

Exploiting Convolution Filter Patterns for Transfer Learning

Mehmet Aygün
Istanbul Technical University
Istanbul, Turkey
aygunme@itu.edu.tr

Yusuf Aytar
MIT
Cambridge, USA
yusuf@csail.mit.edu

Hazım Kemal Ekenel
Istanbul Technical University
Istanbul, Turkey
ekenel@itu.edu.tr

Abstract

In this paper, we introduce a new regularization technique for transfer learning. The aim of the proposed approach is to capture statistical relationships among convolution filters learned from a well-trained network and transfer this knowledge to another network. Since convolution filters of the prevalent deep Convolutional Neural Network (CNN) models share a number of similar patterns, in order to speed up the learning procedure, we capture such correlations by Gaussian Mixture Models (GMMs) and transfer them using a regularization term. We have conducted extensive experiments on the CIFAR10, Places2, and CM-Places datasets to assess generalizability, task transferability, and cross-model transferability of the proposed approach, respectively. The experimental results show that the feature representations have efficiently been learned and transferred through the proposed statistical regularization scheme. Moreover, our method is an architecture independent approach, which is applicable for a variety of CNN architectures.

1. Introduction

The CNN models are found to be successful at various computer vision tasks such as image classification [14, 27, 28], object detection [22], image segmentation [19, 21], and face recognition [20], where large-scale datasets [24, 32, 20] are available. Nevertheless, the performance of CNN models significantly reduces, when training data is limited or the domain of the training set is far from the test set. Today, the most successful and practical solution to address lack of annotated large dataset is training the networks on large-scale annotated datasets like ImageNet [24] and Places [32], then finetuning these pre-trained networks for specific problems. Thanks to community, the pre-trained models of well-known architectures like AlexNet [18], VGG-16 [27], GoogLeNet [28], and ResNet [14] can be found available online. However, when some architectural changes are needed, these pre-trained networks can-

not be used. For such cases, it is necessary to train models on large datasets, then, finetune for the particular problem. Unfortunately, while getting nearly human performances on a lot of applications with these models, training these networks on large datasets is still a significant problem and a very time consuming process.

With recent advances in deep learning such as Residual Learning [14], successful networks have become more and more deep, and training these models become harder in terms of complexity and time. To find a solution for this problem, inspired by *What makes a good detector?* [11], we have investigated *What makes a good CNN filter?* Two successful CNN models, AlexNet [18] and VGG-16 [27], are analyzed from a statistical perspective, and we realized that these models show similar patterns and redundancies. In addition to our findings, in [7], authors show that 95% of weights of neural networks could be predicted without any reduction in accuracy. This leads us to the idea that we can use these redundancies and patterns for learning better representations quickly by transfer learning. Similar to the methods used in [2, 4, 11], we introduce a regularization term for transferring the statistical information to improve the learning scheme. First, the statistical distribution of convolution filters from a well-trained network is learned with a Gaussian Mixture Model. Next, the newly trained model is encouraged to show similar statistics with source models using the regularization term. Extensive experiments on the CIFAR10 [17], Places2 [32], and CMPlaces [4] datasets show that the proposed approach is generalizable and the networks can quickly learn a representation with statistical regularization, which could efficiently be transferred to another task and cross-domains. The overview of our proposed method can be seen in Figure 1. The rest of the paper is organized as follows: in Section 2 related works are summarized, in Section 3 our detailed statistical analysis is reviewed and the regularization term is introduced, in Section 4 experimental results are presented and discussed, and finally in Section 5 the paper is concluded.

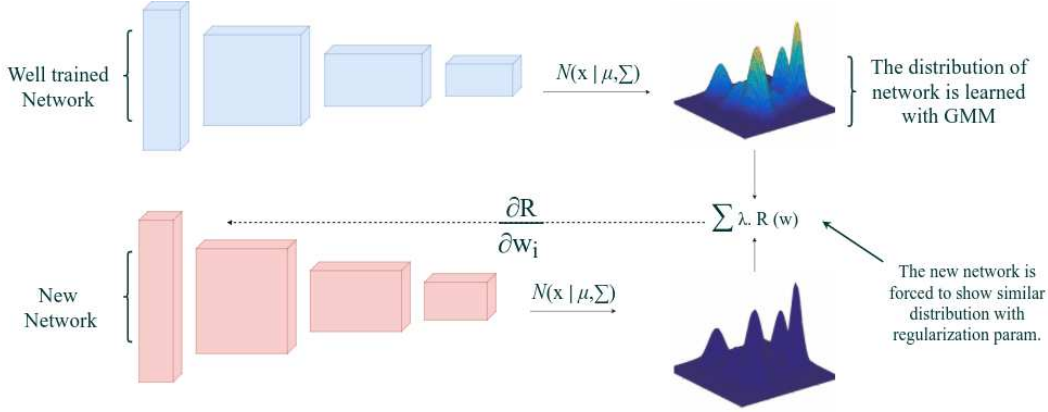


Figure 1: Overview of the proposed system. The blue (top) network is a well-trained CNN where its distribution is learned via GMM. When new red (bottom) network being trained, along with the classification loss, the new regularization term $\sum \lambda \cdot R(w)$ that measures the statistical difference of two distributions are minimized. The statistical knowledge could be transferred from source CNN to target CNN with our regularization term. Best viewed in color.

2. Related Work

Network Distillation: The aim of network distillation approaches are transforming larger networks to smaller ones, while not losing so much information that was learned in the large network. The pioneering work is conducted by Bucila *et al.* [3], where their aim is to compress the ensemble of models to a single model without significant accuracy loss. Later, Hinton *et al.* [15] optimized the smaller network to show similar softmax output of cumbersome model. Then, Romero *et al.* [23] suggest that in addition to softmax outputs, intermediate representations could be used for distilling the network. Recently distillation is also applied in the cross-modal settings [13, 1]. The major drawback of these methods is the necessity to train a large network before using it to train a smaller one. Also these models mainly match the outputs of the networks, whereas we regularize the internal weights of the network.

Domain Adaptation: Domain adaptation is the problem of learning a model that generalizes target domain examples besides source domain ones while learning only from source domain. With the success of CNNs, domain adaptation works have been focused on CNNs and several successful methods have been proposed. For instance, Ganin & Lempitsky [9] introduced *gradient reversal layer* which act like identity transform in forward pass of CNN and change the sign of the gradient and scale in the backward pass. In their work, they added a domain classifier to the end of the feature map and tried to predict the domain of examples. In backpropagation, they changed the gradient using gradient reversal layer and forced the networks to learn domain invariant features by maximizing loss of domain classifier. Tzeng *et al.* [29] added new terms

into objective function of CNN to both increasing domain confusion and transferring inter-class knowledge. The first term *domain confusion loss* forces the networks to learn domain invariant features and *soft label loss* forces the feature of the same class to be similar for both source and target domain. In contradistinction to Ganin & Lempitsky [9], some target labels must be available for optimizing soft label loss. Moreover, recently some other works [10, 12, 30, 26] focused on domain adaptation problems for CNNs. While these methods focus on transferring information about structure of data, our method focuses on transferring more local information.

Statistical Transfer: Statistical transfer is learning statistical properties from a source and to use this statistical knowledge for improving learning procedure. For instance, Aytar & Zisserman [2] proposed part-level transfer regularization which transfers parts of source detectors instead of whole detector. Additionally, they take advantage of part co-occurrence statistics. For example, if there are wheels in the picture, probably another wheel would also appear. They calculated these co-occurrence statistics using the source data and transfer these statistics when a new object detector is learned. Moreover, in the era of hand crafted features, Gao *et al.* [11] analyzed famous HOG [6] templates of successful object detectors and made two observations. Firstly, they observed that activations of individual cell models had some correlations, secondly local neighborhoods of cells also showed the same characteristic. Furthermore, since they wanted to transfer local information in contrast to templates like in [2], they defined their priors such that the correlations from the source model could be transferred to target model without global template align-

ment. In a recent work [4], GMMs are used for aligning cross-modal data. In this approach, statistics of activation maps of different layers for a modality are learned, and the other modalities are forced to show similar statistics in the activations via a regularization term. The proposed method is capable of aligning modalities by using this regularization term when there is no strong alignment between modalities.

Our work is influenced by Gao *et al.* [11] and Castrejon *et al.* [4]. The first work directs us to analyze the weights of convolution filters and find the correlation between the filters, and the other one to use the GMM to enforce them to show similar statistics in a non-convex optimization problem. Different from these works, our regularization forces the weights to show similar statistics and transfers local correlations of convolution filters.

3. Statistical Regularization

In this section, first, we present our analysis on CNN models from a statistical perspective. Next, we describe our approach for capturing statistical knowledge from a CNN and present how to transfer this knowledge via a regularization term.

3.1. Statistical Analysis

For investigating the statistical behavior of a CNN model, we have used VGG-16 [27] and AlexNet [18] models that were trained on ImageNet [24] and Places-365 [32] datasets. Especially, weights of convolution filters are analyzed for the four models. We tried to answer following three questions: (i) are filters separable into clusters?; (ii) how much similar are the filters inside a cluster?; and (iii) how all filters are distributed over clusters?. In order to obtain this information, all 3×3 filters of a model are clustered to ten different clusters using the k-means algorithm and each cluster's covariance matrix and mean value were calculated and visualized. The covariance matrices would provide information about how members of a cluster are correlated. Mean values show whether the sets are similar to each other or not. For instance, it can be seen in Figure 2 (a) the covariance matrices of all the four models show that there is some shared behavior of learned filters across layers and architectures. Also, the cluster centers depict different similarities. Generally, all models have clusters that their mean values are accumulated on the left, right, top, and bottom. Moreover, when we look at the distribution of filters to clusters, the models show different characteristics. While the distributions are Gaussian-like in VGG-16 models, the distributions cannot be fitted to a known distribution in the AlexNet models and most of the clusters have roughly the same number of filters. However, both AlexNet and VGG-16 models show similar distributions by model wise while they are trained on different dataset. The mean values are

shown in Figure 2 (b), and the distributions can be seen in Figure 3.

3.2. The Proposed Method

Since our aim is capturing statistical properties of “good” convolution filters and transferring this statistical knowledge to another network, we modeled the convolution filters' distributions and enforced the network to show similar distributions. Similar to our work, [4] forces the activations of networks to show similar distributions across modalities for aligning cross modal data. In [4], the authors used both mixture and single Gaussian distributions for modelling activations, and the mixture models outperformed single Gaussian model in their problem. Since, our aim is also similar to their work -capturing statistical knowledge and transferring it- we have decided to use mixture models for modeling distributions. The main difference is that we force the weights of filters instead of the activations to show similar distributions across the networks.

Let x_n and y_n be a training image and its corresponding label. We want to minimize

$$\min_w \sum_n L(z(x_n, w), y_n) \quad (1)$$

where $z_n(x_n, w)$ is output of the network. We have added a regularization term R to the loss term that represents the negative log likelihood of a convolution filter to encourage the network to learn the weights that are statistically similar to another network. For filter w_i and distribution P we define R such that,

$$R(w_i) = -\log(P(w_i)) \quad (2)$$

The distribution P is modeled as GMM, therefore P would be,

$$P(w|\pi, \mu, \Sigma) = \sum_{k=1}^K \pi_k \mathcal{N}(w|\mu_k, \Sigma_k) \quad (3)$$

where K is the number of mixtures and $\sum_k \pi_k = 1$, $\pi_k \geq 0 \forall_k$. Therefore, the total negative log likelihood for N convolution filters can be defined as,

$$R(w|\mu, \Sigma) = \sum_{i=1}^N -\log \sum_{k=1}^K \pi_k \mathcal{N}(w_i|\mu_k, \Sigma_k) \quad (4)$$

where \mathcal{N} is specified as,

$$\mathcal{N}(w|\mu, \Sigma) = \frac{1}{(2\pi|\Sigma|)^{-\frac{1}{2}}} \exp(-\frac{1}{2}(w-\mu)^T \Sigma^{-1}(w-\mu)) \quad (5)$$

The derivate of the negative log likelihood must be calculated exactly or analytically, since the network is trained with back-propagation. Still, calculating exact derivative in every iteration would be very expensive during training,




Figure 2: (a) The covariance matrices of ten clusters that are calculated using the weights of convolution filters. The converged VGG-16 and AlexNet models, which are trained on Places and ImageNet datasets, are used for clustering and visualizing. (b) Visualizations of mean values of each clusters.

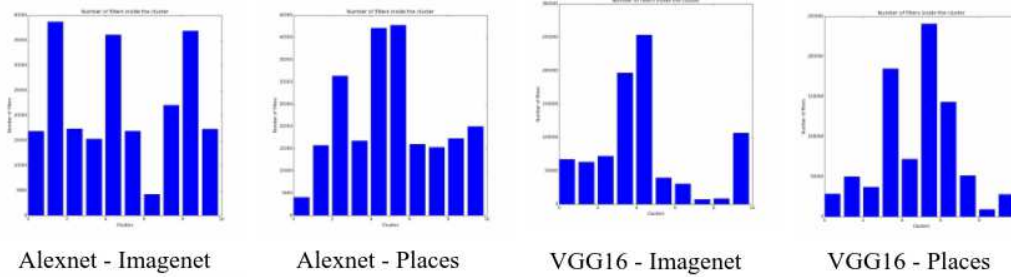


Figure 3: The distributions of the convolution filters over clusters. For both Places and ImageNet, the converged model of VGG and AlexNet show similar distributions.

therefore, we approximately calculate the derivative. To approximately calculate the derivative of a convolution filter, we first pick the mixture component $\mathcal{N}(\mu_s, \Sigma_s)$ that the probability of the convolution filter is maximized. Next, the derivate is calculated using that single Gaussian. The partial derivative of R with respect to a convolution filter w_i would be

$$\frac{\partial R}{\partial w_i} = (w_i - \mu_s) \Sigma_s^{-1} \quad (6)$$

Finally, our complete loss term is defined as

$$\min_w \sum_n L(z(x_n, w), y_n) + \sum \alpha \frac{1}{2} w^2 + \sum \lambda \cdot R(w) \quad (7)$$

where the first and second terms are classification loss and weight decay terms, and the last one is our regularization term. The λ is a hyperparameter that controls the regularization. For experiments in this paper, we have used distribution P , which is learned from convolution filters of VGG-16 model trained on ImageNet using Expectation maximization (EM) algorithm with $K = 1000$ components. To reduce the number of parameters, we assumed that the covariances Σ_k are diagonal. Also, K-Means algorithm was

run on filters for decreasing the convergence time of EM algorithm. While it is known that, different layers, especially first layers, show different characteristics than other layers, we have used a single GMM $-P-$ for all filters. Because, if multiple distributions are used, a layer alignment relationship between layers are needed in transfer time.

4. Experiments and Results

Our hypothesis was that well-trained CNN models show similar statistical patterns, and we could exploit this information in the training phase. Our experiments show that we can learn a better representation faster with the statistical regularization. For example, in section 4.1, we validate that with regularization the convolution filters create more general representations and show less overfitting characteristics than regular trained filter's representations. Also in section 4.2, we see how the regularization helps to learn representations that are successfully transferred to another task. Furthermore, while transferring across task and domains is widely studied, transfer learning for cross modal data is not a well studied problem. We show that our method

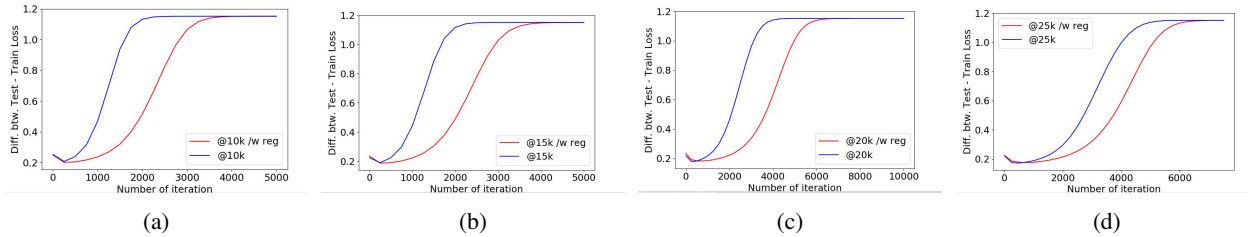


Figure 4: Difference between train and test losses of models in CIFAR-10. After freezing weights the models continued to training. It can be easily seen that later freezing show less overfitting. In the (a),(b),(c), and (d) the regularization effect is compared. With regularization models always overfit slowly (the difference increases slowly) than regular training. Best viewed in color.

could be applied to cross-modal data as well. In section 4.3, we present our experiments applied to a cross-modal dataset CMPlaces [4].

4.1. Generalization

In this section, we evaluate if the statistical regularization helps the generalizability of the learned representation. For this purpose, we have used middle sized CIFAR-10 dataset [17] and CNN architecture described in [25]. Firstly, we trained the network with and without regularization and stopped the training at 10k, 15k, 20k, and 25k iterations. Later on we freeze all 3×3 filters and continue learning with training data. Since only the last layers change during the training and the features extracted from convolution filters are not so generic, the validation loss starts to increase and the validation accuracy is dropped.

We compared the networks, whose training stopped in different iterations with and without regularization. In our experiments, the training loss is started to oscillate in a small interval and does not change much, that is, neither increases nor decreases. However, the validation loss changed. The gap between training and test loss indicates how generalizable the network is.

When we compared the models initialized at 10k, 15k, 20k, and 25k iterations, always the test loss of regularized versions increased slowly compared to those of non-regularized versions. For instance, as can be seen in Figure 4 (a,b,c,d), the difference between training and test loss increases slowly in regularized models than non-regularized models. Interestingly, as regular training time (no freezing) increases, the gap between the regularized and normal training reduces.

4.2. Task Transfer

In this section, we show that the quickly learned representations can also be transferred to another task more successfully with regularization. To validate our claim, we try to transfer filter distributions from ImageNet to Places2 [32] dataset. The trivial solution for a classification problem is

training the network on ImageNet and finetuning the models on the new task afterwards. We follow the same procedure in this section for our experiments, but we want to finish the pre-training stage as early as possible. Also we want to show that our regularization can be used with finetuning. To evaluate our method’s performance, we first train VGG-F model introduced in [5] on ImageNet [24] data with and without regularization. As in Section 4.1, we take snapshots from 10k, 25k, and 50k iterations. Next, we start finetuning on Places2 data [32]. When we compare regularization effect, we see that the regularization can help to learn better representations in the early iterations. For example, as can be seen in Figure 5 (a), the performance of the models that are initialized at the weights learned in pre-training with only 10k iterations, the regularized version shows a better performance than the non-regularized one. When we examine the models initialized at 25k iterations, the performance difference between regularized and non-regularized versions reduces. Finally, in 50k iterations, nearly there is no difference in the performance. The test loss/iterations plots are shown in Figure 5 (a), (b), (c). This experiment shows that as pre-training time increases, the gain obtained from regularization decreases. However, for limited amount of pre-training time, the regularization could increase the efficiency of the pre-training.

4.3. Modality Transfer

While CNNs performances are very good at various computer vision tasks for real-world images, most of the computer vision algorithms fail at non-real images. This shows that generalization performance of the computer vision algorithms is not good for cross modal data. Some recent works [4, 8] focused on this problem and recently Castrejon *et al.* [4] has introduced a new Cross-Modal dataset. In the dataset there are five different modalities for each scene type such as natural image, sketch, clip art, spatial text, and description. As in section 4.2, we firstly train VGG-F models on a large dataset -Places2- and finetune on CMPlaces data. We have used clipart and sketch data to

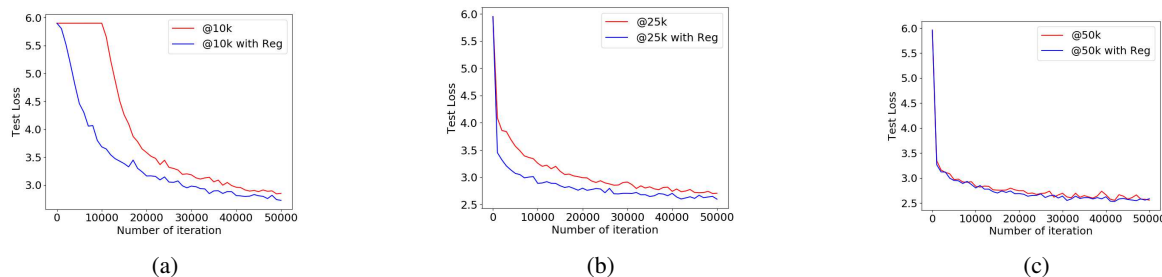


Figure 5: The pre-trained VGG-F models are finetuned on Places2 dataset. Pretraining is stopped at 10k, 25k, and 50k iterations. For (a) and (b) it can be easily seen that with regularization the test loss decreases quickly. Since test losses and accuracies correlated we only provide test losses.

Pre-Training	Top-5 Accuracy
25k	60.45
25k \w Reg.	62.25
50k	63.0
50k \w Reg.	64.5
Converged ImageNet	64.8

Table 1: The first column describes how pretraining is done and the second column shows the top-5 accuracies after finetuning for clipart data.

Pre-Training	Top-5 Accuracy
25k	33.05
25k \w Reg.	40.75
50k	37.65
50k \w Reg.	41.1
Converged ImageNet	53.6

Table 2: The first column describes how pretraining is done and the second column shows the top-5 accuracies after finetuning for sketch data.

evaluate our performance, since we are interested in the visual domain. We take snapshots from 25k and 50k iterations from the regularized and non-regularized networks. Next, we start to finetune on sketch and clipart data and compare their accuracies. Also, we finetune VGG-F model converged on ImageNet data and compare with our pre-trained models. Similar to our experiments, by rising the training pre-training time the performance increases. Also, the regularization helps to learn better representations and increases the performance of pre-training for finetuning. Although, there is a significant gap between the converged ImageNet model and our regularization in the sketch data, there is not a substantial difference between 50k iteration with regularization and the converged ImageNet model on clipart modality. The results for both modalities can be seen in Table 1 and Table 2. These results show that instead of training ImageNet until the model converges, we could train the models using regularization with only few iterations and could employ these pre-trained networks for cross-modality transfer.

4.4. Implementation Details

We have used Caffe [16] deep learning framework in our experiments. Moreover, we have implemented our special convolution layer for applying statistical regularization.

When VGG-F model is trained on ImageNet and Places datasets, stochastic gradient descent with 0.01 learning rate is used for optimization. In the CIFAR experiments, we have used the same parameters described in [25]. Finally, the Gaussian mixture models are learned using the VLFeat library [31].

5. Conclusion

In this paper, we analyzed convolution filters of well-known CNN architectures and found that they share a number of common patterns and redundancies that could be exploited for transfer learning. Gaussian Mixture Models are used for capturing these statistical patterns and a new regularization term is introduced for transferring such patterns to other networks. Our experiments show that we could learn good representations that are transferable to the other tasks and cross-domains quickly with regularization. For instance, we achieved around 25% improvement on the sketch modality in the cross-modal dataset under limited pre-training time. Also, our method gets similar performance on clipart data with converged model that pre-trained on the ImageNet, while pre-training stopped at 50k iterations in our method.

References

- [1] Y. Aytar, C. Vondrick, and A. Torralba. Soundnet: Learning sound representations from unlabeled video. In *Advances in Neural Information Processing Systems*, pages 892–900, 2016.
- [2] Y. Aytar and A. Zisserman. Part level transfer regularization for enhancing exemplar svms. volume 138, pages 114 – 123, 2015.
- [3] C. Bucilu, R. Caruana, and A. Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541. ACM, 2006.
- [4] L. Castrejon, Y. Aytar, C. Vondrick, H. Pirsiavash, and A. Torralba. Learning aligned cross-modal representations from weakly aligned data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2940–2949, 2016.
- [5] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference*, 2014.
- [6] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 886–893. IEEE, 2005.
- [7] M. Denil, B. Shakibi, L. Dinh, N. de Freitas, et al. Predicting parameters in deep learning. In *Advances in Neural Information Processing Systems*, pages 2148–2156, 2013.
- [8] M. Eitz, K. Hildebrand, T. Boubekeur, and M. Alexa. Sketch-based image retrieval: Benchmark and bag-of-features descriptors. *IEEE transactions on visualization and computer graphics*, 17(11):1624–1636, 2011.
- [9] Y. Ganin and V. Lempitsky. Unsupervised domain adaptation by backpropagation. *arXiv preprint arXiv:1409.7495*, 2014.
- [10] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016.
- [11] T. Gao, M. Stark, and D. Koller. What makes a good detector?—structured priors for learning from few examples. *Computer Vision—ECCV 2012*, pages 354–367, 2012.
- [12] M. Ghifary, W. B. Kleijn, M. Zhang, D. Balduzzi, and W. Li. Deep reconstruction-classification networks for unsupervised domain adaptation. In *European Conference on Computer Vision*, pages 597–613. Springer, 2016.
- [13] S. Gupta, J. Hoffman, and J. Malik. Cross modal distillation for supervision transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2827–2836, 2016.
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [15] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [16] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [17] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009.
- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [19] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- [20] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *BMVC*, volume 1, page 6, 2015.
- [21] P. O. Pinheiro, R. Collobert, and P. Dollar. Learning to segment object candidates. In *Advances in Neural Information Processing Systems*, pages 1990–1998, 2015.
- [22] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016.
- [23] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.
- [24] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [25] T. Salimans and D. P. Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in Neural Information Processing Systems*, pages 901–901, 2016.
- [26] O. Sener, H. O. Song, A. Saxena, and S. Savarese. Unsupervised transductive domain adaptation. *arXiv preprint arXiv:1602.03534*, 2016.
- [27] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [28] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015.
- [29] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4068–4076, 2015.
- [30] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial discriminative domain adaptation. *arXiv preprint arXiv:1702.05464*, 2017.
- [31] A. Vedaldi and B. Fulkerson. Vlfeat: An open and portable library of computer vision algorithms. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 1469–1472. ACM, 2010.
- [32] B. Zhou, A. Khosla, A. Lapedriza, A. Torralba, and A. Oliva. Places: An image database for deep scene understanding. *arXiv preprint arXiv:1610.02055*, 2016.