IDIAP RESEARCH REPORT

TRANSFER LEARNING OF VISUAL CONCEPTS ACROSS ROBOTS: A DISCRIMINATIVE APPROACH

Sriram Prasath Elango Tatiana Tommasi
Barbara Caputo

Idiap-RR-06-2012

JANUARY 2012

Centre du Parc, Rue Marconi 19, P.O. Box 592, CH - 1920 Martigny
T +41 27 721 77 11 F +41 27 721 77 12 info@idiap.ch www.idiap.ch

Transfer Learning of Visual Concepts across Robots: a Discriminative Approach

Sriram Prasath Elango, Tatiana Tommasi and Barbara Caputo

***Abstract*—While there is a general consensus that autonomous robots should be able to learn continuously over time, the learning process is traditionally envisioned for each specific robot situated in a given environment. This does not consider the fact that robots performing similar tasks in similar settings would probably learn similar concepts. They would therefore benefit by sharing their prior experience with each other. In this paper we present a transfer learning algorithm that enables robots located in different places to take advantage of each other's experience, boosting the learning process. We specifically assume to have robots equipped with a camera. We do not make any assumption on the type of camera, nor on where it is positioned. We also assume that the robots use the same feature descriptors and learning algorithms. Under these assumptions, we show that one robot can hugely benefit from what has been learned by peer robots performing similar tasks. The advantage concretely means a consistent boost in performance, especially when training data is scarce. Our algorithm is based on Least Square Support Vector Machine, and allows to determine automatically from where to transfer and how much to transfer: this makes it possible to take advantage of the prior when it is useful, while minimizing the risk of negative transfer when the priors are not informative. Thorough experiments on four different publicly available databases show the power of our approach.**

I. INTRODUCTION

Artificial autonomous systems are meant to operate in the real world. However, even the best system we can currently engineer is bound to fail whenever the setting is not heavily constrained. This is because the real world is generally too nuanced, too complicated and too unpredictable to be summarized within a limited set of specifications; there will be inevitably novel situations and the system will always have gaps, conflicts or ambiguities in its own knowledge and capabilities.

To date, research in artificial cognitive systems has approached the problem of learning with two main strategies. The first is the developmental approach, where tabula rasa robots autonomously learn how to deal with their actuators and sensors before constructing the concepts of state and time, the dynamics of sensorimotor coordination for manipulation, and ultimately building and using an increasing number of concepts [1], [2], [3], [4]. The second considers a human supervisor as the key component in the learning process. Implementations of this strategy range from controllers which batch-process large amounts of human-annotated information, to interactive systems which can be tutored online by a human [5], [6], [7]. Consider these two learning strategies in terms of (1) the amount of knowledge available to the system at any point in time, and (2) the cost of accessing this knowledge. When a robot follows the developmental approach, the amount of knowledge available is strictly limited by the environment and the (in)ability of the robot to perceive and interact with objects etc., while the cost of accessing the information that is available is not particularly high. When a robot learns through human supervision, the amount of knowledge available is greater, but the cost of accessing it is commensurately higher. This is particularly true when learning depends on human-robot interaction.

Both strategies envision robots learning in isolation. A third, largely unexplored strategy is that of robots learning from each other. Indeed,

S. Elango, T. Tommasi and B. Caputo are with The Idiap Research Institute, Rue Marconi 19, 1920 Martigny CH {sriram.prasath,ttommasi,bcaputo}@idiap.ch

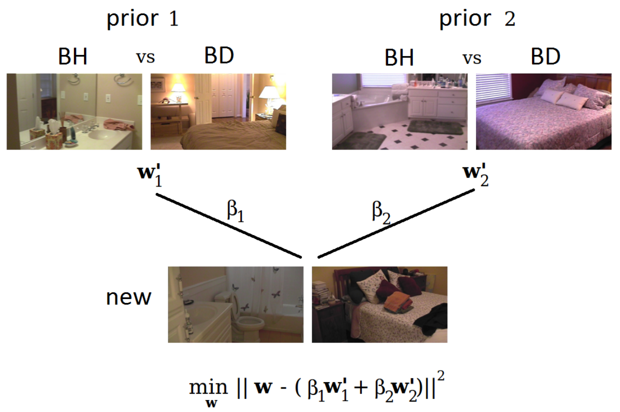

Fig. 1. Schematic example of the proposed transfer learning method. A new robot learns the parameter $\mathbf{w}$ describing the hyperplane which separates bathroom (BH) from bedroom (BD). The optimisation problem is regularized asking $\mathbf{w}$ to be close to a linear combination of known $\mathbf{w}'_j$ associateed with similar problems already solved by other two robots ($j = 1$ prior 1,$j = 2$ prior 2) in different locations. The coefficients β_j indicate how much the new model is close to each of the old ones.

robots situated in similar settings, performing similar tasks, do learn similar perceptual concepts. It seems therefore reasonable that a robot engaged in learning a new task could take advantage of the prior models learned by other robots which have already mastered that same, or a related, task.

We focus here on the learning of visual perceptual tasks, such as learning to recognize visually a place or an object. To enable a robot to perform these tasks while taking advantage of what learned before by peer robots, there is a need for common internal representations of knoweldge. A first thorough attempt in this direction has been the work of Zsolt [8], which builds joint intermediate representations for visual properties like color, texture, shape and so forth (for a comprehensive review of the relevant literature we refer the reader to Section II). The method assumes to have two robots equipped with different cameras, engaged in the same task. As the common representation is built at the sensory level, a key assumption is that both robots perceive the same object at the same time, in the same place. This is a very strong constraint, as it de facto limits the use of this approach to robots physically located in the same environment, at least while learning the joint perceptual representation.

In this paper we propose instead to enable robots to learn from their more experienced peers through joint high level representation, i.e. shared representations at the level of the class decision functions. This preserves the freedom to have each robot possibly equipped with a different camera, positioned at different heights. Our working assumption is that all robots use the same kind of feature descriptors, and the same type of classifier. Although these are non trivial constraints, we argue that they are significantly milder than those proposed in [8]. Furthermore, this strategy does not pose any limitation on the number of agents from which a robot might exploit priors when engaged in a new task, therefore increasing the potential advantages of this learning strategy.

We cast the problem within the transfer learning framework [9], and we take a discriminative approach. Our method is based on Least Square Support Vector Machine (LS-SVM) [10]. A robot learning a new task equates to learning a new class (or set of classes) through adaptation. Knowledge from other k robots is exploited in the form of the classifier hyperplanes $\mathbf{w}'_j$, $j = 1, \ldots, k$. The hyperplane $\mathbf{w}$ learned by the new robot is constrained to be close to a linear combination of those of the priors (Figure 1). Our algorithm learns *from where to transfer* (i.e which prior sources to trust), and *how* *much to transfer* (i.e. how much to trust each prior source), via an optimization process which minimizes the Leave-One-Out (LOO) error on the training set. Determining how much to transfer helps avoiding negative transfer. Therefore, in case of non-informative prior knowledge, transfer might be disregarded completely.

Our working scenario consists of platforms, equipped with the same kind of sensors with similar characteristics, with the same processing capabilities. These robots will be located in indoor office or house environments. We assume that all of them are able to locate themselves semantically in space ("I am in the kitchen", "I am in the corridor"). The knowledge to be shared across platforms will be visual-based semantic place representations, whether at the specific level (Barbara's office, House1 kitchen), or at the categorical level (an office, a kitchen; see Figure 2). We performed experiments on four different public databases (the IDOL database [11], the RobotVision@ImageCLEF2010 database [12], the VPC database [13] and the COLD database [14]) supporting the scenarios described above. Our results clearly indicate that our method significantly boost learning of new perceptual concepts whenever informative priors are available, while it does not affect negatively learning in case of non informative priors.

A preliminary version of this work was first published in [15], within the context of object categorization. Compared to the conference version, in this paper we make the following contributions: (1) we cast for the first time the problem of knowledge sharing across heterogeneous robot platforms within the transfer learning framework, significantly relaxing the requirements over the placement of the robot systems compared to previous work [8]; (2) we extend the algorithm from the binary to the multi-class scenario, obtaining one of the few multi-class transfer learning methods in the literature; (3) we perform a thorough experimental evaluation of our approach over four different publicly available databases, analyzing the behavior of our approach in three different scenarios.

The rest of the paper is organized as follows: section II gives a review of transfer learning and previous work in the robotics community. Section III casts the robot sharing within the transfer learning framework and describes our algorithm. Section IV describes our experimental setup, while Section V reports on our experimental findings and discuss results. We conclude with an overall discussion and possible future directions for research.

II. Related Works

In the following we elaborate on the state of the art on transfer learning from the machine learning and robotics perspective.

A. Transfer Learning: the Machine Learning Perspective

The fundamental motivation for transfer learning in the field of machine learning was first discussed in NIPS-95 workshop on Learning to Learn [16] which focuses on the need for life-long machine learning methods that retain and reuse previous knowledge. A major assumption in many machine learning and data mining systems is that the data distributions between the training and test data are the same, and that the data must be from the same feature representations. However in many real world applications, this assumption does not hold. For example, we sometimes have a classification task in one domain, but we only have sufficient training data in another domain where the data may follow a different distribution, or may be in a different feature space. In these cases, knowledge transfer would greatly benefit learning in our interested domain by avoiding expensive data labelling tasks [9]. To date, there are three main open research issues in transfer learning: *what to transfer*, *how to transfer*, and *when to transfer*.

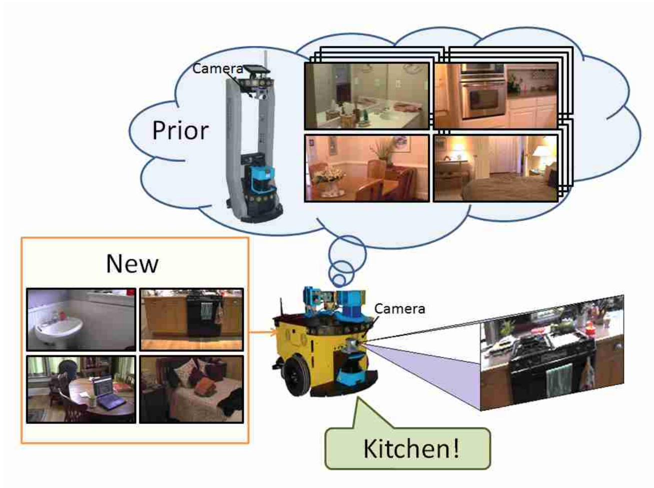

Fig. 2. Schematic example of transfer learning across two similar robots in different house environments (categorical task). Both the agents are asked to solve a four class problem (bathroom, kitchen, diningroom, bedroom). Many image samples are available for the first robot; the discriminative model defined on this prior problem is saved and used by the second robot when learning from few new samples.

The question 'what to transfer?' addresses which knowledge can be used to transfer across domains or tasks. Some knowledge is specific for individual domains, and some knowledge may be common between different domains such that it helps improving performance for the target task [9]. In the literature it is possible to find three answers to this question. One can be referred as *instance-transfer approach*: although the source domain data cannot be reused directly, there are certain parts of the data that can still be considered together with a few labelled data in the target domain. Possible strategies consist in re-weighting the source data or sampling them to remove the misleading training examples [17], [18]. A second solution is defined by *transferring feature representations*. It means learning a common feature structure from different domains that can bridge related tasks [19], [20]. A variant consists in finding suitable kernels for the target data in SVM- based approaches [21]. The third strategy can be described as *parameter-transfer approach*. It assumes that the source task and the target tasks share some parameters or priors of their model. Thus knowledge is formalized e.g. in Gaussian Priors and the parameters describing source tasks are reused to learn the novel target task [22], [23].

Addressing the question 'how to transfer?' corresponds to develop algorithms to transfer knowledge, after discovering which information can be transferred [9]. The answer to this question depends on the application context in which the problem is posed. Inductive Transfer Learning involves extending well known classification and inference algorithms such as neural networks, Bayesian networks and Markov logic networks [24]. In Unsupervised Transfer Learning labelled data in both source and target domains are unavailable, thus the main research issue is to develop an algorithm which transfer knowledge for clustering, dimensionality reduction and other unsupervised learning tasks [24]. Looking at knowledge transfer from the point of view of learning algorithms, we can list works which approach the problem in very different ways. Wu and Dietterich transferred source training examples either as support vectors or as constraints (or both) and demonstrated improved image classification by SVMs [25]. Caruana [26] trained a neural network on several tasks simultaneously as a way to induce efficient internal representation for the target task. Sutton and McCallum [27] demonstrated effective transfer by cascading a class of graphical models with the predictions from one classifier serving as features for the next one in the cascade. Other works involve boosting [17] and k-nearest-neighbor [28].

The issue of 'when to transfer' means in which situations knowledge transfer should be avoided because it would likely affect negatively learning of the new class. The ideal knowledge transfer algorithm should be able to determine automatically if it is worthwhile transferring knowledge or not. A very important aspect of the problem consists in understanding from where to transfer when there exist not just one, but a set of candidate source tasks. Publication focusing on when to transfer evaluate the limit of transfer learning power. Rosenstain et al [29] empirically showed that if two tasks are dissimilar, then brute force transfer hurts the performance of the target task. In [30] an analysis is given for transfer learning using Kolmogorov complexity, where the theoretical bound is proven. In particular, the authors used conditional Kolmogorov complexity to measure relatedness between tasks and transfer the right amount of information in sequential transfer learning in a Bayesian setting. More reacently Eaton et al [31] proposed a model to transfer relationships between tasks to improve inductive learning.

Tommasi et al [15] is one of the first works aiming to address these three aspects at the same time, in a principled framework. The paper face a detection problem where the task is to recognize if a test image belongs to a target object class or not (i.e. belongs to a predefined background class). We proposed a discriminative method based on Least Square Support Vector Machine (LS-SVM) [10] (how to transfer) that learns the new class through adaptation. We defined the prior knowledge as the hyperplanes of the classifiers $\mathbf{w}'_\mathbf{j}$, $j = 1, \ldots, k$ for k classes already learned (what to transfer). Hence knowledge transfer is equivalent to constrain the hyperplanes $\mathbf{w}$ of the $(k+1)$th new category to be close to those of a sub-set of the k classes. We learned the sub-set of classes from where to transfer, and how much to transfer from each of them, via the Leave-One-Out (LOO) error on the training set. Determining how much to transfer helps avoiding negative transfer. Therefore, in case of non-informative prior knowledge, transfer might be disregarded completely (when to transfer). The method was tested successfully in the visual category detection domain.

B. Transfer Learning: The Robotic Perspective

The vast majority of the work on transfer learning in robotics assumes relative homogeneity for the agents, focusing on the problems of conflicting knowledge, differing knowledge levels, and adaptations of the knowledge to the receiving agent's capabilities. This has been studied mostly within the framework of multi-agent systems. Several authors focused on how to determine confidence or trust when interacting with other agents (e.g. [32]). These approaches are based largely on probabilistic representations of previous experience, looking at how to use confidence measures when integrating contributions from different agents.

A more direct attempt to attack the problem of transfer learning across robot platforms was presented in [33], within the context of incremental learning of semantic spatial concepts for mobile platforms. It considered two robot platforms equipped with the same camera, positioned at a different height from the floor, engaged in the same task and with the same processing capability. The approach consisted of a SVM-based framework for modeling concepts, and proposed transfer in the form of stored support vectors from one platform to the other as a way to boost learning. To minimize the risk that transferred knowledge might affect negatively the learning process in the long run, the authors proposed a forgetting strategy to progressively eliminate the transferred vectors, as new data from the actual task become available.

To our knowledge, [8] is the only work so far attempting to study in a coherent framework how robots equipped with different cameras can share experience in order to speed up learning. It focused specifically on differences in sensing and perception, which can be used both for perceptual categorization tasks as well as determining actions based on environmental features. It proposed to abstract raw sensory data into intermediate categories for multiple object features (such as color, texture, shape, etc) by using Gaussian Mixture Models representations, as a way to facilitate effective knowledge transfer. It introduced then a framework to allow robots to build models of their differences with respect to the intermediate representation using joint interaction in the environment, consisting of confusion matrices used to map property pairs between two heterogeneous robots. An information-theoretic metric was introduced to model information loss when going from one robot's representation to another. After this period of joint interaction, the learned models were used to facilitate communication and knowledge transfer in a manner sensitive to the robot's differences. A crucial limitation of this work is that, in order to create abstract internal representations, it assumes that the two robots sharing knowledge are situated in the same place and performing the same task, at the same time. This in practice heavily limits the applicability of the framework.

III. SHARING KNOWLEDGE METHOD: TRANSFER LEARNING ACROSS ROBOTS

Transfer learning focuses on storing knowledge gained while solving some tasks and exploiting it when solving a new, related task. It would be desirable for an autonomous agent, learning how to semantically locate itself in space, to be able to exploit knowledge acquired by an analogous system which already solved a similar problem. Given the difference among the locations the new data will belong to a new probability distribution, in general different from the one previously modeled and stored. Still, as the two robots perform analogous tasks, it is reasonable to expect that the new and the old distributions will be close. It should be possible to use the pre-trained model as starting point when learning on new data.

We propose a method to transfer the knowledge in distinguishing among different locations acquired by robot A to robot B when B starts moving in an environment similar to the one experienced by A. We suppose that robot A is highly reliable in recognizing its position because it worked for a long time acquiring images, classifying them by solving a multiclass problem and storing the obtained classification models. On the other hand, robot B has to define its own place recognition model on the basis of very few images. Instead of starting from scratch, robot B can build its knowledge requiring the new classification model to be close to the one already learned by robot A. The technique should be general enough to allow the transfer of the N-class model of A even when the number of classes seen until now by B is $M < N$. Moreover, if many A robots are available, B should take advantage from all of them in the best possible way according to the similarity between the visited environments.

In the following we summarize the binary transfer learning method on which we build [15] and we present our extension to multiclass using the one-vs-one approach.

A. Model Adaptation

Let us suppose to have k robots each of them perceiving two places in different environments (e.g. robots moving in different hotel rooms where it is possible to recognize two separate areas: bathroom and bedroom). Each robot tackles the place recognition issue as a binary problem acquiring a set of l samples $\{\mathbf{x}_i, y_i\}_{i=1}^{l}$, where $\mathbf{x}_i \in \mathbb{R}^d$ is an input vector describing the i^{th} acquired image and $y_i \in \{-1, 1\}$

is its label. Furthermore each robot learns a linear function $f(\mathbf{x}) = \mathbf{w}' \cdot \phi(\mathbf{x})$ which assigns the correct label to an unseen test sample $\mathbf{x}$. $\phi(\mathbf{x})$ is used to map the input samples to a high dimensional feature space, induced by a kernel function $K(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x}) \cdot \phi(\mathbf{x}')$ [34]. We call $\mathbf{w}'_\mathbf{j}$ the parameter associated to each $j = 1, \dots, k$ robot.

A new robot performing the same task in a new place (e.g. a new hotel room) will learn the corresponding model parameter $\mathbf{w}$ using Least-Square Support Vector Machine (LS-SVM) by solving the following optimisation problem:

$$\min_{\mathbf{w}} \frac{1}{2}\|\mathbf{w}\|^2 + \frac{C}{2}\sum_{i=1}^{l}[y_i - \mathbf{w}\cdot\phi(\mathbf{x}_i)]^2 \ . \tag{1}$$

It can be shown [10] that the optimal $\mathbf{w}$ is expressed by $\mathbf{w} = \sum_{i=1}^{l} \alpha_i \phi(\mathbf{x}_i)$, and $\boldsymbol{\alpha}$ is found solving

$$\left[\mathbf{K} + \frac{1}{C}\mathbf{I}\right]\boldsymbol{\alpha} = \mathbf{y} \ , \tag{2}$$

where $\mathbf{K}$ is the kernel matrix. Let us call $\mathbf{G} = [\mathbf{K} + \frac{1}{C}\mathbf{I}]$, thus the optimisation problem can be solved by simply inverting $\mathbf{G}$. An advantage of the LS-SVM formulation is that it gives the possibility to write the Leave-One-Out (LOO) error in closed form [35]. The LOO error is an unbiased estimator of the classifier generalization error and can be used for model selection [35].

By slightly changing the classical regularization term, it is possible to define a learning method based on adaptation [15]. The idea is to constrain the model learned by the new robot to be close to the set of k pre-trained models [15] already learned by the other robots:

$$\min_{\mathbf{w}} \frac{1}{2}\|\mathbf{w} - \sum_{j=1}^{k}\beta_j\mathbf{w}'_\mathbf{j}\|^2 + \frac{C}{2}\sum_{i=1}^{l}\zeta_i[y_i - \mathbf{w}\cdot\phi(\mathbf{x}_i)]^2 \ , \tag{3}$$

where $\mathbf{w}'_\mathbf{j}$ is the parameter describing each old model and β_j is a scaling factor necessary to control the degree to which the new model is close to the jth old one. To increase robustness to unbalanced distributions of the data, the least-square loss was also weighted with the factor ζ_i depending on the number of positive and negative examples [15]. If we call $\tilde{y}_i$ the LOO prediction, the LOO error in this modified formulation can be written as:

$$r_i^{(-i)} = y_i - \tilde{y}_i = \frac{\alpha_i}{\mathbf{G}_{ii}^{-1}} - \sum_{j=1}^{k}\beta_j\frac{\alpha'_{i(j)}}{\mathbf{G}_{ii}^{-1}} \ , \tag{4}$$

where now $\mathbf{G} = [\mathbf{K} + \frac{1}{C}\mathbf{W}]$ with $\mathbf{W} = diag\{\zeta_1^{-1}, \zeta_2^{-1}, \dots, \zeta_l^{-1}\}$. Here α_i and $\alpha'_{i(j)}$ are respectively elements of the vectors $\boldsymbol{\alpha} = \mathbf{G}^{-1}\mathbf{Y}$ and $\boldsymbol{\alpha_j}' = \mathbf{G}^{-1}\hat{\mathbf{Y}}_\mathbf{j}$ where $\mathbf{Y}$ is the vector of the y_i and $\hat{\mathbf{Y}}_\mathbf{j}$ is the vector of the predictions of the j^{th} known model $\hat{y}_{i(j)} = (\mathbf{w}'_\mathbf{j} \cdot \phi(\mathbf{x}_i))$.

The best vector $\boldsymbol{\beta}$ can be found minimizing the LOO error. To define a convex formulation, it is possible to use the following loss function [15]:

$$\begin{aligned} loss(y_i, \tilde{y}_i) &= \zeta_i \max\left[1 - y_i\tilde{y}_i, 0\right] \\ &= \max\left[y_i\zeta_i\left(\frac{\alpha_i}{\mathbf{G}_{ii}^{-1}} - \sum_{j=1}^{k}\beta_j\frac{\alpha'_{i(j)}}{\mathbf{G}_{ii}^{-1}}\right), 0\right] \ . \end{aligned} \tag{5}$$

This is similar to the hinge loss used in SVM. It is a convex upper bound to the LOO misclassification loss and favours solutions in which $\tilde{y}_i$ has a value of 1, beside having the same sign of y_i. Finally, the objective function is:

$$J = \sum_{i=1}^{l} loss(y_i, \tilde{y}_i) \quad \text{s.t.} \quad \|\boldsymbol{\beta}\|_2 \le 1 \ . \tag{6}$$

B. One-vs-One multiclass extension

Let us assume that k robots perceive N different places (e.g. each robot is driven in one of k distinct apartments, each with N rooms). All the robots are equipped with a learning system based on LS-SVM using the one-vs-one multiclass extension to solve the N-class problem. After an initial stage of data acquisition, the robots learn how to classify the rooms and for each agent all the $N(N-1)/2$ hyperplanes are saved: $\mathbf{w}'_\mathbf{j,c}$ with $j = 1, \dots, k$ and $c = 1, \dots, N(N-1)/2$. A new robot is then asked to solve a similar N-class task and it is allowed to use as prior knowledge the pre-trained models of the other agents. Starting from Equation (3), each classification problem between a pair of classes becomes:

$$\min_{\mathbf{w_c}} \frac{1}{2}\|\mathbf{w_c} - \sum_{j=1}^{k}\beta_{j,c}\mathbf{w}'_\mathbf{j,c}\|^2 + \frac{C}{2}\sum_{i=1}^{l}\zeta_i[y_i - \mathbf{w_c}\cdot\phi(\mathbf{x}_i)]^2 \ , \tag{7}$$

and is solved separately according to the method described in the previous section. If the new robot can access only a subset M of the N classes in the training phase, $M(M-1)/2$ classification models will be defined as above, while for the remaining $(N-M)$ classes, $\mathbf{w_c}$ will be equal to $\sum_{j=1}^{k}\mathbf{w}'_\mathbf{j,c}$. When applied to a test point, one-vs-one classification is performed by max-wins voting strategy. Each classifier assign the instance to one of the two classes and the vote for the assigned class is increased by one; finally the class with most votes determines the instance label.

IV. EXPERIMENTAL SETUP

In this section we describe the experimental setup used for testing our transfer learning algorithm. Section IV-A describes the four databases used; Section IV-B briefly reviews the feature descriptor chosen for all experiments, and Section IV-C describes how we chose the internal parameters of the transfer learning algorithm, as well as the baselines against which we evaluate our results.

A. Databases

We evaluated our transfer learning approach on the semantic place recognition problem and chose four different databases, each representing different scenarios of increasing difficulty.

1) The IDOL 2 Database: The IDOL 2 Database [11] contains 24 image sequences acquired using a perspective camera mounted on two mobile robot platforms (a MobileRobots PeopleBot and a PowerBot). The acquisition was performed within an indoor environment consisting of five rooms of different functionality: One-person Office (OO), Two-person Office (TO), CoRridor (CR), KiTchen (KT) and Printer Area (PA). The sequences were recorded under various weather and illumination conditions (sunny, cloudy and night) and across a time span of six months. Both mobile robot platforms are equipped with the pan-tilt-zoom Canon VC-C4 camera. The cameras on the two robots are mounted at different heights. For each robot platform and for each type of illumination conditions, there are four sequences recorded. The first two were acquired six months before the last two. This means that for each robot and for every illumination condition, there are two sequences acquired under similar conditions and two sequences acquired under very different conditions. In all the experiments we used the sequences acquired under similar conditions, one for training and the other for test. Figure 3 shows exemplar images acquired by the two robot platforms, for all rooms.

2) The RobotVision@ImageCLEF 2010 Database: This database was created for the Robot Vision Task of the ImageCLEF Challenge 2010 [12]. It has been acquired using a MobileRobots Powerbot robot platform equipped with a stereo camera system consisting of two Prosilica GC1380C cameras. The robot was manually driven

Fig. 3. Exemplar images from the IDOL database, acquired by the PeopleBot (top row) and the PowerBot (bottom row).

through several rooms and different floors of a typical indoor office environment under fixed illumination conditions. We considered the robot as a monocular vision system dealing with the images acquired by only one of the two cameras. The model built to classify six rooms (CorRidor(CR), LargeOffice (LO), PrinterArea (PA), RecycleArea (RA), SmallOffice (SO), BatHroom (BH), all frames extracted from the training set of the database) on one floor defines our prior knowledge, which is then exploited when learning the same multiclass problem on a different floor (corresponding to frames extracted from the validation set of the database). Figure 4 shows exemplar images for both floors.

3) The VPC Database : The VPC database [13] consists of room images acquired from 6 different houses using a camcorder (JVC GR-HD1) mounted on a rolling tripod to mimic a mobile robot platform. During the acquisitions the blinds were always closed and artificial light used: this helped to normalize the illumination conditions across homes and times of the day. We consider the five room categories that exist in all homes (BeDroom (BD) , BatHroom (BH), KiTchen (KT), Living Room (LR), and Dining Room (DR)) and we learn a multiclass model separately for five houses. Images coming from all the rooms of the sixth house constitute the new learning task. Figure 5 shows images for all rooms categories, for all houses.

4) The COLD Database: The COLD database [14] contains three separate subsets acquired at different indoor laboratory environments located in three European cities: Ljubljana, Freiburg and Saarbrüchen. The sequences in the database were recorded using three different mobile robots (an ActivMedia PeopleBot, an ActivMedia Pioneer-3 and an iRobot ATRV-Mini), all equipped with two Videre Design MDCS2 digital cameras, one for perspective and one for omnidirectional images. The heights of the cameras varied depending on the robot platform. At each laboratory the image sequences were acquired under different illumination conditions (cloudy, sunny, night) and across several days. Special care was taken in the choice of the rooms to acquire. Therefore, for each laboratory there exists a set of sequences containing 5 rooms with similar functionalities that are also contained in the other two (CorRidor (CR), Two-persons Office (TO), Printer Area (PA), KiTchen (KT), BatHroom (BH)). In the experiments prior knowledge corresponds to the multiclass model learned by the robot in Freiburg using the perspective camera. Two possible new set of data are then considered changing the location (laboratory in Ljubljana, perspective camera) or the visual acquisition system (laboratory in Freiburg, omnidirectional camera). Figure 6 shows sample images for all rooms, different locations and different types of camera.

B. Feature Descriptor

As features, we opted for a state of the art histogram-based global feature in the spatial-pyramid scheme introduced in [36]. This representation scheme was chosen because it combines the structural and statistical approaches: it takes into account the spatial distribution of features over an image, while the local distribution is in turn estimated by mean of histograms; moreover it has proven to be versatile and to achieve higher accuracies in our experiments. Specifically, we used the PHOG features [37] that captures the distribution of edge orientations within an image (computed on the output of a Canny edge detector) and can be extracted in two different variants:

1) with the range of orientations equal to $[0, 180]$ (the sign of the gradient is ignored);
2) with orientations in the range $[0, 360]$.

The orientations range is then quantized into K bins and each edge assigned to the corresponding binned orientation, with a weight proportional to the value of the gradient. Here we used the following parameters: number of pyramid levels $L = 3$, number of histogram bins $K = 40$, angle range $\theta = [0, 360]$. Figure 7 shows two examples of features computed from indoor scene images.

C. Learning Algorithms

For all experiments and all learning algorithms, we used the RBF kernel. The hyperparameters (γ for the RBF kernel and C for the learning problem) were found through cross validation on the prior knowledge.

V. EXPERIMENTAL RESULTS

We present here three set of experiments designed to study the behaviour of our transfer learning method in various scenarios: (1) platforms with *similar* characteristics solving the same *specific* task; (2) platforms with the *same* characteristics solving the same *categorical* task, and (3) platforms with *similar* characteristics solving the same *categorical* task. Figure 8 illustrates these three settings.

We name our transfer learning method **Adapt-Transfer** and we benchmark it against two baselines. **No-Transfer**: it corresponds to learning from scratch just using the new data acquired by the robot. **Prior**: this means to take directly the models learned by the prior knowledge of one (or more) robot(s) and apply them on the new learning agent without any adaptation. We also compare our approach with the one presented in [33], that we call here **SV-Transfer**.

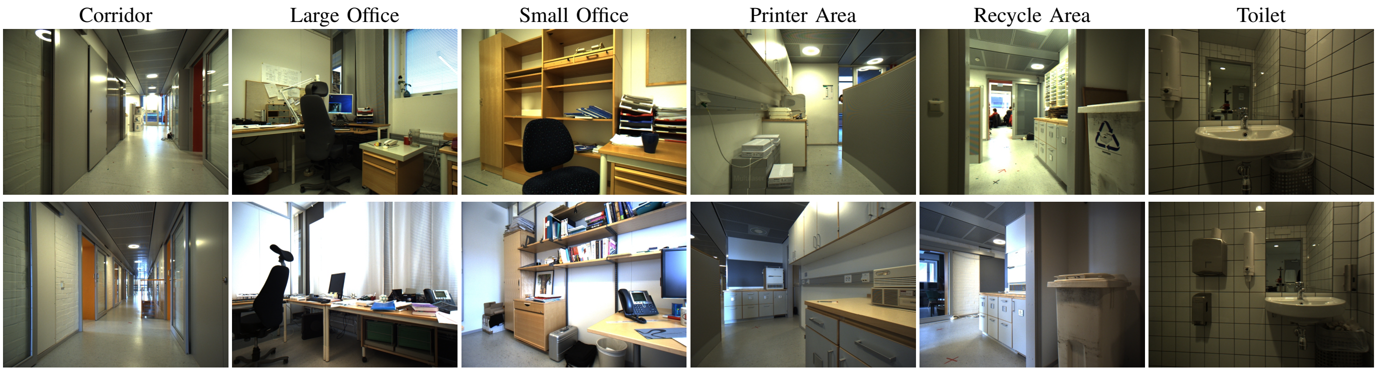

Fig. 4. Sample images from the RobotVision@ImageCLEF 2010 database. Each row corresponds to room on a different floor.

Fig. 5. Exemplar images from the VPC database, for all six houses, one per each row.

A. Transfer learning across platforms with similar characteristics, solving the same specific task

In these experiments we consider to transfer visual knowledge across similar robotic platforms (the only difference is in the height of the camera position), placed in the same environment.

We used the IDOL 2 database and we first chose the intermediate cloudy condition. Four different experiments were run alternating the two robots in the role of prior knowledge and new learning agent, and considering both the sequences recorded with 6 months time distance. The system was updated incrementally in a room by room (i.e. class by class) scenario considering the same room order presented in [33]. As there are five classes in total, training was performed in four steps, starting from two classes, while testing always run on a five-class set.

Fig 9 shows the experimental results obtained at each step. Adapt-Transfer performs better with respect to the other benchmark methods. It is important to remind that Prior corresponds to using the models learned from a sequence acquired under the same illumination and at close time as the training one, but recorded by a similar platform. In [33] the consistency in illumination conditions was the requirement under which all the experiments were run. Moreover the comparison between SV-Transfer and Prior was not shown. Here it is evident that mixing the known Support Vectors with the new training samples does not always guarantee an advantage respect to the direct

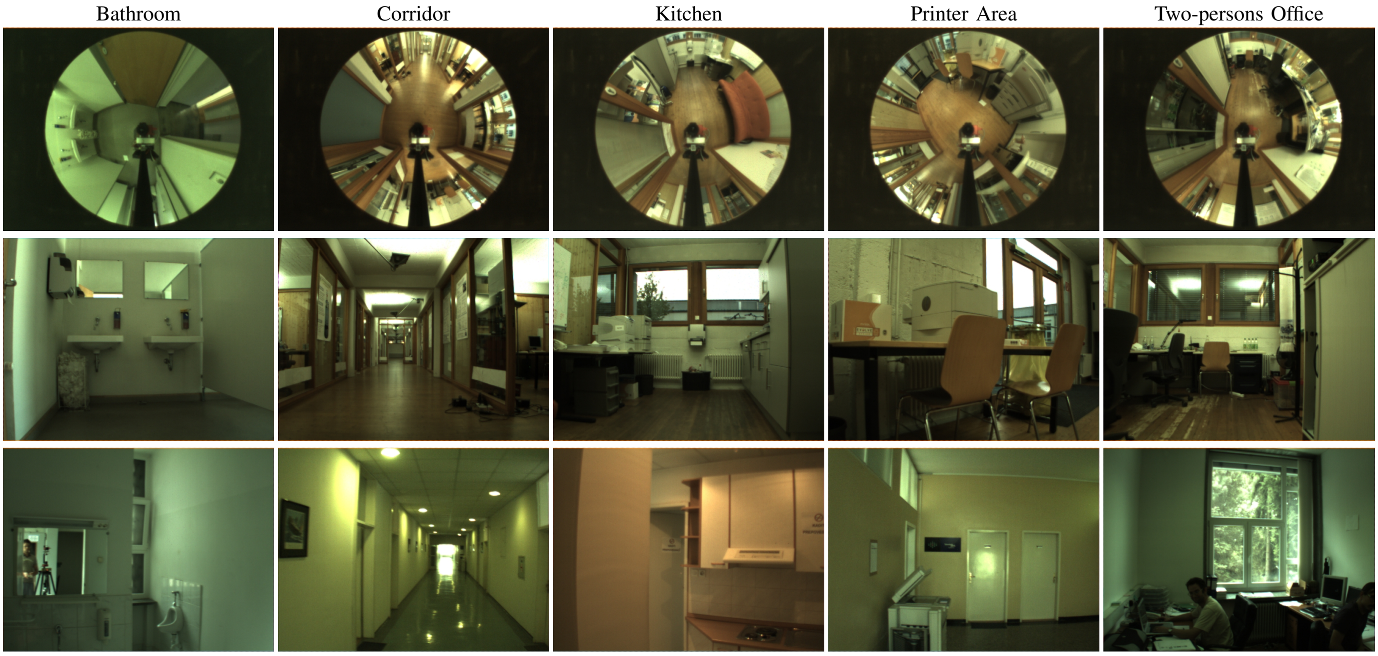

Fig. 6. Exemplar images from the COLD database, acquired using an omnidirectional (top row, Freiburg) and a perspective (middle row, Freiburg; bottom row Ljubljana) camera.

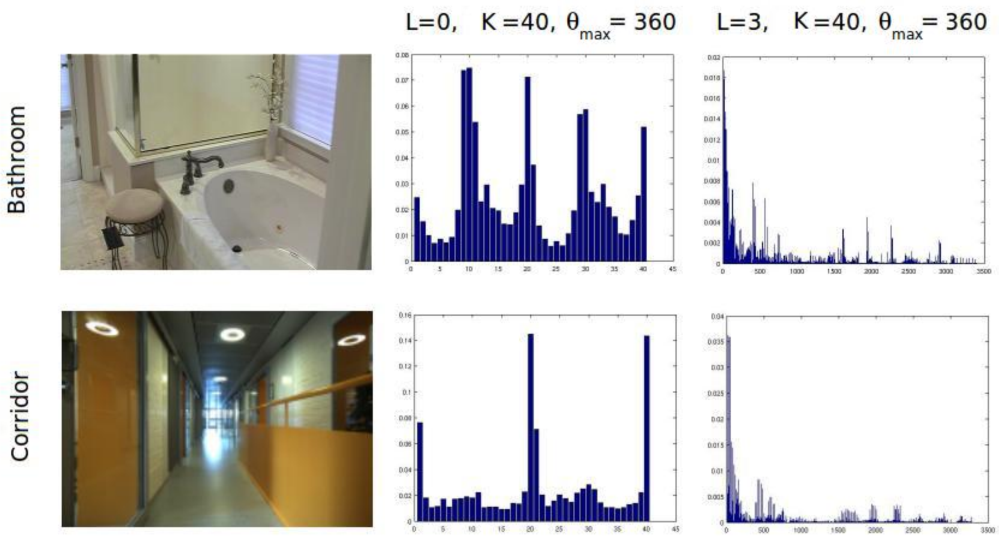

Fig. 7. Examples of features computed at different spatial resolutions, for two different environments.

use of prior knowledge. On the other hand Adapt-Transfer results are always equal or better than Prior.

We also investigated the case where transfer is performed across different illumination conditions. We kept the model learned by one of the robots in the cloudy environment as prior knowledge and we supposed that the new agent tries to exploit this information when learning in sunny weather (Figure 10(left)) and at night (Figure 10(right)). As expected the more the illumination conditions are similar, the higher is the advantage in learning: transferring from cloudy to sunny results in a higher gain in performance with respect to transferring from cloudy to night.

B. Transfer learning across platforms with the same characteristics solving the same categorical task

In this set of experiments we show that the visual knowledge acquired by one robot can be reused when the robot itself is moved in a different location, where it is asked to solve the same categorical task.

We considered the RobotVision database and the problem of learning to classify rooms on one floor when prior knowledge was acquired on a different floor, with rooms belonging to the same semantic categories. We suppose that on the first floor the robot is free to observe all the rooms from many different viewpoints and performing 360 degree turns. When moved on a different floor the robot can either (1) acquire information from one room at the time with the order of the rooms taken randomly or (2) have at

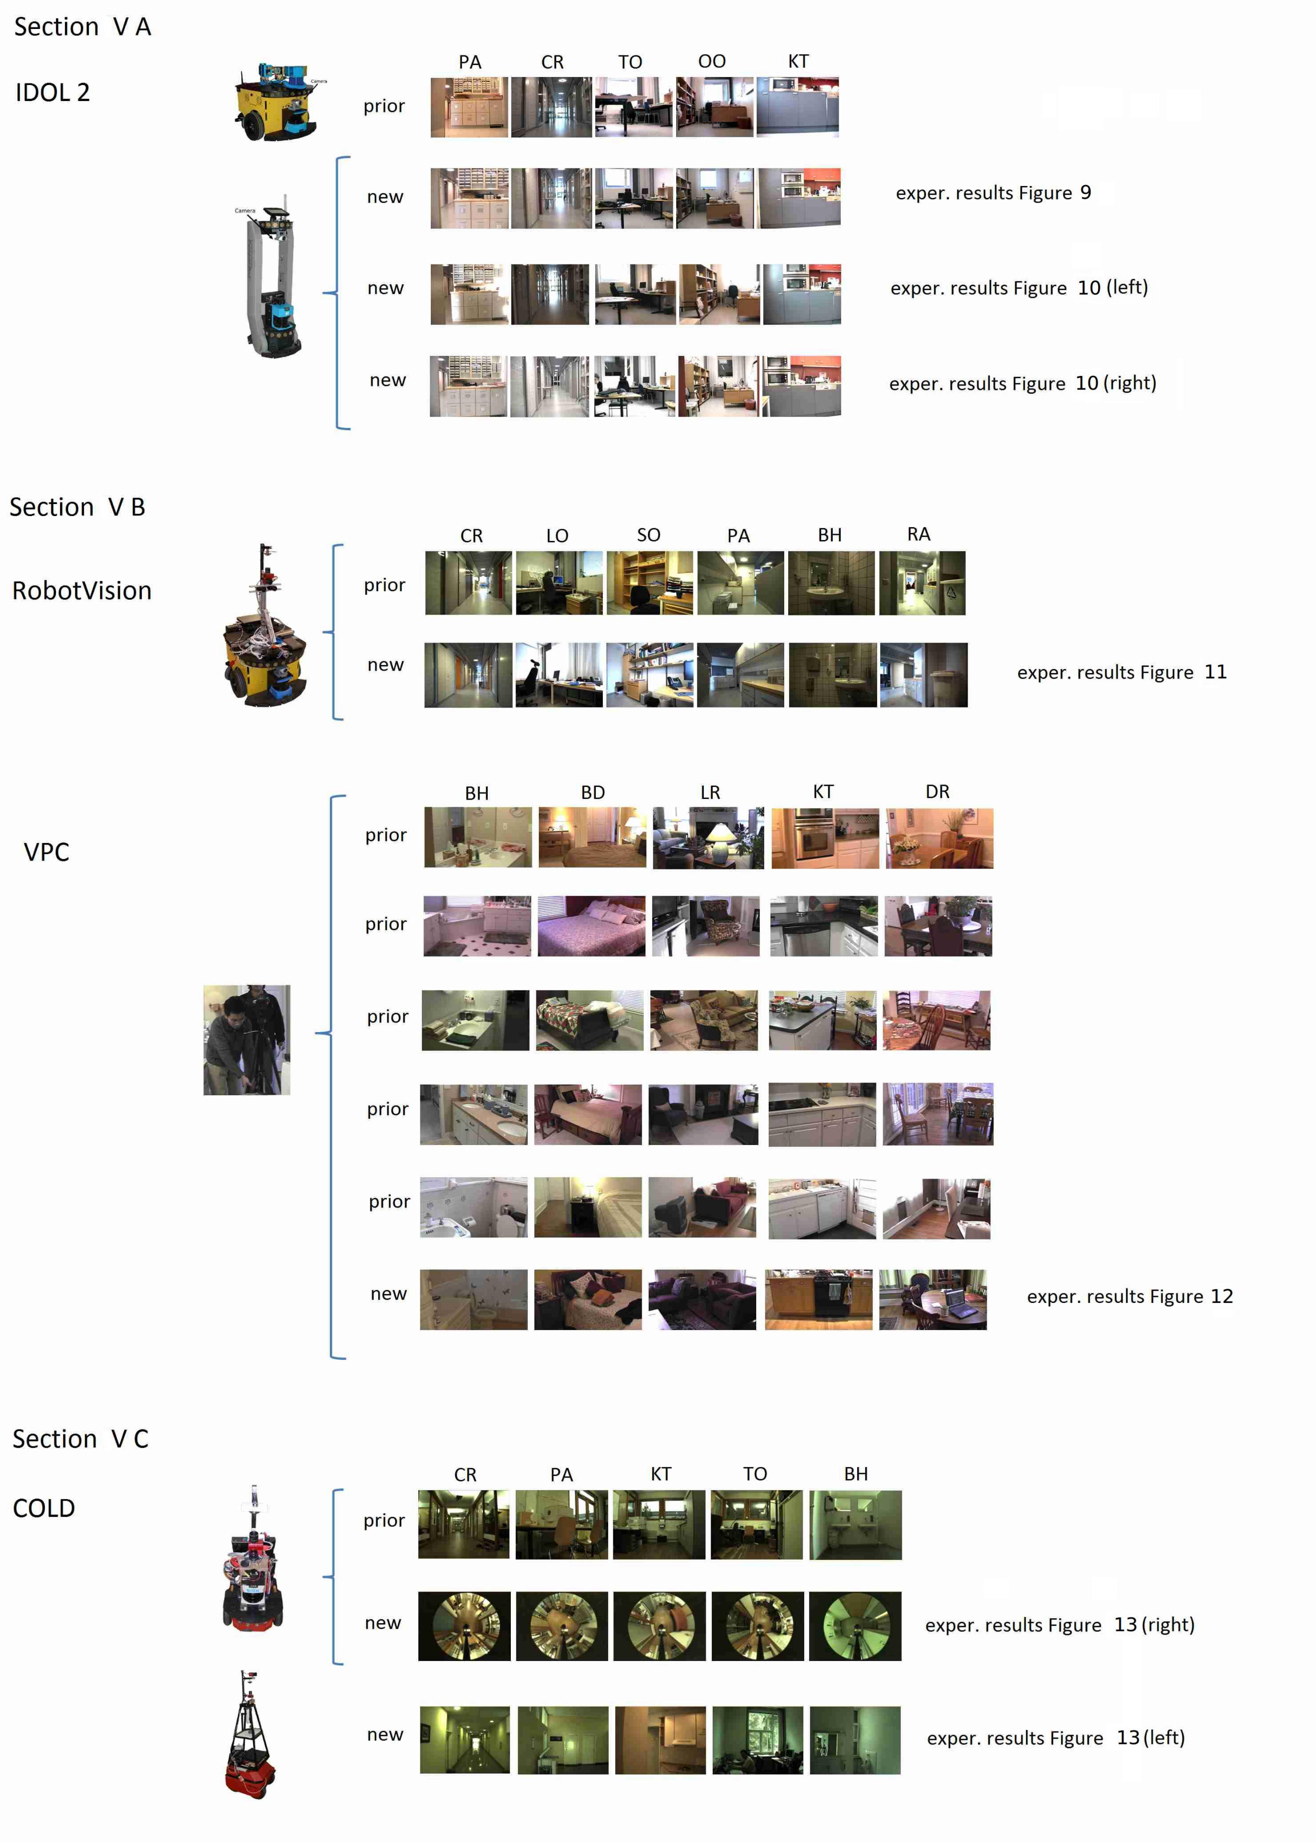

Fig. 8. A schematic representation of the different experimental settings. The rooms are: PA: printer area; CR: corridor; TO: two-person office; OO: one-person office; KT: kitchen; LO: large office; RA: recycle area; SO: small office; BH: bathroom; BD: bedroom; LR: living room; DR: dining room.

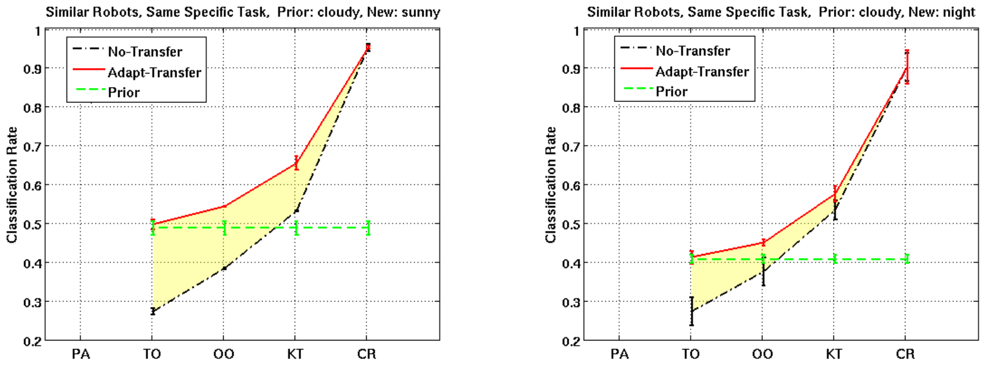

Fig. 10. IDOL2 database, changing illumination conditions. Classification rate at each training step corresponding to images of a new room entering the system. Average $\pm$ standard deviation on two different permutations obtained fixing the acquisition time and alternating the two available sequences of the new robot as training and test set. The area between Adapt-Transfer and No-Transfer is shaded to show the different gain when transferring from cloudy to sunny (left) or night (right).

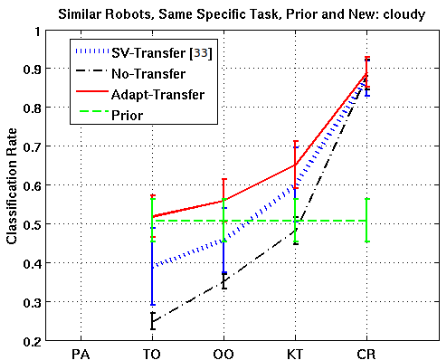

Fig. 9. IDOL2 database, fixed illumination conditions. Classification rate at each training step corresponding to images of a new room entering the system. Average $\pm$ standard deviation on four different permutations, run alternating the role of the two robots and considering different acquisition time.

the beginning a quick look to all the rooms and then it can repeat the whole acquisition collecting more samples all over the floor in other two rounds. Moreover, regardless of the order of the image acquisition, the robot can never make the 360 degree turn inside the rooms of the new floor.

Figure 11(left) reports the results obtained in the setting (1), while Figure 11(right) shows the corresponding results for the setting (2). In both cases Adapt-Transfer clearly outperforms Prior and No-Transfer. We notice that in the room by room experiments, No-Transfer does not follow the same increasing trend which was evident in the results of the previous section. Here exactly the same robot is used both for learning prior knowledge and in the new task, moreover there is a minimal difference between the two floors on which the data are acquired. This means that prior knowledge is extremely relevant for the new learning process which is based on a much smaller number of samples respect to that used to build the prior models (on the other hand the number of samples for the experiments of the previous section was approximately the same for old and new knowledge).

To generalize the described scenario we decided to run another set of experiments on the VPC database. In this case we have five different prior knowledges corresponding to the same robot acquiring images of rooms in five houses. The robot is then moved in a sixth house and required to learn the same categorical task. From the whole dataset we picked house 2, 4 and 6, and we run three experiments each time considering one of them as the new house and all the remaining five as prior knowledge. We reproduced the same setting (1) and (2) described above, the corresponding results are reported respectively in Figure 12(left) and 12(right). The prior knowledge results are comparable to that presented in [13] (Table IV, CENTRIST with filtering) although here we are using a different set of features and a discriminative classification approach without any temporal integration on the frames. Transferring through adaptation produces a small advantage with respect to both learning from scratch or using directly prior knowledge in the room by room experiment. On the other hand, there is no gain in performance in the all rooms experiment. The proposed task is extremely difficult due to the differences among the houses. Although Adapt-Transfer is mixing five different knowledge sources, giving a specific weight to each of them when learning every binary one-vs-one problem, all the weights are small and the average recognition rate is equivalent to No-Transfer.

C. Transfer learning across platforms with similar characteristics solving the same categorical task

In this final set of experiments we want to analyze the case of transferring visual knowledge across similar robot platforms solving the same categorical task.

We used the COLD database and considered two possible cases. In the first one we fix the visual acquisition setting supposing that two robots record images using the same perspective camera mounted at different height, but the acquisitions are performed in two different locations. In the second one, the location is fixed but one of the robot is using a perspective camera while the other an omnidirectional camera. In both cases the illumination condition is the same for the two robots.

Results for the room by room experiments with the order of the rooms taken randomly, are reported respectively in Figure 13(left) and 13(right). In both cases the available prior knowledge is extremely different from what we want to learn as new task. This is the typical condition where applying a blind knowledge transfer can hurt the learning performance. Here the results show that when prior-knowledge is not reliable, Adapt-Transfer performs almost always as learning from scratch, automatically avoiding negative transfer.

D. Discussion

All experiments, for all scenarios, indicate that our algorithm is capable of exploiting the available priors and achieve a boost in performance, whenever the priors contain information which is relevant

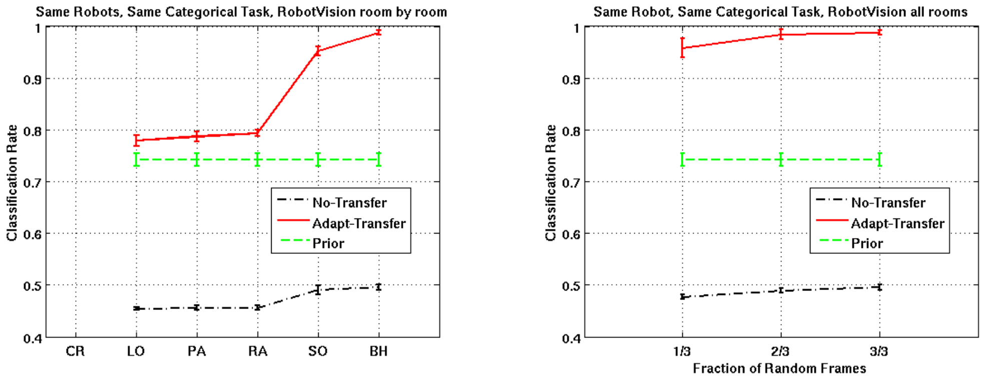

Fig. 11. RobotVision database. Left: classification rate at each training step corresponding to images of a new room entering the system. Right: classification rate when the new sequence is randomized and separated into three parts that progressively enter the system. Average $\pm$ standard deviation on three different runs obtained dividing the RobotVision validation set randomly in 75% for training and 25% for test.

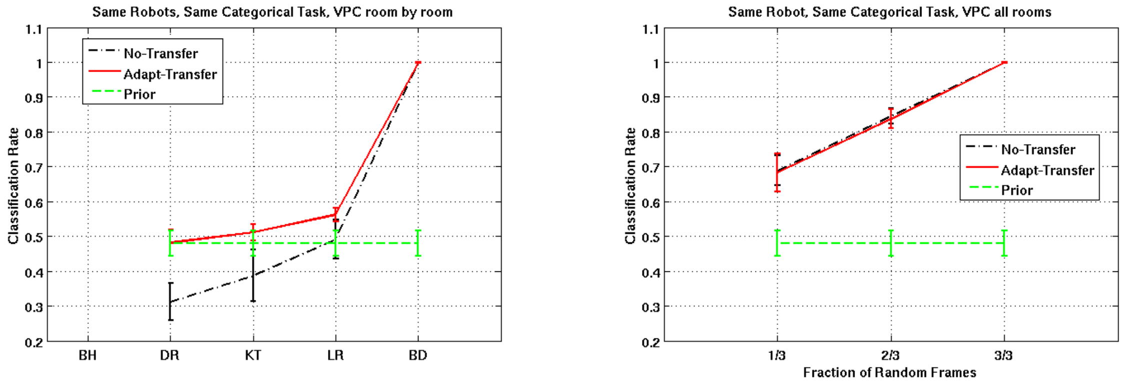

Fig. 12. VPC database. Left: classification rate at each training step corresponding to images of a new room entering the system. Right: classification rate when the new sequence is randomized and separated into three parts that progressively enter the system. Average $\pm$ standard deviation on three different runs obtained considering House 2,4 and 6 as the new tasks and randomly using 75% of the corresponding house image samples for training and 25% for test.

for learning the new task. With the exception of the COLD database (which is probably the most challenging indoor place recognition database publicly available as of today), the gain in performance is always very pronounced at the first learning steps, when the amount of training data is lower. This confirms the importance of transfer learning approaches when learning from few annotated data, a highly relevant scenario for artificial autonomous systems.

When the available priors are not very informative, the transfer learning algorithm tends to behave like the no transfer baseline. This behavior is particularly evident in the experiments on the VPC and the COLD database, where the priors achieve performances around 50%. It is worth stressing that both databases have been created to support research on the indoor place categorization problem, which as of today is still an open research issue.

As a last point, we discuss the computational cost of using our transfer learning algorithm as opposed to the no transfer baseline. For each of the binary problems in the one-vs-one multiclass approach the computational complexity of Adapt-Transfer is $\mathcal{O}(l^3+kl^2)$ with l the number of training examples and k the number of available robots used as prior knowledge. The first term is releated to inverting $\mathbf{G}$ while the second term is the computational complexity of (4). This is to be compared with the computational complexity of a classic LS-SVM which is $\mathcal{O}(l^3)$: the extra computational cost of Adapt-Transfer is linear in the number of prior knowledge robots and for few training samples it is negligible compared to the potential benefit in performance.

VI. SUMMARY AND CONCLUSION

This paper presents an algorithm for transfer learning of visual concepts across robot platforms. Our working assumption is to have robots, possibly located in different settings, that have to solve the same task, whether at the specific or at the categorical level. We assume that all robots are equipped with a camera, but we do not make any assumption on the type of camera, nor on where it is positioned. We also assume that the robots share the same processing capabilities. Our algorithm is based on Least Square Support Vector Machine, and enable one robot to learn a new task by exploiting what the other robots have already learned about that task via adaptation. Very extensive experiments on four different publicly available databases show the power of our approach.

This work can be extended in many ways. First, although here we focused on visual concepts, the algorithm in its current form can handle different types of sensors. Second, we will extend the algorithm so to relax the assumption that all robots share the same features and classifier. This could be achieved by casting the problem within the Multi Kernel Learning framework. Lastly, transfer learning for artificial cognitive systems should be seen as a component of the life long learning strategies necessary to achieve real autonomy. The current version of our algorithm does not allow to automatically stop the transfer process, once the priors have been exploited and/or the new concept has been learned. Preliminary work in this direction shows that the internal parameters of the algorithm might provide useful insights in that direction [38]. Future work will focus in these directions.

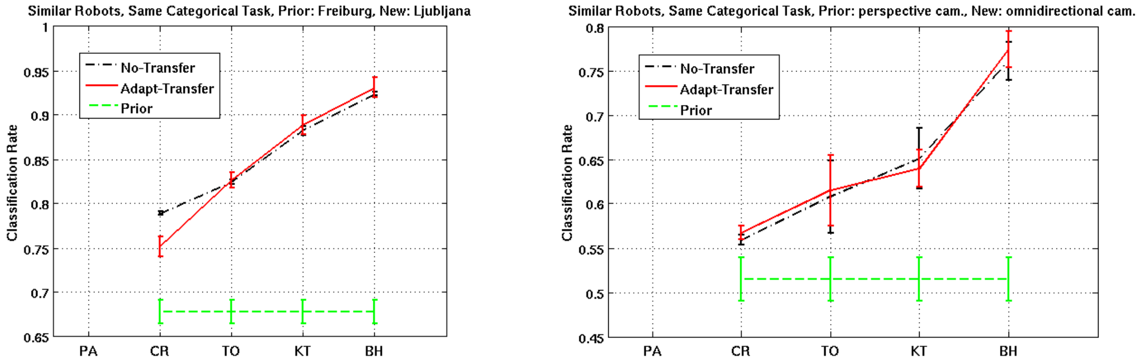

Fig. 13. COLD database. Classification rate at each training step corresponding to images of a new room entering the system. Left: prior and new knowledge from two different laboratories. Right: pior and new knowledge from the same laboratory but registered with different video devices. Average ± standard deviation on two different runs obtained considering two sequences aquired under cloudy illumination conditions and close time, alternated as training and testing set.

REFERENCES

[1] B. Kuipers, P. Beeson, J. Modayil, and J. Provost, "Bootstrap learning of foundational representations," *Connection Science*, vol. 18, no. 2, pp. 145–158, 2006.

[2] L. A. Olsson, C. L. Nehaniv, and D. Polani, "From unknown sensors and actuators to actions grounded in sensorimotor perceptions," *Connection Science*, vol. 18, no. 2, pp. 121–144, 2006.

[3] C. C. Kemp and A. Edsinger, "What can I control?: The development of visual categories for a robot's body and the world that it influences," in *ICDL*, 2006.

[4] J. Modayil, "Robot developmental learning of an object ontology grounded in sensorimotor experience," Ph.D. dissertation, Computer Science Dept., University of Texas at Austin, 2007.

[5] E. A. Topp and H. I. Christensen, "Topological modelling of human augmented mapping," in *Proc. International Conference on Intelligent Robots and Systems (IROS)*, 2006, pp. 2257–2263.

[6] H. I. Christensen, A. Sloman, G. Kruijff, and J. Wyatt, Eds., *Cognitive Systems*. Berlin, DE: Springer Verlag, 2009.

[7] J. Peltason, F. H. Siepmann, T. P. Spexard, B. Wrede, M. Hanheide, and E. A. Topp, "Mixed-initiative in human augmented mapping," in *Proc. International Conference on Robotics and Automation (ICRA)*, 2009.

[8] Z. Kira, "Inter-robot transfer learning for perceptual classification," in *Proc. International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2010, pp. 13–20.

[9] S. J. Pan and Q. Yang, "A survey on transfer learning," Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong, China, Tech. Rep., November 2008.

[10] J. Suykens, T. V. Gestel, J. D. Brabanter, B. D. Moor, and J. Vanderwalle, *Least Squares Support Vector Machines*. World Scientific, 2002.

[11] J. Luo, A. Pronobis, B. Caputo, and P. Jensfelt, "The KTH-IDOL2 database," KTH, CAS/CVAP, Tech. Rep. 304, 2006, available at `http://cogvis.nada.kth.se/IDOL2/`.

[12] A. Pronobis, M. Fornoni, H. I. Christensen, and B. Caputo, "The robot vision track at imageclef 2010," in *CLEF (Notebook Papers/LABs/Workshops)*, 2010, database available at `http://www.imageclef.org/2010/robot`.

[13] J. Wu, H. I. Christensen, and J. M. Rehg, "Visual place categorization: problem, dataset, and algorithm," in *Proc. International Conference on Intelligent Robots and Systems (IROS)*, 2009, pp. 4763–4770, database available at `http://categorizingplaces.com/dataset.html`.

[14] A. Pronobis and B. Caputo, "Cold: The cosy localization database," *International Journal on Robotic Research*, vol. 28, no. 5, pp. 588–594, May 2009.

[15] T. Tommasi, F. Orabona, and B. Caputo, "Safety in numbers: Learning categories from few examples with multi model knowledge transfer," in *Proc. Computer Vision and Pattern recognition (CVPR)*, 2010, pp. 3081–3088.

[16] S. Thrun, "Is learning the n-th thing any easier than learning the first?" in *Advances in Neural Information Processing Systems*. The MIT Press, 1996, pp. 640–646.

[17] Y. Yao and G. Doretto, "Boosting for transfer learning with multiple sources," in *Proceedings of the IEEE Computer Vision and Pattern Recognition (CVPR) Conference*, 2010.

[18] M. Sugiyama, S. Nakajima, H. Kashima, P. von Bünau, and M. Kawanabe, "Direct Importance Estimation with Model Selection and Its Application to Covariate Shift Adaptation," in *Proceedings of the Neural Information Processing Systems (NIPS) Conference*, 2008.

[19] H. Daumé III, "Frustratingly easy domain adaptation," in *ACL*, 2007.

[20] R. Raina, A. Battle, H. Lee, B. Packerand, and A. Y. Ng, "Self-taught learning: transfer learning from unlabeled data," in *Proceedings of the International Conference on Machine Learning (ICML)*, 2007.

[21] A. Quattoni, M. Collins, and T. Darrell, "Transfer learning for image classification with sparse prototype representations," in *Proceedings of the IEEE Computer Vision and Pattern Recognition (CVPR) Conference*, 2008.

[22] M. Stark, M. Goesele, and B. Schiele, "A shape-based object class model for knowledge transfer," in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2009.

[23] L. Fei-Fei, R. Fergus, and P. Perona, "One-shot learning of object categories," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 28, pp. 594–611, 2006.

[24] E. Soria, J. Martin, R. Magdalena, M. Martinez, and A. Serrano, *Handbook of Research on Machine Learning Applications*. IGI Global, 2009, ch. L. Torrey and J. Shavlik, Transfer Learning.

[25] P. Wu and T. G. Dietterich, "Improving svm accuracy by training on auxiliary data sources," in *Proceedings of the International Conference on Machine learning (ICML)*. Morgan Kaufmann, 2004, pp. 871–878.

[26] R. Caruana, "Multitask learning," 1997.

[27] C. Sutton and A. McCallum, "Composition of conditional random fields for transfer learning," in *HLT '05: Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, 2005, pp. 748–754.

[28] K. Saenko, B. Kulis, M. Fritz, and T. Darrell, "Adapting visual category models to new domains," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2010.

[29] M. Rosenstein, Z. Marx, and L. P. Kaelbling, "To transfer or not to transfer," in *Proceedings of the Neural Information Processing Systems (NIPS) Conference, Workshop on Transfer Learning*, 2005.

[30] M. M. Mahmud and S. Ray, "Transfer learning using kolmogorov complexity: Basic theory and empirical evaluations," in *Proceedings of the Neural Information Processing Systems (NIPS) Conference*, 2007.

[31] E. Eaton, M. desJardins, and T. Lane, "Modeling transfer relationships between learning tasks for improved inductive transfer," in *Proceedings of the European Conference on Machine Learning (ECML)*, 2008.

[32] D. C. R. Becker, "Determining confidence when integrating contributions from multiple agents." in *Proc. International Joint conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2007, pp. 449–456.

[33] J. Luo, A. Pronobis, and B. Caputo, "Svm-based transfer of visual knowledge across robotic platforms," in *Proc. of the International Conference in Computer Vision Systems (ICVS)*, 2007.

[34] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines*. Cambridge University Press, 2000.

[35] G. Cawley, "Leave-one-out cross-validation based model selection cri-

teria for weighted LS-SVMs," in *Proc. International Joint Conference on Neural Networks (IJCNN)*, 2006, pp. 1661–1668.

[36] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006, pp. 2169–2178.

[37] A. Bosch, A. Zisserman, X., and Munoz, "Representing shape with a spatial pyramid kernel," in *Proc. ACM international Conference on Image and Video Retrieval (CIVR)*, 2007, pp. 401–408.

[38] T. Tommasi and B. Caputo, "Towards a quantitative measure of rareness," in *DIRAC Workshop at the European Conference on Machine Learning (ECML)*, 2010.