

AdaGraph: Unifying Predictive and Continuous Domain Adaptation through Graphs

Massimiliano Mancini^{1,2}, Samuel Rota Bulò³, Barbara Caputo^{4,5}, Elisa Ricci^{2,6}
¹Sapienza University of Rome, ²Fondazione Bruno Kessler, ³Mapillary Research,
⁴Politecnico di Torino, ⁵Italian Institute of Technology, ⁶University of Trento

mancini@diag.uniroma1.it, samuel@mapillary.com, barbara.caputo@polito.it, eliricci@fbk.eu

Abstract

The ability to categorize is a cornerstone of visual intelligence, and a key functionality for artificial, autonomous visual machines. This problem will never be solved without algorithms able to adapt and generalize across visual domains. Within the context of domain adaptation and generalization, this paper focuses on the predictive domain adaptation scenario, namely the case where no target data are available and the system has to learn to generalize from annotated source images plus unlabeled samples with associated metadata from auxiliary domains. Our contribution is the first deep architecture that tackles predictive domain adaptation and is able to leverage information brought by the auxiliary domains through a graph. Moreover, we present a simple yet effective strategy that allows us to take advantage of the incoming target data at test time, in a continuous domain adaptation scenario. Experiments on three benchmark databases support the value of our approach.

1. Introduction

Over the past years, deep learning has enabled rapid progress in many visual recognition tasks, even surpassing human performance [30]. While deep networks exhibit excellent generalization capabilities, previous studies [8] demonstrated that their performance drops when test data significantly differ from training samples. In other words deep models suffer from the *domain shift* problem, *i.e.* classifiers trained on *source* data do not perform well when tested on samples in the *target* domain. In practice, domain shift arises in many computer vision tasks, as many factors (*e.g.* lighting changes, different view-points, *etc.*) determine appearance variations in visual data.

To cope with this, several efforts focused on developing Domain Adaptation (DA) techniques [33], attempting to reduce the mismatch between source and target data distributions to learn accurate prediction models for the target domain. In the challenging case of unsupervised DA, only source data are labelled while no annotation is pro-

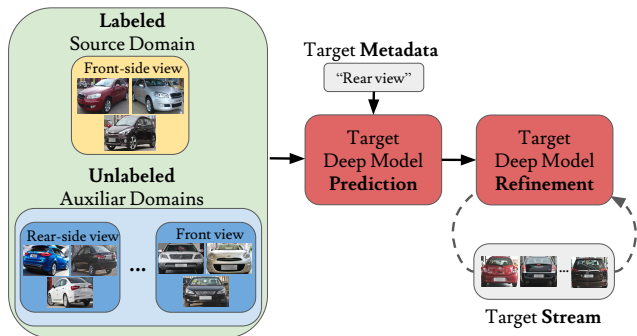


Figure 1. Predictive Domain Adaptation. During training we have access to a labeled source domain (yellow block) and a set of unlabeled auxiliary domains (blue blocks), all with associated metadata. At test time, given the metadata corresponding to the unknown target domain, we predict the parameters associated to the target model. This predicted model is further refined during test, while continuously receiving data of the target domain. Best viewed in color.

vided for target samples. Although it might be reasonable for some applications to have target samples available during training, it is hard to imagine that we can collect data for every possible target. More realistically, we aim for prediction models which can generalize to *new, previously unseen target domains*. Following this idea, previous studies proposed the Predictive Domain Adaptation (PDA) scenario [36], where neither the data, nor the labels from the target are available during training. Only annotated source samples are available, together with additional information from a set of auxiliary domains, in form of unlabeled samples and associated metadata (*e.g.* corresponding to the image timestamp or to camera pose, etc).

In this paper we introduce a deep architecture for PDA. Following recent advances in DA [3, 20, 23], we propose to learn a set of domain-specific models by considering a common backbone network with domain-specific alignment layers embedded into it. We also propose to exploit metadata and auxiliary samples by building a graph which explicitly describes the dependencies among do-

mains. Within the graph, nodes represent domains, while edges encode relations between domains, imposed by their metadata. Thanks to this construction, when metadata for the target domain are available at test time, the domain-specific model can be recovered. We further exploit target data directly at test time by devising an approach for continuously updating the deep network parameters once target samples are made available (Figure 1). We demonstrate the effectiveness of our method with experiments on three datasets: the *Comprehensive Cars* (CompCars) [35], the *Century of Portraits* [12] and the *CarEvolution* datasets [28], showing that our method outperforms state of the art PDA approaches. Finally, we show that the proposed approach for continuous updating of the network parameters can be used for continuous domain adaptation, producing more accurate predictions than previous methods [14, 19].

Contributions. To summarize, the contributions of this work are: (i) we propose the first deep architecture for addressing the problem of PDA; (ii) we present a strategy for injecting metadata information within a deep network architecture by encoding the relation between different domains through a graph; (iii) we propose a simple strategy for refining the predicted target model which exploits the incoming stream of target data directly at test time.

2. Related Works

Unsupervised Deep Domain Adaptation. Previous works on deep DA learn domain invariant representations by exploiting different architectures, such as Convolutional Neural Networks [21, 32, 10, 3, 20, 9, 1], deep autoencoders [37] or GANs [31, 15]. Some methods describe source and target features distributions considering their first and second order statistics and minimize their distance either defining an appropriate loss function [21] or deriving some domain alignment layers [20, 3, 9]. Other approaches rely on adversarial loss functions [32, 11] to learn domain agnostic representations. GAN-based techniques [2, 15, 31] for unsupervised DA focus directly on images and aim at generating either target-like source images or source-like target images. Recent works also showed that considering both the transformation directions is highly beneficial [31].

In many applications multiple source domains may be available. This fact has motivated the study of multi-source DA algorithms [34, 23]. In [34] an adversarial learning framework for multi-source DA is proposed, inspired by [10]. A similar adversarial strategy is also exploited in [38]. In [23] a deep architecture is proposed to discover multiple latent source domains in order to improve the classification accuracy on target data.

Our work performs domain adaptation by embedding into a deep network domain-specific normalization layers as in [20, 3, 9, 29]. However, the design of our layers is different as they are required to guarantee a continuous update

of parameters and to exploit information from the domain graph. Our approach considers information from multiple domains at training time. However, instead of having labeled data from all source domains, we do not have annotations for samples of auxiliary data.

Finally, our work is linked to graph-based domain adaptation methods [6, 5]. Differently from these works however, in our approach a node does not represent a single sample but a whole domain and edges do not link semantically related samples but domains with related metadata.

Domain Adaptation without Target Data. In some applications, the assumption that target data are available during training does not hold. This calls for DA methods able to cope with the domain shift by exploiting either the stream of incoming target samples, or side information describing possible future target domains.

The first scenario is typically referred to as continuous [14] or online DA [22]. To address this problem, in [14] a manifold-based DA technique is employed to model an evolving target data distribution. In [19] Li *et al.* propose to sequentially update a low-rank exemplar SVM classifier as data of the target domain becomes available. In [17], the authors propose to extrapolate the target data dynamics within a reproducing kernel Hilbert space.

The second scenario corresponds to the problem of predictive DA tackled in this paper. PDA is introduced in [36], where a multivariate regression approach is described for learning a mapping between domain metadata and points in a Grassmanian manifold. Given this mapping and the metadata for the target domain, two different strategies are proposed to infer the target classifier.

Other closely related tasks are the problems of zero shot domain adaptation and domain generalization. In zero-shot domain adaptation (ZDDA) [27] the task is to learn a prediction model in the target domain under the assumption that task-relevant source-domain data and task-irrelevant dual-domain paired data are available. We highlight that the PDA problem is related, but different, from ZDDA. ZDDA assumes that the domain shift is known during training from the presence of data of a different task but with the same visual appearance of source and target domains, while in PDA metadata of auxiliary domains is the only available information, and the target metadata is received *only at test time*. For this reason, ZDDA is not applicable to a PDA scenario, and it cannot predict the classification model for a target domain given only the metadata.

Domain generalization methods [25, 18, 7, 24] attempt to learn domain-agnostic classification models by exploiting labeled source samples from multiple domains but without having access to target data. Similarly to Predictive DA in domain generalization, multiple datasets are available during training. However, in PDA data from auxiliary source domains are not labeled.

3. Method

3.1. Problem Formulation

Our goal is to produce a model that is able to accomplish a task in a *target* domain \mathcal{T} for which *no data* are available during training, neither labeled nor unlabeled. The only information we can exploit is a characterization of the content of the target domain in the form of *metadata* $m_{\mathcal{T}}$ plus a set of known domains \mathcal{K} , each of them having associated metadata. All domains in \mathcal{K} carry information about the task we want to accomplish in the target domain. In particular, since in this work we focus on classification tasks, we assume that images from the domains in \mathcal{K} and \mathcal{T} can be classified with semantic labels from a same set \mathcal{Y} . As opposed to standard DA scenarios, the target domain \mathcal{T} does not necessarily belong to the set of known domains \mathcal{K} . Also, we assume that \mathcal{K} can be partitioned into a *labeled source* domain \mathcal{S} and *N unlabeled auxiliary* domains $\mathcal{A} = \{A_1, \dots, A_N\}$.

In the specific, this paper focuses on predictive DA (PDA) problems aimed at regressing the target model parameters using data from the domains in \mathcal{K} . We achieve this objective by (i) interconnecting each domain in \mathcal{K} using the given domain metadata; (ii) building domain-specific models from the data available in each domain in \mathcal{K} ; (iii) exploiting the connection between the target domain and the domains in \mathcal{K} , inferred from the respective metadata, to regress the model for \mathcal{T} .

A schematic representation of the method is shown in Figure 2. We propose to use a graph because of its seamless ability to encode relationships within a set of elements (domains in our case). Moreover, it can be easily manipulated to include novel elements (such as the target domain \mathcal{T}).

3.2. AdaGraph: Graph-based Predictive DA

We model the dependencies between the various domains by instantiating a graph composed of nodes and edges. Each node represents a different domain and each edge measures the relatedness of two domains. Each edge of the graph is weighted, and the strength of the connection is computed as a function of the domain-specific metadata. At the same time, in order to extract one model for each available domain, we employ recent advances in domain adaptation involving the use of domain-specific batch-normalization layers [20, 4]. With the domain-specific models and the graph we are able to predict the parameters for a novel domain that lacks data by simply (i) instantiating a new node in the graph and (ii) propagating the parameters from nearby nodes, exploiting the graph connections.

Connecting domains through a graph. Let us denote the space of domains as \mathcal{D} and the space of metadata as \mathcal{M} . As stated in Section 3.1, in the PDA scenario, we have a set of known domains $\mathcal{K} = \{k_1, \dots, k_n\} \subset \mathcal{D}$ and a bijective mapping $\phi : \mathcal{D} \mapsto \mathcal{M}$ relating domains and metadata. For simplicity, we regard as *unknown* some metadata m that is

not associated to domains in \mathcal{K} , *i.e.* such that $\phi^{-1}(m) \notin \mathcal{K}$.

In this work we structure the domains as a *graph* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} \subset \mathcal{D}$ represents the set of vertices corresponding to domains and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ the set of edges, *i.e.* relations between domains. Initially the graph contains only the known domains so $\mathcal{V} = \mathcal{K}$. In addition, we define an edge weight $\omega : \mathcal{E} \rightarrow \mathbb{R}$ that measures the relation strength between two domains $(v_1, v_2) \in \mathcal{E}$ by computing a distance between the respective metadata, *i.e.*

$$\omega(v_1, v_2) = e^{-d(\phi(v_1), \phi(v_2))}, \quad (1)$$

where $d : \mathcal{M}^2 \rightarrow \mathbb{R}$ is a distance function on \mathcal{M} .

Let Θ be the space of possible model parameters and assume we have properly exploited the domain data from each domain in $k \in \mathcal{K}$ to learn a set of domain-specific models (we will detail this procedure in the next subsection). We can then define a mapping $\psi : \mathcal{K} \mapsto \Theta$, relating each domain to its set of domain-specific parameters. Given some metadata $m \in \mathcal{M}$ we can recover an associated set of parameters via the mapping $\psi \circ \phi^{-1}(m)$ provided that $\phi^{-1}(m) \in \mathcal{K}$. In order to deal with metadata that is unknown, we introduce the concept of *virtual* node. Basically, a virtual node $\tilde{v} \in \mathcal{V}$ is a domain for which no data are available but we have metadata \tilde{m} associated to it, namely $\tilde{m} = \phi(\tilde{v})$. For simplicity, let us directly consider the target domain \mathcal{T} . We have $\mathcal{T} \in \mathcal{D}$ and we know $\phi(\mathcal{T}) = m_{\mathcal{T}}$. Since no data of \mathcal{T} are available, we have no parameters that can be directly assigned to the domain. However, we can estimate parameters for \mathcal{T} by using the domain graph \mathcal{G} . Indeed, we can relate \mathcal{T} to other domains $v \in \mathcal{V}$ using $\omega(\mathcal{T}, v)$ defined in (1) by opportunely extending \mathcal{E} with new edges (\mathcal{T}, v) for all or some $v \in \mathcal{V}$ (*e.g.* we could connect all v that satisfy $\omega(\mathcal{T}, v) > \tau$ for some τ). The extended graph $\mathcal{G}' = (\mathcal{V} \cup \{\mathcal{T}\}, \mathcal{E}')$ with the additional node \mathcal{T} and the new edge set \mathcal{E}' can then be exploited to estimate parameters for \mathcal{T} by *propagating* the model parameters from nearby domains. Formally we regress the parameters $\hat{\theta}_{\mathcal{T}}$ through

$$\hat{\theta}_{\mathcal{T}} = \psi(\mathcal{T}) = \frac{\sum_{(\mathcal{T}, v) \in \mathcal{E}'} \omega(\mathcal{T}, v) \psi(v)}{\sum_{(\mathcal{T}, v) \in \mathcal{E}'} \omega(\mathcal{T}, v)}, \quad (2)$$

where we normalize the contribution of each edge by the sum of the weights of the edges connecting node \mathcal{T} . With this formula we are able to provide model parameters for the target domain \mathcal{T} and, in general, for any unknown domain by just exploiting the corresponding metadata.

We want to highlight that this strategy simply requires extending the graph with a virtual node \tilde{v} and computing the relative edges. While the relations of \tilde{v} with other domains can be inferred from given metadata, as in (1), there could be cases in which no metadata are available for the target domain. In such situations, we can still exploit the incoming target image x to build a probability distribution over

nodes in \mathcal{V} , in order to assign the new data point to a mixture of known domains. To this end, let us define $p(v|x)$ the conditional probability of an image $x \in \mathcal{X}$, where \mathcal{X} is the image space, to be associated with a domain $v \in \mathcal{V}$. From this probability distribution, we can infer the parameters of a classification model for x through:

$$\hat{\theta}_x = \sum_{v \in \mathcal{V}} p(v|x) \cdot \psi(v) \quad (3)$$

where $\psi(v)$ is well-defined for each node linked to a known domain, while it must be estimated with (2) for each virtual domain $\tilde{v} \in \mathcal{V}$ for which $p(\tilde{v}|x) > 0$.

In practice, the probability $p(v|x)$ is constructed from a metadata classifier f^m , trained on the available data, that provides a probability distribution over \mathcal{M} given x , which can be turned into a probability over \mathcal{D} through the inverse mapping ϕ^{-1} .

Extracting node specific models. We have described how to regress model parameters for an unknown domain by exploiting the domain graph. Now, we focus on the actual problem of training domain-specific models using data available from the known domains \mathcal{K} . Since \mathcal{K} entails a labeled source domain \mathcal{S} and a set of auxiliary domains \mathcal{A} , we cannot simply train independent models with data from each available domain due to the lack of supervision on domains in \mathcal{A} for the target classification task. For this reason, we need to estimate the model parameters for the unlabeled domains \mathcal{A} by exploiting DA techniques.

Recent works [20, 3, 4] have shown the effectiveness of applying domain-specific batch-normalization (DA_{BN}) layers to address domain adaptation tasks. In particular, these works rewrite each batch-normalization layer [16] (BN) of the network in order to take into account domain-specific statistics. Given a domain k , a DA_{BN} layer differs from standard BN by including domain-specific information:

$$\text{DA}_{\text{BN}}(x, k) = \gamma \cdot \frac{x - \mu_k}{\sqrt{\sigma_k^2 + \epsilon}} + \beta, \quad (4)$$

where the mean and variance statistics $\{\mu_k, \sigma_k\}$ are estimated from x conditioned on domain k , γ and β are learnable scale and bias parameters, respectively, and ϵ is a small constant used to avoid numerical instabilities. Notice that we have dropped the dependencies on spatial location and channel dimension for simplicity. The effectiveness of this simple DA approach is due to the fact that features of source and target domains are forced to be aligned to the same reference distribution, and this allows to implicitly address the domain shift problem.

In this work we exploit the same ideas to provide each node in the graph with its own BN statistics. At the same time, we depart from [20, 4] since we do not keep scale and bias parameters shared across the domains, but we include also them within the set of domain-specific parameters.

In this scenario, the set of parameters for a domain k $\psi(k) = \theta_k$ is composed of different parts. Formally for each domain we have $\psi(k) = \{\theta^a, \theta_k^s\}$, where θ^a holds the domain-agnostic components and θ_k^s the domain-specific ones. In our case θ^a comprises parameters from standard layers (*i.e.* the convolutional and fully connected layers of the architecture), while θ_k^s comprises parameters and statistics of the domain-specific BN layers.

We start by using the labeled source domain \mathcal{S} to estimate θ^a and initialize θ_k^s . In particular, we obtain $\theta_{\mathcal{S}}$ by minimizing the standard cross-entropy loss:

$$L(\theta_{\mathcal{S}}) = -\frac{1}{|\mathcal{S}|} \sum_{(x,y) \in \mathcal{S}} \log(f_{\theta_{\mathcal{S}}}(y; x)), \quad (5)$$

where $f_{\theta_{\mathcal{S}}}$ is the classification model relative to the source domain, with parameters $\theta_{\mathcal{S}}$.

To extract the domain-specific parameters θ_k^s for each $k \in \mathcal{K}$, we employ 2 steps: the first is a selective forward pass for estimating the domain-specific statistics while the second is the application of a loss to further refine the scale and bias parameters. Formally, we replace each BN layer in the network with a *GraphBN* counterpart (GBN), where the forward pass is defined as follows:

$$\text{GBN}(x, v) = \gamma_v \cdot \frac{x - \mu_v}{\sqrt{\sigma_v^2 + \epsilon}} + \beta_v. \quad (6)$$

Basically in a GBN layer, the set of BN parameters and statistics to apply is conditioned on the node/domain to which x belongs. During training, as for standard BN, we update the statistics by leveraging their estimate obtained from the current batch \mathcal{B} :

$$\hat{\mu}_v = \frac{1}{|\mathcal{B}_v|} \sum_{x \in \mathcal{B}_v} x \quad \text{and} \quad \hat{\sigma}_v^2 = \frac{1}{|\mathcal{B}_v|} \sum_{x \in \mathcal{B}_v} (x - \mu_v)^2, \quad (7)$$

where \mathcal{B}_v is the set of elements in the batch belonging to domain v . As for the scale and bias parameters, we optimize them by means of a loss on the model output. For the auxiliary domains, since the data are unlabeled, we use an entropy loss, while a cross-entropy loss is used for the source domain:

$$L(\Theta^s) = -\frac{1}{|\mathcal{S}|} \sum_{(x,y) \in \mathcal{S}} \log(f_{\theta_{\mathcal{S}}}(y; x)) - \lambda \cdot \sum_{A_i \in \mathcal{A}} \frac{1}{|A_i|} \sum_{x \in A_i} \sum_{y \in \mathcal{Y}} f_{\theta_{A_i}}(y; x) \log(f_{\theta_{A_i}}(y; x)), \quad (8)$$

where $\Theta^s = \{\theta_k^s | k \in \mathcal{K}\}$ represents the whole set of domain-specific parameters and λ is the trade off between the cross-entropy and the entropy loss.

While (8) allows to optimize the domain-specific scale and bias parameters, it does not take into account the presence of the relationship between the domains, as imposed

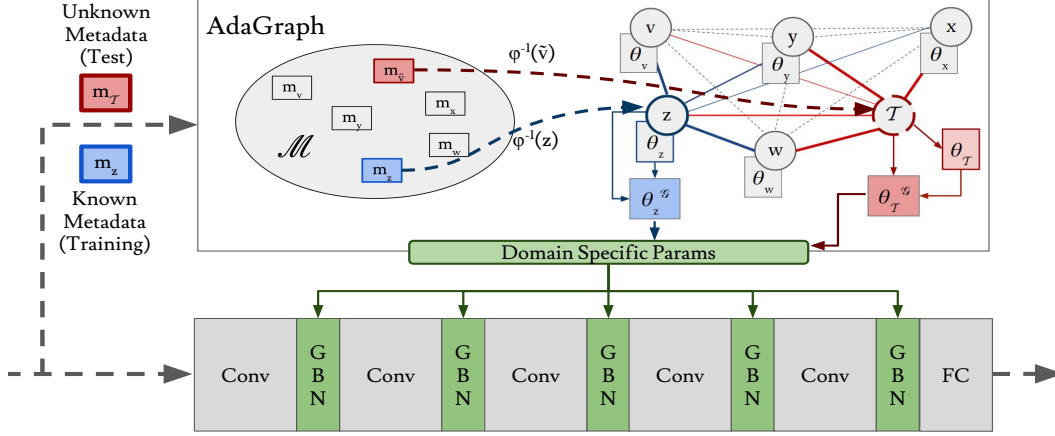


Figure 2. AdaGraph framework (Best viewed in color). Each BN layer is replaced by its GBN counterpart. The parameters used in a GBN layer are computed from a given metadata and the graph. Each domain in the graph (circles) contains its specific parameters (rectangular blocks). During the training phase (blue part), a metadata (*i.e.* m_z , blue block) is mapped to its domain (z). While the statistics of GBN are determined only by the one of z (θ_z), scale and bias are computed considering also the graph edges. During test, we receive the metadata for the target domain (m_T , red block) to which no node is linked. Thus we initialize T and we compute its parameters and statistics exploiting the connection with the other nodes in the graph (θ_T).

by the graph. A way to include the graph within the optimization procedure is to modify (6) as follows:

$$\text{GBN}(x, v, \mathcal{G}) = \gamma_v^{\mathcal{G}} \cdot \frac{x - \mu_v}{\sqrt{\sigma_v^2 + \epsilon}} + \beta_v^{\mathcal{G}}, \quad (9)$$

where we have

$$\nu_v^{\mathcal{G}} = \frac{\sum_{k \in \mathcal{K}} \omega(v, k) \cdot \nu_k}{\sum_{k \in \mathcal{K}} \omega(v, k)}, \quad (10)$$

for $\nu \in \{\beta, \gamma\}$. Basically, we use scale and bias parameters during the forward pass which are influenced by the graph edges, as described in (10).

Taking into account the presence of \mathcal{G} during the forward pass is beneficial for mainly two reasons. First, it allows to keep a consistency between how those parameters are computed at test time and how they are used at training time. Second, it allows to regularize the optimization of γ_v and β_v , which may be beneficial in cases where a domain contains few data. While the same procedure may be applied also for μ_v, σ_v , in our current design we avoid mixing them during training. This choice is related to the fact that each image belongs to a single domain and keeping the statistics separate allows to estimate them more precisely.

At test time, we initialize the domain-specific statistics and parameters of T given metadata m_T using (2), computing the forward pass of each GBN through (9). If no metadata are available, we compute the statistics and parameters through (3), performing the forward pass through (6). In Figure 2, we sketch the behaviour of our method given m_T both at training and test time.

3.3. Model Refinement through Joint Prediction and Adaptation

The approach described in the previous section allows to instantiate GBN parameters and statistics for a novel domain T given the target metadata m_T . However, without any sample of T , we have no way to assess how well the estimated statistics and parameters approximate the real ones of the target domain. This implies that we do not have the possibility to correct the parameters from a wrong initial estimates, a problem which may occur *e.g.* if we have noisy metadata. A possible strategy to solve this issue is to exploit the images we receive at test time to refine the GBN layers. To this extent, we propose a simple strategy for performing continuous domain adaptation [14] within AdaGraph.

Formally, let us define as $X_T = \{x_1, \dots, x_T\}$ the set of images of T that we receive at test time. Without loss of generality, we assume that the images of X_T are processed sequentially, one by one. Given the sequence X_T , our goal is to refine the statistics $\{\mu_T, \sigma_T\}$ and the parameters $\{\gamma_T, \beta_T\}$ of each GBN layer as new data arrives. Following recent works [22, 20], we continuously adapt a model to the target domain by feeding as input to the network batches of target images, updating the statistics as in standard BN. In order to achieve this, we store target samples in a buffer M . The buffer M has a fixed size and stores the samples one by one. Exploiting the buffer, we update the target statistics as follows:

$$\begin{aligned} \mu_T &\leftarrow (1 - \alpha) \cdot \mu_T + \alpha \cdot \mu_M \\ \sigma_T^2 &\leftarrow (1 - \alpha) \cdot \sigma_T^2 + \alpha \cdot \frac{|M|}{|M| - 1} \cdot \sigma_M^2, \end{aligned} \quad (11)$$

where $\{\mu_M, \sigma_M\}$ are computed through (7), replacing \mathcal{B}_v

with M :

$$\mu_M = \frac{1}{|M|} \sum_{x \in M} x \text{ and } \sigma_M^2 = \frac{1}{|M|} \sum_{x \in M} (x - \mu_M)^2. \quad (12)$$

While this allows to update the statistics, using (11) does not produce any refinement on $\{\gamma_{\mathcal{T}}, \beta_{\mathcal{T}}\}$. To this extent, we can easily employ the entropy term in (8):

$$L(\theta_{\mathcal{T}}) = -\frac{1}{|M|} \sum_{x \in M} \sum_{y \in \mathcal{Y}} f_{\theta_{\mathcal{T}}}(y; x) \log(f_{\theta_{\mathcal{T}}}(y; x)). \quad (13)$$

To summarize, with (11) and (13) we define a simple refinement procedure for AdaGraph which allows to recover from bad initialization of the predicted parameters and statistics. The update of statistics and parameters is performed together, each time the buffer is full. To avoid producing a bias during the refinement, we clear the buffer after each update step.

4. Experiments

4.1. Experimental setting

Datasets. We analyze the performance of our approach on three datasets: the *Comprehensive Cars* (CompCars) [35], the *Century of Portraits* [12] and the *CarEvolution* [28].

The *Comprehensive Cars* (CompCars) [35] dataset is a large-scale database composed of 136,726 images spanning a time range between 2004 and 2015. As in [36], we use a subset of 24,151 images with 4 types of cars (MPV, SUV, sedan and hatchback) produced between 2009 and 2014 and taken under 5 different view points (front, front-side, side, rear, rear-side). Considering each view point and each manufacturing year as a separate domain we have a total of 30 domains. As in [36] we use a PDA setting where 1 domain is considered as source, 1 as target and the remaining 28 as auxiliary sets, for a total of 870 experiments. In this scenario, the metadata are represented as vectors of two elements, one corresponding to the year and the other to the view point, encoding the latter as in [36].

Century of Portraits (Portraits) [12] is a large scale collection of images taken from American high school yearbooks. The portraits are taken over 108 years (1905-2013) across 26 states. We employ this dataset in a gender classification task, in two different settings. In the first setting we test our PDA model in a leave-one-out scenario, with a similar protocol to the tests on the *CompCars* dataset. In particular, to define domains we consider spatio-temporal information and we cluster images according to decades and to spatial regions (we use 6 USA regions, as defined in [12]). Filtering out the sets where there are less than 150 images, we obtain 40 domains, corresponding to 8 decades (from 1934 on) and 5 regions (*New England, Mid Atlantic, Mid West, Pacific, Southern*). We follow the same experimental protocol of the *CompCars* experiments, *i.e.* we use

one domain as source, one as target and the remaining 38 as auxiliaries. We encode the domain metadata as a vector of 3 elements, denoting the decade, the latitude (0 or 1, indicating north/south) and the east-west location (from 0 to 3), respectively. Additional details can be found in the supplementary material. In a second scenario, we use this dataset for assessing the performance of our continuous refinement strategy. In this case we employ all the portraits before 1950 as source samples and those after 1950 as target data.

CarEvolution [35] is composed of car images collected between 1972 and 2013. It contains 1008 images of cars produced by three different manufacturers with two car models each, following the evolution of the production of those models during the years. We choose this dataset in order to assess the effectiveness of our continuous domain adaptation strategy. A similar evaluation has been employed in recent works considering online DA [19]. As in [19], we consider the task of manufacturer prediction where there are three categories: Mercedes, BMW and Volkswagen. Images of cars before 1980 are considered as the source set and the remaining are used as target samples.

Networks and Training Protocols. To analyze the impact on performance of our main contributions we consider the ResNet-18 architecture [13] and perform experiments on the Portraits dataset. In particular, we apply our model by replacing each BN layer with its AdaGraph counterpart. We start with the network pretrained on ImageNet, training it for 1 epoch on the source dataset, employing Adam as optimizer with a weight decay of 10^{-6} and a batch-size of 16. We choose a learning rate of 10^{-3} for the classifier and 10^{-4} for the rest of the architecture. We train the network for 1 epoch on the union of source and auxiliary domains to extract domain-specific parameters. We keep the same optimizer and hyper-parameters except for the learning rates, decayed by a factor of 10. The batch size is kept to 16, but each batch is composed by elements of a single pair year-region belonging to one of the available domains (either auxiliary or source). The order of the pairs is randomly sampled within the set of allowed ones.

In order to fairly compare with previous methods we also consider Decaf features [8]. In particular, in the experiments on the CompCars dataset, we use Decaf features extracted at the `f_c7` layer. Similarly, for the experiments on CarEvolution, we follow [19] and use Decaf features extracted at the `f_c6` layer. In both cases, we apply our model by adding either a BN layer or our AdaGraph approach directly to the features, followed by a ReLU activation and a linear classifier. For these experiments we train the model on the source domain for 10 epochs using Adam as optimizer with a learning rate of 10^{-3} , a batch-size of 16 and a weight decay of 10^{-6} . The learning rate is decayed by a factor of 10 after 7 epochs. For CompCars, when training with the auxiliary set, we use the same optimizer, batch size

and weight decay, with a learning rate 10^{-4} for 1 epoch. Domain-specific batches are randomly sampled, as for the experiments on Portraits.

For all the experiments we use as distance measure $d(x, y) = \frac{1}{2\sigma} \cdot \|x - y\|_2^2$ with $\sigma = 0.1$ and set λ equal to 1.0, both in the training and in the refinement stage. At test time, we classify each input image as it arrives, performing the refinement step after the classification. The buffer size in the refinement phase is equal to 16 and we set $\alpha = 0.1$, the same used for updating the GBN components while training with the auxiliary domains.

We implemented¹ our method with the PyTorch [26] framework and our evaluation is performed using a NVIDIA GeForce 1080 Ti GTX GPU.

4.2. Results

In this section we report the results of our evaluation, showing both an empirical analysis of the proposed contributions and a comparison with state of the art approaches.

Analysis of AdaGraph. We first analyze the performance of our approach by employing the Portraits dataset. In particular, we evaluate the impact of (i) introducing a graph to predict the target domain BN statistics (*AdaGraph BN*), (ii) adding scale and bias parameters trained in isolation (*AdaGraph SB*) or jointly (*AdaGraph Full*) and (iii) adopting the proposed refinement strategy (*AdaGraph + Refinement*). As baseline² we consider the model trained only on the source domain and, as an upper bound, a corresponding DA method which is allowed to use target data during training. In our case, the upper bound corresponds to a model similar to the method proposed in [3].

The results of our ablation are reported in Table 1, where we report the average classification accuracy corresponding to two scenarios: *across decades* (considering the same region for source and target domains) and *across regions* (considering the same decade for source and target dataset). The first scenario corresponds to 280 experiments, while the second to 160 tests. As shown in the table, by simply replacing the statistics of BN layers of the source model with those predicted through our graph a large boost in accuracy is achieved (+4% in the *across decades* scenario and +2.4% in the *across regions* one). At the same time, estimating the scale and bias parameters without considering the graph is suboptimal. In fact there is a misalignment between the forward pass of the training phase (*i.e.* considering only domain-specific parameters) and how these parameters will be combined at test time (*i.e.* considering also the connection with the other nodes of the graph). Interestingly, in the *across regions* setting, our full model slightly drops in performance with respect to predicting only the BN statis-

Table 1. Portraits dataset. Ablation study.

Method	Across Decades	Across Regions
Baseline	82.3	89.2
AdaGraph BN	86.3	91.6
AdaGraph SB	86.0	90.5
AdaGraph Full	87.0	91.0
Baseline + Refinement	86.2	91.3
AdaGraph + Refinement	88.6	91.9
DA upper bound	89.1	92.1

tics. This is probably due to how regions are encoded in the metadata (*i.e.* considering geographical location), making it difficult to capture factors (*e.g.* cultural, historical) which can be more discriminative to characterize the population of a region or a state. However, as stated in Section 3.3, employing a continuous refinement strategy allows the method to compensate for prediction errors. As shown in Table 1, with a refinement step (*AdaGraph + Refinement*) the accuracy constantly increases, filling the gap between the performance of the initial model and our DA upper bound.

It is worth noting that applying the refinement procedure to the source model (*Baseline + Refinement*) leads to better performance (about +4% in the *across decades* scenario and +2.1% for *across regions* one). More importantly, the performance of the *Baseline + Refinement* method is always worse than what obtained by *AdaGraph + Refinement*, because our model provides, on average, a better starting point for the refinement procedure.

Figure 3 shows the results associated to the *across decades* scenario. Each bar plot corresponds to experiments where the target domain is associated to a specific year. As shown in the figure, on average, our full model outperforms both *AdaGraph BN* and *AdaGraph SB*, showing the benefit of the proposed graph strategy. The results in the figure clearly also show that our refinement strategy always leads to a boost in performance.

Comparison with the state of the art. Here we compare the performances of our model with state of the art PDA approaches. We use the CompCars dataset and we benchmark against the Multivariate Regression (MRG) methods proposed in [36].

We apply our model in the same setting as [36] and perform 870 different experiments, computing the average accuracy (Table 2). Our model outperforms the two methods proposed in [36] by improving the performances of the Baseline network by 4%. AdaGraph alone outperforms the Baseline model when it is updated with our refinement strategy and target data (*Baseline + Refinement*). When coupled with a refinement strategy, our graph-based model further improves the performances, filling the gap between AdaGraph and our DA upper bound. It is interesting to note that our model is also effective when there are no metadata available in the target domain. In the table, *AdaGraph (images)* corresponds to our approach when, instead of initializing the BN layer for the target exploiting metadata, we employ

¹The code is available at <https://github.com/mancinimassimiliano/adagraph>

²We do not report the results of previous approaches [36] since the code is not publicly available.



Figure 3. Portraits dataset: comparison of different models in the PDA scenario with respect to the average accuracy on a target decade, fixed the same region of source and target domains. The models are based on ResNet-18.

Table 2. CompCars dataset [35]. Comparison with state of the art.

Method	Avg. Accuracy
Baseline [36]	54.0
Baseline + BN	56.1
MRG-Direct [36]	58.1
MRG-Indirect [36]	58.2
AdaGraph (metadata)	60.1
AdaGraph (images)	60.8
Baseline + Refinement	59.5
AdaGraph + Refinement	60.9
DA upper bound	60.9

the current input image and a domain classifier to obtain a probability distribution over the graph nodes, as described in Section 3.2. The results in the Table show that *AdaGraph (images)* is more accurate than *AdaGraph (metadata)*.

Exploiting AdaGraph Refinement for Continuous Domain Adaptation. In Section 3.3, we have shown a way to boost the performances of our model by leveraging the stream of incoming target data and refine the estimates of the target BN statistics and parameters. Throughout the experimental section, we have also demonstrated how this strategy improves the target classification model, with performances close to DA methods which exploit target data during training.

In this section we show how this approach can be employed as a competitive method in the case of continuous domain adaptation [14]. We consider the CarEvolution dataset and compare the performances of our proposed strategy with two state of the art algorithms: the manifold-based adaptation method in [14] and the low-rank SVM strategy presented in [19]. As in [19] and [14], we apply our adaptation strategy after classifying each novel image and compute the overall accuracy. The images of the target domain are presented to the network in a chronological order *i.e.* from 1980 to 2013. The results are shown in Table 3. While the integration of a BN layer alone leads to better performances over the baseline, our refinement strategy produces an additional boost of about 3%. If scale and bias parameters are refined considering the entropy loss, accuracy further increases.

We also test the proposed model on a similar task con-

Table 3. CarEvolution [28]: comparison with state of the art.

Method	Accuracy
Baseline SVM [19]	39.7
Baseline + BN	43.7
CMA+GFK [14]	43.0
CMA+SA [14]	42.7
LLRESVM [19]	43.6
LLRESVM+EDA[19]	44.3
Baseline + Refinement Stats	46.5
Baseline + Refinement Full	47.3

Table 4. Portraits dataset [35]: performances of the refinement strategy on the continuous adaptation scenario

Method	Baseline	Refinement Stats	Refinement Full
Accuracy	81.9	87.3	88.1

sidering the Portraits dataset. The results of our experiments are shown in Table 4. Similarly to what observed on the previous experiments, continuously adapting our deep model as target data become available leads to better performance with respect to the baseline. The refinement of scale and bias parameters contributes to a further boost in accuracy.

5. Conclusions

We present the first deep architecture for Predictive Domain Adaptation. We leverage metadata information to build a graph where each node represents a domain, while the strength of an edge models the similarity among two domains according to their metadata. We then propose to exploit the graph for the purpose of DA and we design novel domain-alignment layers. This framework yields the new state of the art on standard PDA benchmarks. We further present an approach to exploit the stream of incoming target data such as to refine the target model. We show that this strategy itself is also an effective method for continuous DA, outperforming state of the art approaches. Future works will explore methodologies to incrementally update the graph and to automatically infer relations among domains, even in the absence of metadata.

Acknowledgements. We acknowledge financial support from ERC grant 637076 - RoboExNovo and project *DIGIMAP*, grant 860375, funded by the Austrian Research Promotion Agency (FFG). This work was carried out under the "Vision and Learning joint Laboratory" between FBK and UNITN.

References

- [1] Gabriele Angeletti, Barbara Caputo, and Tatiana Tommasi. Adaptive deep learning through visual domain localization. In *ICRA*, 2018. 2
- [2] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *NIPS*, 2016. 2
- [3] Fabio Maria Carlucci, Lorenzo Porzi, Barbara Caputo, Elisa Ricci, and Samuel Rota Bulò. Autodial: Automatic domain alignment layers. In *ICCV*, 2017. 1, 2, 4, 7
- [4] Fabio Maria Carlucci, Lorenzo Porzi, Barbara Caputo, Elisa Ricci, and Samuel Rota Bulò. Just dial: Domain alignment layers for unsupervised domain adaptation. In *ICIAP*, 2017. 3, 4
- [5] Debasmith Das and C. S. George Lee. Graph matching and pseudo-label guided deep unsupervised domain adaptation. In *ICANN*, 2018. 2
- [6] Zhengming Ding, Sheng Li, Ming Shao, and Yun Fu. Graph adaptive knowledge transfer for unsupervised domain adaptation. In *ECCV*, 2018. 2
- [7] Antonio D’Innocente and Barbara Caputo. Domain generalization with domain-specific aggregation modules. In *GCPR*, 2018. 2
- [8] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, 2014. 1, 6
- [9] Geoff French, Michal Mackiewicz, and Mark Fisher. Self-ensembling for visual domain adaptation. *ICLR*, 2018. 2
- [10] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. *ICML*, 2015. 2
- [11] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17(59):1–35, 2016. 2
- [12] Shiry Ginossar, Kate Rakelly, Sarah Sachs, Brian Yin, and Alexei A Efros. A century of portraits: A visual historical record of american high school yearbooks. In *ICCV-WS*, 2015. 2, 6
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [14] Judy Hoffman, Trevor Darrell, and Kate Saenko. Continuous manifold based adaptation for evolving visual domains. In *CVPR*, 2014. 2, 5, 8
- [15] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. 2
- [16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 4
- [17] Christoph H Lampert. Predicting the future behavior of a time-varying probability distribution. In *CVPR*, 2015. 2
- [18] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *ICCV*, 2017. 2
- [19] Wen Li, Zheng Xu, Dong Xu, Dengxin Dai, and Luc Van Gool. Domain generalization and adaptation using low rank exemplar svms. *IEEE T-PAMI*, 40(5):1114–1127, 2018. 2, 6, 8
- [20] Yanghao Li, Naiyan Wang, Jianping Shi, Xiaodi Hou, and Jiaying Liu. Adaptive batch normalization for practical domain adaptation. *Pattern Recognition*, 80:109–117, 2018. 1, 2, 3, 4, 5
- [21] Mingsheng Long and Jianmin Wang. Learning transferable features with deep adaptation networks. In *ICML*, 2015. 2
- [22] Massimiliano Mancini, Hakan Karaoguz, Elisa Ricci, Patric Jensfelt, and Barbara Caputo. Kitting in the wild through online domain adaptation. *IROS*, 2018. 2, 5
- [23] Massimiliano Mancini, Lorenzo Porzi, Samuel Rota Bulò, Barbara Caputo, and Elisa Ricci. Boosting domain adaptation by discovering latent domains. *CVPR*, 2018. 1, 2
- [24] Saeid Motiian, Marco Piccirilli, Donald A. Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *ICCV*, 2017. 2
- [25] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *ICML*, 2013. 2
- [26] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-WS*, 2017. 7
- [27] Kuan-Chuan Peng, Ziyang Wu, and Jan Ernst. Zero-shot deep domain adaptation. In *ECCV*, 2018. 2
- [28] Konstantinos Rematas, Basura Fernando, Tatiana Tommasi, and Tinne Tuytelaars. Does evolution cause a domain shift?, 2013. 2, 6, 8
- [29] Subhankar Roy, Aliaksandr Siarohin, Enver Sangineto, Samuel Rota Bulò, Nicu Sebe, and Elisa Ricci. Unsupervised domain adaptation using feature-whitening and consensus loss. In *CVPR*, June 2019. 2
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015. 1
- [31] Paolo Russo, Fabio Maria Carlucci, Tatiana Tommasi, and Barbara Caputo. From source to target and back: symmetric bi-directional adaptive gan. In *CVPR*, 2018. 2
- [32] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *ICCV*, 2015. 2
- [33] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 2018. 1
- [34] Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *CVPR*, 2018. 2
- [35] Linjie Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. A large-scale car dataset for fine-grained categorization and verification. In *CVPR*, 2015. 2, 6, 8
- [36] Yongxin Yang and Timothy M Hospedales. Multivariate regression on the grassmannian for predicting novel domains. In *CVPR*, 2016. 1, 2, 6, 7, 8

- [37] Xingyu Zeng, Wanli Ouyang, Meng Wang, and Xiaogang Wang. Deep learning of scene-specific classifier for pedestrian detection. In *ECCV*, 2014. 2
- [38] Han Zhao, Shanghang Zhang, Guanhong Wu, Joao P. Costeira, José M. F. Moura, and Geoffrey J. Gordon. Multiple source domain adaptation with adversarial learning. In *ICLR-WS*, 2018. 2