

From source to target and back: Symmetric Bi-Directional Adaptive GAN

Paolo Russo^{1,2}, Fabio M. Carlucci^{1,2}, Tatiana Tommasi² and Barbara Caputo^{1,2}

¹Department DIAG, Sapienza University of Rome, Italy

²Italian Institute of Technology

{Paolo.Russo, Fabio.Carlucci, Tatiana.Tommasi, Barbara.Caputo}@iit.it

Abstract

The effectiveness of GANs in producing images according to a specific visual domain has shown potential in unsupervised domain adaptation. Source labeled images have been modified to mimic target samples for training classifiers in the target domain, and inverse mappings from the target to the source domain have also been evaluated, without new image generation.

In this paper we aim at getting the best of both worlds by introducing a symmetric mapping among domains. We jointly optimize bi-directional image transformations combining them with target self-labeling. We define a new class consistency loss that aligns the generators in the two directions, imposing to preserve the class identity of an image passing through both domain mappings. A detailed analysis of the reconstructed images, a thorough ablation study and extensive experiments on six different settings confirm the power of our approach.

1. Introduction

The ability to generalize across domains is challenging when there is ample labeled data on which to train a deep network (source domain), but no annotated data for the target domain. To attack this issue, a wide array of methods have been proposed, most of them aiming at reducing the shift between the source and target distributions (see Sec. 2 for a review of previous work). An alternative is mapping the source data into the target domain, either by modifying the image representation [10] or by directly generating a new version of the source images [4]. Several authors proposed approaches that follow both these strategies by building over Generative Adversarial Networks (GANs) [13]. A similar but inverse method maps the target data into the source domain, where there is already an abundance of labeled images [39].

We argue that these two mapping directions should

not be alternative, but complementary. Indeed, the main ingredient for adaptation is the ability of transferring successfully the style of one domain to the images of the other. This, given a fixed generative architecture, will depend on the application: there may be cases where mapping from the source to the target is easier, and cases where it is true otherwise. By pursuing both directions in a unified architecture, we can obtain a system more robust and more general than previous adaptation algorithms.

With this idea in mind, we designed SBADA-GAN: Symmetric Bi-directional ADaptive Generative Adversarial Network. Its features are (see Figure 1):

- it exploits two generative adversarial losses that encourage the network to produce target-like images from the source samples and source-like images from the target samples. Moreover, it jointly minimizes two classification losses, one on the original source images and the other on the transformed target-like source images;
- it uses the source classifier to annotate the source-like transformed target images. Such pseudo-labels help regularizing the same classifier while improving the target-to-source generator model by backpropagation;
- it introduces a new semantic constraint on the source images: the *class consistency loss*. It imposes that by mapping source images towards the target domain and then again towards the source domain they should get back to their ground truth class. This last condition is less restrictive than a standard reconstruction loss [41, 17], as it deals only with the image annotation and not with the image appearance. Still, our experiments show that it is highly effective in aligning the domain mappings in the two directions;
- at test time the two trained classifiers are used respectively on the original target images and on their source-like transformed version. The two predictions are integrated to produce the final annotation.

Our architecture yields realistic image reconstructions while competing against previous state-of-the-art classifiers and exceeding them on four out of six different unsupervised adaptation settings. An ablation study showcasing the importance of each component in the architecture, and investigating the robustness with respect to its hyperparameters, sheds light on the inner workings of the approach, while providing further evidence of its value.

2. Related Work

GANs Generative Adversarial Networks are composed of two modules, a generator and a discriminator. The generator synthesizes samples whose distribution closely matches that of real data, while the discriminator distinguishes real from generated samples. GANs are agnostic to the training samples labels, while conditional GAN variants [24] exploit the class annotation as additional information to both the generator and the discriminator. Some works used multiple GANs: in CoGAN [21] two generators and two discriminators are coupled by weight-sharing to learn the joint distribution of images in two different domains without using pair-wise data. CycleGAN [41], Disco-GAN [17] and UNIT [20] encourage the mapping between two domains to be well covered by imposing transitivity: the mapping in one direction followed by the mapping in the opposite direction should arrive where it started. For this image generation process the main performance measure is either a human-based quality control or scores that evaluate the interpretability of the produced images by pre-existing models [31, 41].

Domain Adaptation A widely used strategy consists in minimizing the difference between the source and target distributions [38, 36, 7]. Alternative approaches minimize the errors in target samples reconstruction [12] or impose a consistency condition so that neighboring target samples assigned to different labels are penalized proportionally to their similarity [33]. Very recently, [15] proposed to enforce associations between source and target samples of the same ground truth or predicted class, while [30] assigned pseudo-labels to target samples using an asymmetric tri-training method.

Domain invariance can be also treated as a binary classification problem through an adversarial loss, which encourages mistakes in domain prediction [10]. For all the methods adopting this strategy, the described losses are minimized jointly with the main classification objective function on the source task, guiding the feature learning process towards a domain invariant representation. Only in [39] the two objectives are kept separated and recombined in a second step. In [5] the feature components that

differentiate two domains are modeled separately from those shared among them.

Image Generation for Domain Adaptation In the first style transfer methods [11, 16] new images were synthesized to maintain a specific content while replicating the style of one or a set of reference images. Similar approaches have been used to generate images with different visual domains. In [34] realistic samples were generated from synthetic images and the produced data could work as training set for a classification model on real images. [4] proposed a GAN-based approach that adapts source images to appear as if drawn from the target domain; the classifier trained on such data outperformed several domain adaptation methods. [37] introduced a method to generate source images that resemble the target ones, with the extra consistency constraint that the same transformation should keep the target samples identical. All these methods focus on the source-to-target image generation, not considering the inverse procedure, from target to source, which we show instead to be beneficial.

3. Method

Model We focus on unsupervised cross domain classification. Let us start from a dataset $\mathbf{X}_s = \{\mathbf{x}_s^i, \mathbf{y}_s^i\}_{i=0}^{N_s}$ drawn from a labeled source domain \mathcal{S} , and a dataset $\mathbf{X}_t = \{\mathbf{x}_t^j\}_{j=0}^{N_t}$ from a different unlabeled target domain \mathcal{T} , sharing the same set of categories. The task is to maximize the classification accuracy on \mathbf{X}_t while training on \mathbf{X}_s . To reduce the domain gap, we propose to adapt the source images such that they appear as sampled from the target domain by training a generator model G_{st} that maps any source samples \mathbf{x}_s^i to its target-like version $\mathbf{x}_{st}^i = G_{st}(\mathbf{x}_s^i)$ defining the set $\mathbf{X}_{st} = \{\mathbf{x}_{st}^i, \mathbf{y}_s^i\}_{i=0}^{N_s}$ (see Figure 1, bottom row). The model is also augmented with a discriminator D_t and a classifier C_t . The former takes as input the target images \mathbf{X}_t and target-like source transformed images \mathbf{X}_{st} , learning to recognize them as two different sets. The latter takes as input each of the transformed images \mathbf{x}_{st}^i and learns to assign its task-specific label \mathbf{y}_s^i . During the training procedure for this model, information about the domain recognition likelihood produced by D_t is used adversarially to guide and optimize the performance of the generator G_{st} . Similarly, the generator also benefits from backpropagation in the classifier training procedure.

Besides the source-to-target transformation, we also consider the inverse target-to-source direction by using a symmetric architecture (see Figure 1, top row). Here any target image \mathbf{x}_t^j is given as input to a generator model G_{ts} transforming it to its source-like version $\mathbf{x}_{ts}^j = G_{ts}(\mathbf{x}_t^j)$,

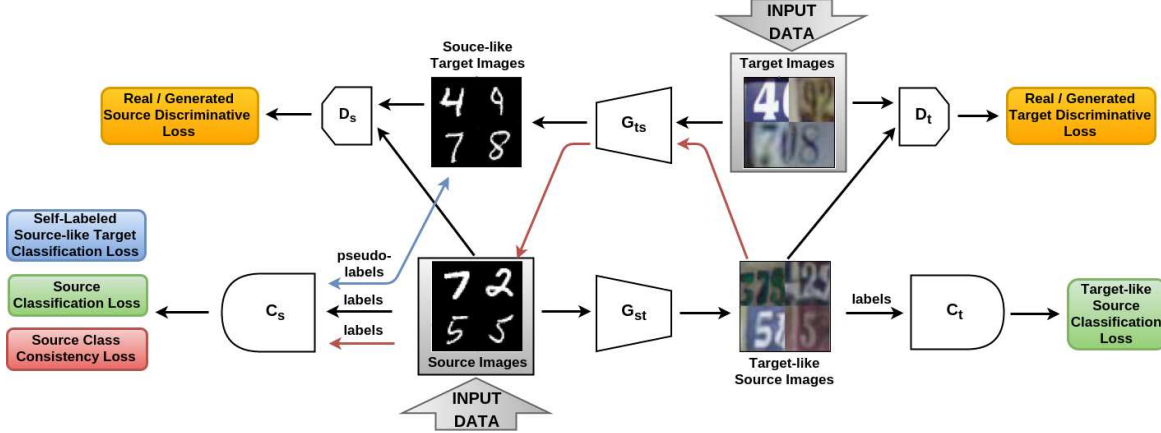


Figure 1: SBADA-GAN, training: the data flow starts from the Input Data arrows. The bottom and top row show respectively the source-to-target and target-to-source symmetric directions. The generative models G_{st} and G_{ts} transform the source images to the target domain and vice versa. D_s and D_t discriminate real from generated images of source and target. Finally the classifiers C_s and C_t are trained to recognize respectively the original source images and their target-like transformed versions. The bi-directional blue arrow indicates that the source-like target images are automatically annotated and the assigned pseudo-labels are re-used by the classifier C_s . The red arrows describe the class consistency condition by which source images transformed to the target domain through G_{st} and back to the source domain through G_{ts} should maintain their ground truth label.

defining the set $\mathbf{X}_{ts} = \{\mathbf{x}_{ts}^j\}_{j=0}^{N_t}$. As before, the model is augmented with a discriminator D_s which takes as input both \mathbf{X}_{ts} and \mathbf{X}_s and learns to recognize them as two different sets, adversarially helping the generator.

Since the target images are unlabeled, no classifier can be trained in the target-to-source direction as a further support for the generator model. We overcome this issue by *self-labeling* (see Figure 1, blue arrow). The original source data \mathbf{X}_s is used to train a classifier C_s . Once it has reached convergence, we apply the learned model to annotate each of the source-like transformed target images \mathbf{x}_{ts}^j . These samples, with the assigned pseudo-labels $y_{ts}^{jself} = \text{argmax}_y(C_s(G_{ts}(\mathbf{x}_t^j)))$, are then used transductively as input to C_s while information about the performance of the model on them is backpropagated to guide and improve the generator G_{ts} . Self-labeling has a long track record of success for domain adaptation: it proved to be effective both with shallow models [6, 14, 26], as well as with the most recent deep architectures [33, 38, 30]. In our case the classification loss on pseudo-labeled samples is combined with our other losses, which helps making sure we move towards the optimal solution: in case of a moderate domain shift, the correct pseudo-labels help to regularize the learning process, while in case of large domain shift, the possible mislabeled samples do not hinder the performance (see Sec. 4.5 for a detailed discussion on the experimental results).

Finally, the symmetry in the source-to-target and target-to-source transformations is enhanced by aligning the two generator models such that, when used in sequence, they bring a sample back to its starting point. Since our main focus is classification, we are interested in preserving the class identity of each sample rather than its overall appearance. Thus, instead of a standard image-based reconstruction condition we introduce a *class consistency* condition (see Figure 1, red arrows). Specifically, we impose that any source image \mathbf{x}_s^i adapted to the target domain through $G_{st}(\mathbf{x}_s^i)$ and transformed back towards the source domain through $G_{ts}(G_{st}(\mathbf{x}_s^i))$ is correctly classified by C_s . This condition helps by imposing a further joint optimization of the two generators.

Learning Here we formalize the description above. To begin with, we specify that the generators take as input a noise vector $\mathbf{z} \in \mathcal{N}(0, 1)$ besides the images, this allows some extra degree of freedom to model external variations. We also better define the discriminators as $D_s(\mathbf{x})$, $D_t(\mathbf{x})$ and the classifiers as $C_s(\mathbf{x})$, $C_t(\mathbf{x})$. Of course each of these models depends from its parameters but we do not explicitly indicate them to simplify the notation. For the same reason we also drop the superscripts i, j .

The source-to-target part of the network optimizes the following objective function:

$$\min_{G_{st}, C_t} \max_{D_t} \alpha \mathcal{L}_{D_t}(D_t, G_{st}) + \beta \mathcal{L}_{C_t}(G_{st}, C_t), \quad (1)$$

where the classification loss \mathcal{L}_{C_t} is a standard *softmax cross-entropy*

$$\mathcal{L}_{C_t}(G_{st}, C_t) = \mathbb{E}_{\substack{\{\mathbf{x}_s, \mathbf{y}_s\} \sim S \\ \mathbf{z}_s \sim \text{noise}}} [-\mathbf{y}_s \cdot \log(\hat{\mathbf{y}}_s)], \quad (2)$$

evaluated on the source samples transformed by the generator G_{st} , so that $\hat{\mathbf{y}}_s = C_t(G_{st}(\mathbf{x}_s, \mathbf{z}_s))$ and \mathbf{y}_s is the one-hot encoding of the class label y_s . For the discriminator, instead of the less robust binary cross-entropy, we followed [23] and chose a *least square* loss:

$$\mathcal{L}_{D_t}(D_t, G_{st}) = \mathbb{E}_{\mathbf{x}_t \sim T} [(D_t(\mathbf{x}_t) - 1)^2] + \mathbb{E}_{\substack{\mathbf{x}_s \sim S \\ \mathbf{z}_s \sim \text{noise}}} [(D_t(G_{st}(\mathbf{x}_s, \mathbf{z}_s)))^2]. \quad (3)$$

The objective function for the target-to-source part of the network is:

$$\min_{G_{ts}, C_s} \max_{D_s} \gamma \mathcal{L}_{D_s}(D_s, G_{ts}) + \mu \mathcal{L}_{C_s}(C_s) + \eta \mathcal{L}_{self}(G_{ts}, C_s), \quad (4)$$

where the discriminative loss is analogous to eq. (3), while the classification loss is analogous to eq. (2) but evaluated on the original source samples with $\hat{\mathbf{y}}_s = C_s(\mathbf{x}_s)$, thus it neither has any dependence on the generator that transforms the target samples G_{ts} , nor it provides feedback to it. The *self* loss is again a classification softmax cross-entropy:

$$\mathcal{L}_{self}(G_{ts}, C_s) = \mathbb{E}_{\substack{\{\mathbf{x}_t, \mathbf{y}_{t_{self}}\} \sim \mathcal{T} \\ \mathbf{z}_t \sim \text{noise}}} [-\mathbf{y}_{t_{self}} \cdot \log(\hat{\mathbf{y}}_{t_{self}})], \quad (5)$$

where $\hat{\mathbf{y}}_{t_{self}} = C_s(G_{ts}(\mathbf{x}_t, \mathbf{z}_t))$ and $\mathbf{y}_{t_{self}}$ is the one-hot vector encoding of the assigned label $y_{t_{self}}$. This loss back-propagates to the generator G_{ts} which is encouraged to preserve the annotated category within the transformation.

Finally, we developed a novel *class consistency* loss by minimizing the error of the classifier C_s when applied on the concatenated transformation of G_{ts} and G_{st} to produce $\hat{\mathbf{y}}_{cons} = (C_s(G_{ts}(G_{st}(\mathbf{x}_s, \mathbf{z}_s), \mathbf{z}_t)))$:

$$\mathcal{L}_{cons}(G_{ts}, G_{st}, C_s) = \mathbb{E}_{\substack{\{\mathbf{x}_s, \mathbf{y}_s\} \sim S \\ \mathbf{z}_s, \mathbf{z}_t \sim \text{noise}}} [-\mathbf{y}_s \cdot \log(\hat{\mathbf{y}}_{cons})]. \quad (6)$$

This loss has the important role of aligning the generators in the two directions and strongly connecting the two main parts of our architecture.

By collecting all the presented parts, we conclude with the complete SBADA-GAN loss:

$$\mathcal{L}_{SBADA-GAN}(G_{st}, G_{ts}, C_s, C_t, D_s, D_t) = \alpha \mathcal{L}_{D_t} + \beta \mathcal{L}_{C_t} + \gamma \mathcal{L}_{D_s} + \mu \mathcal{L}_{C_s} + \eta \mathcal{L}_{self} + \nu \mathcal{L}_{cons}. \quad (7)$$

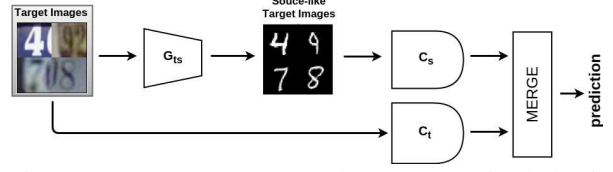


Figure 2: SBADA-GAN, test: the two pre-trained classifiers are applied respectively on the target images and on the transformed source-like target images. Their outputs are linearly combined for the final prediction.

Here $(\alpha, \beta, \gamma, \mu, \eta, \nu) \geq 0$ are weights that control the interaction of the loss terms. While the combination of six different losses might appear daunting, it is not unusual [5]. Here, it stems from the symmetric bi-directional nature of the overall architecture. Indeed each directional branch has three losses as it is custom practice in the GAN-based domain adaptation literature [39, 4]. Moreover, the ablation study reported in Sec. 4.5 indicates that the system is remarkably robust to changes in the hyperparameter values.

Testing The classifier C_t is trained on \mathbf{X}_{st} generated images that mimic the target domain style, and is then tested on the original target samples \mathbf{X}_t . The classifier C_s is trained on \mathbf{X}_s source data, and then tested on \mathbf{X}_{ts} samples, that are the target images modified to mimic the source domain style. These classifiers make mistakes of different type assigning also a different confidence rank to each of the possible labels. Overall the two classification models complement each other. We take advantage of this with a simple ensemble method $\sigma C_s(G_{ts}(\mathbf{x}_t, \mathbf{z}_t)) + \tau C_t(\mathbf{x}_t)$ which linearly combines their probability output, providing a further gain in performance. A schematic illustration of the testing procedure is shown in Figure 2. We set the combination weights σ, τ through cross validation (see Sec. 4.2 for further details).

4. Evaluation

4.1. Datasets and Adaptation Scenarios

We evaluate SBADA-GAN on several unsupervised adaptation scenarios, considering the following widely used datasets and settings:

MNIST \rightarrow MNIST-M: MNIST [19] contains centered, 28×28 pixel, grayscale images of single digit numbers on a black background, while MNIST-M [10] is a variant where the background is substituted by a randomly extracted patch obtained from color photos of BSDS500 [3]. We follow the evaluation protocol of [5, 4, 10].

	MNIST→USPS	USPS→MNIST	MNIST→MNIST-M	SVHN→MNIST	MNIST→SVHN	Synth Signs→GTSRB
Source Only	78.9	57.1 ± 1.7	63.6	60.1 ± 1.1	26.0 ± 1.2	79.0
CORAL [36]	81.7	-	57.7	63.1	-	86.9
MMD [38]	81.1	-	76.9	71.1	-	91.1
DANN [10]	85.1	73.0 ± 2.0	77.4	73.9	35.7	88.7
DSN [5]	91.3	-	83.2	82.7	-	93.1
CoGAN [21]	91.2	89.1 ± 0.8	62.0	not conv.	-	-
ADDA [39]	89.4 ± 0.2	90.1 ± 0.8	-	76.0 ± 1.8	-	-
DRCN [12]	91.8 ± 0.1	73.7 ± 0.1	-	82.0 ± 0.2	40.1 ± 0.1	-
PixelDA [4]	95.9	-	98.2	-	-	-
DTN [37]	-	-	-	84.4	-	-
TRUDA [33]	-	-	86.7	78.8	40.3	-
ATT [30]	-	-	94.2	86.2	52.8	96.2
UNIT [20]	95.9	93.5	-	90.5	-	-
DA _{ass} fix. par. [15]	-	-	89.5	95.7	-	82.8
DA _{ass} [15]	-	-	89.5	97.6	-	97.7
Target Only	96.5	99.2 ± 0.1	96.4	99.5	96.7	98.2
SBADA-GAN C_t	96.7	94.4	99.1	72.2	59.2	95.9
SBADA-GAN C_s	97.1	87.5	98.4	74.2	50.9	95.7
SBADA-GAN	97.6	95.0	99.4	76.1	61.1	96.7
GenToAdapt [32]	92.5 ± 0.7	90.8 ± 1.3	-	84.7 ± 0.9	36.4 ± 1.2	-
CyCADA [1]	94.8 ± 0.2	95.7 ± 0.2	-	88.3 ± 0.2	-	-
Self-Ensembling [2]	98.3 ± 0.1	99.5 ± 0.4	-	99.2 ± 0.3	42.0 ± 5.7	98.3 ± 0.3

Table 1: Comparison against previous work. SBADA-GAN C_t, C_s reports respectively the accuracies produced by the classifier trained in the target domain space, and the results produced by training in the source domain space and testing on the target images mapped to this space. SBADA-GAN reports the results obtained by a weighted combination of the softmax outputs of these two classifiers. Note that all competitors convert SVHN to grayscale, while we deal with the more complex original RGB version. The last three rows report results from online available pre-print papers.

MNIST ↔ USPS: USPS [9] is a digit dataset automatically scanned from envelopes by the U.S. Postal Service containing a total of 9,298 16×16 pixel grayscale samples. The images are centered, normalized and show a broad range of font styles. We follow the evaluation protocol of [4].

SVHN ↔ MNIST: SVHN [27] is the challenging real-world Street View House Number dataset. It contains over 600k 32×32 pixel color samples. Besides presenting a great variety of shapes and textures, images from this dataset often contain extraneous numbers in addition to the labeled, centered one. Most previous works simplified the data by considering a grayscale version, instead we apply our method to the original RGB images. We resize the MNIST images to 32×32 pixels and use the protocol by [5, 12].

Synth Signs → GTSRB: the Synthetic Signs collection [25] contains 100k samples of common street signs obtained from Wikipedia and artificially transformed to

simulate various imaging conditions. The German Traffic Signs Recognition Benchmark (GTSRB, [35]) consists of 51, 839 cropped images of German traffic signs. Both databases contain samples from 43 classes, thus defining a larger classification task than that on the 10 digits. We adopt the protocol proposed in [15].

4.2. Implementation details

We composed SBADA-GAN starting from two symmetric GANs, each with an architecture¹ analogous to that used in [4]. The model is coded² in python and we ran all our experiments in the Keras framework [8]. We use the ADAM [18] optimizer with learning rates for the discriminator and the generator both set to 10^{-4} . The batch size is set to 32 and we train the model for 500 epochs not noticing any overfitting, which suggests that further epochs might be beneficial. The α and γ loss weights (discriminator losses) are set to 1, β and μ (clas-

¹See all the model details in the supplementary material.

²Code available at <https://github.com/engharat/SBADAGAN>

sifier losses) are set to 10, to prevent that generator from indirectly switching labels (for instance, transform 7’s into 1’s). The class consistency loss weight ν is set to 1. All training procedures start with the self-labeling loss weight, η , set to zero, as this loss hinders convergence until the classifier is fully trained. After the model converges (losses stop oscillating, usually after 250 epochs) η is set to 1 to further increase performance. Finally the parameters to combine the classifiers at test time are $\sigma \in [0, 0.1, 0.2, \dots, 1]$ and $\tau = (1 - \sigma)$ chosen on a validation set of 1000 random samples from the target in each different setting.

4.3. Quantitative Results

Table 1 shows results on our evaluation settings. The top of the table reports results by thirteen competing baselines published over the last two years. The Source-Only and Target-Only rows contain reference results corresponding to the no-adaptation case and to the target fully supervised case. For SBADA-GAN, besides the full method, we also report the accuracy obtained by the separate classifiers (C_s, C_t) before the linear combination.

SBADA-GAN improves over the state of the art in four out of six settings. In these cases the advantage with respect to its competitors is already visible in the separate C_s and C_t results and it increases when considering the full combination procedure. Moreover, the gain in performance of SBADA-GAN reaches up to +8 percentage points in the MNIST→SVHN experiment. This setting was disregarded in many previous works: differently from its inverse SVHN→MNIST, it requires a difficult adaptation from the grayscale handwritten digits domain to the widely variable and colorful street view house number domain. Thanks to its bi-directionality, SBADA-GAN leverages on the inverse target to source mapping to produce highly accuracy results.

Conversely, in the SVHN→MNIST case SBADA-GAN ranks eighth out of the thirteen baselines in terms of performance. Our accuracy is on par with ADDA’s [39]: the two approaches share the same classifier architecture, although the number of fully-connected neurons of SBADA-GAN is five time lower. Moreover, compared to DRCN [12], the classifiers of SBADA-GAN are shallower with a reduced number of convolutional layers. Overall here SBADA-GAN suffers of some typical drawbacks of GAN-based domain adaptation methods: although the style of a domain can be easily transferred in the raw pixel space, the generative process does not have any explicit constraint on reducing the overall data distribution shift as instead done by the alternative feature-based domain adaptation approaches. Thus, methods like

Setting	S	T map to S	S map to T	T
MNIST → USPS	0.206	0.219	0.106	0.102
MNIST → MNIST-M	0.206	0.207	0.035	0.032
MNIST → SVHN	0.206	0.292	0.027	0.012
Synth S. → GTSRB	0.105	0.136	0.128	0.154

Table 2: Dataset mean SSIM: this measure of data variability suggests that our method successfully generates images with not only the same style of a chosen domain, but also similar perceptual variability.

DA_{ass} [15], DTN [37] and DSN [5] deal better with the large domain gap of the SVHN→MNIST setting.

Finally, in the Synth Signs → GTSRB experiment, SBADA-GAN is just slightly worse than DA_{ass}, but outperforms all the other competing methods. The comparison remains in favor of SBADA-GAN when considering that its performance is robust to hyperparameter variations (see Sec. 4.5 for more details).

4.4. Qualitative Results

To complement the quantitative evaluation, we look at the quality of the images generated by SBADA-GAN. First, we see from Figure 3 how the generated images mimic the style of the chosen domain, even when going from the simple MNIST digits to the SVHN colorful house numbers.

Visually inspecting the data distribution before and after domain mapping defines a second qualitative evaluation metric. We use t-SNE [22] to project the data from their raw pixel space to a simplified 2D embedding. Figure 4 shows that the transformed dataset tends to replicate faithfully the distribution of the chosen final domain.

A further measure of the quality of the SBADA-GAN generators comes from the diversity of the produced images. Indeed, GAN’s generators may collapse and output a single prototype that maximally fools the discriminators. To evaluate the diversity of samples generated by SBADA-GAN we choose the Structural Similarity (SSIM, [40]) that correlates well with the human perceptual similarity judgments. Its values range between 0 and 1 with higher values corresponding to more similar images. We follow the same procedure used in [28] by randomly choosing 1000 pairs of generated images within a given class. We also repeat the evaluation over all the classes and calculate the average results. Table 2 shows the results of the mean SSIM metric and indicates that the SBADA-GAN generated images not only mimic the same style, but also successfully reproduce the variability of a chosen domain.



Figure 3: Examples of generated digits: we show the image transformation from the original domain to the paired one as indicated under every sub-figure. For each of the (a)-(h) cases, the original/generated images are in the top/bottom row.

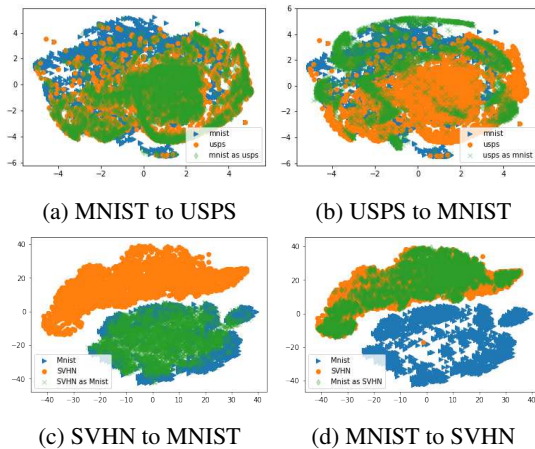


Figure 4: t-SNE visualization of source, target and source mapped to target images. Note how the mapped source covers faithfully the target space both in the (a),(b) case with moderated domain shift and in the more challenging (c),(d) setting.

4.5. Method Analysis

Ablation Study Starting from the core source-to-target GAN module we analyze the effect of adding all the other

S→T GAN		T→S GAN		Class Consist.	Self Label.	Accuracy
\mathcal{L}_{D_t}	\mathcal{L}_{C_t}	\mathcal{L}_{D_s}	\mathcal{L}_{C_s}	\mathcal{L}_{cons}	\mathcal{L}_{self}	MNIST→USPS
✓	✓					94.23
		✓	✓			91.55
✓	✓	✓	✓			94.90
✓	✓	✓	✓	✓		95.45
✓	✓	✓	✓	✓	✓	97.60

Table 3: Analysis of the role of each SBADA-GAN component. We ran experiments by turning on the different losses of the model as indicated by the checkmarks.




Figure 5: G_{ts} outputs (lower line) and their respective inputs (upper line) obtained with: (a) no consistency loss, (b) image-based cycle consistency loss [41, 17], (c) our class consistency loss. In (d) we show some real SVHN samples as a reference.

model parts. At first we add the symmetric target-to-source GAN model. These two parts are then combined and the domain transformation loop is closed by adding the class consistency condition. Finally the model is completed by introducing the target self-labeling procedure. We empirically test each of these model steps on the MNIST→USPS setting and report the results in Table 3. We see the gain achieved by progressively adding the different components, with the largest advantage obtained by the introduction of self-labeling.

An analogous boost due to self-labeling is also visible in all the other experimental settings with the exception of MNIST↔SVHN, where the accuracy remains unchanged if η is equal or larger than zero. A further analysis reveals that here the recognition accuracy of the source classifier applied to the source-like transformed target images is quite low (about 65%, while in all the other settings reaches 80 – 90%), thus the pseudo-labels cannot be considered reliable. Still, using them does not hinder the overall performance.

The crucial effect of the class consistency loss can be better observed by looking at the generated images and comparing them with those obtained in two alternative




Figure 6: CycleGAN [41] vs SBADA-GAN on the Amazon-Webcam experiment of the Office Dataset [29].

cases: setting $\nu = 0$, *i.e.* not using any consistency condition between the two generators G_{st} and G_{ts} , or substituting our class consistency loss with the standard cycle consistency loss [41, 17] based on image reconstruction. For this evaluation we choose the MNIST \rightarrow SVHN case which has the strongest domain shift and we show the generated images in Figure 5. When the consistency loss is not activated, the G_{ts} output images are realistic, but fail at reproducing the correct input digit and provide misleading information to the classifier. On the other hand, using the cycle-consistency loss preserves the input digit but fails in rendering a realistic sample in the correct domain style. Finally, our class consistency loss allows to preserve the distinct features belonging to a category while still leaving enough freedom to the generation process, thus it succeeds in both preserving the digits and rendering realistic samples.

CycleGAN vs SBADA-GAN To further clarify the difference between the two methods, we remind that CycleGAN is unsupervised and works only when transferring style across similarly shaped categories (*e.g.* horses \rightarrow zebras), not across domains. SBADA-GAN instead deals with domains containing multiple categories. The images samples in Figure 5(b) are indeed obtained with CycleGAN: training on them produces an accuracy of 25.5%, much lower than the corresponding 61.1% of SBADA-GAN. Moreover, CycleGAN has a single transformed image as output, while SBADA-GAN exploits a noise vector as input producing multiple outputs for each input image: this is critical for classification as it provides variability through data augmentation, it avoids overfitting and eases generalization. For completeness we also ran an experiment on the challenging Office Dataset [29]: here both the images produced by CycleGAN and SBADA-GAN (see Figure 6) are given as input to a pre-trained AlexNet and the classification accuracy is respectively 52.0% and 50.7%, both lower than the reference 61.6% result produced by the baseline without adaptation. These results confirm the known difficulty

of GAN-based method to deal with domain shifts due to poses and shapes.

Robustness Study SBADA-GAN is robust to the specific choice of the consistency loss weight ν , given that it is different from zero. Changing it in $[0.1, 1, 10]$ induces a maximum variation of 0.6 percentage points in accuracy over the different settings. An analogous evaluation performed on the classification loss weights (β, μ) reveals that changing them in the same range used for ν causes a maximum overall performance variation of 0.2 percentage points. Furthermore SBADA-GAN is minimally sensitive to the batch size used: halving it from 32 to 16 samples while keeping the same number of learning epochs reduces the performance only of about 0.2 percentage points. Such robustness is particularly relevant when compared to competing methods. For instance the most recent DA_{ass} [15] needs a perfectly balanced source and target distribution of classes in each batch, a condition difficult to satisfy in real world scenarios, and halving the originally large batch size reduces by 3.5 percentage points the final accuracy. Moreover, changing the weights of the losses that enforce associations across domains with a range analogous to that used for the SBADA-GAN parameters induces a drop in performance up to 16 accuracy percentage points³.

5. Conclusion

This paper presented SBADA-GAN, an adaptive adversarial domain adaptation architecture that maps simultaneously source samples into the target domain and vice versa with the aim to learn and use both classifiers at test time. To achieve this, self-labeling is exploited to regularize the classifier trained on the source, and we impose a class consistency loss that improves greatly the stability of the architecture, as well as the quality of the reconstructed images in both domains.

We explain the success of SBADA-GAN in several ways. To begin with, thanks to the the bi-directional mapping we avoid deciding a priori which is the best strategy for a specific task. Also, the combination of the two network directions improves performance providing empirical evidence that they are learning different, complementary features. Our class consistency loss aligns the image generators, allowing both domain transfers to influence each other. Finally the self-labeling procedure boost the performance in case of moderate domain shift without hindering it in case of large domain gaps.

Acknowledgments. This work was supported by the ERC grant 637076 - RoboExNovo.

³More details are provided in the supplementary material.

References

- [1] CyCADA: Cycle-Consistent Adversarial Domain Adaptation. In *Blind Submission, International Conference on Learning Representations (ICLR)*, 2018. 5
- [2] Self-ensembling for visual domain adaptation. In *Blind Submission, International Conference on Learning Representations (ICLR)*, 2018. 5
- [3] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):898–916, 2011. 4
- [4] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan. Unsupervised pixel-level domain adaptation with gans. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 4, 5
- [5] K. Bousmalis, G. Trigeorgis, N. Silberman, D. Krishnan, and D. Erhan. Domain Separation Networks. In *Neural Information Processing Systems (NIPS)*, 2016. 2, 4, 5, 6
- [6] L. Bruzzone and M. Marconcini. Domain adaptation problems: A dasvm classification technique and a circular validation strategy. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(5):770–787, 2010. 3
- [7] F. M. Carlucci, L. Porzi, B. Caputo, E. Ricci, and S. Rota Bulò. Autodial: Automatic domain alignment layers. In *International Conference on Computer Vision (ICCV)*, 2017. 2
- [8] F. Chollet. keras. <https://github.com/fchollet/keras>, 2017. 5
- [9] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001. 5
- [10] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. *J. Mach. Learn. Res.*, 17(1):2096–2030, 2016. 1, 2, 4, 5
- [11] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [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 (ECCV)*, 2016. 2, 5, 6
- [13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Neural Information Processing Systems (NIPS)*. 2014. 1
- [14] A. Habrard, J.-P. Peyrache, and M. Sebban. Iterative self-labeling domain adaptation for linear structured image classification. *International Journal on Artificial Intelligence Tools*, 22(5), 2013. 3
- [15] P. Haeusser, T. Frerix, A. Mordvintsev, and D. Cremers. Associative domain adaptation. In *International Conference on Computer Vision (ICCV)*, 2017. 2, 5, 6, 8
- [16] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [17] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim. Learning to discover cross-domain relations with generative adversarial networks. In *International Conference on Machine Learning, ICML*, pages 1857–1865, 2017. 1, 2, 7, 8
- [18] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 5
- [19] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 4
- [20] M.-Y. Liu, T. Breuel, and J. Kautz. Unsupervised image-to-image translation networks. In *Neural Information Processing Systems (NIPS)*, 2017. 2, 5
- [21] M.-Y. Liu and O. Tuzel. Coupled generative adversarial networks. In *Neural Information Processing Systems (NIPS)*. 2016. 2, 5
- [22] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008. 6
- [23] X. Mao, Q. Li, H. Xie, R. Y. Lau, and Z. Wang. Multi-class generative adversarial networks with the l2 loss function. *arXiv preprint arXiv:1611.04076*, 2016. 4
- [24] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*. 2
- [25] B. Moiseev, A. Konev, A. Chigorin, and A. Konushin. Evaluation of traffic sign recognition methods trained on synthetically generated data. In *International Conference on Advanced Concepts for Intelligent Vision Systems (ACIVS)*, 2013. 5
- [26] E. Morvant. Domain adaptation of weighted majority votes via perturbed variation-based self-labeling. *Pattern Recognition Letters*, 51:37–43, 2015. 3
- [27] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, page 5, 2011. 5
- [28] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier GANs. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 70, 2017. 6
- [29] K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting visual category models to new domains. In *European Conference on Computer Vision, (ECCV)*, 2010. 8
- [30] K. Saito, Y. Ushiku, and T. Harada. Asymmetric tri-training for unsupervised domain adaptation. In *International Conference on Machine Learning, (ICML)*, 2017. 2, 3, 5
- [31] T. Salimans, I. J. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *Neural Information Processing Systems (NIPS)*, 2016. 2

- [32] S. Sankaranarayanan, Y. Balaji, C. D. Castillo, and R. Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [33] O. Sener, H. O. Song, A. Saxena, and S. Savarese. Learning transferrable representations for unsupervised domain adaptation. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2110–2118. 2016. 2, 3, 5
- [34] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from simulated and unsupervised images through adversarial training. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [35] J. Stallkamp, M. Schlipsing, J. Salmen, and C. Igel. *The German traffic sign recognition benchmark: a multi-class classification competition*. IEEE, 2011. 5
- [36] B. Sun, J. Feng, and K. Saenko. Return of frustratingly easy domain adaptation. In *Conference of the Association for the Advancement of Artificial Intelligence (AAAI)*, 2016. 2, 5
- [37] Y. Taigman, A. Polyak, and L. Wolf. Unsupervised cross-domain image generation. In *International Conference on Learning Representations (ICLR)*, 2017. 2, 5, 6
- [38] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko. Simultaneous deep transfer across domains and tasks. In *International Conference in Computer Vision (ICCV)*, 2015. 2, 3, 5
- [39] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko. Adversarial discriminative domain adaptation. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 4, 5, 6
- [40] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *Trans. Img. Proc.*, 13(4):600–612, Apr. 2004. 6
- [41] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2, 7, 8