals with too large or too small aspect ratios (≥ 3 or $\leq 1/3$) are ignored. In addition to ImageNet, we add two downstream object detection datasets (MSCOCO and Pasco VOC [7]) as alternative background images. For the same pretraining image, depending on the number of views involved in the contrastive learning, we randomly select the corresponding number of background images to generate the pasted images.

Background invariance is important for object detection, namely, a robust detector can recognize the foreground objects on various backgrounds. We paste the proposals at arbitrary aspect ratios and scales onto various downstream background images. In this process, translation and scale invariance are also considered. The pasted position is treated as the bounding box. Then the bounding box is jittered to contain background images. The jittered boxes are filtered by a proper IoU threshold, we set it to be greater than 0.6. The whole image I and transformed bounding box bb are regarded as inputs of Query network and Key network to conduct contrastive learning.

$$I_q, bb_q = CPJ(I_p, I_{dq}). \quad (1)$$

$$I_{k_i}, bb_{k_i} = CPJ(I_p, I_{dk_i}). \quad (2)$$

Where I_p is the proposal of pretraining images, I_{dq} is the downstream background image for Query network, and I_{dk_i} is the i -th downstream background image for Key network.

To be consistent with existing SSL methods, the default version of our proposed method selects two views for Query network and Key network separately. Actually, multi views help to increase diversity during contrastive learning. We also design a 4 views version, which view 2 to view 4 are inputs of Key network.

3.2. Hierarchical contrastive learning

In our approach, the pipeline of MOCO-V2 [3] is adopted as baseline for learning contrastive representations. We will describe how to achieve structural alignment for MOCO V2 in this section. We select the Resnet50 with FPN as Query network f^q and Key network f^k . FPN is a common component in the object detector, which can fuse different feature maps to reach a better detection performance. To further align with Mask R-CNN, RoIAlign is introduced to extract feature of bounding box from the output of FPN. For Query network, the object-level feature representation v_q is extracted from an image I_q and the corresponding bounding box bb_q as follows. The computation for Key network is similar.

$$v_q = \text{RoIAlign}(f^q(I_q), bb_q). \quad (3)$$

$$v_{k_i} = \text{RoIAlign}(f^k(I_{k_i}), bb_{k_i}). \quad (4)$$

And R-CNN head f^{R-H} is built to obtain embeddings for contrastive learning. The latent embeddings e_q and e_{k_i} for Query network and Key network are as follows:

$$e_q = f^{R-H}(v_q). \quad (5)$$

$$e_{k_i} = f^{R-H}(v_{k_i}). \quad (6)$$

SSL usually selects contrastive loss, i.e. InfoNCE to compute the similarity between views. For Query view and i -th Key view, the loss function is as follows. Notably, The calculation is performed in a hierarchical manner. Because e_q and e_{k_i} can be divided corresponding to the output of FPN ($P2, P3, P4, P5$). Contrastive learning can be carried out at a finer level.

$$L_{q-k_i} = -\log \frac{\exp(e_q \cdot v_{e_i}/\tau)}{\sum_{i=0}^N \exp(e_q \cdot e_{k_i}/\tau)} \quad (7)$$

Table 1. Downstream task performance on COCO by using Mask R-CNN with R50-FPN.

Methods	Epoch	1x Schedule						2x Schedule					
		AP^{bb}	AP_{50}^{bb}	AP_{75}^{bb}	AP^{mk}	AP_{50}^{mk}	AP_{75}^{mk}	AP^{bb}	AP_{50}^{bb}	AP_{75}^{bb}	AP^{mk}	AP_{50}^{mk}	AP_{75}^{mk}
Sepurvised	90	38.9	59.6	42.7	35.4	56.5	38.1	41.3	61.3	45.0	37.3	58.3	40.3
MOCO	200	38.5	58.9	42.0	35.1	55.9	37.7	40.8	61.6	44.7	36.9	58.4	39.7
MOCO v2	200	40.4	60.2	44.2	36.4	57.2	38.9	41.7	61.6	45.6	37.6	58.7	40.5
SimCLR	200	39.6	59.1	42.9	34.6	55.9	37.1	40.8	60.6	44.4	36.9	57.8	39.8
InfoMin	200	40.6	60.6	44.6	36.7	57.7	39.4	42.5	62.7	46.8	38.4	59.7	41.4
InfoMin	800	41.2	61.2	44.8	35.9	57.9	38.4	42.1	62.3	46.2	38.0	59.5	40.8
BYOL	300	40.4	61.6	44.1	37.2	58.8	39.8	42.3	62.6	46.2	38.3	59.6	41.1
SWAV	400	39.6	60.1	42.9	34.7	56.6	36.6	42.3	62.8	46.3	38.2	60.0	41.0
DenseCL	200	40.3	59.9	44.3	36.4	57.0	39.2	41.2	61.9	45.1	37.3	58.9	40.1
InsLoc	200	41.4	61.7	45.0	37.1	58.5	39.6	43.2	63.5	47.5	38.7	60.5	41.9
CoDo	200	41.2	61.3	45.1	36.9	58.2	39.4	42.6	62.6	46.7	38.2	59.7	41.0
CoDo	400	41.9	61.8	45.8	37.4	58.8	40.3	43.1	63.4	47.1	38.8	60.6	41.4
CoDo_m	400	43.1	63.3	47.1	38.3	60.3	41.2	-	-	-	-	-	-

Table 2. The influence of background datasets on COCO by using Mask R-CNN with R50-FPN. Q means Query Network and K means Key Network.

Background Datasets	Epoch	Schedule	AP^{bb}		AP_{50}^{bb}	AP_{75}^{bb}	AP^{mk}	AP_{50}^{mk}	AP_{75}^{mk}
Q:ImageNet K: COCO	200	1x	40.7	-	60.7	44.5	36.5	57.6	39.1
Q/K: ImageNet+COCO	200	1x	40.8	+0.1	60.7	44.7	36.6	57.6	39.3
Q/K: ImageNet+COCO+VOC	200	1x	41.2	+0.5	61.3	45.1	36.9	58.2	39.4
Q: ImageNe K:ImageNet+COCO+VOC	400	1x	41.4	-	61.7	45.1	37.3	58.6	40.1
Q/K: ImageNet+COCO+VOC	400	1x	41.9	+0.5	61.8	45.8	37.4	58.8	40.3

Where τ and N are the temperature and the number of negative samples, respectively.

For multi-view version of our proposed method, we can calculate InfoNCE between view q and all view k . The total loss function is as follows.

$$L_{q-k} = \sum_{j=0}^N (L_{q-k_i})_j \quad (8)$$

Where N is the number of view k .

4. Experiment

4.1. Dataset

The widely used ImageNe-1K with 1.28 million images is adopt as dataset for self-supervised pretraining. MSCOCO is used for finetuning to evaluate the generalization performance on downstream task. Significantly, in the procession of data processing, the training sets of ImageNe-1K, PASCAL VOC0712 and MSCOCO are involved as the background images.

4.2. Setting for pretraining and finetuning

During pretraining, we mainly follow the hyperparameters setting of Moco-v2, the total batch size is set to 1024 over 8 Nvidia A100 GPUs, and the initial learning rate is set to 0.06. The optimization takes 200 and 400 epochs

for the evaluation of downstream tasks, respectively. During pretraining, we employ the data augmentation pipeline of Moco-v2 for pretraing proposals and the background images.

For finetuing, we validate the performance of pre-training representation on downstream tasks based on Detectron2 [24]. On COCO, we adopt the Mask R-CNN with the R50-FPN. The performances of object detection and instance segmentation under 1x and 2x schedules are reported. The batch size is set to 128 over 8 GPUs. The fine tuning iteration step is set to 90000 and 180000 on the 1x schedule and 2x schedule, respectively. The initial learning rate is 0.02. Finally, on the 1x schedule and 2x schedule, the results of AP^{bb} , AP_{50}^{bb} and AP_{75}^{bb} for object detection, AP^{mk} , AP_{50}^{mk} , and AP_{75}^{mk} for instance segmentation are compared with the state-of-the-art methods. R50-FPN (ResNet-50 with FPN) is the common backbone network for Mask R-CNN and Faster R-CNN to evaluate transfer performance.

4.3. Results

We report the experimental results on object detection and instance segmentation with state-of-the-art approaches. Some of these sota methods are designed to focus on classification, such as SimCLR [2], MoCo [10], MoCo v2 [3], BYOL [9], InfoMin [20] and SwAV [1]. While others are specifically suited to object detection, such as

DenseCL [22] and InsLoc [28].

Mask R-CNN on MSCOCO. Table 1 shows the results for Mask R-CNN with R50-FPN. The finetuning follows with the COCO 1× and 2× schedules. We compare the proposed method under 200 and 400 epochs of pretraining. The results show that our method exceeds the above two kinds of methods. On 1× schedule, Our method outperforms the baseline MoCo-v2 by +0.8 AP for the R50-FPN. On 2× schedule, our method exceeds MoCo-v2 by +0.8. The multi-view version of CoDo_m further boosts the performance and reaches a 43.1 AP.

The influence of Background datasets. We analyze the background datasets selection strategy for Query Network and Key Network. Table 2 shows the influence of background datasets on object detection. Firstly, considering the volume gap between ImageNet and object detection datasets, if we only select the object detection datasets as background may cause homogenization. A proper solution is to gather ImageNet and object detection datasets together as the source of the background images. And the diversity of background dataset sources helps to improve performance. Under 200 epochs, this setting contributes a +0.5AP. Secondly, the background datasets selection for Query and Key Network should be same to reduce Disequilibrium. Under 200 epochs, this role also contributes a +0.5AP.

5. Conclusion

In this paper, we noticed that most existing self-supervised learning methods ignore the function of downstream datasets, especially the location ability of foreground objects in downstream scenes. So we propose a new contrastive learning method, CoDo, to achieve the downstream background invariance of the foreground objects. It is implemented by pasting foreground proposals onto various downstream images for contrastive learning. CoDo achieves a strong results on transfer performance for object detection on MSCOCO. The experimental results demonstrate that transfer performance for object detection can be strengthened by considering downstream background invariance.

References

- [1] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020. 1, 4
- [2] 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. 1, 2, 4
- [3] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 3, 4
- [4] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. Debaised contrastive learning. *Advances in neural information processing systems*, 33:8765–8775, 2020. 2
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2
- [6] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015. 2
- [7] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, Jan. 2015. 3
- [8] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. 2
- [9] 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. 1, 2, 4
- [10] 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. 1, 2, 4
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 2
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1
- [13] Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):4037–4058, 2020. 2
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 1
- [15] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 2
- [16] 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. 2

- [17] Songtao Liu, Zeming Li, and Jian Sun. Self-emd: Self-supervised object detection without imagenet. *arXiv preprint arXiv:2011.13677*, 2020. 3
- [18] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering*, 2021. 1
- [19] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016. 2
- [20] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? *Advances in Neural Information Processing Systems*, 33:6827–6839, 2020. 4
- [21] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International journal of computer vision*, 104(2):154–171, 2013. 2
- [22] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3024–3033, 2021. 5
- [23] Fangyun Wei, Yue Gao, Zhirong Wu, Han Hu, and Stephen Lin. Aligning pretraining for detection via object-level contrastive learning. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [24] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 4
- [25] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. 2
- [26] Enze Xie, Jian Ding, Wenhai Wang, Xiaohang Zhan, Hang Xu, Peize Sun, Zhenguo Li, and Ping Luo. Detco: Unsupervised contrastive learning for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8392–8401, 2021. 2
- [27] Jiahao Xie, Xiaohang Zhan, Ziwei Liu, Yew Ong, and Chen Change Loy. Unsupervised object-level representation learning from scene images. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [28] Ceyuan Yang, Zhirong Wu, Bolei Zhou, and Stephen Lin. Instance localization for self-supervised detection pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3987–3996, 2021. 2, 5
- [29] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021. 2
- [30] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016. 2
- [31] Nanxuan Zhao, Zhirong Wu, Rynson WH Lau, and Stephen Lin. Distilling localization for self-supervised representation learning. *arXiv preprint arXiv:2004.06638*, 2020. 3
- [32] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020. 1