

Artistic Object Recognition by Unsupervised Style Adaptation

Christopher Thomas and Adriana Kovashka

University of Pittsburgh, Pittsburgh PA 15260, USA
{chris,kovashka}@cs.pitt.edu

Abstract. Computer vision systems currently lack the ability to reliably recognize artistically rendered objects, especially when such data is limited. In this paper, we propose a method for recognizing objects in artistic modalities (such as paintings, cartoons, or sketches), without requiring any labeled data from those modalities. Our method explicitly accounts for stylistic domain shifts between and within domains. To do so, we introduce a complementary training modality constructed to be similar in artistic style to the target domain, and enforce that the network learns features that are invariant between the two training modalities. We show how such artificial labeled source domains can be generated automatically through the use of style transfer techniques, using diverse target images to represent the style in the target domain. Unlike existing methods which require a large amount of unlabeled target data, our method can work with as few as ten unlabeled images. We evaluate it on a number of cross-domain object and scene classification tasks and on a new dataset we release. Our experiments show that our approach, though conceptually simple, significantly improves the accuracy that existing domain adaptation techniques obtain for artistic object recognition.

1 Introduction

Clever design of convolutional neural networks, and the availability of large-scale image datasets for training [23, 29], have greatly increased the performance of object recognition systems. However, models trained on one dataset often do not perform well on another [28]. For example, training a model on photographs and applying it to an artistic modality such as cartoons is unlikely to yield acceptable results given the large differences in object appearance across domains [21], as illustrated in Fig. 1 (a).

One possible solution to this problem is to obtain labeled training data from the domain that the model will be applied on (i.e. the *target* modality). Unfortunately, obtaining sufficient data for training today’s deep convolutional networks on each modality of interest may be prohibitively expensive or impossible. In our scenario of recognition in artistic modalities, this problem is especially pronounced: a given artist usually does not produce limitless samples for a particular object class, yet each artist may have a unique style. This problem has led researchers to explore *domain adaptation* techniques: a model is trained on




Fig. 1. (a) Illustration of the visual variation between instances of the same object across domains. (b) Overview of our method: We transform our training data so it resembles the style of the target domain. Left: We illustrate how we use style modification to create a synthetic source modality that looks like the target modality, causing the learned features to be more target-appropriate. Right: Training on any single modality causes learned features to be *too* domain specific, so we train on both domains.

a domain where data is plentiful (the *source* domain) while ensuring that what the model learned can be applied to a disjoint *target* domain. Usually these approaches rely on a small amount of labeled data in the target domain [35] or on seeing some *unlabeled* target data and learning a feature space where the distributions on the source and target domains are similar [1, 2, 6, 8, 11, 13, 25, 36].

In contrast to the more general domain adaptation methods above, we focus on a *specific type of domain difference*. We consider domains which exhibit differences due to artistic style [7, 17, 19], such as paintings, cartoons and sketches. Note that artistic domains often contain limited data—e.g. the PACS dataset of [21] contains on average less than 2,500 images per domain.

We propose two variants of our method, both of which learn style-invariant representations for recognition, one using as few as ten unlabeled target images. Our methods construct a new, *synthetic source modality* from a single photorealistic modality. The synthetic modality bears the *style* of the target modality, and is obtained for free (without human annotations) through the use of style transfer techniques [17, 19]. We then train a network using both source modalities, ensuring that the features the model learns produce similar activations in the two source domains. Unlike existing methods which also create synthetic training data for domain adaptation [1, 36], our method is easier to train, and requires orders of magnitude less target data (e.g. ten vs ten thousand). This translates to more accurate performance for domains where even unlabeled target data is sparse. We illustrate our method in Fig. 1 (b).

Our method of generating synthetic data in the same artistic style as the target domain is applicable to any setting where the primary difference between the domains is the style of the content, e.g. cartoons, paintings, sketches, and line drawings. It is not however applicable to problems where the source and target domain differences extend beyond artistic style, e.g. RGB-D data.

We evaluate our approach on multiple datasets (both new and from prior work) and show it outperforms the state-of-the-art domain adaptation methods.

The main contributions of our work are as follows:

- We propose a framework for constructing synthetic source modalities useful for learning style-invariant classifiers through style transfer.
- We develop a novel style selection component, which enables accurate adaptation with just ten target images.
- We release a new large dataset of photographs and artistic images across four domains: natural photographs, cartoons, paintings, and sketches. The dataset is available at http://www.cs.pitt.edu/~chris/artistic_objects/.
- The two versions of our approach outperform recent domain adaptation techniques on the challenging task of object recognition in artistic modalities, which exhibits larger domain shift than traditional adaptation scenarios.
- We conduct rich evaluation against many baselines in different settings (also see our supplementary file).

2 Related Work

Recognition on artistic modalities. There is limited work in performing object recognition on artistic modalities. [9] show that as paintings become more abstract, the performance of person detection degrades. [14, 38] benchmark existing methods on artistic objects, but only propose improvements over older, non-convolutional techniques, which are greatly outperformed by neural networks. [37] provide a large dataset of artistic domains with per-image object annotations, but we found the labels were too coarse (due to multiple objects in the same image and no bounding box annotations) and the human annotations too sparse, to reliably perform recognition. [21] publish a dataset of objects in artistic modalities, and propose a domain *generalization* approach. We use this dataset but instead focus on domain *adaptation*, where we target a specific unlabeled modality, rather than attempting to do well on an unseen modality. Further, [21] require the availability of multiple original (non-synthetic) source modalities, while we only require one real modality.

Cross-domain retrieval. There has been interest in retrieving samples across domain boundaries, e.g. retrieving photographs that match a *sketch*. [31] find a joint embedding for sketches and photographs that ensures a paired sketch and photo are more similar in the learned space than a photo and sketch that are not paired. We utilize the dataset of [31] but perform recognition and not retrieval, and do not assume availability of labeled data from the target domain. [3] retrieve scene types across e.g. sketches and clipart. They propose a supervised approach (which we modify to an unsupervised setting for comparison) that encourages source and target neural activations to fit the same distribution.

Domain adaptation. Domain adaptation techniques can be broadly divided into two categories: 1) semi-supervised [24, 35] in which a limited amount of labeled target data is available, and 2) unsupervised techniques [1, 6, 8, 11, 25]

(such as this paper) where only unlabeled target data is available. Older domain adaptation techniques either modify a classifier or ensure that the feature space in which we attempt to classify a novel modality are similar to the features for the source modality [11]. Recently, researchers have tackled the same problem in the context of training convolutional neural networks. [6, 35] show how to learn domain-invariant features with CNNs by confusing domains, i.e. ensuring that a classifier cannot guess to which domain an image belongs based on its activations. Other approaches [2, 8] train an autoencoder which captures features specific to the target domain, then train a classifier on the source and apply it to the target. While these approaches attempt to bring the *source and target* domains closer, we encourage feature invariance between our *two source* domains, while explicitly bringing the source and *target* closer via style transfer. We compare against a number of recent, state-of-the-art unsupervised domain adaptation approaches and find that our approach consistently outperforms them.

Several works use *multiple source domains* for domain adaptation but these assume multiple real human-labeled sources are available [13, 21] or can be discovered in the given data [10]. Recent works [1, 15, 36] have tackled unsupervised domain adaptation using generative adversarial networks (GANs) [12] to generate synthetic target data. The networks are trained to produce novel data such that a discriminator network cannot tell it apart from the real target dataset. [15] extends the idea further by training a GAN which adapts source data at both the image and feature level. Our method has several important advantages over GAN-based approaches. We exploit the fact that our target modality’s domain gap with respect to the source can be bridged by controlling for artistic style. Since we know we want to model style differences, we can explicitly extract and distill them, rather than requiring a GAN to learn these in an unsupervised way. This allows our method to use *orders of magnitude less target data*. The ability to leverage limited target data is essential in applications where target data is extremely limited, such as artistic domains [21]. Second, training generative adversarial networks is challenging [30] as the training is unstable and dataset-dependent. Our image translation networks are *easier to train* and require no tuning. Our approach outperforms [1] in experiments.

Other domain adaptation works have also attempted to model style. [40] propose a framework for handwriting recognition from different writers which projects each writer’s writing into a “style-free” space and uses it to train a classifier. Another classic work [33] separates style and content using bilinear models and can generate a character in a certain writing style. Unlike our proposed methods, neither of these works use CNNs or perform object recognition.

Domain generalization. Domain generalization methods [21, 22, 26, 27, 39] attempt to leverage knowledge from one or more source domains during training, and apply it on completely unseen target domains. [21] proposes a method of domain generalization on a dataset of paintings, cartoons, and sketches. In order to generalize to cartoons, for example, they assume data is available from all other modalities. In contrast, we perform adaptation towards a known, but unlabeled target modality, without requiring multiple source modalities.

Style transfer. Style transfer methods [4, 5, 7, 16, 17, 19] modify an image or video to have the style of some single target image. For example, we might modify a portrait to have the same artistic style as “Starry Night” by Van Gogh. [7] use a pixel-by-pixel loss, whereas [19] use a more holistic perceptual loss to ensure that both content and style are preserved while an input image is transformed. While [7] use a backpropagation technique to iteratively transform an image to another style, [19] train a CNN to directly perform the style transfer in a feed-forward network. More recently, [17] propose an encoder-decoder framework which allows mimicking the style of a target image without retraining a network. Earlier work [34] models style at the object part and scene level. We use [17, 19] as the style transfer component of our framework. Importantly, the transformed data retains the classification labels from the photograph source data.

3 Approach

Our method explicitly controls for the stylistic gap between the source and target domains. While our method is based on existing style transfer approaches, its novelty is two-fold. First, we explicitly encourage variation in the styles towards which the source modality is modified, which we call *style selection*. In Sec. 4.3 and our supplementary file, we show that this novel style selection is crucial for attaining good performance in practice. Second, while other methods proposed to use synthetic data for adaptation, we apply our approach to a new problem: object recognition in *artistic* domains. This problem is characterized by larger domain shift compared to prior adaptation benchmarks, as shown in [21]; and by sparse unlabeled target data (see dataset sizes in Sec. 4.1).

The main intuition behind our approach is that we know a large part of the domain gap between photographs and our target modalities is their difference in artistic style. The key observation is that most of the images within an artistic modality exhibit one of a few “representative” styles. For example, some paintings may be cubist, expressionist, or photorealistic. Sketches may be different due to low-level detail, but look similar in terms of coarse visual appearance. Other unsupervised methods [1, 8] require a large dataset of target images to learn this shared appearance. Because we specifically focus on artistic style variations, we can explicitly distill “representative” styles from clusters of artistically similar images, and then transfer those styles onto our labeled training data. We train on both original and transformed photos, and add a constraint to our method to encourage the features on both domains to be indistinguishable. Because the only difference between the real and synthetic domains is their style, this constraint, perforce, *factors out artistic style* from our feature representation and bridges the domain gap. Below we first describe how we obtain additional source modalities, and then how we learn style-invariant features during training.

Our approach automatically transforms photographs into the representative styles found in the target modality. Because the photos are labeled, the transformed photos retain their original labels. We consider two techniques for style transfer. While they are comparable in terms of perceptual quality overall,

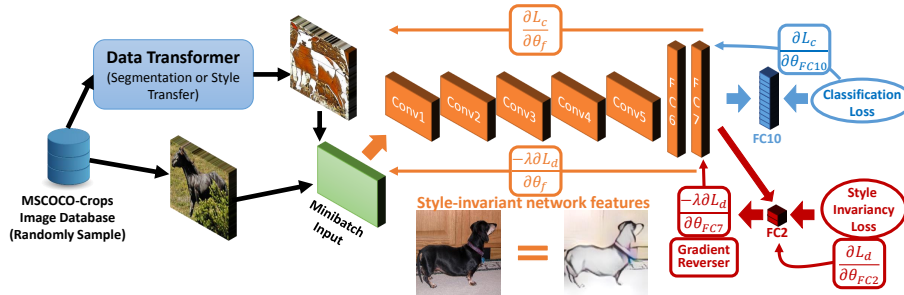


Fig. 2. Training with multiple modalities and style-invariance constraint. We train networks on real and synthetic data. We show an example of style transfer transforming photos into labeled synthetic cartoons. The style-invariance loss trains the FC2 layer to predict which modality the image came from. During backpropagation, we reverse its gradient before propagating it to the layers used by both classifiers. This encourages those layers to learn style-invariant features.

one [19] produces much higher-quality sketches (see the supplementary file), but is computationally less efficient, which limits the number of transformations we can create. We enable this method to produce rich, albeit few, transformations using a novel representative *style selection* technique (Sec. 3.1). The second style transfer technique [17] can be used to create more transformations as it is more efficient, but it produces less stable results.

3.1 Style transfer via Johnson [19]

Johnson et al. [19] present an artistic style transfer technique which trains a feed-forward network to perform transfer towards *a single style*. Thus, performing transfer towards additional styles requires training multiple networks (which is computationally expensive). The style transformation network is trained to minimize two losses: the source loss \mathcal{S} , which preserves the content of a given source image, and the target loss \mathcal{T} , which preserves the style of a given (unlabeled) target image. Let $\Phi_i(I)$ be the neural activations of layer i of a pretrained VGG-16 [32] network on input image I . Let I_s denote a source domain image, I_t a target domain image, and \hat{I}_s a source domain image transformed by the style transfer network. The shape of $\Phi_i(I)$ is given by $C_i \times H_i \times W_i$, which represents the number of channels, height, and width of the activations of layer i respectively. The source loss \mathcal{S} is thus given by:

$$\mathcal{S}^\Phi = \sum_i \frac{1}{C_i H_i W_i} \left\| \Phi_i(\hat{I}_s) - \Phi_i(I_s) \right\|_2^2 \quad (1)$$

In practice, we use only a single VGG layer to compute our source loss, $i = \text{relu3_3}$, based on the experimental results of [19].

The target loss, on the other hand, maintains the “style” of the target image, preserving colors and textures. Style spatial invariance is achieved by taking




Fig. 3. Examples of COCO images transformed towards two modalities: cartoons (left) and sketches (right). Background information is lost, which might be useful in non-photorealistic modalities such as cartoons that may not even have a background.

the distance between the correlations of different filter responses, by computing the Gram matrix $G^{\Phi_i}(I)$, where the matrix entry $G_{a,b}^{\Phi_i}(I)$ is the inner product between the flattened activations of channel C_{i_a} and C_{i_b} from $\Phi_i(I)$. Our target loss is then just the normalized difference between the Gram matrices of the transformed source image and the target image:

$$\mathcal{T}^{\Phi} = \sum_i \left\| G^{\Phi_i}(\hat{I}_s) - G^{\Phi_i}(I_t) \right\|_F^2 \quad (2)$$

where F represents the Frobenius norm. Following [19]’s results, we use $i \in \{\text{relu1.2}, \text{relu2.2}, \text{relu3.3}, \text{relu4.3}\}$ for our target loss. The final loss formulation for our style transformation network is then $\mathcal{L} = \lambda_s \mathcal{S}^{\Phi} + \lambda_t \mathcal{T}^{\Phi}$, where $\lambda_s = 1$ and $\lambda_t = 5$. We omit the total variation regularizer term for simplicity.

Style selection. In order to capture the artistic style variation *within* our target dataset, we choose a set of representative style images from our dataset to represent clusters with the same artistic style. We run the target data through layers in the style loss network that are used in computing the style reconstruction loss. We then compute the Gram matrix on the features to remove spatial information. Because of the high dimensionality of the features, we perform PCA and select the first thousand components. We cluster the resulting vectors using k -means, with $k = 10$. From each “style cluster” we select the image closest to its centroid as its style representative. We then train one style transfer network for each style representative, as described above.

3.2 Style transfer via Huang [17]

Because style transfer via the previous approach requires training a network for every style representative, it is too computationally expensive for more than

a few representative styles. We thus also explore a second, more recent style transfer technique by Huang and Belongie [17]. The model follows an hourglass, encoder-decoder architecture. The key takeaway of the approach is that style transfer can be performed on the output of the encoder (i.e. in feature space) by transferring channel-wise mean and variance across spatial locations. Using the same notation as above, style transfer is performed by transforming the encoder’s output as follows, where $\widehat{\Phi}_i(I_s)$ represents the transformed source image features, and μ and σ represent the mean and variance across all spatial locations in the given feature map:

$$\widehat{\Phi}_i(I_s) = \sigma(\Phi_i(I_t)) \left(\frac{\Phi_i(I_s) - \mu(\Phi_i(I_s))}{\sigma(\Phi_i(I_s))} \right) + \mu(\Phi_i(I_t)) \quad (3)$$

Notice that after training the decoder reconstruction network (the encoder is a pre-trained VGG network, as in Johnson [19]), performing style transfer towards any arbitrary image only requires one to compute $\Phi_i(I_s)$ and $\Phi_i(I_t)$, transform the source image features, and provide $\widehat{\Phi}_i(I_s)$ to the decoder to get the style-transferred image. This enables style transfer to be performed towards any arbitrary target image in a feedforward, efficient manner, without any retraining.

3.3 Creating the synthetic source modality

Formally, let $\mathbf{I}_s = \{I_s^i, y_s^i\}_{i=1}^{N_s}$ denote the dataset of labeled source data and $\mathbf{I}_t = \{I_t^i\}_{i=1}^{N_t}$ the dataset of unlabeled target data, where N_s and N_t represent the number of images in each dataset. Our style transfer network $\Psi(I_s^i, \theta_t^j) \rightarrow \widehat{I}_s^i$, transforms each source domain image to a synthetic image \widehat{I}_s^i which appears in the style of the single target image j captured by the parameters θ_t^j . In the case of Johnson’s method [19], we train one style-transfer network for each of ten “style representatives,” and apply each learned network to modify all source images. In the case of Huang’s method [17], we randomly select ten target images *for each source image* I_s^i as our target styles. The transformed dataset is thus more stylistically diverse for our adaptation approach via Huang’s method. For each target image sampled for either method, we obtain a transformed dataset in the style of that target, $\widehat{\mathbf{I}}_s^j = \{\Psi(\mathbf{I}_s, \theta_t^j), \mathbf{y}_s\}$, where \mathbf{y}_s denote object labels from the source data. We emphasize the targets are only used for their style, *not their content* (i.e. not their object labels). Examples of style modification for cartoons and sketches are shown in Fig. 3.

3.4 Training with two modalities

We train convolutional neural networks which are exposed to two modalities of data. We modify the training procedure such that each minibatch passed to the network during training contains data from both real and synthetic domains. In order to explicitly control for style differences between the source and target

domains in our learning process, we impose a style-invariance constraint on the network. The constraint causes the network to learn a feature representation such that it is difficult to predict which of the two source modalities an image came from, given the network’s features. In other words, the network learns to *confuse or perform transfer between the source modalities*.

The key observation is that because we have deliberately constructed the synthetic modality such that it differs from the real modality *only in its artistic style*, in order to learn modality-indistinguishable features across both modalities, the network is forced to *factor-out* style from its representation and thus learn style-invariant features for recognition. A similar criterion is used in prior work [6,35] but to confuse the *source and target modalities*. This is an important distinction: because the source and target modalities may differ in several ways, learning modality-invariant features between them may cause the network to disable filters which ultimately are useful for classification on the target domain. In our case of source-to-source invariance, because our source domains contain the same images except in different styles, the network is only forced to disable filters which depend on artistic style. We illustrate our method in Fig. 2.

For each modality, we train a CNN using our source and transformed data to perform the following two classification tasks:

$$T\left(\left\{\mathbf{I}_s, \widehat{\mathbf{I}}_s\right\}; \theta_Y; \theta_F; \theta_D\right) \rightarrow \{\mathbf{y}, \mathbf{d}\} \quad (4)$$

where $\widehat{\mathbf{I}}_s$ denotes the collection of transformed data output by our style transfer networks. The first classification task (predicting \mathbf{y}) predicts the object class of the image using the parameters θ_Y . The second task is to predict from which modality a given image came from (denoted by \mathbf{d}) using the parameters θ_D . Both classifiers make their predictions using the intermediate network features θ_F . Thus, we seek the following parameters of the network: θ_F which *maximizes* the loss of the modality classifier while simultaneously minimizing the loss of the object classifier, θ_Y which *minimizes* the object classification loss, and θ_D which *minimizes* the modality classifier’s loss. Our optimization process seeks the saddle point of the following adversarial functional:

$$\min_{\theta_F, \theta_Y} \max_{\theta_D} \alpha \mathcal{L}_d\left(\left\{\mathbf{I}_s, \widehat{\mathbf{I}}_s\right\}, \mathbf{d}\right) + \beta \mathcal{L}_y\left(\left\{\mathbf{I}_s, \widehat{\mathbf{I}}_s\right\}, \mathbf{y}\right) \quad (5)$$

where α and β are weighting parameters. In practice, we set α to be 1/10 the value of β . We use a multinomial logistic loss to train both classifiers. \mathcal{L}_d represents our domain classification loss, while \mathcal{L}_y is the object classifier loss which trains our network to recognize the object categories of interest.

Implementation Details: We experiment exclusively with the AlexNet architecture in this paper, but we include results with ResNet in our supplementary. We train all CNNs using the Caffe [18] framework. We use a batch size of 256 for AlexNet and a learning rate of $\alpha = 1e-4$; higher learning rates caused training to diverge. We used SGD with momentum $\mu = 0.9$ and learning rate decay of $\gamma = 0.1$ every 25K iterations. We trained all networks for a maximum of 150K iterations, but found the loss typically converged much sooner.

4 Experimental Validation

4.1 Datasets

Most of our experiments focus on object recognition, using the **PACS** dataset of [21] and the **Sketchy Database** of [31] as our target domains. We complement these with our own dataset, described below. To test our method with a broader set of categories, we also test on **Castrejon** et al.’s [3] large dataset containing 205 scene categories in different modalities. Our method also uses a photorealistic modality as the original source modality; this is either the photo domain from the respective dataset, or COCO [23] (for our new dataset).

PACS [21] contains four domains, three artistic (sketches, cartoons, and paintings) and one photorealistic, which we use as the original, non-synthetic source domain. The dataset contains seven object categories (e.g. “dog”, “elephant”, “horse”, “guitar”) and 9,991 images in total. The authors demonstrate the much larger shift exhibited in PACS compared to prior adaptation datasets, by showing KL divergence and cross-domain accuracy loss statistics.

The Sketchy Database [31] contains 75,471 human-drawn sketches and accompanying photographs, from 125 common object categories in ImageNet.

Castrejon [3] extends the Places dataset [41] containing 205 scene categories by providing four new modalities: sketches, clipart, text descriptions, and spatial text images. We use the two *visual* modalities, sketches and clipart.

New dataset: CASPA. Since the PACS dataset [21] is small, we complemented it by assembling images in the same three modalities as PACS. We use this additional dataset to test our conclusions on more test data. We call the dataset CASPA: (Cartoons, Sketches, Paintings). We assembled a dataset of 5,047 cartoons by querying Google Image Search with each of the ten animal categories from COCO: “bear”, “bird”, “cat”, “cow”, “dog”, “elephant”, “giraffe”, “horse”, “sheep”, and “zebra”. We chose these categories because sketch data of these animal classes was available. We also collected a new dataset of how different painters paint objects. We downloaded 1,391 paintings¹ which cover eight of our ten categories (except “giraffe” and “zebra”), and annotated these images with *2,834 bounding boxes*. Note that another painting dataset exists [14] but it only contains four of the ten classes we consider. To maintain the same set of domains as those used in PACS, we also include 12,008 sketches from the Sketchy Database [31]. PACS uses seven categories from [31], while we use twenty categories and collapse them down into coarser categories (i.e. different types of birds became the coarse category “bird”). In total, the dataset contains 18,446 images, almost twice more than PACS’ 9,991. The dataset is available at http://www.cs.pitt.edu/~chris/artistic_objects/. As our photorealistic source domain, we use COCO [23]. Because images in COCO contain multiple objects, we use the provided bounding boxes to produce cropped images containing a single object, and eliminate crops which do not contain the ten animal categories listed above.

¹ from [http://www.arab-painting.com/pic/Oil Painting Styles on Canvas/Animals](http://www.arab-painting.com/pic/Oil%20Painting%20Styles%20on%20Canvas/Animals)

4.2 Baselines

We compare against several recent unsupervised domain adaptation methods; these do not require any labeled data on the target domain.

- GONG et al. [11] is a pre-CNN method which shows how statistical differences between features extracted from two domains can be corrected, by learning the geodesic flow between the two domains. We use FC7 features.
- GANIN and Lempitsky [6] make CNNs modality-invariant by learning features such that the CNN is unable to distinguish which modality the image came from, via a domain confusion loss. Since no labeled data is available from the target, we only use [6]’s domain confusion loss but not the classification loss.
- GHIFARY et al. [8] train an encoder-decoder architecture to compute the reconstruction loss with the target domain’s data, and use the resulting features to train a classifier on the labeled source data.
- LONG et al. [25] propose to train separate classifiers for the source and target domain which differ only by a residual function. The source and target classifiers can still preserve domain-specific cues useful for classifying in that domain, but the source classifier is expressed as a residual function of the target classifier, and its residual layers capture the source-only features.
- CASTREJON et al. [3] learn a representation across scenes, by fixing the higher layers of the network while the lower layers are allowed to learn a domain-dependent representation. However, unlike [3], we cannot learn modality-specific lower-layer features without target labels. We train on ImageNet and fine-tune the higher layers on our photo labeled data, but we skip the second part of [3]’s “Method A” since we do not have labels on the modality data; we instead use generic pretrained ImageNet lower-level features.
- BOUSMALIS et al. [1] use a generative adversarial network (GAN) to transform labeled source data to appear as if it were target data. The generator adapts source images to fool the discriminator, while the discriminator attempts to distinguish between the generator’s output and real target images. A classifier is jointly trained to classify the original and translated source images.

We also compare against PHOTO-ALEXNET, a standard AlexNet [20], trained on 1.3M ImageNet samples, fine-tuned on our photorealistic domain (COCO / Places, respectively). See our supplementary file for results using ResNet.

Note that we do not compare to the domain *generalization* approach in [21] since it assumes the availability of *multiple* human-labeled source modalities at training time, while we only require one. For BOUSMALIS et al. and GHIFARY et al. we use data generated by their methods instead of our style-transferred data, then proceed with training a network with two modalities and domain confusion loss. We found this made these methods more competitive.

4.3 Results

In Table 1, we show the results of our comparison to state-of-the-art unsupervised domain adaptation methods. We first show the result of *not* performing any

Method	Target Domain									
	PACS				CASPA				Sketchy	
	Paint	Cart	Sketch	AVG	Paint	Cart	Sketch	AVG	Sketch	
PHOTO-ALEXNET	0.560	0.276	0.328	0.388	0.663	0.222	0.398	0.428	0.093	
GONG [11]	0.487	0.310	0.305	0.367	0.309	0.200	0.463	0.324	0.112	
GANIN [6]	0.624	<i>0.638</i>	0.351	0.538	0.652	0.251	0.202	0.368	0.112	
GHIFARY [8]	0.453	0.561	0.371	0.462	0.631	0.388	0.418	0.479	0.225	
LONG [25]	0.614	0.668	0.417	0.566	0.628	<i>0.425</i>	0.497	0.517	<i>0.303</i>	
CASTREJON [3]	0.580	0.293	0.328	0.400	0.628	0.231	0.458	0.439	0.085	
BOUSMALIS [1]	0.609	0.472	0.559	0.547	0.666	0.408	0.482	0.519	0.284	
OURS-JOHNSON	0.624	0.485	<i>0.653</i>	<i>0.587</i>	<i>0.677</i>	0.406	0.625	0.569	0.326	
OURS-HUANG	0.619	0.480	0.689	0.596	0.698	0.464	<i>0.501</i>	<i>0.554</i>	0.234	
UPPER BOUND	0.863	0.927	0.921	0.904	0.842	0.741	0.917	0.833	0.822	

Table 1. Our method outperforms all other domain adaptation techniques on average. The best method is in **bold** (excluding upper bound), and the second-best in *italics*.

adaptation (PHOTO-ALEXNET). We then show the performance of six domain adaptation techniques (from 2012, 2015, 2016, and 2017). We next show the two versions of our method, using style transfer techniques from Johnson (Sec. 3.1) and Huang (Sec. 3.2). Finally, we show the performance of an UPPER BOUND method which “cheats” by having access to labeled target data at training time, which none of the baselines nor our methods see. We split the data from each modality into 90%/10% training/test sets (for the sake of the upper bound which requires target data). All results show top-1 accuracy.

Our results. We see that one of our methods is always the best performer per dataset (average columns for **PACS** and **CASPA**, and the single column for **Sketchy**), and the other is second-best in two of three cases. For cartoons and paintings, OURS-HUANG is the stronger of our two methods, or our two methods perform similarly. On the full **Sketchy** dataset, which **PACS-Sketch** and **CASPA-Sketch** are subsets of, OURS-JOHNSON performs much better (by about 40%). This is because [17] and thus OURS-HUANG has trouble generating high-quality sketch transfer. This tendency holds for the subset chosen in **CASPA** (as [19] is better than [17] by about 25%), but on the subset in **PACS**, [17] does somewhat better (6%). We believe this is due to the choice of categories in each sketch subset; we illustrate these results in our supplementary. Note **Sketchy** contains many more classes hence the lower performance overall.

While we observe *some* qualitative correlation between visual quality of the style transferred images and utility of the data for training our classifiers, we note that our key concern in this work is utility, not visual quality. We do observe, however, that in some cases loss of detail in synthetic sketches may be causing our classifiers trained on such data to struggle with classification, particularly for OURS-HUANG. We explore this problem in detail in our supplementary material.

Baseline results. Our style transfer approach outperforms all other domain adaptation methods on all types of data except **PACS-Painting** (where it ties

for best performance with GANIN), and on **PACS-Cartoon**. Our method outperforms the strongest baseline by 23% for **PACS-Sketch** (wrt BOUSMALIS), 5% for **CASPA-Painting** (wrt BOUSMALIS), 9% for **CASPA-Cartoon** (wrt LONG), and 25% for **CASPA-Sketch** (wrt LONG). On average, OURS-JOHNSON and OURS-HUANG outperform the best baseline by 4%/5% respectively on **PACS** (wrt LONG), and 10%/7% on **CASPA** (wrt LONG). OURS-JOHNSON outperforms LONG by 8% for **Sketchy**. To underscore the importance of style adaptation as opposed to naïve data transformations, we show using style transferred data outperforms edge maps (which are sketch-like) on sketches in our supp. file.

Interestingly, on **CASPA-Painting**, only three domain adaptation methods outperform the no-adaptation photo-only model, but our two are the only ones that do so by a significant margin. As shown in [21], paintings are closer to photographs than cartoons and sketches are, so it is not surprising that several domain adaptation methods fail to improve upon the source model’s performance. By explicitly accounting for the style of the target modality, we improve the photo model’s accuracy by 5%.

Data requirements. Overall, the strongest baseline is LONG. The most recent method, BOUSMALIS [1], performs second-best. The performance of this GAN model is limited by the amount of training data available. While [1] uses all unlabeled target samples, that set of target samples is too limited to adequately train a full generative adversarial network. As a point of comparison, in their original work, [1] use 1000 *labeled* target images just to verify their hyperparameters. This is 100 times the amount of target data we use for domain adaptation in OURS-JOHNSON. This result underscores one of the key strengths of our approach: its ability to explicitly distill and control for style while requiring little target data. Notice that for all domain adaptation methods, there is still a large gap between the unsupervised recognition performance, and the supervised upper bound’s performance.

Ablations. Our method is a framework for domain adaptation that explicitly accounts for the artistic style variations between domains. Within this framework, we show how to make the best use of limited target data, via the style selection technique in Sec. 3.1. In Sec. 3.4, we encourage style-invariant features to be learned, via a domain confusion loss over the source domains. We briefly demonstrate the benefit of these two techniques here; see the supplementary file for the full ablation study.

Style selection of ten diverse target style images has a great advantage over randomly choosing ten images for transfer. On **PACS-Sketches**, style selection achieves a 15% improvement over randomly choosing the style images (0.653 vs 0.569) and it achieves 14% improvement on **PACS-Cartoons** (0.485 vs 0.425). While style selection does require access to the target dataset to *choose* the representative styles from, our technique still significantly improves performance when a single set of ten images are selected at random (i.e. when the target data pool is limited to ten images); see our supp. for more details.

Encouraging style invariance via the domain confusion loss also boosts performance, by 28% (0.384 for style transfer using random targets and domain

Method	Castrejon-Clipart	Castrejon-Sketches
PHOTO-ALEXNET	0.0689	0.0213
LONG [25]	0.0727	0.0402
BOUSMALIS [1]	0.0839	<i>0.0464</i>
OURS-JOHNSON	<i>0.0847</i>	0.0479
OURS-HUANG	0.0885	0.0456
UPPER BOUND	0.5937	0.3120

Table 2. Scene classification results using [3]’s data.

confusion loss, vs 0.299 without that loss) for **CASPA-Cartoons**, and 16% for **CASPA-Sketches** (0.594 vs 0.513).

Results on Castrejon’s dataset. In Table 2, we evaluate our methods against the two strongest baselines from Table 1, this time on the scene classification task in [3]’s dataset. Note that unlike [3], we do not train on target labels, and we only train with 250K Places images. Performance is low overall on this 205-way task. While our methods’ performance is closely followed by BOUSMALIS, our methods perform best. In particular, OURS-JOHNSON outperforms both BOUSMALIS and LONG on both modalities. This experiment confirms our conclusions from Table 1 that explicitly controlling for style variation is key when the domain shift is caused by artistic style.

Supplementary file. Please see our supplementary for extensive additional experiments showing the performance of ablations of our method (comparing style selection vs. selecting random style images), methods that use a single source modality (including the synthetic modalities generated by our method), domain-agnostic style transfer using edge maps, a comparison of methods using the ResNet architecture, and a detailed exploration of the differences between the Huang [17] and Johnson [19] style transfer methods which help determine which method is appropriate for a given target domain.

5 Conclusion

We addressed the problem of training CNNs to recognize objects in unlabeled artistic modalities. We found the introduction of style-mimicking training modalities generated for free results in a domain-invariant representation. We confirm in extensive experiments that our method outperforms a variety of state-of-the-art and standard baselines. Compared to recent works which create synthetic training data, our method requires far less target data (as few as ten unlabeled target images) and is easier to train. We also release a large dataset in photographs and images showing objects across multiple artistic domains. Currently our method cannot handle large, exaggerated shape differences found in some cartoons or sketches. In future work, we will investigate making our domain adaptation method even more robust by performing more aggressive transformations, so we can recognize objects in even more abstract artistic renderings.

Acknowledgement: This material is based upon work supported by the National Science Foundation under NSF CISE Award No. 1566270.

References

1. Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., Krishnan, D.: Unsupervised pixel-level domain adaptation with generative adversarial networks. In: CVPR (IEEE) (2017)
2. Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., Erhan, D.: Domain separation networks. In: NIPS (2016)
3. Castrejon, L., Aytar, Y., Vondrick, C., Pirsiavash, H., Torralba, A.: Learning aligned cross-modal representations from weakly aligned data. In: CVPR (IEEE) (2016)
4. Chen, D., Liao, J., Yuan, L., Yu, N., Hua, G.: Coherent online video style transfer. In: ICCV (IEEE) (2017)
5. Chen, D., Yuan, L., Liao, J., Yu, N., Hua, G.: Stylebank: An explicit representation for neural image style transfer. In: CVPR (IEEE) (2017)
6. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: ICML (2015)
7. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: CVPR (IEEE) (2016)
8. Ghifary, M., Kleijn, W.B., Zhang, M., Balduzzi, D., Li, W.: Deep reconstruction-classification networks for unsupervised domain adaptation. In: ECCV (Springer) (2016)
9. Ginosar, S., Haas, D., Brown, T., Malik, J.: Detecting people in cubist art. In: ECCV Workshops (Springer) (2014)
10. Gong, B., Grauman, K., Sha, F.: Reshaping visual datasets for domain adaptation. In: NIPS (2013)
11. Gong, B., Shi, Y., Sha, F., Grauman, K.: Geodesic flow kernel for unsupervised domain adaptation. In: CVPR (IEEE) (2012)
12. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS (2014)
13. Gupta, S., Hoffman, J., Malik, J.: Cross modal distillation for supervision transfer. In: CVPR (IEEE) (2016)
14. Hall, P., Cai, H., Wu, Q., Corradi, T.: Cross-depiction problem: Recognition and synthesis of photographs and artwork. *Computational Visual Media* (Springer) **1**(2), 91–103 (2015)
15. Hoffman, J., Tzeng, E., Park, T., Zhu, J., Isola, P., Saenko, K., Efros, A.A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: ICML (2018)
16. Huang, H., Wang, H., Luo, W., Ma, L., Jiang, W., Zhu, X., Li, Z., Liu, W.: Real-time neural style transfer for videos. In: CVPR (IEEE) (2017)
17. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV (IEEE) (2017)
18. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding. In: Proceedings of the ACM International Conference on Multimedia (2014)
19. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV (Springer) (2016)

20. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
21. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: ICCV (IEEE) (2017)
22. Li, W., Xu, Z., Xu, D., Dai, D., Van Gool, L.: Domain generalization and adaptation using low rank exemplar svms. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017)
23. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (Springer) (2014)
24. Long, M., Cao, Y., Wang, J.: Learning transferable features with deep adaptation networks. In: ICML (2015)
25. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Unsupervised domain adaptation with residual transfer networks. In: NIPS (2016)
26. Muandet, K., Balduzzi, D., Schölkopf, B.: Domain generalization via invariant feature representation. In: ICML (2013)
27. Niu, L., Li, W., Xu, D.: Visual recognition by learning from web data: A weakly supervised domain generalization approach. In: CVPR (IEEE) (2015)
28. Oquab, M., Bottou, L., Laptev, I., Sivic, J.: Learning and transferring mid-level image representations using convolutional neural networks. In: CVPR (IEEE) (2014)
29. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* (Springer) (2015)
30. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. In: NIPS (2016)
31. Sangkloy, P., Burnell, N., Ham, C., Hays, J.: The sketchy database: learning to retrieve badly drawn bunnies. *ACM Transactions on Graphics (TOG)* (2016)
32. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)
33. Tenenbaum, J.B., Freeman, W.T.: Separating style and content. NIPS (1997)
34. Thomas, C., Kovashka, A.: Seeing behind the camera: Identifying the authorship of a photograph. In: CVPR (IEEE) (2016)
35. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous deep transfer across domains and tasks. In: ICCV (IEEE) (2015)
36. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: CVPR (IEEE) (2017)
37. Wilber, M.J., Fang, C., Jin, H., Hertzmann, A., Collomosse, J., Belongie, S.: Bam! the behance artistic media dataset for recognition beyond photography. In: CVPR (IEEE) (2017)
38. Wu, Q., Cai, H., Hall, P.: Learning graphs to model visual objects across different depictive styles. In: ECCV (Springer) (2014)
39. Xu, Z., Li, W., Niu, L., Xu, D.: Exploiting low-rank structure from latent domains for domain generalization. In: ECCV (Springer) (2014)
40. Zhang, X.Y., Liu, C.L.: Writer adaptation with style transfer mapping. *IEEE transactions on pattern analysis and machine intelligence* **35**(7), 1773–1787 (2013)
41. Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., Oliva, A.: Learning deep features for scene recognition using places database. In: NIPS (2014)