
Node-Adapt, Path-Adapt and Tree-Adapt: Model-Transfer Domain Adaptation for Random Forest

Azadeh S. Mozafari
Department of Computer Engineering
Sharif University of Technology
Tehran, Iran
amozaffari@ce.sharif.edu

David Vazquez
Computer Vision Center
Universitat Autònoma de Barcelona
Barcelona, Spain
dvazquez@cvc.uab.es

Mansour Jamzad
Department of Computer Engineering
Sharif University of Technology
Tehran, Iran
jamzad@sharif.edu

Antonio M. Lopez
Computer Vision Center
Universitat Autònoma de Barcelona
Barcelona, Spain
antonio@cvc.uab.es

Abstract

Random Forest (RF) is a successful paradigm for learning classifiers due to its ability to learn from large feature spaces and seamlessly integrate multi-class classification, as well as the achieved accuracy and processing efficiency. However, as many other classifiers, RF requires domain adaptation (DA) provided that there is a mismatch between the training (source) and testing (target) domains which provokes classification degradation. Consequently, different RF-DA methods have been proposed, which not only require target-domain samples but revisiting the source-domain ones, too. As novelty, we propose three inherently different methods (Node-Adapt, Path-Adapt and Tree-Adapt) that only require the learned source-domain RF and a relatively few target-domain samples for DA, i.e. source-domain samples do not need to be available. To assess the performance of our proposals we focus on image-based object detection, using the *pedestrian detection* problem as challenging proof-of-concept. Moreover, we use the RF with *expert* nodes because it is a competitive *patch-based* pedestrian model. We test our Node-, Path- and Tree-Adapt methods in standard benchmarks, showing that DA is largely achieved.

1 Introduction

Random Forest (RF) is a paradigm that allows to learn an ensemble of randomized decision-tree classifiers [1, 2]. RF is very successful due to its ability to learn from large feature spaces and seamlessly integrate multi-class classification, as well as the accuracy and processing efficiency achieved by the ensemble. Note that the evaluation of the trees can be parallelized (their training too) or a sequential soft-cascade used [3]; moreover, evaluating a relatively well-balanced tree shows logarithmic complexity with the number of nodes. We can find recent paradigmatic applications such as its use for pose recognition in depth images (Kinect) [4], or as classification layer in the GoogLeNet deep CNN architecture for large-scale image classification [5].

Usually, when learning any classifier, including RF, it is assumed that the training and testing samples follow the same probability distribution. However, in practice, there are application-dependent reasons that make this assumption to fail. The consequence is that the classifier performs worse than expected. In this case, we have to augment the machine learning formalism with the concepts of

source and *target* domains. The training samples are acquired in the former, and the testing ones in the latter. Then, the challenge is to define a domain adaptation (DA) method which allows to adapt the source-domain-trained classifier for performing well in the target domain, with a minimum effort. By *minimum effort* we refer to several assumptions. For instance, it is assumed little data labeling effort, i.e. either we have relatively few labeled target-domain samples to perform the (*supervised*) DA, or we have eventually many target-domain samples but without labels (*unsupervised* DA). Moreover, DA methods that use the source-domain classifier and target-domain samples (*model-transfer*) are preferred over those relying on the source-domain samples (*feature-transfer*). Note that in some cases it may not be possible to have access to the source-domain samples (e.g. a trained classifier can be available for public use, while the data used for its training may be confidential), and in any case the DA used to be slower and more memory demanding.

In this paper, we focus on supervised model-transfer DA for RF classifiers. In fact, as we will summarize in Sect. 2, there are previous DA methods for RF; however, they require revisiting the source-domain samples. Thus, aiming model-transfer DA for RF is a novelty of this paper. Actually, after quickly reviewing the specific RF paradigm used here (Sect. 3), we will propose three inherently different methods (Sect. 4), termed as *Node-Adapt*, *Path-Adapt* and *Tree-Adapt*, according to the level at which adaptation is performed.

Model-transfer DA is especially relevant when dealing with complex data such as images. Therefore, in this paper, to assess the performance of our proposals, we focus on image-based object detection, using the *pedestrian detection* problem as challenging proof-of-concept. Moreover, we use the RF with *local expert* splitting nodes (RF-LE) [6–8] because it allows to learn a *patch-based* pedestrian model which is competitive even with more elaborated deep CNNs [9]. As we will see (Sect. 5), we test our Node-, Path- and Tree-Adapt methods in standard benchmarks, showing that DA is largely achieved. Interestingly, we will show how a pedestrian detector based on a RF-LE trained with synthetic data (no manual labeling required) together with a relatively low number of real-world labels, performs not too far than an equivalent detector based on the labeling of 10-20 times more pedestrians. Finally, Sect. 6 draws our conclusions.

2 Related works

Several methods have been proposed to face the DA problem for RF. Goussies et al. [10] consider supervised transfer learning for gesture and character recognition. They propose a method for tuning the splitting function parameters of each node, so that the information gain of the target and source-domain samples decrease simultaneously. Danielsson and Aghazadeh [11] focus on unsupervised DA to adapt a pose estimator learned in synthetic data to operate in real scenarios. They propose an optimization function which selects the best parameters for the splitting function in each node. This is done by minimizing the information gain of the source-domain samples while minimizing the chi-square distance between the source and target feature histograms. Vezhnevets and Buhmann [12] performed DA for the so-called *semantic texton forest* used for semantic segmentation. They consider the same theoretical setting as [11] but using L1-norm for calculating the distance between two histograms instead of the chi-square.

All these RF-DA methods require the source-domain samples, i.e. they follow the feature-transfer paradigm. In contrast, model-transfer DA methods are eventually more interesting since the source data is not always available and they tend to be more computationally efficient; thus, opening the door to hierarchical and online adaptation strategies. In fact, there are already model-transfer DA proposal for popular successful classifiers such as SVM [13–15], structural SVM (SSVM) [16], deep CNNs [17], SSVM in hierarchical target domains [18], and SSVM in an online setting [19]. In this paper, as novelty, we aim at proposing and evaluating three complementary model-transfer DA methods for the RF-LE.

3 Random forest of local experts

Random Forest is a paradigm that allows to learn an ensemble of randomized decision-tree classifiers from a set of training samples. The classification of a testing sample depends on the average classification probabilities over all the trees. Randomness is the core to produce non-correlated decision trees which together lead to a highly accurate ensemble. In this paper, we focus on the

Algorithm 1 Training steps for the j^{th} node of the Random Forest of Local Experts (details in [7])

- 1) K feature selectors (patches) $\phi_1(\vec{v}), \dots, \phi_K(\vec{v})$ are selected randomly.
 - 2) For $k = 1, \dots, K$ do:
 - 2.1) Let S_{j,ϕ_k} be the transformed set of samples, $S_{j,\phi_k} = \{\phi_k(\vec{v}) : \vec{v} \in S_j\}$.
 - 2.2) Learn the linear SVM weights $\vec{\psi}_k$ with S_{j,ϕ_k} .
 - 2.3) Let $\theta = (\phi_k, \vec{\psi}_k, \tau)$ and $h(\vec{v}; \theta) = \{\vec{\psi}_k \cdot \phi_k(\vec{v}) + \tau \geq 0 : l; r\}$.
Let $\tau_k = \arg \max_{\tau} I(S_{j,\phi_k}; \theta_k)$ according to Eq. (1) and then $\theta_k = (\phi_k, \vec{\psi}_k, \tau_k)$.
 - 3) Let $j = \arg \max_{k \in \{1, \dots, K\}} I(S_{j,\phi_k}; \theta_k)$. Then $\theta_j = (\phi_j, \vec{\psi}_j, \tau_j)$.
-

RF-LE proposed in [7] for binary classification (i.e. *pedestrian vs background*), which is competitive for pedestrian detection [9]. In this case the training samples come in the form of image windows with a canonical size and the label that indicates if the window contains either a pedestrian or background. Randomness is introduced at splitting nodes. In particular, many possible rectangular sub-windows within the canonical window are considered (*patches*), but just a randomly chosen subset is selected. The samples reaching a splitting node are evaluated according to such selected patches. For each patch a combination of its HOG and LBP features is feed into a linear SVM learning process. The patch that allows to learn the *best* linear SVM is chosen as the *local expert* that remains attached to the splitting node. Here, *best* is defined in terms of the purity of the two sample partitions (considering only the samples reaching the node) induced by applying each linear SVM with an optimum score threshold. In the following we summarize this procedure more formally.

Let S_j be the set of samples reaching the splitting j^{th} node of the tree under training. Let $h(\vec{v}; \theta_j) \in \{l, r\}$ be the splitting parametric function at node j^{th} , being θ_j the set of parameters and $\vec{v} \in S_j$ a feature vector (representing a sample). The optimal θ_j is learned as $\theta_j = \arg \max_{\theta} I(S_j; \theta)$, where $I(S_j; \theta)$ is the information gain defined by:

$$I(S_j; \theta) = H(S_j) - ((|S_{j,\theta}^l|/|S_j|)H(S_{j,\theta}^l) + (|S_{j,\theta}^r|/|S_j|)H(S_{j,\theta}^r)), H(S) \equiv \text{ClassEntropyOf}(S), \quad (1)$$

where θ induces a two-subset split of S_j ($S_j = S_{j,\theta}^l \cup S_{j,\theta}^r$) by evaluating S_j with $h(\vec{v}; \theta)$; i.e. $\forall \vec{v} \in S_j$, \vec{v} will go to $S_{j,\theta}^c$ if $h(\vec{v}; \theta) = c, c \in \{l, r\}$. More precisely, $\theta = (\phi, \vec{\psi}, \tau)$, so that $h(\vec{v}; \theta) = \{\vec{\psi} \cdot \phi(\vec{v}) + \tau \geq 0 : l; r\}$. Following [7], \vec{v} contains the HOG+LBP features representing a given canonical window (i.e. a sample), $\phi(\vec{v})$ contains the HOG+LBP features of a selected sub-window within the canonical window, and $\vec{\psi} \cdot \phi(\vec{v})$ consists in applying a linear SVM ($\vec{\psi}$ is the vector of weights, and " \cdot " means scalar product). All θ_j are fixed at testing time, at training time randomness is introduced while looking for the best $\vec{\psi}$ at each splitting node. Algorithm 1 summarizes the learning procedure for any splitting node. A node becomes a leaf node if either a maximum tree depth is reached, or the number of samples reaching the node is below a threshold (no sense to split them more), or if the percentage of samples belonging to a class is above a threshold. The percentage of positive and negative samples (e.g. pedestrians and backgrounds) reaching a leaf node determines the posterior probability distribution attached to the node, i.e. the node $P(y|\vec{v})$, being y the class label (*pedestrian or background*). The different trees of the forest are trained independently.

A sample \vec{v} under classification is processed by all the forest' trees (but soft-cascade can be used for speeding up the process [7]), which is done by thresholding the following posterior probability:

$$P(y|\vec{v}) = \frac{1}{T} \sum_{i=1}^T P_i(y|\vec{v}), \quad (2)$$

where T is the number of trees in the forest, and $P_i(y|\vec{v})$ is the posterior probability that tree i assigns to sample \vec{v} . For obtaining $P_i(y|\vec{v})$, \vec{v} is routed through the tree i by the splitting nodes until reaching a leaf node, and the $P(y|\vec{v})$ at this leaf node is taken as the final $P_i(y|\vec{v})$ of the tree.




Figure 1: Node- and Path-Adapt rely on the structure of source trees, which are reshaped during DA. Eq. (4) is derived for the highlighted paths of the corresponding source- and target-domain trees.

4 Domain adaptation for random forest

We assume the so-called *covariate shift* commonly considered for object detection problems [20]; i.e. the marginal distribution of the source- and target-domain samples are different ($P^s(\vec{v}^s) \neq P^t(\vec{v}^t)$) while their posterior distributions are similar ($P^s(y|\vec{v}^s) \simeq P^t(y|\vec{v}^t)$). Then, in model-transfer DA, we can search for a classifier in the target-domain hypothesis space which can discriminate the available target-domain samples properly, while being similar to the source classifier which acts as regularizer (prior knowledge). This regularization is important since we assume a supervised DA setting where the number of target-domain samples is relatively low. In particular, we have a set S^{t_a} of m target-domain labeled samples and a RF-LE (F^s) trained on n source-domain samples, $m \ll n$. Overall, our final goal is to learn an accurate RF-LE (F^{t_a}), using S^{t_a} and F^s as regularizer. To achieve this, we present three novel and inherently different supervised model-transfer DA methods: Node-Adapt (Sect. 4.1), Path-Adapt (Sect. 4.2), and Tree-Adapt (Sect. 4.3).

In Node- and Path-Adapt we build a forest F^{t_a} only using S^{t_a} and forcing it to match as much as possible the structure of F^s . Briefly, the number of trees is the same (T), and the feature selectors at nodes (ϕ_j) as well as the maximum depth of each path are not changed. In Node-Adapt such adaptation is done at node level, both the linear SVM weights ($\vec{\psi}_j$) and thresholds (τ_j) are adapted. In Path-Adapt, the adaptation is done simultaneously for all the nodes of each root-to-leaf path, but only the thresholds (τ_j) are adjusted. Tree-Adapt consists in a *reforestation*. A ratio C of randomly selected trees of F^s is replaced by the same number of trees trained with S^{t_a} from the scratch.

4.1 Node-Adapt

Node-Adapt applies a node level adaptation separately for each of the trees cloned from F^s ; thus, we can focus on how this is done for the i^{th} tree by using S^{t_a} . The method proceeds as follows (see Fig. 1). The first node of i to be adapted is the root, whose local expert is parameterized by $\theta_j^s = (\phi_j^s, \vec{\psi}_j^s, \tau_j^s)$ for $j = 1$, where the upper script reminds that it was learned in the source domain. The adaptation will act on $\vec{\psi}_j^s$ and τ_j^s but the feature selector ϕ_j^s does not change, so we can just denote it as ϕ_j . In particular, Node-Adapt applies the Adaptive SVM (A-SVM) [13], a model-transfer convex DA method, to obtain an adapted SVM hyper-plane $\vec{\psi}_j^{t_a}$ by using $S_j^{t_a}$ and $\vec{\psi}_j^s$. This is done by solving the following optimization problem formulated for any node j (for the root $j = 1 \wedge S_1^{t_a} = S^{t_a}$):

$$\begin{aligned} \vec{\psi}_j^{t_a} &= \arg \min_{\vec{\psi}_j} \frac{1}{2} \|\vec{\psi}_j - C_1 \vec{\psi}_j^s\|^2 + C_2 \sum_{k=1}^{|S_j^{t_a}|} \epsilon_k, \\ \text{s.t. : } \quad &\epsilon_k \geq 0 \quad \wedge \quad y_k \cdot \vec{\psi}_j \cdot \phi_j(\vec{v}_k^{t_a}) \geq 1 - \epsilon_k \quad \forall k \in \{1, \dots, |S_j^{t_a}|\}, \end{aligned} \quad (3)$$

where y_k is the class label of the target-domain feature vector $\vec{v}_k^{t_a}$. Note that class labels do not change from source to target domains, so we do not need to distinguish domains for them. The main

goal of Eq. (3) is to obtain a hyper-plane ($\vec{\psi}_j^{t_a}$) which can separate the different classes properly in target domain while, according to covariate shift at node level, being similar to the source-domain hyper-plane ($\vec{\psi}_j^s$). C_1 and C_2 are the coefficients which control the amount of similarity to the source-domain hyper-plane and the amount of classification correctness on the target domain, respectively. Once we have $\vec{\psi}_j^{t_a}$, the new $\tau_j^{t_a}$ is obtained similarly than in step 2.2) of Algorithm 1, i.e. $\tau_j^{t_a} = \arg \max_{\tau} I(S_j^{t_a}; \phi_j, \vec{\psi}_j^{t_a}, \tau)$. Now, we have the new parameters defining the adapted root node, i.e. $\theta_j^{t_a} = (\phi_j, \vec{\psi}_j^{t_a}, \tau_j^{t_a})$ ($j = 1$). This means that we can apply $h(\vec{v}_k^{t_a}; \theta_j^{t_a}) \forall \vec{v}_k^{t_a} \in S_j^{t_a}$; therefore, we obtain the split $S_j^{t_a} = S_j^{l,t_a} \cup S_j^{r,t_a}$. Next, we apply the same procedure for the child nodes of the root using S_j^{l,t_a} and S_j^{r,t_a} as input for the left and right nodes, respectively. The process continues until the leaves of the i^{th} tree are reached. If the conditions for establishing a leaf node (the same than for learning source-domain trees) fulfill in a splitting node of the i^{th} tree, then it becomes a leaf node and its corresponding posterior density is computed using the target-domain samples that reached it. Note that this can happen because the original structure of the i^{th} tree was learned with n source-domain samples, while the adaptation is based on the m target-domain samples, and we assume $m \ll n$. Thus, the adapted tree is a reshaped version of the original one (Fig. 1). In addition, since we apply A-SVM at node level, we need to show that the covariate shift assumption is valid between two corresponding nodes of the source and target domains, which we prove as follows:

$$\begin{aligned}
P^s(y|\vec{v}) &\simeq P^t(y|\vec{v}) \\
&\Rightarrow P_i^s(y|\phi_1(\vec{v})) \simeq P_i^t(y|\phi_1(\vec{v})) \\
&\Rightarrow [\vec{\psi}_1^s, \tau_1^s] \simeq [\vec{\psi}_1^t, \tau_1^t] \quad (\text{SVM and } I(S_1^t; \phi_1, \vec{\psi}_1^t, \tau) \text{ are related to the posteriors in the domains}) \\
&\Rightarrow P_i^s(y|\vec{\psi}_1^s \cdot \phi_1(\vec{v}) + \tau_1^s) \simeq P_i^t(y|\vec{\psi}_1^t \cdot \phi_1(\vec{v}) + \tau_1^t) \\
&\Rightarrow P_i^s(y|\vec{\psi}_1^s \cdot \phi_1(\vec{v}) + \tau_1^s \geq 0) \simeq P_i^t(y|\vec{\psi}_1^t \cdot \phi_1(\vec{v}) + \tau_1^t \geq 0) \quad (\text{posteriors of the samples which reach to the second node in the source and target domains are similar}) \\
&\Rightarrow [\vec{\psi}_2^s, \tau_2^s] \simeq [\vec{\psi}_2^t, \tau_2^t] \\
&\Rightarrow \dots \\
&\Rightarrow P_i^s(y|\vec{\psi}_1^s \cdot \phi_1(\vec{v}) + \tau_1^s, \dots, \vec{\psi}_d^s \cdot \phi_d(\vec{v}) + \tau_d^s) \simeq P_i^t(y|\vec{\psi}_1^t \cdot \phi_1(\vec{v}) + \tau_1^t, \dots, \vec{\psi}_d^t \cdot \phi_d(\vec{v}) + \tau_d^t) \\
&\Rightarrow P_i^s(y|\vec{\psi}_1^s \cdot \phi_1(\vec{v}) + \tau_1^s \geq 0, \dots, \vec{\psi}_d^s \cdot \phi_d(\vec{v}) + \tau_d^s \geq 0) \simeq \\
&\quad P_i^t(y|\vec{\psi}_1^t \cdot \phi_1(\vec{v}) + \tau_1^t \geq 0, \dots, \vec{\psi}_d^t \cdot \phi_d(\vec{v}) + \tau_d^t \geq 0).
\end{aligned} \tag{4}$$

Of course, this reasoning can be extended to all other paths in the tree.

4.2 Path-Adapt

Path-Adapt performs adaptation for each path of each source-domain tree instead of for each node individually; i.e. for all splitting nodes in a given path, the adaptation is done simultaneously. Therefore, we need to show that the covariate shift assumption holds at path level, i.e. beyond node level as required for Node-Adapt. In order to do that, as for Node-Adapt, we assume that the structure of the source-domain forest, F^s , is similar to the target-domain forest we are searching for, F^t . Thus, according to Eq. (4), we can show the relation $P_i^s(y|\vec{\psi}_1^s \cdot \phi_1(\vec{v}) + \tau_1^s, \vec{\psi}_2^s \cdot \phi_2(\vec{v}) + \tau_2^s, \dots, \vec{\psi}_d^s \cdot \phi_d(\vec{v}) + \tau_d^s) \simeq P_i^t(y|\vec{\psi}_1^t \cdot \phi_1(\vec{v}) + \tau_1^t, \vec{\psi}_2^t \cdot \phi_2(\vec{v}) + \tau_2^t, \dots, \vec{\psi}_d^t \cdot \phi_d(\vec{v}) + \tau_d^t)$, is established between each path of the i^{th} tree of F^t and its corresponding path of the i^{th} tree of F^s , where d is the index of the last splitting node in the p^{th} path of the i^{th} tree. Now, we define $Q_p^s = [\vec{\psi}_1^s, \tau_1^s; \vec{\psi}_2^s, \tau_2^s; \dots; \vec{\psi}_d^s, \tau_d^s]$ and $Q_p^t = [\vec{\psi}_1^t, \tau_1^t; \vec{\psi}_2^t, \tau_2^t; \dots; \vec{\psi}_d^t, \tau_d^t]$ as two mapping functions for projecting the feature set $\Phi_p = [\phi_1(\vec{v}), 1; \phi_2(\vec{v}), 1; \dots; \phi_d(\vec{v}), 1]$ to a new feature space. Let $[\vec{W}_p^s, b_p^s]$ be the SVM hyper-plane that we could learn in the $Q_p^s \cdot \Phi_p(\vec{v}) = (\vec{\psi}_1^s \cdot \phi_1(\vec{v}) + \tau_1^s, \dots, \vec{\psi}_d^s \cdot \phi_d(\vec{v}) + \tau_d^s)$ d -dimensional feature space, using sufficient source-domain data. Let $[\vec{W}_p^t, b_p^t]$ the analogous case using $Q_p^t \cdot \Phi_p(\vec{v})$ and target-domain data. Then, $[\vec{W}_p^s, b_p^s]$ and $[\vec{W}_p^t, b_p^t]$ should be similar since, according to Eq. (4), the relations $Q_p^s \simeq Q_p^t$ and $P_i^s(y|Q_p^s \cdot \Phi_p(\vec{v})) \simeq P_i^t(y|Q_p^t \cdot \Phi_p(\vec{v}))$ exist for every path of the tree (i.e. the covariate shift assumption in path level).

Accordingly, the strategy of Path-Adapt starts by searching for a \tilde{Q}_p^t function so that, after projecting the target-domain samples in S^{t_a} to the new feature space, they can be class-discriminated successfully

by $[\vec{W}_p^s, b_p^s]$. In particular, using the same structure than F^s (same T , maximum depth of each path, and $\phi_j(\vec{v})$ feature selectors), as well as S^{t_a} , a new forest \tilde{F}^t is learned. Note that in this process there is no randomness, \tilde{F}^t is a reshaped version of F^s (leaves change, remind $m \ll n$). Now, we can define $\tilde{Q}_p^t = [\tilde{\psi}_1^t, \tilde{\tau}_1^t; \tilde{\psi}_2^t, \tilde{\tau}_2^t; \dots; \tilde{\psi}_d^t, \tilde{\tau}_d^t]$ for the p^{th} path of \tilde{F}^t . At this point, Path-Adapt uses S^{t_a} and $[\vec{W}_p^s, b_p^s]$ to adapt \tilde{Q}_p^t . However, the dimensionality of \tilde{Q}_p^t can be eventually too high depending on the length of the path and the particular feature selectors in it, while at the same time the number of samples in S^{t_a} that eventually are considered along the path can be too low. Therefore, we propose to adapt only the thresholds. In particular, we propose the following optimization problem:

$$B^{t_a} = \underset{B^t}{\operatorname{argmin}} \frac{1}{2} \|B^t - \tilde{B}^t\|^2 + C \sum_{p=1}^{R_i} \sum_{k=1}^m \epsilon_{k,p}; \quad (5)$$

$$s.t.: \epsilon_{k,p} \geq 0 \wedge y_k(\vec{W}_p^s \cdot (\tilde{Q}_p^t \cdot \Phi_p(\vec{v}_k^{t_a}) + b_p^s)) \geq -\epsilon_{k,p} \quad \forall k \in \{1, \dots, m\}, p \in \{1, \dots, R_i\},$$

where $B^t = [\tau_1^t, \tau_2^t, \dots, \tau_N^t]$, and $\tilde{B}^t = [\tilde{\tau}_1^t, \tilde{\tau}_2^t, \dots, \tilde{\tau}_N^t]$. N is equal to the number of nodes and R_i stands for the number of paths in the i^{th} tree. C is a coefficient that controls the amount of similarity to the source model. We also restrict the adapted thresholds (B^{t_a}) to be similar to the target thresholds (\tilde{B}^t) to make balance between similarity to the source model and correctness of classification on target domain. Eq. (5) is a QP convex optimization problem that can be solved with standard optimization packages. Finally, after obtaining the new thresholds for each tree, S^{t_a} is used again for reshaping and computing the posterior probabilities at the new leaves, so obtaining a final adapted forest F^{t_a} .

The hyper-plane $[\vec{W}_p^s, b_p^s]$ is computed after the training of F^s . In other words, for further domain adaptation, Path-Adapt requires to learn the different $[\vec{W}_p^s, b_p^s]$ as source-domain work. In fact, for each tree of F^s , and for each path in a given tree, we need to learn and store different $[\vec{W}_p^s, b_p^s]$ since we cannot know a priori which splitting nodes in F^s will become leaves in \tilde{F}^t and, therefore, which will be the final depth of a given path p . In particular, if d is the length a path of a tree in F^s , then it is required to learn d linear SVMs for such a path, using the source-domain samples. However, this is not necessarily quite time consuming since we even expect $\dim(Q_p^s \cdot \Phi_p) \ll \dim(\phi_j(\vec{v}))$ for each node j in the path. Moreover, note that during Path-Adapt, the $\vec{\psi}_j^s$ are not used. Intuitively, such a collection of hyper-planes acts as a compact proxy of the source-domain data and the local experts $\vec{\psi}_j^s$ during the Path-Adapt process.

4.3 Tree-Adapt

According to the covariate shift assumption between source and target domains, in Tree-Adapt we replace a ratio C of randomly selected trees in F^s , by an equivalent number of trees trained with S^{t_a} from the scratch. If we name as $F_m^{t_a}$ the forest formed by such target-oriented set of trees, then joining $F_m^{t_a}$ and F^s , gives rise to the adapted forest F^{t_a} of T trees. Thus, F^{t_a} has a posterior distribution based on F^s and $F_m^{t_a}$, balanced by C :

$$P(y|\vec{v}) = \frac{1}{T} \left(\sum_{i=1}^{CT} P_i^{t_a}(y|\vec{v}) + \sum_{i=1}^{(1-C)T} P_i^s(y|\vec{v}) \right), \text{ where } 0 < C \leq 1. \quad (6)$$

5 Experimental results

In this section we show the effectiveness of the proposed DA methods for RF-LE. We present quantitative results on vision-based pedestrian detection as a challenging proof-of-concept object detection problem. We select four widely used pedestrian datasets, namely Virtual [21], INRIA [22], Daimler [23] and KITTI [24], to evaluate the RF-DA methods. In the case of KITTI we follow the train/test partition proposed in [?]. Pedestrians and backgrounds from these datasets can be seen in Fig. 2.



Figure 2: Pedestrians and backgrounds of four different datasets. The samples of these datasets show differences due to the camera (resolution, B/W-vs-color, sensor sensitivity, focal length, lens distortion); image acquisition conditions (usual viewpoint, scene illumination, backgrounds); and even the typical human poses and clothes.

Table 1: The average miss rate (AMR) for 11 points in the interval of $[0.01, 1]$ FPPI is calculated and shown for different baselines (Src, Tar100%, TarX% where $X = 10$ for INRIA and KITTI and $X=5$ for Daimler) and RF-DA methods (Node-Adapt, Path-Adapt and Tree-Adapt). Where V = Virtual, I = INRIA, D = Daimler, K = KITTI. Each "A \rightarrow B" experiment is run five times to provide a mean AMR and corresponding standard deviation; thus the lower the better. Note how, Node-Adapt, Path-Adapt and Tree-Adapt show successful adaptation in most combinations (comparing with Src and TarX%, even with Tar100).

	Src	Tar100%	Tar X%	Node-Adapt	Path-Adapt	Tree-Adapt
V \rightarrow I	19.06 \pm 1.77	11.09 \pm 1.12	17.74 \pm 2.17	16.37 \pm 1.84	16.70 \pm 1.47	14.59 \pm 1.39
V \rightarrow D	31.95 \pm 1.35	25.57 \pm 1.36	30.98 \pm 3.06	27.95 \pm 3.70	29.41 \pm 2.70	25.39 \pm 1.94
V \rightarrow K	51.07 \pm 2.40	36.02 \pm 2.28	53.06 \pm 3.93	45.55 \pm 3.98	38.95 \pm 2.36	41.59 \pm 2.20
I \rightarrow D	23.28 \pm 1.65	25.57 \pm 1.36	30.77 \pm 3.08	26.72 \pm 3.69	26.69 \pm 4.46	22.87 \pm 2.30
I \rightarrow K	44.41 \pm 1.79	36.02 \pm 2.28	53.06 \pm 3.93	44.88 \pm 3.07	38.84 \pm 4.37	37.24 \pm 1.95
D \rightarrow I	21.98 \pm 1.29	11.09 \pm 1.12	17.74 \pm 2.17	14.77 \pm 1.58	14.66 \pm 2.27	13.92 \pm 1.17
D \rightarrow K	58.96 \pm 1.85	36.02 \pm 2.28	53.06 \pm 3.93	44.91 \pm 2.42	41.58 \pm 2.97	42.10 \pm 2.25
K \rightarrow I	48.19 \pm 1.57	11.09 \pm 1.12	17.74 \pm 2.17	17.74 \pm 3.21	16.22 \pm 2.19	18.17 \pm 1.92
K \rightarrow D	44.87 \pm 2.37	25.57 \pm 1.36	30.77 \pm 3.08	26.41 \pm 3.43	27.41 \pm 2.81	26.36 \pm 2.10

5.1 Experimental setup & implementation details

The Virtual dataset is used only as source domain. The rest are considered both as source and target domains thanks to the use of different training and testing sets per dataset. In order to run a DA experiment we select a relatively low number of the samples available in the training set of the target-domain dataset. For INRIA and KITTI this is $\sim 10\%$, while for Daimler we use $\sim 5\%$ since it is a larger training dataset.

Regarding the training of the RF-LE, we follow the instructions reported in [7], e.g. we use 100 trees and the maximum depth for the trees is set to 7, and HOG+LBP features are used as image descriptors.

The average miss rate (AMR) vs false positives per image (FPPI) [25] to report detection accuracy. To set the hyper-parameters of the different adaptation methods (i.e. the C 's), we used the Virtual-INRIA DA problem. These values are then used for the rest of experiments.

5.2 Results

In order to show the effectiveness of the proposed RF-DA methods, we run experiments for different combinations of source and target domains. Each train-test experiment is repeated five times and we report the mean and standard deviation of AMR versus FPPI in the interval of $[0.01, 1]$. We compare Node-Adapt, Path-Adapt and Tree-Adapt methods in each experiment with the RF trained only with source data (Src), the RF trained with all samples of the target domain (Tar100%), and with the RF trained only on the limited number of target-domain samples (TarX%).

Table 2: Analogous to Table 1 but for a more challenging situation in terms of available target-domain data. In particular, $X = 5$ for INRIA and KITTI and $X = 3$ for Daimler.

	Src	Tar100%	Tar X%	Node-Adapt	Path-Adapt	Tree-Adapt
V \rightarrow I	19.06 \pm 1.77	11.09 \pm 1.12	24.97 \pm 3.85	19.29 \pm 1.38	18.77 \pm 2.72	16.65 \pm 2.15
V \rightarrow D	31.95 \pm 1.35	25.57 \pm 1.36	31.01 \pm 3.63	28.55 \pm 3.14	29.56 \pm 4.80	25.77 \pm 1.97
V \rightarrow K	51.07 \pm 2.40	36.02 \pm 2.28	57.96 \pm 3.69	44.24 \pm 4.26	38.34 \pm 2.03	43.36 \pm 2.28
I \rightarrow D	23.28 \pm 1.65	25.57 \pm 1.36	31.71 \pm 3.23	27.12 \pm 3.34	28.65 \pm 3.78	23.71 \pm 1.79
I \rightarrow K	44.41 \pm 1.79	36.02 \pm 2.28	58.34 \pm 2.89	44.33 \pm 2.41	40.52 \pm 2.23	39.79 \pm 2.87
D \rightarrow I	21.98 \pm 1.29	11.09 \pm 1.12	20.62 \pm 5.53	17.66 \pm 2.32	16.33 \pm 2.73	15.23 \pm 2.29
D \rightarrow K	58.96 \pm 1.85	36.02 \pm 2.28	56.41 \pm 2.35	47.39 \pm 2.56	45.48 \pm 2.75	45.20 \pm 2.50
K \rightarrow I	48.19 \pm 1.57	11.09 \pm 1.12	24.34 \pm 5.41	21.02 \pm 2.84	19.08 \pm 3.43	19.31 \pm 3.18
K \rightarrow D	44.87 \pm 2.37	25.57 \pm 1.36	30.92 \pm 3.23	27.53 \pm 2.97	27.81 \pm 3.16	26.36 \pm 3.44

In this case, DA is successful if reduces the AMR of both Src and TarX%, and is close to the AMR of Tar100%. Therefore, we can see in Table 1, that Node-Adapt, Path-Adapt and Tree-Adapt provide successful results. Notably, Tree-Adapt, the simplest strategy, is more successful than the other two.

Node-Adapt and Path-Adapt report similar performance. However, when the number of available samples in target domain is limited, the probability of error propagation becomes higher for Node-Adapt (too few samples to train the local experts) and, a priori, Path-Adapt should be more successful since it uses all samples for adapting to adapt less parameters. Tree-Adapt can suffer the same problem that Node-Adapt since the learned target-domain trees can be too shallow. To assess this point, we repeat the same experiments with half of available target-domain samples. Table 2 shows the corresponding results¹.

Indeed, Path-Adapt is more effective than Node-Adapt in this demanding scenario. However, Tree-Adapt still is the most effective method. On the other hand, for low memory adaptive embedded systems (e.g. a domestic cleaning robot), Path-Adapt is interesting in the sense that it requires scarce information from the source domain since the data and node hyper-planes are not required, just the source tree topology and the compact path SVMs are needed. So the system can come with a sort of common knowledge and be (self-)adapted for the particular environment. After all, Path-Adapt results are just scarcely worse than those of Tree-Adapt.

Overall, we conclude that among the three tested strategies, Tree-Adapt is the simplest to implement and the one that offers better results. Interestingly, Table 1 shows that the pedestrian detector based on the RF-LE trained with synthetic data (no manual labeling required) and adapted with a relatively low number of real-world labels using Tree-Adapt, performs not too far than an equivalent detector based on the labeling of 10-20 times more pedestrians (Tar100%). In addition, as we have mentioned, Path-Adapt can be a good alternative depending on the application.

6 Conclusions

We have proposed three novel supervised model-transfer DA methods for RF-LE models, namely Node-Adapt, Path-Adapt and Tree-Adapt. These methods are inherently different since they perform the adaption at different levels of the RF-LE. We have investigated the validity of the covariate shift assumption for them to support their approaches of adaptation. We have assessed their performance by facing the challenge of vision-based pedestrian detection. In particular, we have considered four well-known benchmark datasets as different domains. The conclusion is that, although these three proposals achieve DA, Tree-Adapt is the most successful method to adapt RF-LE models, and opens the door to the use of virtual worlds as source domain for reducing the number of object annotations. As future work we plan several working lines. First, we want to assess the possibility of combining the core ideas behind the proposed methods so that adaptation is further improved; alternatively they can be used to self-annotate training data for progressively improving its accuracy (each method acting as a different oracle). Second, we want to assess how Tree-Adapt can be used in an online learning setting.


¹The information in Table 1 and in Table 2 is plotted as AMR-vs-FPPI curves in the supplementary material.

References


- [1] L. Breiman, “Random forests,” *ML*, vol. 45, no. 1, pp. 5 – 32, 2001.
- [2] A. Criminisi, J. Shotton, and E. Konukoglu, “Decision forests for classification, regression, density estimation, manifold learning and semi-supervised learning,” *F&T-CGV*, vol. 7, no. 2 - 3, pp. 81 – 227, 2012.
- [3] L. Bourdev and J. Brandt, “Robust object detection via soft cascade,” in *CVPR*, (San Diego, CA, USA), 2005.
- [4] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake, “Real-time human pose recognition in parts from a single depth image,” in *CVPR*, (Colorado Springs, CO, USA), 2011.
- [5] P. Kotschieder, M. Fiterau, A. Criminisi, and S. Rota, “Deep neural decision forests,” in *ICCV*, (Santiago de Chile), 2015.
- [6] B. Yao, A. Khosla, and L. Fei-Fei, “Combining randomization and discrimination for fine-grained image categorization,” in *CVPR*, (Colorado Springs, CO, USA), 2011.
- [7] J. Marin, D. Vázquez, A. López, J. Amores, and B. Leibe, “Random forests of local experts for pedestrian detection,” in *ICCV*, (Sydney, Australia), 2013.
- [8] M. Ristin, M. Guillaumin, J. Gall, and L. Van Gool, “Incremental learning of random forests for large-scale image classification,” *T-PAMI*, vol. 38, no. 3, pp. 490 – 503, 2016.
- [9] H. Fukui, T. Yamashita, Y. Yamauchi, H. Fujiyoshi, and H. Murase, “Pedestrian detection based on deep convolutional neural network with ensemble inference network,” in *IV*, (Seoul, South Korea), 2015.
- [10] N. Goussies, S. Ubalde, and M. Mejail, “Transfer learning decision forests for gesture recognition,” *JMLR*, vol. 15, no. 1, pp. 3667 – 3690, 2014.
- [11] O. Danielsson and O. Aghazadeh, “Human pose estimation from RGB input using synthetic training data.” arXiv:1405.1213, 2014.
- [12] A. Vezhnevets and J. Buhmann, “Agnostic domain adaptation,” in *DAGM*, (Frankfurt, Germany), 2011.
- [13] J. Yang, R. Yan, and A. Hauptmann, “Cross-domain video concept detection using adaptive SVMs,” in *ACM-MM*, (Augsburg, Germany), 2007.
- [14] Y. Aytar and A. Zisserman, “Tabula rasa: model transfer for object category detection,” in *ICCV*, 2011.
- [15] A. Mozafari and M. Jamzad, “A SVM-based model-transferring method for heterogeneous domain adaptation,” *PR*, vol. 56, no. C, pp. 142 – 158, 2016.
- [16] J. Xu, S. Ramos, D. Vázquez, and A. López, “Domain adaptation of deformable part-based models,” *T-PAMI*, vol. 36, pp. 2367 – 2380, 2014.
- [17] J. Hoffman, S. Guadarrama, E. Tzeng, R. Hu, J. Donahue, R. Girshick, T. Darrell, and K. Saenko, “LSDA: large scale detection through adaptation,” in *NIPS*, (Montréal, Canada), 2014.
- [18] J. Xu, S. Ramos, D. Vázquez, and A. López, “Hierarchical adaptive structural SVM for domain adaptation,” *IJCV*, February 2016.
- [19] J. Xu, D. Vázquez, A. López, and K. Mikolajczyk, “Hierarchical online domain adaptation of deformable part-based models,” in *ICRA*, (Stockholm, Sweden), 2016.
- [20] J. Jiang, “A literature survey on domain adaptation of statistical classifiers,” tech. rep., School of Information Systems, Singapore Management University, 2008.
- [21] D. Vazquez, A. López, J. Marín, D. Ponsa, and D. Gerónimo, “Virtual and real world adaptation for pedestrian detection,” *T-PAMI*, vol. 36, pp. 797 – 809, 2014.
- [22] N. Dalal and B. Triggs, “Histograms of oriented gradients for human detection,” in *CVPR*, (San Diego, CA, USA), 2005.
- [23] M. Enzweiler and D. Gavrilu, “Monocular pedestrian detection: survey and experiments,” *T-PAMI*, vol. 31, pp. 2179 – 2195, 2009.
- [24] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the KITTI vision benchmark suite,” in *CVPR*, (Providence, RI, USA), 2012.
- [25] P. Dollár, C. Wojek, B. Schiele, and P. Perona, “Pedestrian detection: an evaluation of the state of the art,” *T-PAMI*, vol. 34, pp. 743 – 761, 2012.

Supplementary Materials


1. Results of Table 1 as plots



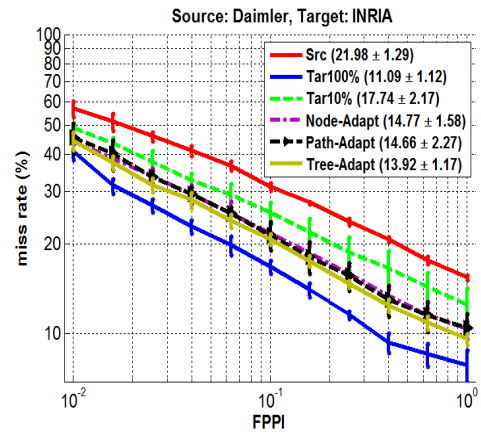
(a)




(b)



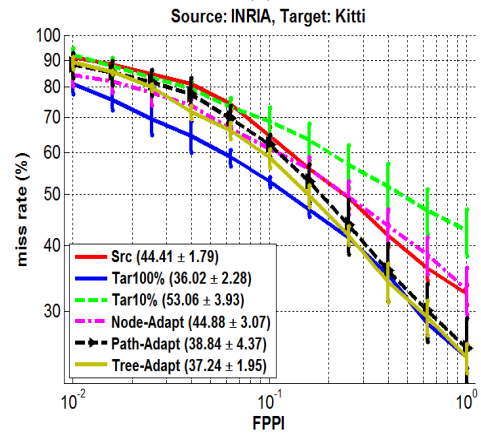
(c)



(d)



(e)



(f)








Figure 3: Average miss rate vs FPPI curve in the interval of $[0.01, 1]$ is plotted for Node-Adapt, Path-Adapt and Tree-Adapt in comparison to different baselines: Src, Tar100%, Tar10% (when the target domains are Kitti or INRIA) and Tar5% (when the target domain is Daimler). The experiment is repeated for all combination of the source and target domains. As can be seen, the proposed DA methods can achieve lower miss rate than Src and Tar10% or Tar5% baselines which indicates successful adaptation. Tree-Adapt is the most successful one.


2. Results of Table 2 as plots




(a)




(b)




(c)



(d)



(e)



(f)

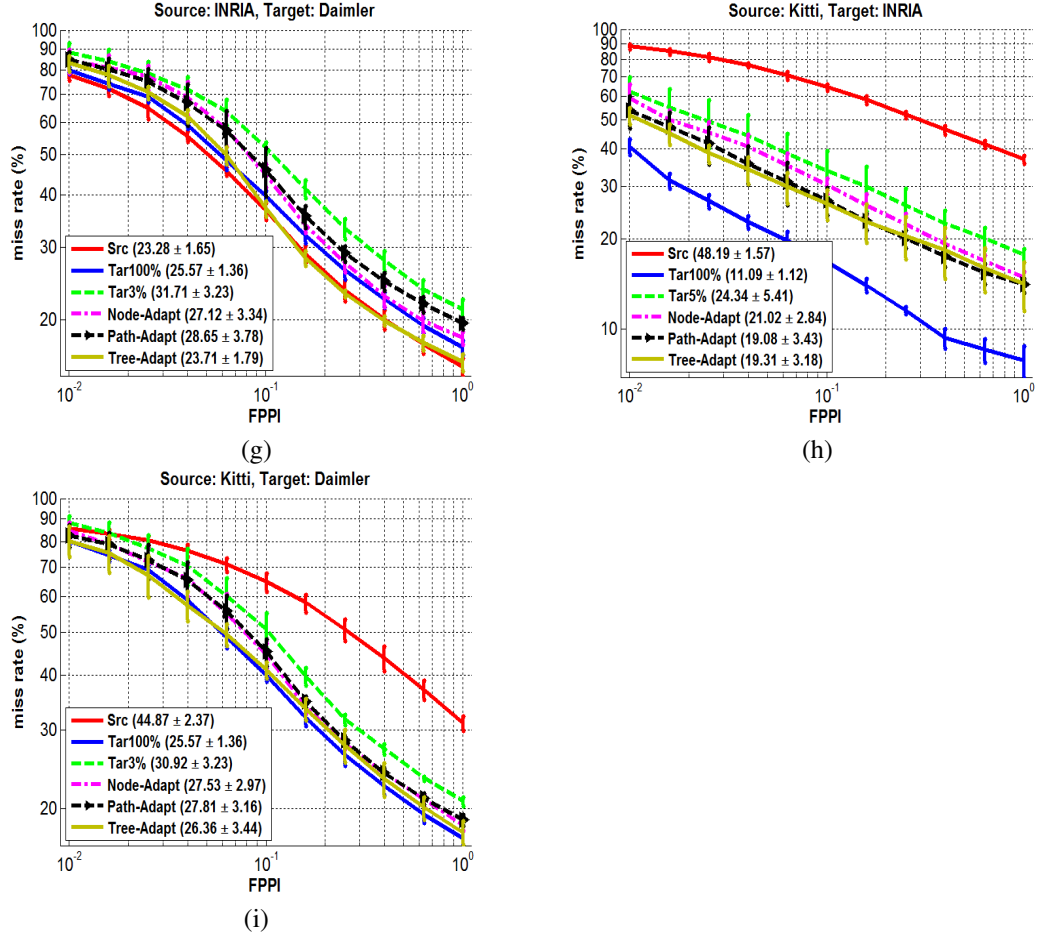


Figure 4: Average miss rate vs FPPI curve in the interval of $[0.01, 1]$ is plotted for Node-Adapt, Path-Adapt and Tree-Adapt in comparison to different baselines: Src, Tar100%, Tar5% (for Kitti or INRIA as the target domain) and Tar3% (for Daimler as the target domain) when the number of available target samples is very low. Comparing to Figure1, Node-Adapt obtains higher miss rate than Path-Adapt which indicates that the robustness of Path-Adapt is more than Node-Adapt in the few training target samples situation.