

Unsupervised Representation Learning for Binary Networks by Joint Classifier Learning

Dahyun Kim^{1,2}

¹Upstage AI Research

kdahyun@upstage.ai

Jonghyun Choi^{2,3,†}

²NAVER AI Lab.

³Yonsei University

jc@yonsei.ac.kr

Abstract

Self-supervised learning is a promising unsupervised learning framework that has achieved success with large floating point networks. But such networks are not readily deployable to edge devices. To accelerate deployment of models with the benefit of unsupervised representation learning to such resource limited devices for various downstream tasks, we propose a self-supervised learning method for binary networks that uses a moving target network. In particular, we propose to jointly train a randomly initialized classifier, attached to a pretrained floating point feature extractor, with a binary network. Additionally, we propose a feature similarity loss, a dynamic loss balancing and modified multi-stage training to further improve the accuracy, and call our method BURN. Our empirical validations over five downstream tasks using seven datasets show that BURN outperforms self-supervised baselines for binary networks and sometimes outperforms supervised pretraining. Code is available at <https://github.com/naver-ai/burn>.

1. Introduction

Self-supervised learning (SSL) has achieved great success with floating point (FP) networks in recent years [4, 5, 7, 9, 13, 14, 16, 18, 20, 28, 41, 42, 44, 50]. Models learned by SSL methods perform *on par* with or even outperform the ones learned by supervised pretraining by the help of large scale unlabeled data in a number of downstream tasks such as image classification [1, 5], semi-supervised fine-tuning [5, 7, 18] and object detection [20]. While recent works [5, 7, 18, 20] from resourceful research groups have shown that the gains from SSL scale up with model size and/or dataset size used for pretraining, there is little work where the resulting pretrained models are small in size, *i.e.*, quantized. SSL for such small models is important since it could expedite the AI deployment for a wide range of appli-

This work was done while DK and JC were an intern and an AI technical advisor at NAVER AI Lab., respectively. † indicates corresponding author.

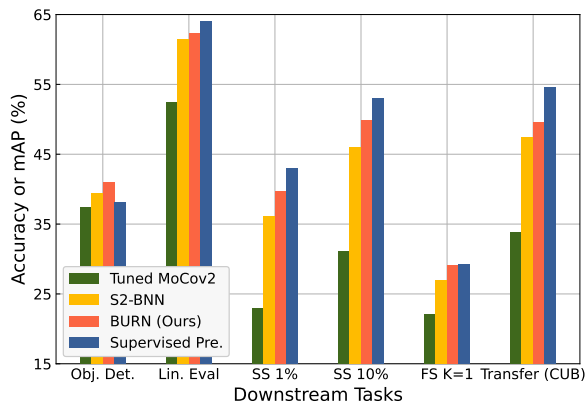


Figure 1. Comparison of various representation learning methods on multiple downstream tasks (pretrained with ImageNet). ‘Obj. Det.’ refers to object detection, ‘Lin. Eval’ refers to linear evaluation, ‘SS 1/10%’ refers to semi-supervised fine-tuning with 1 or 10% data respectively, ‘FS K=1’ refers to few-shot learning with 1 shot, and ‘Transfer (CUB)’ means transfer learning to the CUB dataset. ‘Tuned MoCov2’ and ‘S2-BNN’ are SSL methods from [39]. Proposed BURN outperforms all comparable methods in various tasks, and even the Supervised Pre. in certain tasks.

cations onto models with high efficiency in computational and memory costs, and energy consumption [12]. At the extreme of resource constrained scenarios, binary networks exhibit superior efficiency and the accuracy is being significantly improved [2, 3, 23, 31–34, 38]. Thus, developing an SSL method for binary networks could further accelerate the deployment of models to edge devices for various downstream tasks, yet is seldom explored.

Providing additional supervisory signals with a pre-trained FP network by using the KL divergence loss between the softmax outputs from the classifiers of the FP target network and binary network, which we denote as ‘*supervised KL div.*’, has become a popular and effective method for training binary networks [2, 3, 32, 34]. Recently, [39] propose an unsupervised representation learning method for binary networks based on the supervised KL div. method. To extract meaningful softmax probabilities from the FP network, they pretrain the classifier as well as the feature

extractor using SSL. Then, the FP network is completely frozen when used as the target network, which could lead to stale targets [18] or be dependent on the pretraining dataset used for the fixed FP network being similar to the dataset used for training the binary network.

Thus, to avoid the potential pitfalls of a fixed target, we are motivated to develop an SSL method for binary networks that uses a moving FP network as the target, similar to other SSL methods [8, 9, 18, 20], and call our method **Binary Unsupervised Representation** learning or **BURN**. Specifically, we first construct the FP target network by combining a fixed FP feature extractor pretrained in an SSL manner and a randomly initialized FP classifier. We then use the outputs of the randomly initialized FP classifier as targets for the binary network and *jointly optimize both the FP classifier and the binary network*, using the KL divergence loss, to keep updating the FP network overtime. But the gradients provided by the randomly initialized FP classifier could have unexpectedly large magnitudes, especially during early training phase. To alleviate this problem, we additionally propose to enforce feature similarity across both precision, providing stable gradients that bypass the randomly initialized classifier. As relative importance of the feature similarity loss decreases as the FP classifier gets jointly trained to provide less random targets, we further propose to *dynamically balance* the KL divergence term and the feature similarity term in the loss function. Finally, we modify the multi-stage training scheme [34] for BURN to further improve performance.

We conduct extensive empirical validations with a wide variety of downstream tasks such as object detection on Pascal VOC, linear evaluation on ImageNet, semi-supervised fine-tuning on ImageNet with 1% and 10% labeled data, SVM classification and few-shot SVM classification on Pascal VOC07, and transfer learning to various datasets such as CIFAR10, CIFAR100, CUB-200-2011, Birdsnap, and Places205. In the validations, the binary networks trained by our method outperforms other SSL methods by large margins (see Fig. 1 and Sec. 4.1).

We summarize our contributions as follows:

- We propose a novel SSL method for binary networks that uses a jointly trained FP classifier to obtain targets that can adapt overtime to the current training scenario.
- We propose to use a feature similarity loss and dynamic balancing with modified multi-stage training to significantly improve the accuracy.
- Our BURN outperforms prior arts by large margins on a wide variety of downstream tasks.
- We analyze our proposed BURN by in-depth investigations.

2. Related Work

2.1. Self-Supervised Representation Learning

To reduce the annotation cost for representation learning, self-supervised representation learning (SSL) methods including [6, 7, 9, 16, 20, 42–44, 50] and many more have been shown to be effective, with the Info-NCE loss [36] being a popular choice for many works. These methods use the instance discrimination task as the pretext task which aims to pull instances of the same image closer and push instances of different images farther apart [36, 48]. Different to these methods, [1, 5, 15, 18, 29, 47] use feature regression with an EMA target [18], matching cluster assignments [5, 29], or matching similarity score distributions [1, 15, 47] as the pretext task. We compare with BYOL [18] and SWAV [5] as they show high performance and have similarities with other SSL methods [7, 9, 44]. However, while these methods show promising results for large FP models, they do not consider resource constrained scenarios which are more practical, *e.g.*, quantized models with smaller complexity.

2.2. Binary Networks

At the extreme of quantized models, numerous works on binary networks [2, 3, 19, 23, 24, 30–35, 37, 38] have been proposed. These include searching architectures for binary networks [2, 22, 23], using specialized activation functions [32], and object detection using the information bottleneck principle [46]. Note that previous works mostly focused on the supervised training set-up.

Among many proposals, two stand out as the state-of-the-art binary network backbones due to their strong empirical performance: ReActNet [32] and High-Capacity Expert Binary Networks (HCEBN) [3]. [32] propose to learn thresholds for binarization by the RSign and RReLU activation functions. [3] use multiple experts for conditional computing and increase the representation capacity of binary networks without increasing the operation count.

Recently, both ‘supervised KL div.’ method and the multi-stage training scheme [32, 34] have become popular for training binary networks. The supervised KL div. method uses a pretrained FP network to provide targets for the KL div. loss in training binary networks. The multi-stage training scheme trains a binary network in multiple stages, where more parts of the network are binarized. The most related work is the recently published S2-BNN [39] which utilize the supervised KL div. method to train binary networks in an SSL manner. They use unlabeled data when training the binary network with only the KL divergence loss between the binary and the pretrained and frozen FP networks. S2-BNN shows good performance but the frozen FP network as a target may be limiting *i.e.*, many works in the SSL literature suggest that a *changing* target is effective [8, 9, 18, 20]. We extensively compare with S2-BNN.

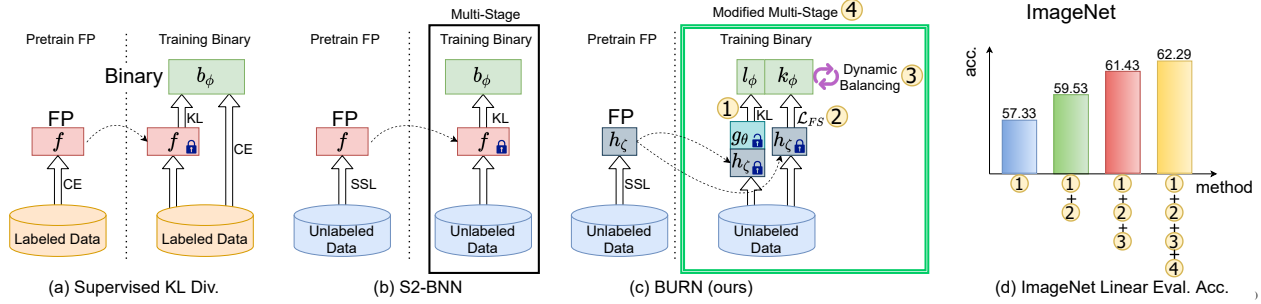


Figure 2. Illustrative comparison of Supervised KL div. method [32, 34], S2-BNN [39] and proposed BURN, and linear evaluation accuracy of ablated models on ImageNet. f is the FP network, h_ζ is the FP feature extractor, g_θ is the trainable classifier, and k_ϕ, l_ϕ are the binary feature extractor and classifier decoupled from the binary network b_ϕ . f and h_ζ are initialized with pretrained weights as indicated by the dotted arrow. The ‘locked’ icons in f and h_ζ indicate they are not trained while the ‘unlocked’ icon in g_θ indicates that it is trained. Our baseline (①) already achieves a relatively high top-1 accuracy of 57.33% on ImageNet and our proposed components provide noticeable gains. For details, please see Sec. 3.1 for ①, Sec. 3.2 for ②, Sec. 3.3 for ③, and Sec. 3.4 for ④.

In contrast, we aim to develop an unsupervised representation learning method for binary networks which uses a FP network that is changing as the target.

3. Approach

The supervised KL div. method is an effective method to train binary networks [32, 34] that utilizes a FP network pretrained with labeled data. But, as we are interested in the self-supervised learning with no access to labeled data at any time during training, the supervised KL div. is not applicable. Recently, S2-BNN [39] propose to use the supervised KL div. for unsupervised learning of binary networks. They pretrain the classifier and the feature extractor of the FP network to obtain meaningful softmax probabilities and use a completely fixed FP network as the target. In contrast, we propose an unsupervised representation learning method for binary networks that uses a changing FP network as the target such that the FP network can adapt to the current dataset and binary network to provide more useful targets overtime. We illustrate the supervised KL div. method [32, 34], S2-BNN [39], and our proposal in Fig. 2.

Specifically, instead of using softmax outputs from a fixed pretrained FP network [39], we propose to use softmax outputs from a *randomly initialized classifier* attached to a pretrained FP feature extractor, and jointly train the classifier with the binary network using the KL divergence loss. As the supervision from the untrained classifier makes gradients with unexpectedly high magnitudes, we subdue gradients by proposing an additional feature similarity loss across precision. We propose to use a dynamic balancing scheme between the loss terms to better balance the KL divergence and feature similarity losses and employ a modified multi-stage training [34] to improve learning efficacy.

3.1. Joint Classifier Training as Moving Targets

Grill *et al.* [18] show that even when a randomly initialized exponential moving average (EMA) network is used as

the target network, the online network improves by training with it. One possible reason for the improvement is that the randomly initialized target network is also updated in an EMA manner during training, improving it gradually. Motivated by this, we conjecture whether a randomly initialized classifier combined with a pretrained FP feature extractor can be used as a *moving target network* for training binary networks. To gradually improve the target network, we jointly train the classifier of the target network and the binary network. Note that training just the classifier can improve the target network as is shown in the SSL literature [6, 7, 9, 16, 20, 42–44, 50]. We discuss other moving targets, *e.g.*, the EMA target [18] or the momentum encoder [20] for binary networks in Sec. 4.2.

The joint training of the randomly initialized classifier is depicted in ① in Fig. 2-(b). Specifically, instead of a fixed FP network $f(\cdot)$, the randomly initialized and trainable classifier $g_\theta(\cdot)$ and the pretrained and fixed FP feature extractor $h_\zeta(\cdot)$ are combined to create the target network. Then, we use the outputs of $g_\theta(\cdot)$ as targets for training the binary network $b_\phi(\cdot)$. Our objective is to minimize the KL divergence between the outputs of $g_\theta(\cdot)$ and $b_\phi(\cdot)$ as:

$$\min_{\theta, \phi} \mathbb{E}_{x \sim \mathcal{D}} [\mathcal{L}_{KL}(g_\theta(h_\zeta(x)), b_\phi(x))], \quad (1)$$

where x is a sample from the dataset \mathcal{D} and $\mathcal{L}_{KL} = D_{KL}(\cdot, \cdot)$ is the KL divergence between the outputs of $g_\theta(\cdot)$ and $b_\phi(\cdot)$. However, the softmax outputs from the classifier would be close to random early on. Thus, using the random outputs as the only target for the binary network, especially in early training, could result in noisy gradients.

3.2. Stabilize Gradients by Feature Similarity Across Precision

To alleviate the issue of unreliable gradients from the randomly initialized classifier being the only target, particularly in the early the training, we propose an additional loss term that enforces feature similarity between the target




Figure 3. Gradient magnitude for the binary classifier (a) and the binary feature extractor (b) during early training with and without \mathcal{L}_{FS} for pretraining on ImageNet. With only KL, the gradients of the classifier is extremely large and this carries over to the feature extractor. Additionally, we observe intermediate spikes for both the classifier and the feature extractor. The addition of \mathcal{L}_{FS} significantly lowers the gradient magnitudes of the classifier as well as the feature extractor at early iterations. Additionally, the surges in gradient magnitudes are also subdued.

and the binary network. Specifically, $g_\theta(\cdot)$ is largely updated (due to the fixed feature extractor) in the early phase of the joint training. As the binary classifier uses the quickly changing $g_\theta(\cdot)$ as a target to transfer knowledge from, the binary classifier might receive large gradients. To address the potentially undesirably large gradients caused by using the randomly initialized classifier as the only target, we propose to augment an additional loss term that bypasses the classifier. We call it the *feature similarity loss*.

Specifically, we use the cosine distance between the feature vectors from the FP and binary feature extractors as the feature similarity loss; $\mathcal{L}_{FS}(v_1, v_2) = 1 - \frac{\langle v_1, v_2 \rangle}{\|v_1\|_2 \cdot \|v_2\|_2}$ for smoothness and a bounded nature to prevent large gradients. The cosine distance (or $1 - \text{cosine similarity}$) is widely used in representation learning [6, 18, 20, 49] (discussion on other choices for \mathcal{L}_{FS} is in Sec. 4.2).

Augmenting the cosine distance to the KL divergence loss, we can write our new optimization problem as:

$$\min_{\theta, \phi} \mathbb{E}_{x \sim \mathcal{D}} [(1 - \lambda) \mathcal{L}_{KL}(g_\theta(h_\zeta(x)), l_\phi(k_\phi(x))) + \lambda \mathcal{L}_{FS}(h_\zeta(x), k_\phi(x))], \quad (2)$$

where the binary network $b_\phi(\cdot)$ is decoupled into $k_\phi(\cdot)$ the binary feature extractor and the classifier $l_\phi(\cdot)$, λ is a static balancing factor, and $\mathcal{L}_{FS}(\cdot, \cdot)$ is the feature similarity loss.

The new loss provides additional supervisory signals from the feature extractor of the FP network. Since the feature extractor of the FP network is pretrained and fixed, it provides stationary and stable targets as opposed to the randomly initialized classifier. Empirically, we observe the gradient of the binary classifier and feature extractor with and without \mathcal{L}_{FS} in Fig 3. Note that with only KL, the gradients of the binary classifier are extremely large; it starts at roughly 20,000 then drops to roughly 3,000 for some iterations and finally drops to a small value around roughly

9,000 iterations. In addition, there is a surge in gradient magnitude around 7,500 iterations. The binary feature extractor also shows a similar trend where the gradients exhibit a sudden spike at around 7,500 iterations. Both very high magnitudes of the gradients at the start and the sudden spikes, occurring after some iterations, would harm training stability [10, 51]. However, as shown in the figures, addition of the proposed $\mathcal{L}_{FS}(\cdot, \cdot)$ significantly reduces the gradient magnitudes of the binary classifier and the feature extractor at early iterations as well as the surges throughout the training, which leads to better training efficacy and accuracy.

3.3. Dynamic Balancing of λ

As g_θ is gradually updated, it provides more meaningful targets and \mathcal{L}_{FS} becomes less important. Thus, we propose a temporally *dynamic balancing* strategy to replace the static balancing factor λ in Eq. 2 by a smooth cosine annealing similar to how [18] annealed the momentum value as:

$$\lambda(t) = \lambda_{T_{max}} - (\lambda_{T_{max}} - \lambda_0) \cdot (\cos(\pi t / T_{max}) + 1) / 2, \quad (3)$$

where λ_0 and $\lambda_{T_{max}}$ are the initial and final values of $\lambda(t)$, T_{max} is the maximum training iteration and t is the current training iteration. Thus, $\lambda(t)$ will start at λ_0 then gradually decay to $\lambda_{T_{max}}$, emphasizing the cosine distance more at the beginning and less as the learning progresses. Discussion on other choices of $\lambda(t)$ are in Sec. 4.2.

Finally, our optimization problem can be rewritten as:

$$\min_{\theta, \phi} \mathbb{E}_{x \sim \mathcal{D}} [(1 - \lambda(t)) \mathcal{L}_{KL}(g_\theta(h_\zeta(x)), l_\phi(k_\phi(x))) + \lambda(t) \mathcal{L}_{FS}(h_\zeta(x), k_\phi(x))]. \quad (4)$$

3.4. Modified Multi-Stage Training for BURN

The multi-stage training [2, 32, 34] is known to be effective in training binary networks. It trains the network with only binarized activations in the first stage. Then, it uses the trained weights of the partially binarized network as initial values for training the fully binarized network, *i.e.*, binarized weights and activations, in the second stage. Unfortunately, we cannot use this strategy as the binary networks converge quickly thanks to the good initial values learned in the first stage [34] whereas the randomly initialized FP classifier g_θ does not converge as quickly as the binary network. This discrepancy in the convergence speeds of the binary network and the FP classifier harms training efficacy.

To apply the multi-stage training to BURN, we modify it to give good initial points to the FP classifier as well as the binary network. Specifically, we initialize g_θ in the second stage with the weights of g_θ obtained in the first stage, similar to the binary network. As a result, g_θ starts from a good initial point and converges quickly to provide useful targets.

We describe the full algorithm of **BURN** in Alg. 1

Algorithm 1 Binary Unsupervised Representation learning (BURN)

```
1: function BURN( $\mathcal{D}, t, \zeta, h_\zeta, g_\theta, k_\phi, l_\phi$ )
2:    $\theta, \phi \leftarrow$  Pretrain( $\mathcal{D}, t, \zeta, h_\zeta, g_\theta, k_\phi, l_\phi$ , STAGE1)
3:    $W \leftarrow \{\zeta\} \cup \{\theta, \phi\}$ 
4:    $\theta, \phi \leftarrow$  Pretrain( $\mathcal{D}, t, W, h_\zeta, g_\theta, k_\phi, l_\phi$ , STAGE2)
5:   return  $k_\phi$ 
6: end function
7: function PRETRAIN( $\mathcal{D}, t, W, h_\zeta, g_\theta, k_\phi, l_\phi$ , F)
8:   if F is STAGE1 then
9:      $k_\phi, l_\phi \leftarrow$  Binarize Activations
10:     $h_\zeta \leftarrow W$  ▷ Load pretrained weights
11:   else
12:      $k_\phi, l_\phi \leftarrow$  Binarize Activations and Weights
13:      $h_\zeta, g_\theta, k_\phi, l_\phi \leftarrow W$  ▷ Load pretrained weights
14:   end if
15:    $x =$  RandomSelect( $\mathcal{D}$ ) ▷ Sample  $x \sim \mathcal{D}$ 
16:    $v_1, v_2 = h_\zeta(x), k_\phi(x)$  ▷ Feature vectors  $v_1, v_2$ 
17:    $p_1, p_2 = g_\theta(v_1), l_\phi(v_2)$  ▷ Softmax Probabilities  $p_1, p_2$ 
18:    $\mathcal{L}_{\zeta, \theta, \phi} =$  AugmentedLoss( $v_1, v_2, p_1, p_2, t$ )
19:    $\theta \leftarrow$  Optimizer( $\nabla_\theta \mathcal{L}_{\zeta, \theta, \phi}, \eta$ ) ▷ Update  $\theta$ 
20:    $\phi \leftarrow$  Optimizer( $\nabla_\phi \mathcal{L}_{\zeta, \theta, \phi}, \eta$ ) ▷ Update  $\phi$ 
21:   return  $\theta, \phi$ 
22: end function
23: function AUGMENTEDLOSS( $v_1, v_2, p_1, p_2, t$ )
24:    $\mathcal{L}_{KL} = \mathcal{D}_{KL}(p_2 || p_1)$  ▷ KL Divergence
25:    $\mathcal{L}_{FS} = 1 - \frac{\langle v_1, v_2 \rangle}{\|v_1\|_2 \cdot \|v_2\|_2}$  ▷ Cosine Distance
26:    $\lambda(t) = \lambda_T - (\lambda_T - \lambda_0) \cdot (\cos(\pi t/T) + 1)/2$  ▷ Eq. 3
27:    $\mathcal{L} = (1 - \lambda(t)) \cdot \mathcal{L}_{KL} + \lambda(t) \cdot \mathcal{L}_{aug}$  ▷ Eq. 4
28:   return  $\mathcal{L}$ 
29: end function
```

4. Experiments

Experimental Details. Following [11, 21, 27, 45, 49, 52], we use ImageNet [26] for pretraining. We use 1) object detection, 2) linear evaluation, 3) semi-supervised fine-tuning, 4) full-shot and few-shot image classification using SVM, and 5) transfer learning via linear evaluation for downstream tasks. We strictly follow the SSL evaluation protocols of the downstream tasks [6, 7, 17, 20]. Downstream task and implementation details are in the supplement. Experiments were partly based on NAVER Smart Machine Learning (NSML) platform [25, 40]. Code is available at <https://github.com/naver-ai/burn>.

Baselines. We pretrain the ReActNet-A backbone with BYOL [18], SWAV [5], tuned MoCov2 [39], and S2-BNN [39] as our SSL baselines. We also show supervised pretraining, *i.e.*, ‘Supervised Pre.’. More comparisons to SimCLRv2 [7] and InfoMin [44] are in the supplement. We pretrain the models for 200 epochs in all methods.

Method	mAP (%)	AP50 (%)	AP75 (%)
Supervised Pre.	38.22	68.53	37.65
SWAV [5]	37.22	67.47	35.91
BYOL [18]	36.92	67.13	35.65
Tuned MoCov2 [39]	37.42	67.30	36.37
S2-BNN [39]	39.50	70.09	39.15
BURN (Ours)	41.00	70.91	41.45

Table 1. Object detection (mAP, AP50 and AP75) on Pascal VOC after pretraining. BURN outperforms all the compared methods including supervised pretraining (‘Supervised Pre.’).

4.1. Results on Downstream Tasks

We evaluate our method along with prior arts in various downstream tasks. We denote the best results except for supervised pretraining in each table in **bold**.

Object Detection. We first conduct object detection (mAP (%), AP50 (%) and AP75 (%)) on Pascal VOC and summarize the results in Table 1. Once the feature extractor is pretrained, we use the pretrained weights as initial weights for fine-tuning a detection pipeline. BURN outperforms all other methods, including supervised pretraining, in all three metrics. We believe one of the reasons for the performance of BURN is that it utilizes a FP network trained in an SSL manner that mostly learned low- and mid-level features [52] which would help object detection.

Linear Evaluation. We then conduct linear evaluation (top-1) on ImageNet and summarize the results in Table 2. Once the binary feature extractor is pretrained, it is frozen and only the attached classifier is trained for classification. As shown in the table, BURN outperforms other SSL methods by up to +13.04% top-1 accuracy, possibly because it utilizes the knowledge from the FP network. Interestingly, BURN even outperforms past supervised ImageNet classification by binary networks *e.g.*, XNOR-Net (51.20%) [38].

Semi-Supervised Fine-Tuning. We now conduct semi-supervised fine-tuning (top-1 and top-5) and summarize the results in Table 2. We fine-tune the entire network on the labeled subset (1% or 10%) from ImageNet. BURN outperforms other SSL baselines by large margins across all metrics; at least +3.67% top-1 accuracy and +5.30% top-5 accuracy on the 1% labels setting and +3.98% top-1 accuracy and +4.41% top-5 accuracy on the 10% labels setting, respectively. Interestingly, BURN seems to outperform other SSL methods by larger amounts in this task than the linear evaluation, implying that BURN may be more beneficial in tasks with limited supervision as also discussed by [16].

SVM Image Classification. We conduct SVM classification (mAP (%)) and summarize results for both the few-shot and full-shot (‘Full’) settings on VOC07 in Table 3. For the few-shot results, the results are averaged over 5 runs. The number of shots k is varied from 1 to 96.

Method	Linear Eval.		Semi-Supervised Fine-tuning		
	Top-1 (%)	1% Labels		10% Labels	
		Top-1 (%)	Top-5 (%)	Top-1 (%)	Top-5 (%)
Supervised Pre.	64.10	42.96	69.10	53.07	77.40
SWAV	49.41	24.66	46.57	33.83	57.81
BYOL	49.25	23.05	43.90	34.66	58.78
Tuned MoCov2	52.50	22.96	45.12	31.18	55.64
S2-BNN	61.50	36.08	61.83	45.98	71.11
BURN (Ours)	62.29	39.75	67.13	49.96	75.52

Table 2. Linear evaluation (top-1) and semi-supervised fine-tuning (1% labels or 10% labels) on ImageNet after pretraining. BURN outperforms all other SSL methods by large margins across for both the linear evaluation and semi-supervised fine-tuning.

Method	k = 1	k = 2	k = 4	k = 8	k = 16	k = 32	k = 64	k = 96	Full
Supervised Pre.	29.28±0.94	36.46±2.97	49.67±1.20	56.99±0.67	64.68±0.89	70.08±0.58	73.49±0.53	74.96±0.17	77.47
SWAV	22.97 ± 1.21	27.91±2.37	37.91±1.11	44.5±1.51	52.79±0.81	59.15±0.62	64.38±0.59	66.72±0.19	71.23
BYOL	23.45 ± 0.76	28.04±2.40	38.09±1.07	44.69±1.66	51.5±0.90	57.44±0.24	62.07±0.28	64.37±0.13	69.16
Tuned MoCov2	22.12 ± 0.74	27.45 ± 2.06	36.81 ± 0.82	43.19 ± 1.4	51.93 ± 0.84	57.95 ± 0.62	63.07 ± 0.43	65.15 ± 0.05	69.73
S2-BNN	27.00±1.54	33.39 ± 2.72	46.31 ± 2.11	54.14 ± 1.32	61.86 ± 1.14	68.01 ± 0.41	71.89 ± 0.44	73.55±0.29	76.49
BURN (Ours)	29.20±1.51	36.14 ± 2.15	48.49 ± 1.08	55.12 ± 1.59	62.36 ± 1.01	68.10 ± 0.3	72.1 ± 0.39	74.06 ± 0.18	77.49

Table 3. SVM classification (mAP) for the few-shot and full-shot settings on VOC07 after pretraining. BURN outperforms all other SSL methods by large margins and performs on par with supervised pretraining on both settings. The number of shots (k) is varied from 1 to 96. We report the averaged performance over 5 runs with the standard deviation.

Method	Object-Centric				Scene-Centric
	CIFAR10	CIFAR100	CUB-200-2011	Birdsnap	Places205
Supervised Pre.	78.30	57.82	54.64	36.90	46.38
SWAV	75.78	56.78	36.11	25.54	46.90
BYOL	76.68	58.18	38.80	27.11	44.62
Tuned MoCov2	78.29	57.56	33.79	23.37	44.90
S2-BNN	82.70	61.90	47.50	34.10	46.58
BURN (Ours)	84.60	61.99	49.62	34.48	47.22

Table 4. Transfer learning (top-1) on either object-centric or scene-centric datasets after pretraining. CIFAR10, CIFAR100, CUB-200-2011, and Birdsnap are used as the object-centric datasets while Places205 is used as the scene-centric dataset. BURN outperforms all other SSL baselines on the object-centric datasets and on Places205.

For the few-shot setting, BURN outperforms all other SSL methods by roughly +1% to +10% mAP depending the number of shots. Noticeably, BURN performs very close to the supervised pretraining regardless of the number of shots. This is consistent with the semi-supervised fine-tuning results; BURN shows strong performance in tasks with limited supervision such as the few-shot classification [16]. In the full-shot setting, BURN outperforms other SSL methods by up to +8.33% mAP and performs very similarly to supervised pretraining. In both settings, representations learned with ImageNet by BURN is still effective on a different dataset such as VOC07, potentially due to BURN using a FP network to obtain targets that are generally useful *i.e.*, low to mid level representations [52].

Transfer Learning. While [39] use VOC07, we use Places205 instead of VOC07 to show more diverse transfer scenarios as we already show transfer learning results

to VOC07 using SVM in Table 3. Given that we use ImageNet (object-centric) for pretraining, to evaluate the transferability of learned representations across domains, we use two types of datasets for the transferability experiments, *i.e.*, object-centric and scene-centric datasets. Specifically, we use CIFAR10, CIFAR100, CUB-200-2011 and Birdsnap as the ‘object-centric datasets’, and use Places205 as the scene-centric dataset. Once we pretrain the binary feature extractor with ImageNet, the feature extractor is frozen and only the attached classifier is trained on the target datasets.

As shown in the table, BURN outperforms all SSL methods on the object-centric datasets with large margins in CIFAR10, CUB-200-2011 and Birdsnap. It implies that the representations learned using BURN transfers well across multiple object-centric datasets. For the scene-centric dataset (Places205), we observed that the transfer learning results for methods vary less. It is quite expected

Method	① Rand. Init.Cls.	② Feat. Sim. Loss	③ Dyn. Bal.	④ Multi-Stage	Top-1 (%)
①	✓	✗	✗	✗	57.33
①+②	✓	✓	✗	✗	59.53
①+②+③	✓	✓	✓	✗	61.43
①+②+③+④ (=BURN)	✓	✓	✓	✓	62.29

Table 5. Ablation studies on the proposed components of BURN using linear evaluation (top-1) on ImageNet. ① refers to using a randomly initialized classifier as targets. ② ‘Feat. Sim.’ refers to feature similarity loss (Eq. 2). ③ ‘Dyn. Bal.’ refers to using the dynamic balancing. ④ refers to using the modified multi-stage training. Each step of improving BURN contribute to a non-trivial performance gain as the evaluation is done with the ImageNet dataset. Also, using only ① already outperforms all other SSL baselines except S2-BNN [39].

since ImageNet is object-centric, thus transferring knowledge to a scene-centric dataset may suffer from domain gap and performance marginally differ across methods.

4.2. Further Analyses

We further investigate our method using linear evaluation (top-1) on ImageNet for detailed analyses.

Training Time. A comparison of the training time vs. linear evaluation on ImageNet is shown in Fig. 4. Without multi-stage (‘MST’), BURN shows much higher accuracy than other baselines with similar or less training time. With MST, BURN shows higher accuracy than S2-BNN but comparable training time.




Figure 4. Training time and linear evaluation accuracy.

Ablation Study. We ablate the model and summarize the results in Table 5. We number each components, following the convention in Fig. 2. Our baseline method (①) alone achieves 57.33% top-accuracy, which is higher than all other SSL baselines except S2-BNN (see Table 2). In addition, every component in BURN contributes to non-trivial gains. In particular, both the feature similarity loss and dynamic balancing provide noticeable improvements. We believe the reason is that the addition of \mathcal{L}_{FS} stabilizes the gradients (see Sec. 3.2) and a dynamic balancing of \mathcal{L}_{FS} captures the changing importance of \mathcal{L}_{FS} stabilizing the gradients even more effectively, resulting in improved performance. The modified multi-stage training also contributes to the accuracy by a non-trivial margin.

Choice of Moving Targets for Binary Networks. Note that the downstream task performance of BYOL [18] and Tuned MoCov2 [39], which use the EMA target and momentum encoder respectively, are worse than BURN (see Sec. 4.1), which uses jointly trained classifier as moving targets. In addition to the above quantitative comparison, we provide intuition for why the EMA target or the momentum encoder may be less effective choices as the moving target




Figure 5. The L2 distance between the *Backbone* and *Target* vectors averaged over 10 independent runs. The shaded areas denote the region of 1 std. The L2 distance is larger and has higher variance in the binary case, indicating that the the EMA target is more different from the backbone *i.e.*, less effective as targets.

for binary networks. The EMA target and the momentum encoder are both based on momentum or EMA updates of the target network and incur differences between the target and the backbone networks. We conjecture that EMA updates may not be well-suited for binary networks because the differences of the target and backbone network from the EMA update could be amplified by the binarization process.

To elaborate our intuition, we design an experiment to compare how different the backbone and EMA target networks are where the backbone and the EMA target are simplified to 100-dimensional vectors denoted as *Backbone* and *Target*, respectively. We then simulate the training of the backbone and the EMA update of the target as:

$$\begin{aligned} \text{Backbone} &\leftarrow \text{Backbone} + \eta \cdot \Delta n, \\ \text{Target} &\leftarrow \tau \cdot \text{Target} + (1 - \tau) \cdot \text{Backbone}, \end{aligned} \quad (5)$$

where *Backbone* and *Target* are initialized with the same random values, η is the learning rate set as 4.8 [18], Δn is a random perturbation drawn from a standard normal distribution, $\mathcal{N}(0, 1)$, to simulate the gradient update, and τ is the momentum used in the EMA update set as 0.99. We perform the update described in Eq. 5 for 100 iterations.

We compare the L2 distance between the *Backbone* and *Target* vectors in Fig. 5. The L2 distance between the *Backbone* and *Target* is much larger with higher variance for the binary case, indicating that the EMA target may differ heavily from the backbone with binary networks. Not

\mathcal{L}_{FS}	Bounded	Top-1 (%)
L_1	✗	51.46
L_2	✗	50.28
Cosine	✓	62.29

Table 6. L_1 , L_2 , and the cosine distances are compared. The cosine distance is by far the best choice amongst the three, as is supported by our intuition that a bounded loss term would be better as the feature similarity loss.




Figure 6. Gradient magnitude of (a) binary classifier and (b) the binary feature extractor during early training for various choices of \mathcal{L}_{FS} such as the cosine, L_1 , and L_2 distances. Both L_1 and L_2 distances show very high gradients at the beginning in the classifier, especially L_2 . Moreover, L_1 and L_2 distances exhibit potential gradient explosions in the feature extractor. Cosine distance shows none of these trends that harm training efficacy.

only that, in the binary case, the L_2 distance does not decrease over time. These results are in-line with our intuition and empirical results that the EMA target could be more different for binary networks and hence be less effective.

Choice of Feature Similarity Loss. We further investigate the choices of the feature similarity loss (\mathcal{L}_{FS} in Eq. 4) in Table 6. Besides the cosine distance, *i.e.*, $1 - \frac{\langle v_1, v_2 \rangle}{\|v_1\|_2 \cdot \|v_2\|_2}$, used in BURN, we compare the L_1 , *i.e.*, $\|v_1 - v_2\|_1$, and L_2 , *i.e.*, $\|v_1 - v_2\|_2$, distances. We believe that as both the L_1 and L_2 distances are not bounded, they may potentially cause problems such as gradient explosion, leading to the worse performance unlike the cosine distance.

Cosine distance outperforms L_1 and L_2 by large margins. In Fig. 6, we illustrate gradient magnitudes of both the classifier and feature extractor when using cosine, L_1 , or L_2 distances as \mathcal{L}_{FS} . L_1 and L_2 distances show very high gradients early on in the classifier, especially L_2 where the gradients start at 1×10^6 . Even more importantly, L_1 and L_2 distances show signs of gradient explosions in the feature extractor, *i.e.*, the gradients keep increasing as the iterations proceed, with L_2 exhibiting more severe trends. In contrast, cosine distance exhibits small and subdued gradients for both the classifier and the feature extractor.

$\lambda(t)$	Top-1 (%)
Constant	55.83
Heaviside Step: $H(-t + T_{max}/2)$	55.60
Eq. 3	62.29

Table 7. Comparison of dynamic balancing functions. Please refer to the supplement for a plot comparing the choices of $\lambda(t)$. The Eq. 3 (smooth annealing) is the best amongst the three choices. The constant function does not capture the dynamic nature of the balancing factor and the Heaviside step function disrupts the training midway due to the discontinuity.

Choice of the Dynamic Balancing Function. We finally investigate the choice of the dynamic balancing function. Specifically, we consider (1) a constant function, $\lambda(t) = 0.7$, (2) a Heaviside step function that is shifted and horizontally reflected, $H(-t + T_{max}/2)$, and (3) a smooth cosine annealing function (Eq. 3) used in our BURN, and compare their respective accuracy in Table 7.

The constant function fails to capture that the importance of \mathcal{L}_{FS} may change as learning progresses, leading to poor results. The Heaviside step function abruptly changes the balancing factor mid-training. This disrupts the training and leads to poor performance. In contrast, smooth annealing (Eq. 3) captures the dynamic nature of the importance of the feature similarity loss while smoothly varying the balancing factor, resulting in the best performance among the options.

5. Conclusion

We propose BURN, an unsupervised representation learning framework for binary networks that jointly trains the FP classifier and the binary network. We propose a feature similarity loss, dynamic loss balancing, and a modified multi-stage training to improve BURN. We conduct extensive empirical validations with five downstream tasks and seven datasets. In all downstream tasks, BURN consistently outperforms existing SSL baselines by large margins and sometimes even the supervised pretraining. We also provide detailed analysis on various aspects of BURN.

Limitations. One limitation of BURN is that it requires a pretrained FP feature extractor. While using pretrained FP networks is a commonly used approach for binary networks [2, 32, 34, 39], it would be more efficient to design a framework that does not require a pretrained FP model.

Potential Negative Societal Impact. We aim to improve the representation power learned by binary networks, which would facilitate AI on edge devices with vision sensors. Consequently, AI surveillance systems could become prevalent and may lead to the monitoring of mass populations for private information and personal attributes. Although the authors have *no intention* to allow such consequences, such negative effects could manifest. Efforts to prevent such effects would include gating the code and pretrained models behind security and ethical screenings.

Acknowledgement. The authors thank Jung-Woo Ha for the valuable discussions. This work was partly supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.2022R1A2C4002300) and Institute for Information & communications Technology Promotion (IITP) grants funded by the Korea government (MSIT) (No.2020-0-01361-003 and 2019-0-01842, Artificial Intelligence Graduate School Program (Yonsei University, GIST), and No.2021-0-02068 Artificial Intelligence Innovation Hub).

References

- [1] Soroush Abbasi Koohpayegani, Ajinkya Tejankar, and Hamed Pirsiavash. Compress: Self-supervised learning by compressing representations. In *NeurIPS*, 2020. 1, 2
- [2] Adrian Bulat, Brais Martínez, and Georgios Tzimiropoulos. Bats: Binary architecture search. In *ECCV*, 2020. 1, 2, 4, 8
- [3] Adrian Bulat, Brais Martinez, and Georgios Tzimiropoulos. High-capacity expert binary networks. In *ICLR*, 2021. 1, 2
- [4] Zhaowei Cai, Avinash Ravichandran, Subhransu Maji, Charles Fowlkes, Zhuowen Tu, and Stefano Soatto. Exponential moving average normalization for self-supervised and semi-supervised learning. In *CVPR*, pages 194–203, June 2021. 1
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 1, 2, 5
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2, 3, 4, 5
- [7] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. In *NeurIPS*, 2020. 1, 2, 3, 5
- [8] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning, 2020. 2
- [9] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, pages 15750–15758, 2021. 1, 2, 3
- [10] Xiangyi Chen, Steven Z Wu, and Mingyi Hong. Understanding gradient clipping in private sgd: a geometric perspective. *NeurIPS*, 33, 2020. 4
- [11] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. Debiasing contrastive learning. In *NeurIPS*, 2020. 5
- [12] Devin Coldewey. Xnor’s saltine-sized, solar-powered ai hardware redefines the edge. <https://techcrunch.com/2019/02/13/xnors-saltine-sized-solar-powered-ai-hardware-redefines-the-edge/>. 1
- [13] Linus Ericsson, Henry Gouk, and Timothy M. Hospedales. How well do self-supervised models transfer? In *CVPR*, pages 5414–5423, June 2021. 1
- [14] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sanginetto, and Nicu Sebe. Whitening for self-supervised representation learning. In *ICML*, pages 3015–3024. ICML, 2021. 1
- [15] Zhiyuan Fang, Jianfeng Wang, Lijuan Wang, Lei Zhang, Yezhou Yang, and Zicheng Liu. {SEED}: Self-supervised distillation for visual representation. In *ICLR*, 2021. 2
- [16] Priya Goyal, Mathilde Caron, Benjamin Lefaudeaux, Min Xu, Pengchao Wang, Vivek Pai, Mannat Singh, Vitaliy Liptchinsky, Ishan Misra, Armand Joulin, et al. Self-supervised pretraining of visual features in the wild. *arXiv preprint arXiv:2103.01988*, 2021. 1, 2, 3, 5, 6
- [17] Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. In *ICCV*, 2019. 5
- [18] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning. In *NeurIPS*, 2020. 1, 2, 3, 4, 5, 7
- [19] Kai Han, Yunhe Wang, Yixing Xu, Chunjing Xu, Enhua Wu, and C. Xu. Training binary neural networks through learning with noisy supervision. In *ICML*, 2020. 2
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2, 3, 4, 5
- [21] Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. In *NeurIPS*, 2020. 5
- [22] Dahyun Kim and Jonghyun Choi. Bnas v2: Learning architectures for binary networks with empirical improvements. *arXiv preprint https://arxiv.org/abs/2110.08562*, 2021. 2
- [23] Dahyun Kim, Kunal Pratap Singh, and Jonghyun Choi. Learning architectures for binary networks. In *ECCV*, 2020. 1, 2
- [24] Hyungjun Kim, Kyungsu Kim, Jinseok Kim, and Jae-Joon Kim. Binaryduo: Reducing gradient mismatch in binary activation network by coupling binary activations. In *ICLR*, 2020. 2
- [25] Hanjoo Kim, Minkyu Kim, Dongjoo Seo, Jinwoong Kim, Heungseok Park, Soeun Park, Hyunwoo Jo, KyungHyun Kim, Youngil Yang, Youngkwan Kim, et al. Nsm1: Meet the mlaas platform with a real-world case study. *arXiv preprint arXiv:1810.09957*, 2018. 5
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012. 5
- [27] Kibok Lee, Yian Zhu, Kihyuk Sohn, Chun-Liang Li, Jinwoo Shin, and Honglak Lee. i-mix: A strategy for regularizing contrastive representation learning. In *ICLR*, 2021. 5
- [28] Bin Li, Yin Li, and Kevin W. Eliceiri. Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In *CVPR*, pages 14318–14328, June 2021. 1
- [29] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *ICLR*, 2021. 2
- [30] Mingbao Lin, Rongrong Ji, Zihan Xu, Baochang Zhang, Yan Wang, Yongjian Wu, Feiyue Huang, and Chia-Wen Lin. Rotated binary neural network. In *NeurIPS*, 2020. 2
- [31] Xiaofan Lin, Cong Zhao, and Wei Pan. Towards accurate binary convolutional neural network. In *NeurIPS*, 2017. 1, 2
- [32] Zechun Liu, Zhiqiang Shen, Marios Savvides, and Kwang-Ting Cheng. Reactnet: Towards precise binary neural network with generalized activation functions. In *ECCV*, 2020. 1, 2, 3, 4, 8

- [33] Zechun Liu, Baoyuan Wu, Wenhan Luo, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *ECCV*, 2018. 1, 2
- [34] Brais Martinez, Jing Yang, Adrian Bulat, and Georgios Tzimiropoulos. Training binary neural networks with real-to-binary convolutions. In *ICLR*, 2020. 1, 2, 3, 4, 8
- [35] Xiangming Meng, Roman Bachmann, and Mohammad Emtiyaz Khan. Training binary neural networks using the bayesian learning rule. In *ICML*, 2020. 2
- [36] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2
- [37] Haotong Qin, Ruihao Gong, Xianglong Liu, Mingzhu Shen, Ziran Wei, Fengwei Yu, and Jingkuan Song. Forward and backward information retention for accurate binary neural networks. In *CVPR*, 2020. 2
- [38] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *ECCV*, 2016. 1, 2, 5
- [39] Zhiqiang Shen, Zechun Liu, Jie Qin, Lei Huang, Kwang-Ting Cheng, and Marios Savvides. S2-bnn: Bridging the gap between self-supervised real and 1-bit neural networks via guided distribution calibration. In *CVPR*, pages 2165–2174, 2021. 1, 2, 3, 5, 6, 7, 8
- [40] Nako Sung, Minkyu Kim, Hyunwoo Jo, Youngil Yang, Jingwoong Kim, Leonard Lausen, Youngkwan Kim, Gayoung Lee, Donghyun Kwak, Jung-Woo Ha, et al. Nsm1: A machine learning platform that enables you to focus on your models. *arXiv preprint arXiv:1712.05902*, 2017. 5
- [41] Yuandong Tian, Xinlei Chen, and Surya Ganguli. Understanding self-supervised learning dynamics without contrastive pairs. *ICML*, 2021. 1
- [42] Yonglong Tian, Olivier J Henaff, and Aaron van den Oord. Divide and contrast: Self-supervised learning from uncurated data. *arXiv preprint arXiv:2105.08054*, 2021. 1, 2, 3
- [43] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, 2020. 2, 3
- [44] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? In *NeurIPS*, 2020. 1, 2, 3, 5
- [45] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *ICML*, 2020. 5
- [46] Ziwei Wang, Ziyi Wu, Jiwen Lu, and Jie Zhou. Bidet: An efficient binarized object detector. In *CVPR*, 2020. 2
- [47] Chen Wei, Huiyu Wang, Wei Shen, and Alan Yuille. {CO}2: Consistent contrast for unsupervised visual representation learning. In *ICLR*, 2021. 2
- [48] Zhirong Wu, Yuanjun Xiong, X Yu Stella, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 2
- [49] Tete Xiao, Xiaolong Wang, Alexei A Efros, and Trevor Darrell. What should not be contrastive in contrastive learning. In *ICLR*, 2021. 4, 5
- [50] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, 2021. 1, 2, 3
- [51] Jingzhao Zhang, Tianxing He, Suvrit Sra, and Ali Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. *arXiv preprint arXiv:1905.11881*, 2019. 4
- [52] Nanxuan Zhao, Zhirong Wu, Rynson W. H. Lau, and Stephen Lin. What makes instance discrimination good for transfer learning? In *ICLR*, 2021. 5, 6