

Switchable Representation Learning Framework with Self-compatibility

Shengsen Wu^{1,2}, Yan Bai², Yihang Lou³, Xiongkun Linghu⁴, Jianzhong He³, Ling-Yu Duan^{1,2,*}
¹ Peng Cheng Laboratory, China; ² Peking University, China;
³ Huawei Technologies; ⁴ Tsinghua University, China

{sswu, yanbai, lingyu}@pku.edu.cn,
 {louyihang1, jianzhong.he}@huawei.com, lhxk20@mails.tsinghua.edu.cn

Abstract

Real-world visual search systems involve deployments on multiple platforms with different computing and storage resources. Deploying a unified model that suits the minimal-constrain platforms leads to limited accuracy. It is expected to deploy models with different capacities adapting to the resource constraints, which requires features extracted by these models to be aligned in the metric space. The method to achieve feature alignments is called “compatible learning”. Existing research mainly focuses on the one-to-one compatible paradigm, which is limited in learning compatibility among multiple models. We propose a Switchable representation learning Framework with Self-Compatibility (SFSC). SFSC generates a series of compatible sub-models with different capacities through one training process. The optimization of sub-models faces gradients conflict, and we mitigate this problem from the perspective of the magnitude and direction. We adjust the priorities of sub-models dynamically through uncertainty estimation to co-optimize sub-models properly. Besides, the gradients with conflicting directions are projected to avoid mutual interference. SFSC achieves state-of-the-art performance on the evaluated datasets.

1. Introduction

Visual search systems are widely deployed, which recall the nearest neighbors in gallery features according to their distances to the query feature. Real-world visual search systems consist of multiple models deployed on different platforms [17, 38, 45] (e.g. clouds, mobiles, smart cameras), where different platforms interact with each other by visual features. Typically, as for person re-identification systems, images are captured and processed into features on edge sides. And such features are sent to the cloud side to compare with the database features for identifying

a specific person by feature similarities. As diverse platforms meet different computing and storage resource limitations, deploying the unified model which suits the minimal-constrain platforms leads to a waste of resources on other platforms and limited accuracy. To make better use of resources and achieve higher accuracy, it is expected to deploy models with different capacities which adapt to the resource limitations. Such a solution requires compatibility among the models to satisfy their interaction requirements, which means that the similar data processed by different models are close to each other in the feature space while dissimilar data are far apart.

To achieve compatibility, compatible learning methods have been proposed. Existing research focuses on the one-to-one compatible paradigm [6, 49, 62], which constrains the learned features of a latter (learnable) model to be compatible with its previous (fixed) version. They are limited in achieving many-to-many compatibilities, which means any two in a series of models are compatible with each other.

In this work, we propose a Switchable representation learning Framework with Self-Compatibility (SFSC) to achieve many-to-many compatibilities. As shown in Figure 1, SFSC can deploy different sub-models to suit the diverse computing and storage resource limitation of different platforms. The compatibility between any two sub-models can be achieved, termed as “self-compatibility”. However, there are gradients conflicts between sub-models as they are co-optimized during the training process. We summarize such conflicts into the gradient magnitude and gradient direction. Specifically, a gradient of a large magnitude will dominate the optimization process, which may lead to the overfitting of its corresponding sub-model and impair the improvements of other sub-models. To solve this problem, we estimate the uncertainty of different sub-models to weight the gradient, denoting the optimization priorities. Besides, as the sub-models may produce gradients in various directions, there is mutual interference between them. Thus the improvement in different sub-models may be overestimated

*Corresponding author.

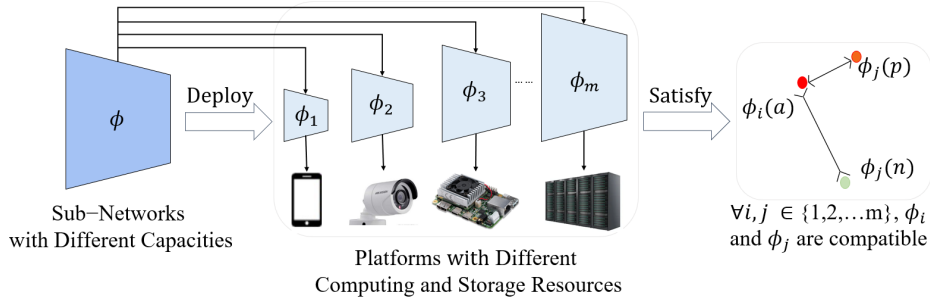


Figure 1. An example of SFSC in multi-platform deployments. The switchable neural network generates a series of sub-networks $\{\phi_1, \phi_2, \dots, \phi_m\}$ with different capacities for different computing resources. $\forall i, j \in \{1, 2, \dots, m\}, |\phi_i(a) - \phi_j(p)| < |\phi_i(a) - \phi_j(n)|$, where (a, p) is a pair of samples with the same class ID, while (a, n) is a pair of samples with different class IDs.

or underestimated. To tackle such conflict, the gradients are projected onto planes that are orthogonal to each other.

The contributions of this paper are summarized as follows:

- We propose a switchable representation learning framework with self-compatibility (SFSC). SFSC generates a series of feature-compatible sub-models with different capacities for visual search systems, which can be deployed on different platforms flexibly.
- We resolve the conflicts between sub-models from the aspect of gradient magnitude and gradient direction. A compatible loss based on uncertainty estimation is proposed to guide optimization priorities and alleviate the imbalance of gradient magnitude between sub-models. An aggregation method based on the gradient projection is proposed to avoid mutual interference and find a generic optimal direction for all sub-models.
- SFSC achieves state-of-the-art performance on the evaluated benchmark datasets. Compared to deploying a unified model, adopting SFSC to obtain different sub-models can achieve 6% to 8% performance improvements on three different datasets.

2. Related Work

2.1. One-to-one Compatible Learning Paradigm

Given a fixed old model, the one-to-one compatible learning methods aim at aligning the learnable models with the fixed old one in metric space, termed as “backward compatible learning”. Existing methods focus on the optimization difficulties caused by the old model. Yan *et al.* [63] propose Focal Distillation to reduce the “negative flips” which means a test sample is correctly classified by the old model but incorrectly classified by the new one. Wu *et al.* [62] start from neighborhood structures and classification entropy to eliminate the effect of outliers in old features. Shen *et al.* [48] propose a pseudo classifier and enhance it with a

random walk algorithm to get rid of the old classifier. Different from these one-to-one compatible learning methods, this paper studies a new compatible learning paradigm that aims to learn many-to-many compatibility among multiple learnable models.

2.2. Dynamic Neural Networks

Dynamic neural networks are able to adapt their structures or parameters to the input during inference [19]. In recent years, research in this field has mainly focused on dynamic depth [5, 18, 26, 54, 55, 58], dynamic width [3, 4, 11, 27, 29, 36], dynamic routing [22, 25, 28, 64], dynamic parameters [12, 15, 21, 51, 69], etc. In addition to hand-crafted methods, Neural Architecture Search (NAS) method is also adopted to train and optimize multiple sub-nets with various capacities in one-time training [7, 8, 44, 53]. Dynamic neural networks can allocate computations on demand for different platforms by selectively activating model components. They have been used in classification [41, 42, 59], segmentation [31, 34, 43], and detection tasks [12, 56, 69]. In general, current dynamic neural networks output with definite semantic information (*e.g.* class id, detection box) which are naturally interoperable. However, for retrieval tasks, the networks output with features, which requires feature compatibility.

2.3. Gradient-based Multi-task Optimization

From the gradient perspective, multi-task optimization is to aggregate gradients corresponding to different tasks and find a descent direction that decreases all objectives. Existing methods treat it as a nonlinear programming problem and take the multi-objective Karush-Kuhn-Tucker [33] (KKT) conditions to perform gradient projection and rescaling under different assumptions on tasks, *e.g.* multiple tasks which are equally important [10, 46, 66], main tasks with some auxiliary tasks [13, 57], a series of the tasks which are learned sequentially [14, 35, 37]. This paper treats optimizations of multiple sub-models as multiple tasks. As

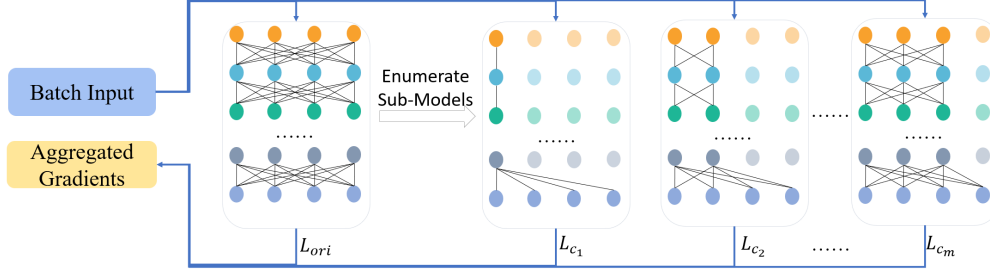


Figure 2. Overview of SFSC. We first transform the full model ϕ into the sub-models with various capacities. Then we calculate the compatible learning loss $\{L_{c_1}, L_{c_2}, \dots, L_{c_m}\}$ on different sub-models with various capacities. Finally, we optimize the sub-models by aggregating losses.

the sub-models vary greatly in their capacity, the optimization of different models converges at different speeds. The challenge is to dynamically adjust the priorities of different sub-models to avoid the overfitting of someones as well as underfitting of others.

2.4. Evidential Deep Learning

Evidential deep learning [47] is used to model the uncertainty of deterministic neural networks. Compared with the Bayesian neural networks [32, 40], evidential deep learning methods have the advantage of low computational complexity, as there is no need for complex posterior probability estimates and multiple Mento Carlo samplings. Based on the DST [16] (Dempster-Shafer Theory of Evidence), evidential deep learning treats the output of the neural networks as subjective opinions and formalizes them as Dirichlet distribution. It has been used in classification tasks to overcome the overconfidence problem of softmax cross entropy [30], action recognition task to identify unknown samples [2], and multi-view learning task [20] to dynamically integrate different views. This paper treats the output features of different sub-models as evidence to access uncertainty and relates the priority of different sub-models with the obtained uncertainty.

3. Methods

In this section, we will describe the proposed SFSC method. We first introduce the construction of switchable neural networks. Then, we analyze the difficulties for optimizations and propose our compatible loss as well as the gradient aggregation method.

3.1. Overview

SFSC involves the training process of multiple sub-models whose features are compatible. Sub-models with different capacities are transformed from the full model ϕ according to the crop ratio list W . We termed the crop ratio list of each sub-model as $W = \{\gamma_1, \gamma_2, \dots, \gamma_m\}$, where γ_j

denotes the ratio of the width of the sub-model ϕ_j to the full model ϕ .

As shown in Figure 2, for each batch input, SFSC first calculates the original loss L_{ori} . The compatible loss $\{L_{c_1}, L_{c_2}, \dots, L_{c_m}\}$ is calculated on different sub-models. By aggregating losses $\{L_{ori}, L_{c_1}, L_{c_2}, \dots, L_{c_m}\}$, SFSC computes the general gradient and updates the full model ϕ by gradient descent methods.

3.2. Switchable Neural Network for Sub-models Construction

We first convert the traditional convolutional neural network into its corresponding switchable neural network. As the platforms generally perform model inference layer by layer, the resource limitation is more related to the widths of layers than the number of layers. Here we perform channel pruning on ϕ with crop ratio γ_j to obtain sub-model ϕ_j . Specifically, the classic convolutional neural network is mainly composed of some basic modules, such as convolution, fully connected layer, and batch normalization. We modify these basic modules as follows:

1) For convolution (Conv) and fully connected (FC) layers, if their output is not the final feature for retrieval, we perform pruning in both input and output channels. Otherwise, pruning is only performed in input channels.

2) For batch normalization (BN), we assign an independent BN to the sub-model ϕ_j . As the intermediate features extracted by different sub-models vary greatly in distributions [65], channel pruning can not be performed on BN. Fortunately, the parameters of BN only occupy a small part of all parameters of the entire network (shown in Table 1). Such an operation only results in a slight increase in storage resources in training progress but does not incur additional deployment costs.

Through such design, we obtain the sub-model ϕ_j with the following characteristics: 1) The features extracted by different sub-models have the same dimensions. 2) The ratio of layer widths in the sub-model ϕ_j to that in the full model ϕ is around γ_j . 3) The sub-model ϕ_j reuses the pa-

NetWorks	ResNet-18 [23]	ResNet-50 [23]	MobileNet V2 [45]	ShuffleNet [67]
Conv and FC	11,679,912 99.918%	25,503,912 99.792%	3,470,870 99.027%	5,402,826 99.102%
BN	9,600 0.0820%	53,120 0.208%	34,112 0.973%	48,960 0.898%

Table 1. Number and proportion of different modules’ parameters in different network structures

rameters from the full model ϕ except for BN.

3.3. Aggregation Method based on the Gradients Projection

As the gradients corresponding to different sub-models may have different directions, simply aggregating the gradients by summation may cause mutual interference, which means the improvements of sub-models are overestimated or underestimated, as shown in Figure 3.

Inspired by multi-task learning [1, 46, 66], we tackle such a problem by projecting conflicting gradients and keeping the components perpendicular to each other. Given a pair of gradients (g_i, g_j) corresponding to sub-model ϕ_i and ϕ_j , if the cosine of their direction $\cos(g_i, g_j) < 0$, they are conflict with each other. For a pair of conflicting gradients (g_i, g_j) , we project them to the orthogonal plane of each other, respectively, as shown in Equation 1.

$$g'_j = g_j - \frac{g_j * g_i}{\|g_i\|^2} * g_i. \quad (1)$$

Such projection results in a direction that improves both the i -th and j -th sub-model simultaneously.




Figure 3. An example of gradient conflict between pair (g_i, g_j) . Simply taking the summation aggregation results in the elimination of the improvement of the i -th sub-model. Our aggregation module ensures the generic direction to improve both the i -th and j -th sub-model.

3.4. Compatible Loss based on the Uncertainty Estimation

The aggregation method aims to optimize all sub-models coordinately. However, it may not work as expected when facing the imbalance of gradient magnitude between sub-models, as shown in Figure 4. Therefore, it is critical to design a compatible loss.

As sub-models are co-optimized, constraining them with existing compatible loss [6, 49] faces conflicts in the magnitude of gradients. As the sub-models vary in capacities, their optimization process converges at different speeds. The small-capacity sub-models will converge faster. Therefore, the small-capacity sub-model may have reached its




Figure 4. An example of gradient conflict between pair (g_i, g_j) when there is an imbalance of gradient magnitude. The i -th sub-model contributes little to the optimization even using gradient projection.

performance ceiling while still producing losses of large magnitude, resulting in the overfitting of such sub-model and weakening of other sub-models. Intuitively, a sub-model with low uncertainty tends to reach its performance ceiling. Therefore, it is expected to adaptively adjust the optimization priorities of each sub-models through their uncertainty dynamically during the training process.

Here we propose a compatible loss based on the uncertainty estimation. Inspired by Dempster-Shafer Theory of evidence [16] (DST), we measure the uncertainty through belief, where high uncertainty corresponds to the model with small capacities. Typically, the belief is measured by a set of mass $B = \{b_1, b_2, \dots, b_C\}$ and an overall uncertainty mass u , where b_c denotes the belief degree belonging to the class c . B and u are related to a Dirichlet distribution, which satisfy Equation 2.

$$u + \sum_{c=1}^C b_c = 1, u \geq 0, b_c \geq 0. \quad (2)$$

The belief is estimated by the features extracted by sub-models. Given sample i , its features extracted by the j -th sub-model termed $\phi_j(i)$. The evidence $E_i^j = \{e_{i1}^j, e_{i2}^j, \dots, e_{iC}^j\}$ is transformed by $\phi_j(i)$ through linear projection and activation function (*i.e.*, $\exp(\cdot)$), where the linear projection layers share parameters with the classifier of the $1.0\times$ model ϕ . The belief of sample i on the j -th sub-model is computed by Equation 3, where $\alpha_i^j = \{\alpha_{i1}^j, \alpha_{i2}^j, \dots, \alpha_{iC}^j\}$ is the Dirichlet intensity on all classes.

$$\alpha_{ic}^j = e_{ic}^j + 1, \quad S_i^j = \sum_{c=1}^C \alpha_{ic}^j, b_{ic}^j = \frac{\alpha_{ic}^j - 1}{S_i^j}, w_i^j = \frac{C}{S_i^j}. \quad (3)$$

With α_i^j , the Dirichlet distribution $D(p_i | \alpha_i^j)$ is obtained, where p_i is the class assignment probabilities on a simplex. Inspired by the Cross-Entropy loss, we integrate the probability p_i on $D(p_i | \alpha_i^j)$, as shown in Equation 4.

$$\begin{aligned} \mathcal{L}(\alpha_i^j) &= \int \left[\sum_{c=1}^C -y_{ic}^j \log(p_{ic}^j) \right] \frac{1}{\beta(\alpha_i^j)} \prod_{c=1}^C p_{ic}^{(\alpha_{ic}^j - 1)} dp_i^j \\ &= \sum_{c=1}^C y_{ic}^j (\psi(S_i^j) - \psi(\alpha_{ic}^j)), \end{aligned} \quad (4)$$

where $\psi(\cdot)$ is the *digamma* function. Since Equation 4 may converge to a trivial solution $D(p_i|1)$ during optimization, we penalize the distance between $D(p_i|1)$ and $D(p_i|\alpha_i^j)$ as the regularization. The complete compatible loss L_{c_j} is shown in Equation 5. L_{c_j} adaptively adjusts the magnitude of gradients through their corresponding uncertainty.

$$KL [D(p_i|\alpha_i^j)||D(p_i|1)] = \log \left(\frac{\sum_{c=1}^C \Gamma(\alpha_{ic}^j)}{\Gamma(C) \prod_{c=1}^C (\alpha_{ic}^j)} \right) + \sum_{c=1}^C (\alpha_{ic}^j - 1) \left[\psi(\alpha_{ic}^j) - \psi \left(\sum_{c=1}^C \alpha_{ic}^j \right) \right], \quad (5)$$

$$L_{c_j}(i) = \mathcal{L}(\alpha_i^j) + \lambda * KL [D(p_i|\alpha_i^j)||D(p_i|1)].$$

We further conduct a gradient analysis for the proposed compatible loss. The Cross-Entropy loss is written as:

$$p_{ic} = \frac{e^{z_{ic}}}{\sum_{k=1}^C e^{z_{ik}}}, L_{ce} = \sum_{c=1}^C -y_{ic} \log(p_{ic}), \quad (6)$$

where z_i is the classification vector of sample i outputted by the model. The corresponding gradient of Equation 6 can be computed as:

$$\frac{\partial L_{ce}}{\partial z_{ic}} = \frac{e^{z_{ic}}}{\sum_{k=1}^C e^{z_{ik}}} - y_{ic}. \quad (7)$$

The corresponding gradient of the proposed compatible loss in 5 can be computed as:

$$\frac{\partial L_{c_j}(\alpha_i^j)}{\partial \alpha_{ic}^j} = \psi'(\alpha_{ic}^j) [\lambda(\alpha_{ic}^j - 1) - y_{ic}] + \psi'(S_i^j) [-\lambda(S_i^j - C) + 1]. \quad (8)$$

We take $\frac{1}{\alpha_{ic}^j}$ as the approximation of $\psi'(\alpha_{ic}^j)$ and obtain Equation 9.

$$\begin{aligned} \frac{\partial L_{c_j}(\alpha_i^j)}{\partial \alpha_{ic}^j} &\approx \frac{1}{\alpha_{ic}^j} [\lambda(\alpha_{ic}^j - 1) - y_{ic}] \\ &+ \frac{1}{S_i^j} [-\lambda(S_i^j - C) + 1] \\ &= -\left(\frac{\lambda + y_{ic}}{\alpha_{ic}^j} + \frac{\lambda C + 1}{S_i^j} \right), \end{aligned} \quad (9)$$

where the first term corresponds to the Cross-Entropy loss in Equation 7, and the second term corresponds to the uncertainty estimation. As the uncertainty $u_i^j = \frac{C}{S_i^j} \propto \frac{1}{S_i^j}$, thus higher u_i^j would result in a gradient with larger magnitude through $\frac{\lambda C + 1}{S_i^j}$. Therefore, the magnitude of gradients corresponding to different sub-models is dynamically adjusted to avoid someone dominating the optimization process.

Such compatible loss can work well with gradient projection, as shown in Algorithm 1.

Algorithm 1 Training Process of SFSC

Require: batch input \mathcal{B} , the crop list $W = \{\gamma_0, \gamma_1, \dots, \gamma_m\}$, the full model ϕ

- 1: $g_{ori} \leftarrow \frac{\delta L_{ori}(\phi, \mathcal{B})}{\delta \phi}$
- 2: **for** $\gamma_j \in W$ **do**
- 3: Switch to the sub-model ϕ_j
- 4: $g_j \leftarrow \frac{\delta L_{c_j}(\phi_j, \mathcal{B})}{\delta \phi}$
- 5: **end for**
- 6: $G, G^p \leftarrow \{g_{ori}, g_0, \dots, g_m\}$
- 7: **for** $g_a \in G^p$ **do**
- 8: Shuffle(G)
- 9: **for** $g_b \in G$ **do**
- 10: $d = g_a * g_b$
- 11: **if** $d < 0$ **then** //have conflicts
- 12: $g_a \leftarrow g_a - \frac{g_b * d}{\|g_b\|^2}$
- 13: **end if**
- 14: **end for**
- 15: **end for**
- 16: **return** Update ϕ by $\Delta\theta = -\sum_{g_a \in G^p} g_a$

In general, SFSC aims to achieve the Pareto optimality [9] among the optimization of sub-models. The proposed compatible loss adaptively adjusts the magnitude of gradients through uncertainty to control the optimization priorities of sub-models dynamically. Besides, the aggregation method performs the gradient projection to find a generic direction to improve all sub-models.

4. Experiments

4.1. Experimental Configuration

Datasets. We evaluate the proposed SFSC on person ReID datasets, *i.e.*, Market-1501 [68], MSMT17 [61], and one vehicle ReID dataset VeRi-776 [60]. Market1501 consists of 32,668 annotated images of 1,501 identities shot from 6 cameras. MSMT17 consists of 126,441 bounding boxes of 4,101 identities taken by 15 cameras. VeRi-776 contains over 50,000 images of 776 vehicles captured by 20 cameras.

Comparison methods. For the one-to-one compatible learning paradigm, BCT [49] and Asymmetric [6] are adopted, as they are representatives of compatible learning through classification losses and embedding losses. BCT and Asymmetric can only obtain sub-models in a one-to-one compatible way. Specifically, given the fixed largest capacity model which is trained without compatible regularization, we train a series of small-capacity models to be compatible with it, respectively. Besides, we also incorporate BCT/Asymmetric loss in our proposed self-compatible learning paradigm for a fair comparison. We take the compatible loss in BCT/Asymmetric to constrain each sub-model compatible with the full model, and use summation to aggregate losses from sub-models, termed as ‘‘BCT-S’’ and ‘‘Asymmetric-S’’.

Evaluation metrics. Mean average precision (mAP)

$M(\phi_q, \phi_g) \backslash \phi_g$	$\phi_{\frac{1}{16}}$	$\phi_{\frac{1}{4}}$	$\phi_{\frac{9}{16}}$	ϕ	$\phi_{\frac{1}{16}}$	$\phi_{\frac{1}{4}}$	$\phi_{\frac{9}{16}}$	ϕ
ϕ_q	Unified Model				Ours			
$\phi_{\frac{1}{16}}$	55.25	-	-	-	58.19	62.88	63.48	69.43
$\phi_{\frac{1}{4}}$	-	67.48	-	-	61.24	70.74	71.37	76.37
$\phi_{\frac{9}{16}}$	-	-	71.25	-	62.43	71.25	72.06	77.26
ϕ	-	-	-	80.91	68.03	74.76	75.67	81.43
BCT-S				Asymmetric-S				
$\phi_{\frac{1}{16}}$	54.40	58.10	60.22	61.29	48.79	51.96	54.09	56.31
$\phi_{\frac{1}{4}}$	55.71	66.38	68.29	69.71	52.24	62.45	64.98	67.23
$\phi_{\frac{9}{16}}$	56.50	68.30	68.42	71.99	54.44	64.74	67.97	71.28
ϕ	61.79	69.73	71.93	73.61	56.54	67.30	71.91	78.38
BCT				Asymmetric				
$\phi_{\frac{1}{16}}$	55.55	56.64	59.40	60.35	55.83	50.85	53.71	55.39
$\phi_{\frac{1}{4}}$	55.83	65.74	67.10	68.54	52.66	67.04	62.08	66.59
$\phi_{\frac{9}{16}}$	55.66	68.02	67.49	70.93	54.48	62.33	66.33	70.14
ϕ	58.69	69.48	70.04	-	56.08	66.66	69.71	-

Table 2. Baseline performance comparison on Market1501 (mAP). “Unified Model” are models without any compatible regularization.

	$M(\phi_{\frac{1}{16}}, \phi)$		$M(\phi_{\frac{1}{4}}, \phi)$		$M(\phi_{\frac{9}{16}}, \phi)$		$M(\phi, \phi)$	
	R1	mAP	R1	mAP	R1	mAP	R1	mAP
MSMT17								
Unified Model	28.82	11.69	48.65	22.86	57.56	30.06	70.22	43.06
BCT-S	46.49	20.11	56.22	28.13	58.11	30.14	58.76	30.99
Asymmetric-S	32.64	13.34	49.62	24.04	57.18	30.04	66.08	40.39
Ours	56.22	28.16	63.29	35.32	65.13	37.74	70.38	43.89
VERI-776								
Unified Model	81.07	44.79	85.30	53.63	88.69	55.42	92.20	66.50
BCT-S	86.25	49.38	87.08	57.03	90.42	58.20	86.73	58.46
Asymmetric-S	68.75	42.16	84.29	56.55	88.04	61.37	89.82	65.35
Ours	79.76	55.04	89.94	62.28	90.48	62.72	92.32	66.55

Table 3. Performance comparison on MSMT17 and VeRi-776 datasets.

	$M(\phi_{\frac{1}{16}}, \phi)$		$M(\phi_{\frac{1}{4}}, \phi)$		$M(\phi_{\frac{9}{16}}, \phi)$		$M(\phi, \phi)$	
	R1	mAP	R1	mAP	R1	mAP	R1	mAP
Circle Loss								
Unified Model	70.25	50.79	72.66	60.43	78.21	64.66	87.38	74.06
BCT-S	63.18	42.45	72.74	52.04	77.38	59.49	87.53	74.58
Asymmetric-S	69.66	48.21	77.85	57.02	80.52	61.81	86.90	73.08
Ours	73.99	54.61	80.58	63.46	82.39	67.02	88.95	77.05
Softmax + Triplet								
Unified Model	77.55	55.82	83.97	65.30	87.89	70.69	90.94	78.30
BCT-S	81.62	59.58	87.86	69.05	88.75	71.38	90.11	74.71
Asymmetric-S	75.56	52.60	85.36	64.21	86.91	68.83	90.45	77.34
Ours	87.36	66.21	90.94	76.64	91.21	77.93	92.64	81.46

Table 4. Performance comparison between different loss functions on Market1501.

and top-1 accuracy (R1) are adopted to evaluate the performances of retrieval. Define $M(\phi_a, \phi_b)$ as the performances of retrieval where the query features are extracted by model ϕ_a and gallery features are extracted by model ϕ_b . We set ϕ_a and ϕ_b as different sub-models to evaluate the perfor-

mance of compatible learning.

Implementation details. The implementations of all compared methods and SFSC are based on the FastReID [24]. The default configuration named “bag of tricks” [39] is adopted. λ is set to 0.2 for all experiments.

	$M(\phi_{\frac{1}{16}}, \phi)$		$M(\phi_{\frac{1}{4}}, \phi)$		$M(\phi_{\frac{9}{16}}, \phi)$		$M(\phi, \phi)$	
	R1	mAP	R1	mAP	R1	mAP	R1	mAP
ResNet-50								
Unified Model	85.27	65.44	88.84	72.07	90.8	76.40	92.22	82.13
BCT-S	89.31	72.09	91.00	77.22	91.36	78.27	91.86	79.40
Asymmetric-S	85.04	67.39	90.26	76.28	91.60	78.53	91.95	82.79
Ours	90.41	76.40	91.03	78.20	91.66	78.76	93.44	83.66
MobileNet V2								
Unified Model	69.00	41.14	89.58	72.85	91.54	77.99	92.28	79.50
BCT-S	80.46	55.17	88.48	70.68	89.88	73.72	89.96	73.87
Asymmetric-S	67.28	43.64	86.97	68.95	89.31	73.82	90.74	78.45
Ours	83.11	63.82	90.53	76.52	91.18	78.30	92.61	81.34
ShuffleNet								
Unified Model	82.54	60.12	89.85	64.45	91.95	72.96	84.38	79.13
BCT-S	85.51	65.76	88.42	66.85	89.73	70.59	86.67	74.41
Asymmetric-S	82.30	62.50	84.41	57.43	88.42	66.43	78.92	74.92
Ours	89.25	73.09	89.93	74.86	89.93	75.14	92.81	81.97

Table 5. Performance comparison between different model architectures on Market1501.

4.2. Baseline Comparisons

We conduct baseline experiments on Market1501. We take “Softmax + Cross Entropy” as the loss function, ResNet-18 [23] as the backbone, and $W = \{0.25\times, 0.50\times, 0.75\times\}$ as the crop ratio list of the sub-models. Notes that for a model with crop ratio γ_i , the ratio of its parameters to the full model is γ_i^2 . Therefore, the ratio of the number of parameters in the $0.25\times$, $0.50\times$ and $0.75\times$ sub-model to that in the full model are $\frac{1}{16}$, $\frac{1}{4}$, $\frac{9}{16}$, respectively.

Comparison methods. As for one-to-one compatible learning methods BCT and Asymmetric, the compatibility does not outperform that achieved by their corresponding self-compatible methods especially when the models are not constrained to be compatible directly. For self-compatible learning methods, BCT-S and Asymmetric-S, the conflicts among sub-models lead to a limited average accuracy.

Our Method (SFSC). As we resolve the conflicts among different sub-models, our methods can achieve many-to-many compatibility among all the sub-models, which significantly outperforms all comparison methods.

4.3. Performances under Different Settings

Different datasets. We evaluate each self-compatible learning method on the other person ReID datasets MSMT17, and the vehicle ReID dataset VeRi-776 under baseline configurations, as shown in Table 3. The results on the two datasets with larger scales are consistent with that on the Market1501. Our methods outperform others significantly for the small capacity sub-model (e.g. the $0.25\times$ sub-model $\phi_{\frac{1}{4}}$).

Different loss functions. We replace the original loss

function L_{ori} from “Softmax” to “Softmax + Triplet” and “Circle Loss” [52] to evaluate the effectiveness of each method, while Triplet loss is a kind embedding loss and circle loss is a kind of classification loss. The results in Table 4 show that the accuracy of BCT-S and Asymmetric-S drops greatly when training with circle loss, as Our method still stably achieves the best performance.

Different model architectures. We also study whether each method can be stably applied to various model architectures, i.e., ResNet-50 [23], MobileNet V2 [45] and ShuffleNet [67]. ShuffleNet and MobileNet are both designed especially for edge devices, which meet the application scenarios of self-compatible learning. It is shown in Table 5 that BCT-S and Asymmetric-S do not achieve higher accuracy than deploying a unified model with ShuffleNet and MobileNet (e.g. for the $0.75\times$ sub-model $\phi_{\frac{9}{16}}$). And our methods always improve the accuracy in all cases.

4.4. Ablation Study

We replace the proposed compatible loss and aggregation module with existing compatible loss BCT [50] and summation aggregation to test the proposed two methods.

	$M(\phi_{\frac{1}{16}}, \phi)$	$M(\phi_{\frac{1}{4}}, \phi)$	$M(\phi_{\frac{9}{16}}, \phi)$	$M(\phi, \phi)$
Unified Model	55.25	67.48	71.25	80.91
w/o Ours Loss	66.41	74.58	76.86	77.87
w/o Ours Aggregation	68.08	73.76	74.84	75.38
SFSC	69.43	76.37	77.26	81.43

Table 6. Ablation Study on Market1501.

Table 6 shows that the proposed compatible loss mainly improves the retrieval accuracy of the small capacity sub-models while the proposed aggregation method can ensure the performance of the full model.

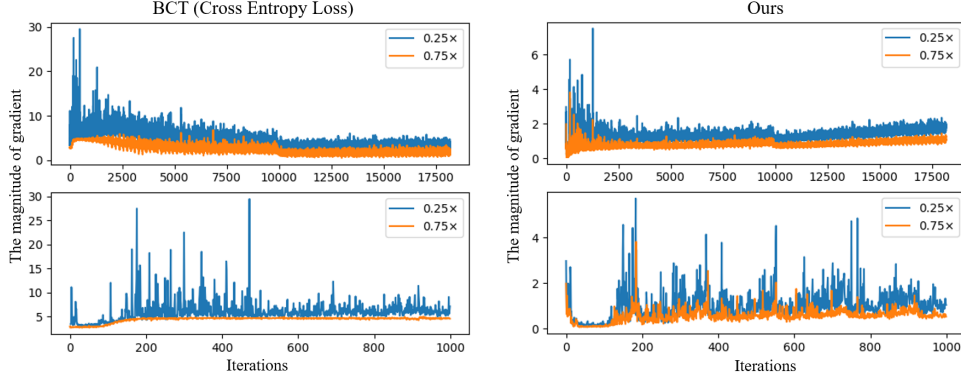


Figure 5. Gradient visualization on Market1501-R18 baseline. As for the original cross-entropy loss (BCT [49]), the magnitude of the $0.25\times$ sub-model is much larger than that of the $0.75\times$ sub-model, especially at the early stage of the training process (0~1000 iterations). Therefore, the $0.25\times$ sub-model dominates the training progress and the improvements of the $0.75\times$ sub-models are eliminated. While our method ensures co-optimization through uncertainty estimation.

4.5. Hyperparameter Analysis

Here we verify the impact of hyperparameter settings on the performance of the proposed method. The key hyperparameter is the crop ratio list for sub-models W . Intuitively, if the number of W increases or the minimal crop ratio decreases, the compatibility constraints on model ϕ would be more strict, thereby increasing the difficulty of compatible learning. Therefore, we take the baseline configuration and set W with different values.

	$M(\phi_{\frac{1}{16}}, \phi)$	$M(\phi_{\frac{81}{400}}, \phi)$	$M(\phi_{\frac{169}{400}}, \phi)$	$M(\phi_{\frac{289}{400}}, \phi)$	$M(\phi, \phi)$
Unified Model	55.25	62.54	69.54	72.57	80.27
BCT-S	61.45	68.63	73.42	72.80	75.39
Asymmetric-S	62.87	69.45	75.20	73.14	74.60
Ours	69.63	75.81	77.10	77.36	80.59

Table 7. Performance comparison with $W=\{0.25\times, 0.45\times, 0.65\times, 0.85\times\}$ on Market1501 (mAP).

	$M(\phi_{\frac{1}{100}}, \phi)$	$M(\phi_{\frac{1}{16}}, \phi)$	$M(\phi_{\frac{9}{16}}, \phi)$	$M(\phi, \phi)$
Unified Model	40.40	55.25	71.25	80.27
BCT-S	42.34	59.75	71.73	76.39
Asymmetric-S	42.58	58.20	70.50	75.42
Ours	45.44	68.18	78.00	81.46

Table 8. Performance comparison with $W=\{0.10\times, 0.25\times, 0.75\times\}$ on Market1501 (mAP).

Table 7 shows the performances with the number of W increases, while Table 8 shows the performances with the minimum among W decreases. The result of the comparison methods has fluctuated to varying degrees. SFSC still achieves the best retrieval accuracy.

4.6. Visualization

We visualize the norm of the gradient during the training progress on Market1501-R18 baseline, as shown in Figure 5. We also visualize the cumulative number of gradient pairs which are conflict with each other during the training process. As shown in Figure 6, our method has found a generic direction to improve all sub-models. Therefore, there are almost no conflict pairs after 5000 iterations. As

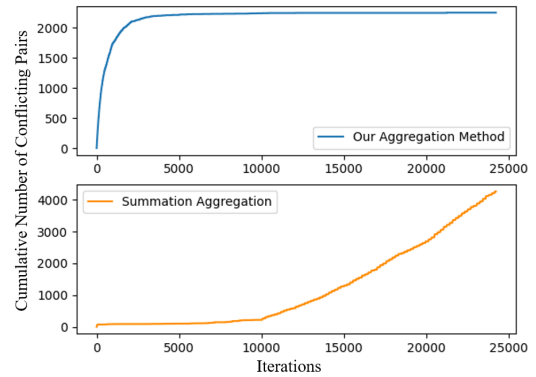


Figure 6. The cumulative number of conflict pairs during the training process.

for the summation aggregation, it faces fierce conflicts especially after 10000 iterations, resulting in the improvement of different sub-models overestimated or underestimated.

5. Conclusion

This paper proposes a Switchable Representation Learning Framework with Self Compatibility (SFSC) for multi-platform model collaboration. SFSC enables us to obtain models with various capacities to fit different computing and storage resource constraints on diverse platforms. It finds a generic direction to improve all sub-models by uncertainty estimation and gradient projection. SFSC achieves state-of-the-art performances and shows convincing robustness in different case studies.

6. Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant 62088102, the Basic and Frontier Research Project of PCL, and in part by the PKU-NTU Joint Research Institute (JRI) sponsored by a donation from the Ng Teng Fong Charitable Foundation.

References

- [1] Lu Bai, Yew-Soon Ong, Tiantian He, and Abhishek Gupta. Multi-task gradient descent for multi-task learning. *Memetic Computing*, 12(4):355–369, 2020. 4
- [2] Wentao Bao, Qi Yu, and Yu Kong. Evidential deep learning for open set action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13349–13358, 2021. 3
- [3] Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. Conditional computation in neural networks for faster models. *arXiv preprint arXiv:1511.06297*, 2015. 2
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 2
- [5] Tolga Bolukbasi, Joseph Wang, Ofer Dekel, and Venkatesh Saligrama. Adaptive neural networks for efficient inference. In *International Conference on Machine Learning*, pages 527–536. PMLR, 2017. 2
- [6] Mateusz Budnik and Yannis Avrithis. Asymmetric metric learning for knowledge transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8228–8238, 2021. 1, 4, 5
- [7] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791*, 2019. 2
- [8] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. *arXiv preprint arXiv:1812.00332*, 2018. 2
- [9] Yair Censor. Pareto optimality in multiobjective problems. *Applied Mathematics and Optimization*, 4(1):41–59, 1977. 5
- [10] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, pages 794–803. PMLR, 2018. 2
- [11] Kyunghyun Cho and Yoshua Bengio. Exponentially increasing the capacity-to-computation ratio for conditional computation in deep learning. *arXiv preprint arXiv:1406.7362*, 2014. 2
- [12] Jifeng Dai, Haozhi Qi, Yuwen Xiong, and Yi Li. Guodong, zhang, han hu, and yichen wei. *Deformable convolutional networks*. In, *ICCV*, 2(3):7, 2017. 2
- [13] Yunshu Du, Wojciech M Czarnecki, Siddhant M Jayakumar, Mehrdad Farajtabar, Razvan Pascanu, and Balaji Lakshminarayanan. Adapting auxiliary losses using gradient similarity. *arXiv preprint arXiv:1812.02224*, 2018. 2
- [14] Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li. Orthogonal gradient descent for continual learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3762–3773. PMLR, 2020. 2
- [15] Hang Gao, Xizhou Zhu, Steve Lin, and Jifeng Dai. Deformable kernels: Adapting effective receptive fields for object deformation. *arXiv preprint arXiv:1910.02940*, 2019. 2
- [16] Jean Gordon and Edward H Shortliffe. The dempster-shafer theory of evidence. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, 3:832–838, 1984. 3, 4
- [17] Benjamin Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12259–12269, 2021. 1
- [18] Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016. 2
- [19] Yizeng Han, Gao Huang, Shiji Song, Le Yang, Honghui Wang, and Yulin Wang. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [20] Zongbo Han, Changqing Zhang, Huazhu Fu, and Joey Tianyi Zhou. Trusted multi-view classification. *arXiv preprint arXiv:2102.02051*, 2021. 3
- [21] Adam W Harley, Konstantinos G Derpanis, and Iasonas Kokkinos. Segmentation-aware convolutional networks using local attention masks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5038–5047, 2017. 2
- [22] Hussein Hazimeh, Natalia Ponomareva, Petros Mol, Zhenyu Tan, and Rahul Mazumder. The tree ensemble layer: Differentiability meets conditional computation. In *International Conference on Machine Learning*, pages 4138–4148. PMLR, 2020. 2
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 4, 7
- [24] Lingxiao He, Xingyu Liao, Wu Liu, Xinchun Liu, Peng Cheng, and Tao Mei. Fastreid: A pytorch toolbox for general instance re-identification. *arXiv preprint arXiv:2006.02631*, 6(7):8, 2020. 6
- [25] Thomas M Hehn, Julian FP Kooij, and Fred A Hamprecht. End-to-end learning of decision trees and forests. *International Journal of Computer Vision*, 128(4):997–1011, 2020. 2
- [26] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens Van Der Maaten, and Kilian Q Weinberger. Multi-scale dense networks for resource efficient image classification. *arXiv preprint arXiv:1703.09844*, 2017. 2
- [27] Gao Huang, Shichen Liu, Laurens Van der Maaten, and Kilian Q Weinberger. Condensenet: An efficient densenet using learned group convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2752–2761, 2018. 2
- [28] Yani Ioannou, Duncan Robertson, Darko Zikic, Peter Kotschieder, Jamie Shotton, Matthew Brown, and Antonio Criminisi. Decision forests, convolutional networks and the models in-between. *arXiv preprint arXiv:1603.01250*, 2016. 2
- [29] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991. 2

- [30] Taejong Joo, Uijung Chung, and Min-Gwan Seo. Being bayesian about categorical probability. In *International Conference on Machine Learning*, pages 4950–4961. PMLR, 2020. 3
- [31] Shu Kong and Charless Fowlkes. Pixel-wise attentional gating for scene parsing. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1024–1033. IEEE, 2019. 2
- [32] Igor Kononenko. Bayesian neural networks. *Biological Cybernetics*, 61(5):361–370, 1989. 3
- [33] Harold W Kuhn and Albert W Tucker. Nonlinear programming. In *Traces and emergence of nonlinear programming*, pages 247–258. Springer, 2014. 2
- [34] Xiaoxiao Li, Ziwei Liu, Ping Luo, Chen Change Loy, and Xiaoou Tang. Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3193–3202, 2017. 2
- [35] Sen Lin, Li Yang, Deliang Fan, and Junshan Zhang. Trgp: Trust region gradient projection for continual learning. *arXiv preprint arXiv:2202.02931*, 2022. 2
- [36] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE international conference on computer vision*, pages 2736–2744, 2017. 2
- [37] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017. 2
- [38] Yihang Lou, Ling-Yu Duan, Shiqi Wang, Ziqian Chen, Yan Bai, Changwen Chen, and Wen Gao. Front-end smart visual sensing and back-end intelligent analysis: A unified infrastructure for economizing the visual system of city brain. *IEEE Journal on Selected Areas in Communications*, 37(7):1489–1503, 2019. 1
- [39] Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. 6
- [40] David JC MacKay. Bayesian neural networks and density networks. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 354(1):73–80, 1995. 3
- [41] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. *Advances in neural information processing systems*, 27, 2014. 2
- [42] Amir Rosenfeld and Shimon Ullman. Visual concept recognition and localization via iterative introspection. In *Asian Conference on Computer Vision*, pages 264–279. Springer, 2016. 2
- [43] Abhijit Guha Roy, Nassir Navab, and Christian Wachinger. Concurrent spatial and channel ‘squeeze & excitation’ in fully convolutional networks. In *International conference on medical image computing and computer-assisted intervention*, pages 421–429. Springer, 2018. 2
- [44] Manas Sahni, Shreya Varshini, Alind Khare, and Alexey Tumanov. Compofa: Compound once-for-all networks for faster multi-platform deployment. *arXiv preprint arXiv:2104.12642*, 2021. 2
- [45] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 1, 4, 7
- [46] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31, 2018. 2, 4
- [47] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty. *Advances in Neural Information Processing Systems*, 31, 2018. 3
- [48] Yantao Shen, Fanzi Wu, and Ying Shan. Non-inherent feature compatible learning. 2020. 2
- [49] Yantao Shen, Yuanjun Xiong, Wei Xia, and Stefano Soatto. Towards backward-compatible representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6368–6377, 2020. 1, 4, 5, 8
- [50] Yantao Shen, Yuanjun Xiong, Wei Xia, and Stefano Soatto. Towards backward-compatible representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6368–6377, 2020. 7
- [51] Hang Su, Varun Jampani, Deqing Sun, Orazio Gallo, Erik Learned-Miller, and Jan Kautz. Pixel-adaptive convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11166–11175, 2019. 2
- [52] Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. Circle loss: A unified perspective of pair similarity optimization. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6398–6407, 2020. 7
- [53] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2820–2828, 2019. 2
- [54] Surat Teerapittayanon, Bradley McDanel, and Hsiang-Tsung Kung. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 2464–2469. IEEE, 2016. 2
- [55] Andreas Veit and Serge Belongie. Convolutional networks with adaptive inference graphs. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018. 2
- [56] Thomas Verelst and Tinne Tuytelaars. Dynamic convolutions: Exploiting spatial sparsity for faster inference. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2320–2329, 2020. 2
- [57] Xinyi Wang, Ankur Bapna, Melvin Johnson, and Orhan Firat. Gradient-guided loss masking for neural machine translation. *arXiv preprint arXiv:2102.13549*, 2021. 2
- [58] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E Gonzalez. Skipnet: Learning dynamic routing in

- convolutional networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 409–424, 2018. 2
- [59] Yulin Wang, Kangchen Lv, Rui Huang, Shiji Song, Le Yang, and Gao Huang. Glance and focus: a dynamic approach to reducing spatial redundancy in image classification. *Advances in Neural Information Processing Systems*, 33:2432–2444, 2020. 2
- [60] Zhongdao Wang, Luming Tang, Xihui Liu, Zhuliang Yao, Shuai Yi, Jing Shao, Junjie Yan, Shengjin Wang, Hongsheng Li, and Xiaogang Wang. Orientation invariant feature embedding and spatial temporal regularization for vehicle re-identification. *Proceedings of the IEEE International Conference on Computer Vision*, pages 379–387, 2017. 5
- [61] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 79–88, 2018. 5
- [62] Shengsen Wu, Liang Chen, Yihang Lou, Yan Bai, Tao Bai, Minghua Deng, and Lingyu Duan. Neighborhood consensus contrastive learning for backward-compatible representation. *arXiv preprint arXiv:2108.03372*, 2021. 1, 2
- [63] Sijie Yan, Yuanjun Xiong, Kaustav Kundu, Shuo Yang, Siqi Deng, Meng Wang, Wei Xia, and Stefano Soatto. Positive-congruent training: Towards regression-free model updates. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14299–14308, 2021. 2
- [64] Zhicheng Yan, Hao Zhang, Robinson Piramuthu, Vignesh Jagadeesh, Dennis DeCoste, Wei Di, and Yizhou Yu. Hd-cnn: hierarchical deep convolutional neural networks for large scale visual recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 2740–2748, 2015. 2
- [65] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. *arXiv preprint arXiv:1812.08928*, 2018. 3
- [66] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33:5824–5836, 2020. 2, 4
- [67] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6848–6856, 2018. 4, 7
- [68] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. *Proceedings of the IEEE international conference on computer vision*, pages 1116–1124, 2015. 5
- [69] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9308–9316, 2019. 2