

# CURL: Image Classification using Co-training and Unsupervised Representation Learning

Simone Bianco<sup>a</sup>, Gianluigi Ciocca<sup>a</sup>, Claudio Cusano<sup>b</sup>

<sup>a</sup>*Department of Informatics, Systems and Communications, University of Milano-Bicocca,  
Milano, 20126 Italy  
e-mail: {bianco, ciocca}@disco.unimib.it.*

<sup>b</sup>*Department of Electrical, Computer and Biomedical Engineering, University of Pavia,  
Pavia, 27100 Italy  
e-mail: claudio.cusano@unipv.it*

---

## Abstract

In this paper we propose a strategy for semi-supervised image classification that leverages unsupervised representation learning and co-training. The strategy, that is called CURL from Co-training and Unsupervised Representation Learning, iteratively builds two classifiers on two different views of the data. The two views correspond to different representations learned from both labeled and unlabeled data and differ in the fusion scheme used to combine the image features.

To assess the performance of our proposal, we conducted several experiments on widely used data sets for scene and object recognition. We considered three scenarios (inductive, transductive and self-taught learning) that differ in the strategy followed to exploit the unlabeled data. As image features we considered a combination of GIST, PHOG, and LBP as well as features extracted from a Convolutional Neural Network. Moreover, two embodiments of CURL are investigated: one using Ensemble Projection as unsupervised representation learning coupled with Logistic Regression, and one based on LapSVM. The results show that CURL clearly outperforms other supervised and semi-supervised learning methods in the state of the art.

### *Keywords:*

Image classification, machine learning algorithms, pattern analysis, semi-supervised learning.

---

## 1. Introduction

Semi-supervised learning [1] consists in taking into account both labeled and unlabeled data when training machine learning models. It is particularly effective when there is plenty of training data, but only a few instances are labeled. In the last years, many semi-supervised learning approaches have been proposed [2] including generative methods [3, 4], graph-based methods [5, 6], and methods based on Support Vector Machines [7, 8]. Co-training is another example of semi-supervised technique [9]. It consists in training two classifiers independently which, on the basis of their level of confidence on unlabeled data, co-train each other through the identification of good additional training examples. The difference between the two classifiers is that they work on different *views* of the training data, often corresponding to two feature vectors. Pioneering works on co-training identified the conditional independence between the views as the main reason of its success. More recently, it has been observed that conditional independence is a sufficient, but not necessary condition, and that even a single view can be considered, provided that different classification techniques are used [10].

In this work we propose a semi-supervised image classification strategy which exploits unlabeled data in two different ways: first two image representations are obtained by unsupervised representation learning (URL) on a set of image features computed on all the available training data; then co-training is used to enlarge the labeled training set of the corresponding co-trained classifiers (C). The difference between the two image representations is that one is built on the combination of all the image features (*early fusion*), while the other is the combination of sub-representations separately built on each feature (*late fusion*). This strategy has the advantages that only a set of features must be defined and in the co-training the different fusion schemes are used as different

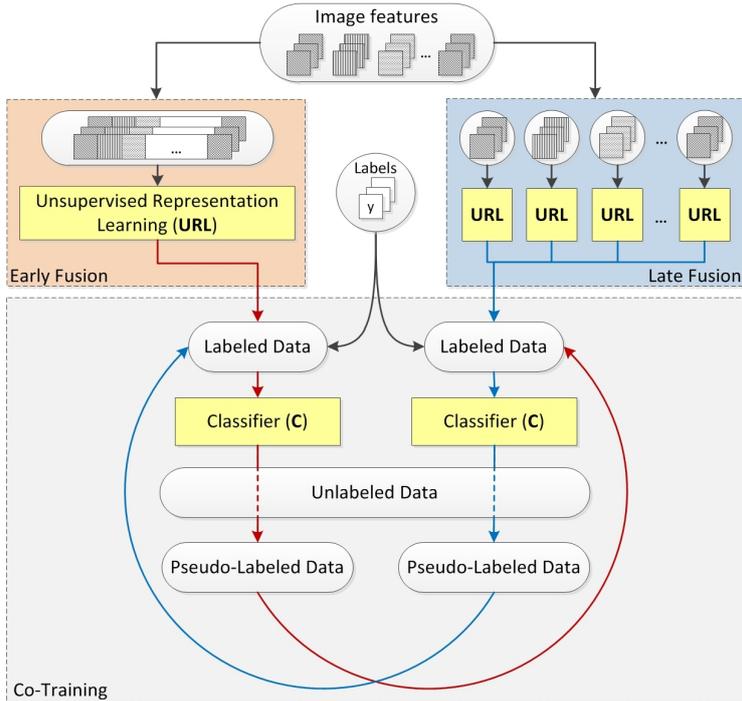


Figure 1: Schema of the proposed strategy.

views. Moreover, by exploiting both representation learning and co-training, classification can be effectively performed when very few labeled data are available. We call the proposed strategy CURL from the combination of C and URL components. The schema of CURL is illustrated in Fig. 1.

In standard co-training each classifier is built on a single view, often corresponding to a single feature. However, the combination of multiple features is often required to recognize complex visual concepts [11, 12, 13]. Both the classifiers built by CURL exploit all the available image features in such a way that these concepts can be accurately recognized. We argue that the use of two different fusion schemes together with the non-linear transformation produced by the unsupervised learning procedure, makes the two image representations uncorrelated enough to allow an effective co-training of the classifiers.

The proposed strategy is built on two base components: the unsupervised representation learning, and the classifier used in co-training. By changing these two components we can have different embodiments of CURL that can be experimented and evaluated.

To assess the merits of our proposal we conducted several experiments on widely used data sets: the 15-scene data set, the Caltech-101 object classification data set, and the ILSVCR 2012 data set which contains 1000 different classes. We considered a variety of scenarios including *transductive learning* (i.e. unlabeled test data available during training), *inductive learning* (i.e. test data not available during training), and *self-taught learning* (i.e. test and training data coming from two different data sets). In order to verify the efficacy of the CURL classification strategy, we also tested two embodiments: one that uses Ensemble Projection unsupervised representation coupled with Logistic Regression classification, and one based on LapSVM semi-supervised classification. Moreover different variants of the embodiments are evaluated as well. The results show that CURL clearly outperforms other semi-supervised learning methods in the state of the art.

Summarizing, the contributions of this work are: the proposal of a new classification strategy based on unsupervised representation learning; the use of a single set of visual features to create two new image representations; the use of Early and Late feature fusion schemes in a co-training workflow to allow effective classification in presence of few labeled data; extensive experiments to demonstrate the effectiveness of the proposed classification strategy under inductive, transductive and self-taught scenarios on different image datasets of heterogeneous contents and cardinalities.

## 2. Related Work

There is a large literature on semi-supervised learning. For the sake of brevity, we discuss only the paradigms involved in the proposed strategy. More information about these and other approaches to semi-supervised learning can be found in the book by Chapelle *et al.* [1].

### 2.1. Co-training

Blum and Mitchell proposed co-training in 1998 [9] and verified its effectiveness for the classification of web pages. The basic idea is that two classifiers are trained on separate views (features) and then used to train each other. More precisely, when one of the classifiers is very confident in making a prediction for unlabeled data, the predicted labels are used to augment the training set of the other classifier. The concept has been generalized to three [14] or more views [15, 16]. Co-training has been used in several computer vision applications including video annotation [17], action recognition [18], traffic analysis [19], speech and gesture recognition [20], image annotation [21], biometric recognition [22], image retrieval [23], image classification [24, 25], object detection [19, 26], and object tracking [27].

According to Blum and Mitchell, a sufficient condition for the effectiveness of co-training is that, beside being individually accurate, the two classifiers are conditionally independent given the class label. However, conditional independence is not a necessary condition. In fact, Whang and Zhou [28] showed that co-training can be effective when the diversity between the two classifiers is larger than their errors; their results provided a theoretical support to the success of single-view co-training variants [29, 30, 31] (the reader may refer to an updated study from the same authors [32] for more details about necessary and sufficient conditions for co-training).

### 2.2. Unsupervised representation learning

New image representations can be efficiently learned if a large amount of labeled data is available [33, 34, 35]. For example, the features extracted from the ConvNets [34] are widely used in different application scenarios. In the last years, as a consequence of the success of deep learning frameworks we observed an increased interest in methods that make use of unlabeled data to automatically learn new representations. In fact, these has been demonstrated to be very effective for the pre-training of large neural networks [36, 37]. Restricted Boltzmann Machines [38] and auto-encoder networks [39] are notable examples of this kind of methods. The tutorial by Bengio covers in detail this family of approaches [40].

A conceptually simpler approach consists in using clustering algorithms to identify frequently occurring patterns in unlabeled data that can be used to define effective representations. The K-means algorithm has been widely used for this purpose [41]. In computer vision this approach is very popular and lead to the many variants of bag-of-visual-words representations [42, 43, 44, 45]. Briefly, clustering on unlabeled data is used to build a vocabulary of visual words. Given an image, multiple local features are extracted and for each of them the most similar visual word is searched. The final representation is a histogram counting the occurrences of the visual words. Sparse coding can be seen as an extension of this approach, where each local feature is described as a sparse combination of multiple words of the vocabulary [46, 47, 48]. Spectral clustering can build a more discriminating dissimilarity measure between data points by using kernel functions e.g [49, 50, 51].

Another strategy for unsupervised feature learning is represented by Ensemble Projection (EP) [52]. From all the available data (labeled and unlabeled) Ensemble Projection samples a set of prototypes. Discriminative learning is then used to learn projection functions tuned to the prototypes. Since a single set of projections could be too noisy, multiple sets of prototypes are sampled to build an ensemble of projection functions. The values computed according to these functions represent the components of the learned representations.

LapSVM [8] can be seen as an unsupervised representation learning method as well. In this case the learned representation is not explicit but it is implicitly embedded in a kernel learned from unlabeled data.

### 2.3. Fusion schemes

Combining multimodal information is an important issue in pattern recognition. The fusion of multimodal inputs can bring complementary information from various sources, useful for improving the quality of the image retrieval and classification performance [53]. The problem arises in defining how these modalities are to be combined or fused. In general, the existing fusion approaches can be categorized as early and late fusion approaches, which refers to their relative position from the feature comparison or learning step in the whole processing chain. Early fusion usually refers to the combination of the features into a single representation before comparison/learning. Late fusion refers to the combination, at the last stage, of the responses obtained after individual features comparison or learning [54, 55]. There is no universal conclusion as to which strategy is the preferred method for a given task. For example, Snoek et al. [54] found that late fusion is better than early fusion for the TRECVID 2004 semantic indexing task, while Ayache et al. [56] stated that early fusion gets better results than late fusion on the TRECVID 2006 semantic indexing task. A combination of these approaches can also be exploited as hybrid fusion approach [57].

Another form of data fusion is Multiple Kernel Learning (MKL). MKL has been introduced by Lanckriet et al. [58] as extension of the support vector machines (SVMs). Instead of using a single kernel computed on the image representation as in standard SVMs, MKL learns distinct kernels. The kernels are combined with a linear or non linear function and the function’s parameters can be determined during the learning process. MKL can be used to learn different kernels on the same image representation or by learning different kernels each one on a different image representation [59]. The former corresponds to have different notion of similarity, and to choose the most suitable one for the problem and representation at hand. The latter corresponds to have multiple representations each with a, possibly, different definition of similarity that must be combined together. This kind of data fusion, in [55], is termed *intermediate fusion*.

### 3. The Proposed strategy: CURL

In the semi-supervised image classification setup the training data consists of both labeled examples  $\{\mathcal{X}_l, \mathcal{Y}\} = \{(\mathbf{x}_i, y_i)\}_{i=1}^L$  and unlabeled ones  $\mathcal{X}_u = \{\mathbf{x}_i\}_{i=L+1}^{L+U}$ , where  $\mathbf{x}_i$  denotes the feature vector of image  $i$ ,  $y_i \in \{1, \dots, K\}$  is its label, and  $K$  is the number of classes.

In this work, for each image  $i$  a set of  $S$  different image features  $\mathbf{x}_i^{(s)}$ ,  $s = 1, \dots, S$  is considered. Two views are then generated by using two different fusion strategies: early and late fusion. In case of Early Fusion (EF), the image features are concatenated and then used to learn a new representation  $\mathbf{x}_i^{EF} = \varphi([\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(S)}])$  in an unsupervised way, where  $\varphi(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^m$  is a projection function. In case of Late Fusion (LF), an unsupervised representation  $\varphi_s(\mathbf{x}_i^{(s)})$  is independently learned for each image feature and then the representations are concatenated to obtain  $\mathbf{x}_i^{LF} = [\varphi_1(\mathbf{x}_i^{(1)}), \dots, \varphi_S(\mathbf{x}_i^{(S)})]$ .

Using the learned EF and LF unsupervised representations, the two views are built:  $\mathcal{X}_l^{EF} = \{\mathbf{x}_i^{EF}\}_{i=1}^L$ ,  $\mathcal{X}_u^{EF} = \{\mathbf{x}_i^{EF}\}_{i=L+1}^{L+U}$  and  $\mathcal{X}_l^{LF} = \{\mathbf{x}_i^{LF}\}_{i=1}^L$ ,  $\mathcal{X}_u^{LF} = \{\mathbf{x}_i^{LF}\}_{i=L+1}^{L+U}$ . Furthermore, two label sets  $\mathcal{Y}^{EF}$  and  $\mathcal{Y}^{LF}$  are initialized equal to  $\mathcal{Y}$ .

Once the two views are generated, our method iteratively co-trains two classifiers  $\phi_{EF}$  and  $\phi_{LF}$  on them [9]. SVMs, logistic regressions, or any other similar technique can be used to obtain them. The idea of iterative co-training is that one can use a small labeled sample to train the initial classifiers over the respective views (i.e.  $\phi_{EF} : \mathcal{X}_l^{EF} \mapsto \mathcal{Y}^{EF}$  and  $\phi_{LF} : \mathcal{X}_l^{LF} \mapsto \mathcal{Y}^{LF}$ ), and then iteratively bootstrap by taking unlabeled examples for which at least one of the classifiers is very confident. The confident classifier determines pseudo-labels [60] that are then used as if they were true labels to improve the other classifier [61]. The confidence scores depend on the classification framework; for instance posterior probabilities can be used for this purpose if the classifiers provide them.

Given the classifier confidence scores  $\mathbf{w}_i^{EF} = \phi_{EF}(\mathbf{x}_i^{EF})$  and  $\mathbf{w}_i^{LF} = \phi_{LF}(\mathbf{x}_i^{LF})$ , the pseudo-labels  $\hat{y}_i^{EF}$  and  $\hat{y}_i^{LF}$  are respectively obtained as:

$$\hat{y}_i^{EF} = \arg \max_{j=1, \dots, K} \mathbf{w}_i^{EF}[j] \quad (1)$$

$$\hat{y}_i^{LF} = \arg \max_{j=1, \dots, K} \mathbf{w}_i^{LF}[j] \quad (2)$$

In each round of co-training, the classifier  $\phi_{LF}$  chooses some examples in  $\mathcal{X}_u^{EF}$  to pseudo-label for  $\phi_{EF}$ , and vice versa. For each class  $k$ , let us call  $\mathcal{X}_*$  the set of candidate unlabeled examples to be pseudo-labeled for  $\phi_{EF}$ . Each  $\mathbf{x}_* \in \mathcal{X}_*$  must belong to the unlabeled set, i.e.  $\mathbf{x}_* \in \mathcal{X}_u^{EF}$ , has not to be already used for training, i.e.  $\mathbf{x}_* \notin \mathcal{X}_l^{EF}$ , and its pseudo-label has to be  $\hat{y}_*^{LF} = k$ . Furthermore,  $\phi_{LF}$  should be more confident on the classification of  $\mathbf{x}_*$  than  $\phi_{EF}$ , and its confidence should be higher than a fixed threshold  $t_1$ :

$$\forall \mathbf{x}_* \in \mathcal{X}_* : \mathbf{w}_*^{EF}[k] < \mathbf{w}_*^{LF}[k], \mathbf{w}_*^{LF}[k] > t_1 \quad (3)$$

If no  $\mathbf{x}_*$  satisfying Eq. 3 are found, then the constraints are relaxed:

$$\forall \mathbf{x}_* \in \mathcal{X}_* : \mathbf{w}_*^{LF}[k] > t_2, \text{ with } t_2 < t_1 \quad (4)$$

Non-maximum suppression is applied to add one single pseudo-labeled example for each class by extracting the most confident  $\mathbf{x}_* \in \mathcal{X}_*$ :

$$\text{find } \mathbf{x}_* \in \mathcal{X}_* : \mathbf{w}_*^{LF}[k] = \operatorname{argmax}_j \mathbf{w}_j^{LF}[k] \quad (5)$$

The selected  $\mathbf{x}_*$  and its corresponding pseudo-label  $\hat{y}_*$  are added to  $\mathcal{X}_l^{LF}$  and  $\mathcal{Y}^{LF}$  respectively. If no  $\mathbf{x}_*$  satisfying Eq. 4 are found, then nothing is added to  $\mathcal{X}_l^{LF}$  and  $\mathcal{Y}^{LF}$ .

Similarly, the classifier  $\phi_{EF}$  chooses some examples in  $\mathcal{X}_u^{LF}$  to pseudo-label for  $\phi_{LF}$ . At the next co-training round, two new classifiers  $\phi_{EF}$  and  $\phi_{LF}$  are trained on the respective views, that now contain both labeled and pseudo-labeled examples. The complete procedure of the CURL method is outlined in Algorithms 1-3.

The Unsupervised Representation Learning is done once, and its computational cost is dominated by that of co-training which, in turn, mostly depends on the cost of the supervised training steps. The number of these steps is linear in the number of co-training rounds, and their individual cost increases with the number of pseudo-labels assigned.

## 4. Experiments

CURL is parametric with respect to the projection function  $\varphi$  used in the unsupervised representation learning URL, and the supervised classification technique C used during to co-train  $\phi_{EF}$  and  $\phi_{LF}$ . As first embodiment of CURL, we used Ensemble Projection [52] for the former and logistic regression for the latter. Another embodiment, based on LapSVM [8] is presented in Section 5.4.

### 4.1. Data sets

We evaluated our method on two data sets: Scene-15 (S-15) [43], and Caltech-101 (C-101) [62]. Scene-15 data set contains 4485 images divided into 15 scene categories with both indoor and outdoor environments. Each category has 200 to 400 images. Caltech-101 contains 8677 images divided into 101 object categories, each having 31 to 800 images. Furthermore, we collected a set of random images by sampling 20,000 images from the ImageNet data set [63] to evaluate our method on the task of self-taught image classification. Since the current version of ImageNet has 21841 synsets (i.e. categories) and a total of more than 14 millions images, there is a small probability that the random images and images in the two considered data sets come from the same distribution.

### 4.2. Image features

In our experiments we used the following three features: GIST [64], Pyramid of Histogram of Oriented Gradients (PHOG) [65], and Local Binary Patterns (LBP) [66]. GIST was computed on the rescaled images of  $256 \times 256$  pixels, at 3 scales with 4, 8 and 8 orientations respectively. PHOG was computed with a 2-layer pyramid and in 8 directions. Uniform LBP with radius equal to 1, and 8 neighbors was used.

In Section 5.3 we also investigate the use of features extracted from a CNN [67] in combination with the previous ones.

---

**Algorithm 1: CURL**

---

**Data:** Labeled data  $\{\mathcal{X}_l, \mathcal{Y}\}$ , unlabeled data  $\mathcal{X}_u$

**Result:** Classifiers  $\phi_{EF}(\cdot)$  and  $\phi_{LF}(\cdot)$

**begin**

$[\mathcal{X}_l^{EF}, \mathcal{X}_u^{EF}, \mathcal{X}_l^{LF}, \mathcal{X}_u^{LF}] = \text{computeURL}(\mathcal{X}_l, \mathcal{X}_u)$   
 $\mathcal{Y}^{EF} = \mathcal{Y}^{LF} = \mathcal{Y}$

train classifier  $\phi_{EF} : \mathcal{X}_l^{EF} \mapsto \mathcal{Y}^{EF}$

train classifier  $\phi_{LF} : \mathcal{X}_l^{LF} \mapsto \mathcal{Y}^{LF}$

**for** *co-training round*  $c = 1 : C$  **do**

initialize  $\mathcal{W}^{EF} = \mathcal{W}^{LF} = \hat{\mathcal{Y}}^{EF} = \hat{\mathcal{Y}}^{LF} = \emptyset$

**foreach**  $\mathbf{x}_i^{EF} \in \mathcal{X}_u^{EF}$  **do**

add  $\mathbf{w}_i^{EF} = \phi_{EF}(\mathbf{x}_i^{EF})$  to  $\mathcal{W}^{EF}$

add  $\hat{y}_i^{EF} = \arg \max_{j=1, \dots, K} \mathbf{w}_i^{EF}[j]$  to  $\hat{\mathcal{Y}}^{EF}$

**foreach**  $\mathbf{x}_i^{LF} \in \mathcal{X}_u^{LF}$  **do**

add  $\mathbf{w}_i^{LF} = \phi_{LF}(\mathbf{x}_i^{LF})$  to  $\mathcal{W}^{LF}$

add  $\hat{y}_i^{LF} = \arg \max_{j=1, \dots, K} \mathbf{w}_i^{LF}[j]$  to  $\hat{\mathcal{Y}}^{LF}$

**for** *class number*  $k = 1 : K$  **do**

**for**  $(v_1, v_2) \in \{(EF, LF), (LF, EF)\}$  **do**

find  $\{\mathcal{X}_*, \hat{\mathcal{Y}}_*^{v_2}\}$  with

$\mathcal{X}_* \subset \mathcal{X}_u^{v_1}$ ,  $\mathcal{X}_* \cap \mathcal{X}_l^{v_1} = \emptyset$ ,  $\hat{\mathcal{Y}}_*^{v_2} \subset \hat{\mathcal{Y}}^{v_2}$  s.t.:

$\forall \mathbf{x}_* \in \mathcal{X}_*$  and  $\forall \hat{y}_*^{v_2} \in \hat{\mathcal{Y}}_*^{v_2}$ , hold:

$\hat{y}_*^{v_2} = k$ ,  $\mathbf{w}_*^{v_1}[k] < \mathbf{w}_*^{v_2}[k]$ ,  $\mathbf{w}_*^{v_2}[k] > t_1$

**if**  $\mathcal{X}_* = \emptyset$  **then**

find  $\{\mathcal{X}_*, \hat{\mathcal{Y}}_*^{v_2}\}$  with

$\mathcal{X}_* \subset \mathcal{X}_u^{v_1}$ ,  $\mathcal{X}_* \cap \mathcal{X}_l^{v_1} = \emptyset$ ,  $\hat{\mathcal{Y}}_*^{v_2} \subset \hat{\mathcal{Y}}^{v_2}$  s.t.:

$\forall \mathbf{x}_* \in \mathcal{X}_*$  and  $\forall \hat{y}_*^{v_2} \in \hat{\mathcal{Y}}_*^{v_2}$ , hold:

$\hat{y}_*^{v_2} = k$ ,  $\mathbf{w}_*^{v_2}[k] > t_2$

$[\mathcal{X}_*, \hat{\mathcal{Y}}_*^{v_2}] = \text{nonMaxSuppr}(\mathcal{X}_*, \hat{\mathcal{Y}}_*^{v_2}, \mathcal{W}^{v_2})$

$\mathcal{X}_l^{v_1} = \mathcal{X}_l^{v_1} \cup \mathcal{X}_*$

$\mathcal{Y}^{v_1} = \mathcal{Y}^{v_1} \cup \hat{\mathcal{Y}}_*^{v_2}$

train classifier  $\phi_{EF} : \mathcal{X}_l^{EF} \mapsto \mathcal{Y}^{EF}$

train classifier  $\phi_{LF} : \mathcal{X}_l^{LF} \mapsto \mathcal{Y}^{LF}$

---

---

**Algorithm 2: compute URL**

---

**Data:** Labeled data  $\mathcal{X}_l$  and unlabeled data  $\mathcal{X}_u$

**Result:** Unsup. representations  $\mathcal{X}_l^{EF}, \mathcal{X}_u^{EF}, \mathcal{X}_l^{LF}, \mathcal{X}_u^{LF}$

**begin**

Learn EF representation  $\varphi$  on  $\{[\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(S)}]\}_{i=1}^{L+U}$

$\mathcal{X}_l^{EF} = \varphi(\{[\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(S)}]\}_{i=1}^L)$

$\mathcal{X}_u^{EF} = \varphi(\{[\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(S)}]\}_{i=L+1}^{L+U})$

Learn LF representations  $\varphi_s$  on  $\{\mathbf{x}_i^{(s)}\}_{i=1}^{L+U}$   $\mathcal{X}_l^{LF} = \{[\varphi_1(\mathbf{x}_i^{(1)}), \dots, \varphi_S(\mathbf{x}_i^{(S)})]\}_{i=1}^L$

$\mathcal{X}_u^{LF} = \{[\varphi_1(\mathbf{x}_i^{(1)}), \dots, \varphi_S(\mathbf{x}_i^{(S)})]\}_{i=L+1}^{L+U}$

---

---

**Algorithm 3:** non-maximum suppression

---

**Data:**  $\mathcal{X}, \hat{\mathcal{Y}}, \mathcal{W}, k$ **Result:**  $\mathcal{X}_*, \hat{\mathcal{Y}}_*$ **begin**
$$\left[ \begin{array}{l} \text{find } \{\mathcal{X}_*, \hat{\mathcal{Y}}_*\} \text{ with } \mathcal{X}_* \in \mathcal{X}, \hat{\mathcal{Y}}_* \in \hat{\mathcal{Y}} \text{ s.t.:} \\ \mathbf{w}_*[k] = \operatorname{argmax}_j \mathbf{w}_j[k], \text{ with } \mathbf{w}_j \in \mathcal{W} \end{array} \right.$$

---

#### 4.3. Ensemble projection

Differently from others semi-supervised methods that train a classifier from labeled data with a regularization term learned from unlabeled data, Ensemble Projection [52] learns a new image representation from all known data (i.e. labeled and unlabeled data), and then trains a plain classifier on it.

Ensemble Projection learns knowledge from  $T$  different prototype sets  $\mathcal{P}^t = \{(s_i^t, c_i^t)\}_{i=1}^{rn}$ , with  $t \in \{1, \dots, T\}$  where  $s_i^t \in \{1, \dots, L + U\}$  is the index of the  $i$ -th chosen image,  $c_i^t \in \{1, \dots, r\}$  is the pseudo-label indicating to which prototype  $s_i^t$  belong to.  $r$  is the number of prototypes in  $\mathcal{P}^t$ , and  $n$  is the number of images sampled for each prototype. For each prototype set,  $m$  hypotheses are randomly sampled, and the one containing images having the largest mutual distance is kept.

A set of discriminative classifiers  $\phi^t(\cdot)$  is learned on  $\mathcal{P}^t$ , one for each prototype set, and the projected vectors  $\phi^t(\mathbf{x}_i)$  are obtained. The final feature vector is obtained by concatenating these projected vectors.

Following [52] we set  $T = 300$ ,  $r = 30$ ,  $n = 6$ ,  $m = 50$ , using Logistic Regression (LR) as discriminative classifier  $\phi^t(\cdot)$  with  $C = 15$ .

Within CURL, Ensemble Projection is used to learn both Early Fusion and Late Fusion unsupervised representations. In the case of Early Fusion (EF), the feature vector  $\mathbf{x}_i$  is obtained concatenating the  $S$  different features available  $\mathbf{x}_i = [\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(S)}]$ ,  $s = 1, \dots, S$ . In the case of Late Fusion (LF), the feature vector  $\mathbf{x}_i$  is made by considering just one single feature at time  $\mathbf{x}_i = \mathbf{x}_i^{(s)}$ . For both EF and LF, the same number  $T$  of prototypes is used in order to assure that the unsupervised representations have the same size. The posterior probabilities computed by the LR classifier on the feature vector  $\mathbf{x}_i$  are used as the confidence scores  $\mathbf{w}_i$ .

#### 4.4. Experimental settings

We conducted two kinds of experiments: (1) comparison of our strategy with competing methods for semi-supervised image classification; (2) evaluation of our method at different number of co-training rounds. We considered three scenarios corresponding to three different ways of using unlabeled data. In the *inductive learning* scenario 25% of the unlabeled data is used together with the labeled data for the semi-supervised training of the classifier; the remaining 75% is used as an independent test set. In the *transductive learning* scenario all the unlabeled data is used during both training and test. In the *self-taught learning* scenario the set of unlabeled data is taken from an additional data set featuring a different distribution of image content (i.e. the 20,000 images from ImageNet); all the unlabeled data from the original data set is used as an independent test set.

As evaluation measure we followed [52] and used the multi-class average precision (MAP), computed as the average precision over all recall values and over all classes. Different numbers of training images per class were tested for both Scene-15 and Caltech-101 (i.e. 1, 2, 3, 5, 10, and 20). All the reported results represent the average performance over ten runs with random labeled-unlabeled splits.

The performance of the proposed strategy are compared with those of other supervised and semi-supervised baseline methods. As supervised classifiers we considered Support Vector Machines (SVM). As semi-supervised classifiers, we used LapSVM [68, 8]. LapSVM extend the SVM framework including a smoothness penalty term defined on the Laplacian adjacency graph built from both labeled and unlabeled data. For both SVM and LapSVM we experimented with the linear, RBF and  $\chi^2$  kernels computed on the concatenation of the three available image features as in [52]. The parameters of SVM and LapSVM have been determined by a greedy search with

a three-fold cross validation on the training set. We also compared the present embodiment of CURL against Ensemble Projection coupled with a logistic regression classifier (EP+LR) as in [52].

## 5. Experimental results

As a first experiment we compared CURL against EP+LR, and against SVMs and LapSVMs with different kernels. Specifically, we tested the two co-trained classifiers operating on early-fused and late-fused representations, both employing EP for URL and LR as classifier C, that we call CURL-EF(EP+LR) and CURL-LF(EP+LR) respectively. We also included a variant of the proposed method. It differs in the number of pseudo-labeled examples that are added at each co-training round. The variant skips the non-maximum suppression step, and at each round, adds all the examples satisfying Eq. 3. We denote the two co-trained classifiers of the variant as CURL-EF<sub>n</sub>(EP+LR) and CURL-LF<sub>n</sub>(EP+LR).

Fig. 2 shows the classification performance with different numbers of labeled training images per class, in the three learning scenarios for both the Scene-15 and Caltech-101 data sets. For the CURL-based methods we considered five co-training rounds, and the reported performance correspond to the last round. For SVM and LapSVM only the results using  $\chi^2$  kernel are reported, since they consistently showed the best performance across all the experiments. Detailed results

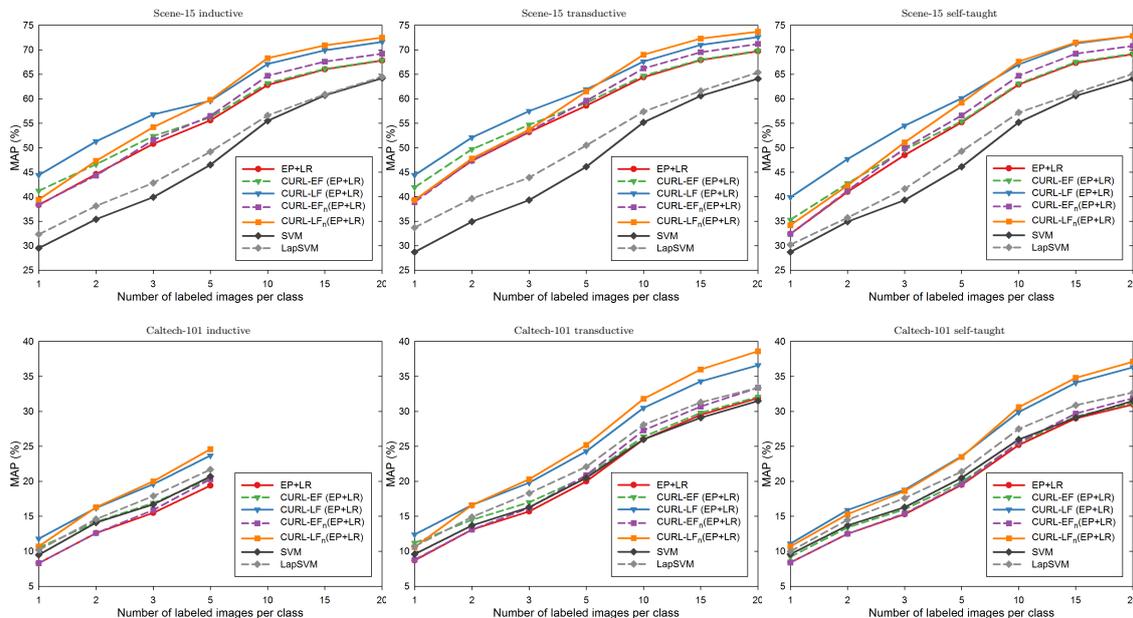


Figure 2: Mean Average Precision (MAP) varying the number of labeled images per class, obtained on the Scene-15 data set (first row), and on the Caltech-101 data set (second row). Three scenarios are considered: inductive learning (left column), transductive learning (middle column) and self-taught learning (third column). Note that inductive learning on the Caltech-101 data set is limited to 5 labeled images per class because otherwise for some classes there wouldn't be enough unlabeled data left for both training and evaluation.

for all the tested baseline methods, and for the CURL variants across the co-training rounds are available in Tables 1, 2 and 3.

The behavior of the methods is quite stable with respect to the three learning scenarios, with slightly lower MAP obtained in the case of self-taught learning. It is evident that our strategy outperformed the other methods in the state of the art included in the comparison across all the data sets and all the scenarios considered. Among the variants considered, CURL-LF(EP+LR) demonstrated to be the best in the case of a small number of labeled images, while CURL-LF<sub>n</sub>(EP+LR) obtained the best results when more labeled data is available. Classifiers obtained on early-fused representations performed generally worse than the corresponding ones obtained on late-fused representations, but they are still uniformly better than the original EP+LR Ensemble Projection

Table 1: Mean Average Precision (MAP) of the baseline algorithms, varying the number of labeled images per class in the three learning scenarios considered: inductive (IND), transductive (TRD), and self-taught (ST).

# img	method	Scene-15			# img	method	Caltech-101		
		IND	TRD	ST			IND	TRD	ST
1	SVM <sub>lin</sub>	22.7	22.3	22.3	1	SVM <sub>lin</sub>	6.0	6.3	6.3
	SVM <sub>rbf</sub>	26.5	25.8	25.8		SVM <sub>rbf</sub>	7.1	7.3	7.3
	SVM <sub><math>\chi^2</math></sub>	29.5	28.7	28.7		SVM <sub><math>\chi^2</math></sub>	9.5	9.6	9.6
	LapSVM <sub>lin</sub>	28.1	29.1	26.4		LapSVM <sub>lin</sub>	9.2	9.6	9.1
	LapSVM <sub>rbf</sub>	28.8	29.8	26.7		LapSVM <sub>rbf</sub>	9.8	10.2	9.6
	LapSVM <sub><math>\chi^2</math></sub>	32.3	33.7	30.2		LapSVM <sub><math>\chi^2</math></sub>	10.2	10.7	10.0
	EP+LR	38.3	39.3	32.4		EP+LR	8.3	8.7	8.4
2	SVM <sub>lin</sub>	27.4	26.9	26.9	2	SVM <sub>lin</sub>	9.1	9.2	9.2
	SVM <sub>rbf</sub>	32.9	31.3	31.3		SVM <sub>rbf</sub>	10.1	9.8	9.8
	SVM <sub><math>\chi^2</math></sub>	35.4	34.9	34.9		SVM <sub><math>\chi^2</math></sub>	14.1	13.7	13.7
	LapSVM <sub>lin</sub>	33.7	34.9	31.2		LapSVM <sub>lin</sub>	12.3	12.8	12.1
	LapSVM <sub>rbf</sub>	34.6	35.7	32.5		LapSVM <sub>rbf</sub>	12.7	13.3	12.4
	LapSVM <sub><math>\chi^2</math></sub>	38.1	39.6	35.7		LapSVM <sub><math>\chi^2</math></sub>	14.6	14.9	14.5
	EP+LR	44.6	47.3	41.0		EP+LR	12.6	13.1	12.5
3	SVM <sub>lin</sub>	30.0	30.2	30.2	3	SVM <sub>lin</sub>	10.7	10.8	10.8
	SVM <sub>rbf</sub>	36.5	36.7	36.7		SVM <sub>rbf</sub>	11.7	11.6	11.6
	SVM <sub><math>\chi^2</math></sub>	39.9	39.3	39.3		SVM <sub><math>\chi^2</math></sub>	16.7	16.3	16.3
	LapSVM <sub>lin</sub>	37.6	38.6	36.4		LapSVM <sub>lin</sub>	13.8	14.3	13.5
	LapSVM <sub>rbf</sub>	37.7	38.9	37.2		LapSVM <sub>rbf</sub>	14.0	14.6	13.9
	LapSVM <sub><math>\chi^2</math></sub>	42.8	43.9	41.6		LapSVM <sub><math>\chi^2</math></sub>	17.9	18.3	17.6
	EP+LR	50.8	53.2	48.5		EP+LR	15.5	15.7	15.3
5	SVM <sub>lin</sub>	35.5	35.2	35.2	5	SVM <sub>lin</sub>	13.4	13.3	13.3
	SVM <sub>rbf</sub>	43.5	42.8	42.8		SVM <sub>rbf</sub>	14.9	14.6	14.6
	SVM <sub><math>\chi^2</math></sub>	46.5	46.1	46.1		SVM <sub><math>\chi^2</math></sub>	20.7	20.5	20.5
	LapSVM <sub>lin</sub>	43.4	44.5	43.5		LapSVM <sub>lin</sub>	16.3	16.6	16.0
	LapSVM <sub>rbf</sub>	43.8	44.7	44.0		LapSVM <sub>rbf</sub>	16.8	17.1	16.6
	LapSVM <sub><math>\chi^2</math></sub>	49.2	50.5	49.3		LapSVM <sub><math>\chi^2</math></sub>	21.7	22.1	21.4
	EP+LR	55.6	58.6	55.2		EP+LR	19.4	20.0	19.5
10	SVM <sub>lin</sub>	43.8	43.7	43.7	10	SVM <sub>lin</sub>	-	17.3	17.3
	SVM <sub>rbf</sub>	51.9	51.3	51.3		SVM <sub>rbf</sub>	-	19.2	19.2
	SVM <sub><math>\chi^2</math></sub>	55.5	55.2	55.2		SVM <sub><math>\chi^2</math></sub>	-	26.0	26.0
	LapSVM <sub>lin</sub>	51.4	52.5	52.3		LapSVM <sub>lin</sub>	-	20.5	20.1
	LapSVM <sub>rbf</sub>	52.4	53.1	53.0		LapSVM <sub>rbf</sub>	-	21.6	21.2
	LapSVM <sub><math>\chi^2</math></sub>	56.6	57.4	57.2		LapSVM <sub><math>\chi^2</math></sub>	-	28.1	27.5
	EP+LR	62.8	64.4	62.9		EP+LR	-	26.0	25.2
15	SVM <sub>lin</sub>	50.1	50.3	50.3	15	SVM <sub>lin</sub>	-	19.7	19.7
	SVM <sub>rbf</sub>	57.4	57.1	57.1		SVM <sub>rbf</sub>	-	22.1	22.1
	SVM <sub><math>\chi^2</math></sub>	60.7	60.6	60.6		SVM <sub><math>\chi^2</math></sub>	-	29.1	29.1
	LapSVM <sub>lin</sub>	55.1	56.0	55.5		LapSVM <sub>lin</sub>	-	22.9	22.5
	LapSVM <sub>rbf</sub>	57.8	58.7	58.3		LapSVM <sub>rbf</sub>	-	24.0	23.5
	LapSVM <sub><math>\chi^2</math></sub>	60.9	61.6	61.2		LapSVM <sub><math>\chi^2</math></sub>	-	31.3	30.9
	EP+LR	66.0	67.9	67.3		EP+LR	-	29.5	29.0
20	SVM <sub>lin</sub>	54.0	53.5	53.5	20	SVM <sub>lin</sub>	-	21.5	21.5
	SVM <sub>rbf</sub>	60.4	60.3	60.3		SVM <sub>rbf</sub>	-	24.2	24.2
	SVM <sub><math>\chi^2</math></sub>	64.2	64.1	64.1		SVM <sub><math>\chi^2</math></sub>	-	31.5	31.5
	LapSVM <sub>lin</sub>	58.3	58.7	58.6		LapSVM <sub>lin</sub>	-	24.6	24.0
	LapSVM <sub>rbf</sub>	60.8	61.4	61.2		LapSVM <sub>rbf</sub>	-	26.8	26.1
	LapSVM <sub><math>\chi^2</math></sub>	64.5	65.4	65.1		LapSVM <sub><math>\chi^2</math></sub>	-	33.4	32.7
	EP+LR	67.8	69.7	69.1		EP+LR	-	31.9	31.0

Table 2: Mean Average Precision (MAP) of the CURL variants, in the (EP+LR) embodiment, varying the number of labeled images per class at the different co-training rounds obtained on the Scene-15 data set in the three learning scenarios considered: inductive (left), transductive (middle), and self-taught (right). For clarity, the (EP+LR) suffixes have been omitted.

# img	method	# co-train round					# img	method	# co-train round					# img	method	# co-train round				
		0	1	2	3	4			5	0	1	2	3			4	5	0	1	2
1	CURL-EF	38.3	41.0	41.1	41.1	41.1	41.9	41.9	41.9	41.9	41.9	41.9	32.4	34.8	35.2	35.3	35.3	35.3	35.3	35.3
	CURL-LF	40.0	42.4	43.8	44.0	44.3	44.4	44.4	44.4	44.4	44.4	44.4	36.7	38.6	39.6	39.7	39.8	39.9	39.9	
	CURL-EF&LF	-	43.3	43.7	43.8	44.0	44.0	-	-	-	-	-	36.7	34.9	34.8	32.8	32.8	32.3	34.6	34.2
	CURL-EF <sub>n</sub>	38.3	38.4	38.5	38.4	38.3	38.2	39.3	38.9	38.9	38.9	38.7	38.7	32.4	32.4	32.8	32.8	32.8	32.8	32.3
	CURL-LF <sub>n</sub>	40.0	39.9	39.9	39.7	39.6	39.4	39.8	39.9	39.9	39.7	39.6	39.3	36.7	34.9	34.8	34.8	34.8	34.6	34.2
2	CURL-EF	44.6	46.3	46.4	46.5	46.6	46.6	47.3	49.2	49.5	49.6	49.7	49.7	41.0	42.4	42.5	42.6	42.6	42.7	
	CURL-LF	48.5	50.2	50.7	50.9	51.1	51.3	48.7	50.6	51.5	51.9	52.0	52.1	45.4	47.1	47.6	47.5	47.7	47.7	
	CURL-EF&LF	-	49.7	50.0	50.1	50.2	50.3	-	51.2	51.7	51.9	52.0	52.1	46.4	46.5	46.6	46.7	46.7	46.7	
	CURL-EF <sub>n</sub>	44.6	44.3	44.7	44.7	44.5	44.2	47.3	47.4	47.8	47.5	47.1	46.9	41.0	41.2	42.7	42.7	43.1	41.4	
	CURL-LF <sub>n</sub>	48.5	48.8	48.8	48.4	48.0	47.3	48.7	49.3	49.0	48.8	48.5	47.8	45.4	44.3	44.2	43.7	43.1	42.3	
3	CURL-EF	50.8	52.1	52.2	52.3	52.4	52.4	53.2	54.4	54.5	54.6	54.7	54.7	48.5	49.4	49.5	49.5	49.5	49.6	
	CURL-LF	54.1	55.2	56.2	56.4	56.6	56.8	54.8	55.8	56.6	57.2	57.4	57.5	53.4	54.0	54.5	54.4	54.5	54.5	
	CURL-EF&LF	-	55.3	55.7	55.9	56.0	56.1	-	56.4	56.8	57.2	57.3	57.4	53.5	53.6	53.6	53.6	53.6	53.6	
	CURL-EF <sub>n</sub>	50.8	51.6	52.5	52.6	52.9	51.8	53.2	53.4	53.7	53.5	52.7	52.2	48.5	49.9	50.6	50.6	49.8	49.3	
	CURL-LF <sub>n</sub>	54.1	55.3	55.3	55.1	54.7	54.2	54.8	55.4	55.4	54.9	54.4	53.7	53.4	52.0	52.3	52.0	51.6	51.1	
5	CURL-EF	55.6	56.1	56.2	56.2	56.2	56.2	58.6	59.1	59.2	59.2	59.2	59.2	55.2	55.4	55.5	55.5	55.5	55.4	
	CURL-LF	58.3	59.0	59.3	59.5	59.5	59.6	60.3	61.3	61.6	61.7	61.9	61.9	59.4	59.8	60.0	60.1	60.1	60.1	
	CURL-EF&LF	-	59.0	59.1	59.2	59.3	59.4	-	61.5	61.6	61.7	61.8	61.8	59.4	59.3	59.4	59.4	59.4	59.4	
	CURL-EF <sub>n</sub>	55.6	56.5	57.7	57.7	57.7	57.6	58.6	59.6	60.4	60.3	60.0	59.7	55.2	56.6	57.5	57.5	57.2	57.1	
	CURL-LF <sub>n</sub>	58.3	59.7	59.8	60.0	59.8	59.8	60.3	61.8	62.0	61.9	61.7	61.5	59.4	59.5	59.6	59.6	59.4	59.2	
10	CURL-EF	62.8	63.1	63.2	63.2	63.2	63.2	64.4	64.6	64.6	64.6	64.7	64.7	62.9	63.1	63.1	63.1	63.1	63.1	
	CURL-LF	66.2	66.5	66.6	66.9	67.0	67.1	66.6	67.0	67.2	67.3	67.5	67.6	66.0	66.4	66.7	66.7	66.9	67.0	
	CURL-EF&LF	-	66.3	66.4	66.5	66.6	66.7	-	67.1	67.1	67.2	67.3	67.4	66.6	66.7	66.7	66.7	66.8		
	CURL-EF <sub>n</sub>	62.8	64.7	65.3	65.5	65.8	65.8	64.4	66.2	67.0	67.3	67.4	67.5	62.9	64.7	65.7	66.0	66.1	66.3	
	CURL-LF <sub>n</sub>	66.2	67.6	67.8	68.2	68.2	68.3	66.6	68.5	68.8	69.0	69.0	69.0	66.0	67.1	67.3	67.5	67.6	67.6	
15	CURL-EF	66.0	66.1	66.1	66.1	66.1	66.1	67.9	68.0	68.0	68.0	68.0	68.0	67.3	67.5	67.5	67.5	67.5	67.5	
	CURL-LF	69.2	69.4	69.6	69.8	69.8	69.9	70.1	70.5	70.6	70.7	70.8	71.0	70.5	70.8	70.9	71.1	71.2	71.3	
	CURL-EF&LF	-	69.3	69.3	69.4	69.4	69.5	-	70.5	70.6	70.6	70.7	70.8	71.0	71.0	71.1	71.2	71.2		
	CURL-EF <sub>n</sub>	66.0	67.6	68.3	68.5	68.7	68.8	67.9	69.5	70.4	70.7	70.8	70.9	67.3	69.2	69.8	70.1	70.2	70.2	
	CURL-LF <sub>n</sub>	69.2	70.6	70.9	71.0	71.0	70.9	70.1	71.7	71.9	72.2	72.3	72.3	70.5	70.8	71.1	71.3	71.5	71.5	
20	CURL-EF	67.8	67.9	67.9	67.9	67.9	67.9	69.7	69.8	69.8	69.8	69.8	69.8	69.1	69.2	69.2	69.2	69.2	69.2	
	CURL-LF	71.1	71.3	71.3	71.5	71.5	71.6	72.1	72.2	72.3	72.4	72.5	72.6	72.3	72.4	72.4	72.5	72.6	72.8	
	CURL-EF&LF	-	71.2	71.2	71.2	71.2	71.3	-	72.3	72.3	72.4	72.4	72.4	72.6	72.6	72.6	72.7	72.8		
	CURL-EF <sub>n</sub>	67.8	69.2	69.8	70.1	70.2	70.3	69.7	71.2	71.7	72.1	72.2	72.3	69.1	70.8	71.5	71.7	71.8	71.9	
	CURL-LF <sub>n</sub>	71.1	72.0	72.4	72.5	72.5	72.5	72.1	73.3	73.6	73.7	73.7	73.7	72.3	72.4	72.7	72.8	72.8		

Table 3: Mean Average Precision (MAP) of the CURL variants, in the (EP+LR) embodiment, varying the number of labeled images per class at the different co-training rounds obtained on the Caltech-101 data set in the three learning scenarios considered: inductive (left), transductive (middle), and self-taught (right). For clarity, the (EP+LR) suffixes have been omitted.

# img	method	# co-train round					# img	method	# co-train round					# img	method	# co-train round				
		0	1	2	3	4			5	0	1	2	3			4	5	0	1	2
1	CURL-EF	8.3	10.4	10.5	10.5	10.5	10.5	8.7	11.1	11.2	11.2	11.2	11.2	8.4	9.3	9.3	9.2	9.2	9.2	
	CURL-LF	10.1	11.6	11.8	11.8	11.8	11.8	10.5	12.2	12.4	12.4	12.4	12.4	10.7	11.4	11.3	11.2	11.1	11.1	
	CURL-EF&LF	-	11.5	11.6	11.6	11.6	11.6	-	12.1	12.2	12.2	12.2	12.2	10.7	10.9	10.8	10.7	10.7	10.7	
	CURL-EF <sub>n</sub>	8.3	8.3	8.5	8.5	8.8	8.8	8.7	8.8	8.8	8.8	8.9	8.9	8.4	8.4	8.4	8.4	8.4	8.4	
	CURL-LF <sub>n</sub>	10.1	10.3	10.3	10.6	10.7	10.7	10.5	10.5	10.5	10.6	10.6	10.6	10.7	10.7	10.7	10.7	10.7	10.7	
2	CURL-EF	12.6	14.2	14.3	14.3	14.3	14.2	13.1	14.4	14.5	14.6	14.5	14.5	12.5	13.3	13.3	13.4	13.4	13.4	
	CURL-LF	15.3	16.3	16.3	16.3	16.2	16.2	15.7	16.5	16.6	16.6	16.6	16.6	15.4	16.0	16.0	15.9	15.9	15.9	
	CURL-EF&LF	-	15.9	15.9	15.8	15.8	15.8	-	16.0	16.1	16.1	16.1	16.0	15.7	15.5	15.4	15.4	15.4		
	CURL-EF <sub>n</sub>	12.6	12.6	13.4	13.5	13.9	14.1	13.1	13.1	13.8	13.9	14.1	14.2	12.5	12.5	12.5	12.5	12.5		
	CURL-LF <sub>n</sub>	15.3	15.7	15.9	16.1	16.2	16.3	15.7	16.2	16.4	16.5	16.6	16.6	15.4	15.3	15.3	15.4	15.4		
3	CURL-EF	15.5	16.8	16.9	16.8	16.9	16.9	15.7	16.8	16.9	16.9	17.0	17.0	15.3	15.9	16.0	16.0	16.0	16.0	
	CURL-LF	18.8	19.4	19.6	19.6	19.7	19.6	19.0	19.5	19.8	19.8	19.8	19.8	18.5	18.8	18.9	18.8	18.8	18.8	
	CURL-EF&LF	-	18.9	19.0	18.9	18.9	18.9	-	18.9	19.0	19.0	19.0	19.0	18.4	18.4	18.4	18.4	18.4		
	CURL-EF <sub>n</sub>	15.5	15.9	16.8	16.9	17.1	17.1	15.7	16.2	17.2	17.3	17.5	17.5	15.3	15.4	15.4	15.4	15.5		
	CURL-LF <sub>n</sub>	18.8	19.5	19.8	20.0	20.0	20.0	19.0	20.0	20.2	20.3	20.3	20.3	18.5	18.5	18.5	18.6	18.6		
5	CURL-EF	19.4	20.2	20.4	20.4	20.4	20.4	20.0	20.6	20.7	20.8	20.7	20.7	19.5	19.9	19.9	19.9	19.9		
	CURL-LF	23.2	23.4	23.7	23.7	23.7	23.7	23.7	24.0	24.2	24.3	24.3	24.3	23.3	23.4	23.6	23.6	23.6		
	CURL-EF&LF	-	22.9	23.0	23.0	23.0	23.0	-	23.2	23.3	23.4	23.4	23.4	22.9	23.0	22.9	22.9	22.9		
	CURL-EF <sub>n</sub>	19.4	20.3	21.0	21.2	21.3	21.3	20.0	20.9	21.8	21.9	22.0	22.0	19.5	19.6	19.7	19.8	19.8		
	CURL-LF <sub>n</sub>	23.2	24.2	24.4	24.6	24.6	24.6	23.7	24.9	25.1	25.2	25.2	25.2	23.3	23.4	23.5	23.6	23.6		
10	CURL-EF	-	-	-	-	-	-	26.0	26.3	26.4	26.4	26.4	26.4	25.2	25.4	25.5	25.5	25.5		
	CURL-LF	-	-	-	-	-	-	30.2	30.2	30.4	30.4	30.5	30.5	29.9	29.9	29.9	29.9	29.9		
	CURL-EF&LF	-	-	-	-	-	-	29.6	29.6	29.6	29.7	29.7	29.7	29.3	29.3	29.3	29.3			
	CURL-EF <sub>n</sub>	-	-	-	-	-	-	26.0	27.3	27.8	27.9	27.9	27.9	25.2	25.6	25.9	26.0	26.1		
	CURL-LF <sub>n</sub>	-	-	-	-	-	-	30.2	31.4	31.8	31.8	31.8								

which can be considered as their non-co-trained version. SVMs and LapSVMs performed poorly on the Scene-15 data set, but they outperformed EP+LR and some of the CURL variants on the Caltech-101 data set.

Co-training allows to make good use of the early fusion representations that otherwise lead to worse results than late fusion representations. In our opinion this happens because the two views capture different relationships among data. This fact is visible in Fig. 3, which shows 2D projections obtained by applying the t-SNE [69] method to GIST, PHOG, LBP features, their concatenation, and their learnt early- and late-fused representations. Unsupervised representation

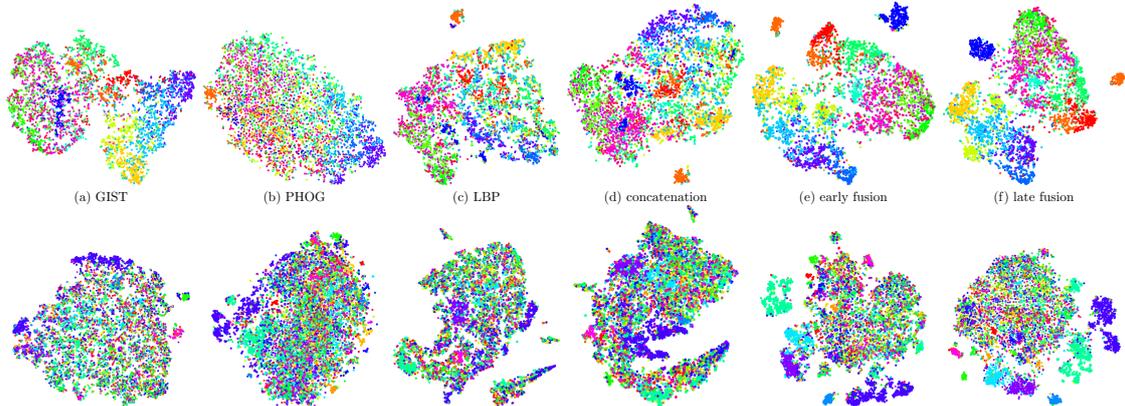


Figure 3: t-SNE 2D projections for the different features used. They are relative to the Scene-15 (top row) and Caltech-101 (bottom row) data sets. Different classes are represented in different colors, and the same class with the same color across the row.

learning allows t-SNE to identify groups of images of the same class. Moreover, representations based on early and late fusion induce different relationships among the classes. For instance, in the second row of Fig. 3f the blue and the light green classes have been placed close to each other on the bottom right; in Fig. 3e, instead, the two classes are well separated. The difference in the two representations explains the effectiveness of co-training and justifies the difference in performance between CURL-EF(EP+LR) and CURL-LF(EP+LR).

As further investigation, we also combined the two classifiers produced by the co-training procedure obtaining two other variants of CURL that we denoted as CURL-EF&LF(EP+LR) and CURL-EF&LF<sub>n</sub>(EP+LR). However, in our experiments, these variants did not caused any significant improvement when compared to CURL-LF(EP+LR).

### 5.1. Performance across co-training rounds

Here we analyze in more details the performance of our strategy across the five co-training rounds. Results are reported in Fig. 4 with lines of increasing color saturation corresponding to rounds one to five. CURL-LF(EP+LR) is reported in red lines, while CURL-LF<sub>n</sub>(EP+LR) in blue. Results are reported in terms of MAP improvements with respect to EP+LR, which, we recall, corresponds to CURL-EF(EP+LR) with zero co-training rounds. For CURL-LF(EP+LR), performances always increase with the number of rounds. For CURL-LF<sub>n</sub>(EP+LR), this is not true on the Scene-15 data set with a small number of labeled examples. In CURL-LF<sub>n</sub>(EP+LR) each round of co-training adds all the promising unlabeled samples, with a high chance of including some of them with the wrong pseudo-label. This may result in a ‘concept drift’, with the classifiers being pulled away from the concepts represented by the labeled examples. This risk is lower on the Caltech-101 (which tends to have more homogeneous classes than Scene-15) and when there are more labeled images. The original CURL-LF(EP+LR) is more conservative, since each of its co-training rounds adds a single image per class. As a result, increasing the rounds usually increases MAP and never decreases it by an appreciable amount.

We observed the same behavior for CURL-EF(EP+LR) and CURL-EF<sub>n</sub>(EP+LR). We omit the relative figures for sake of brevity.

The plots confirm that CURL-LF(EP+LR) is better suited for small sets of labeled images, while CURL-LF<sub>n</sub>(EP+LR) is to be preferred when more labeled examples are available. The

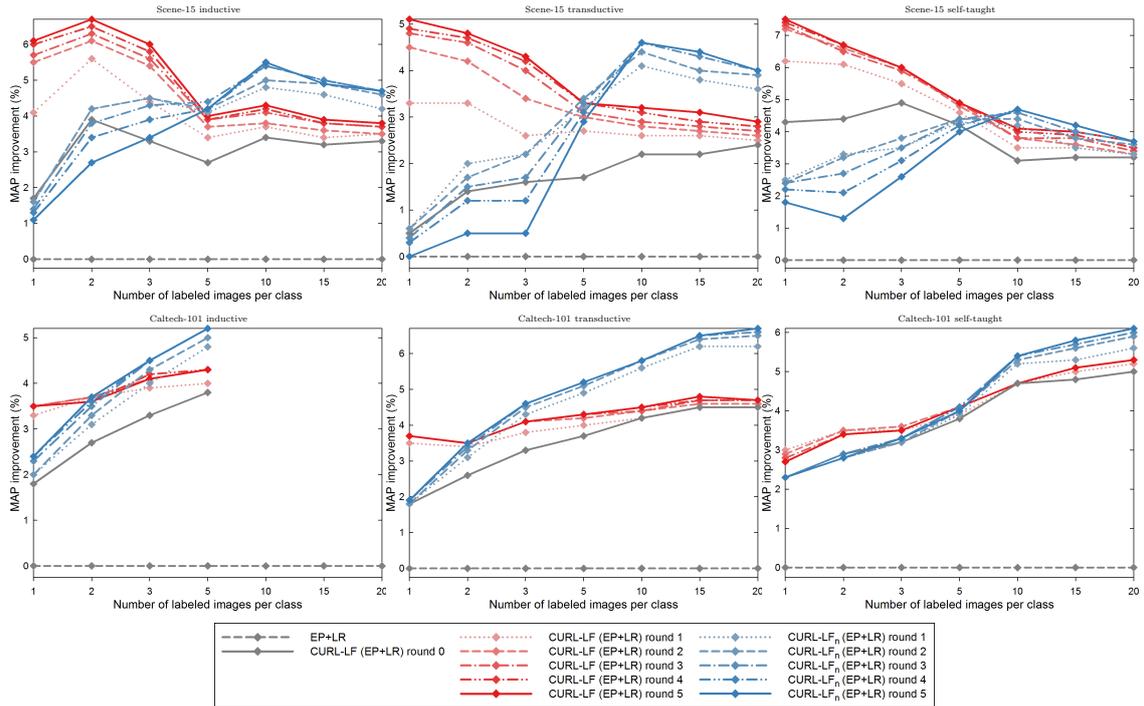


Figure 4: Performance obtained by CURL-LF(EP+LR) and CURL-LF<sub>n</sub>(EP+LR) varying the number of co-training rounds. Performance are reported in terms of MAP improvement with respect to Ensemble Projection. Due to the small cardinality of some classes, inductive learning on the Caltech-101 has been limited to five labeled images per class.

representation learned from late fused features explains part of the effectiveness of CURL. In fact, even CURL-LF(EP+LR) without co-training (zero rounds) outperforms the baseline represented by Ensemble Projection.

### 5.2. Comparison with standard co-training

In this section we compare the proposed classification strategy with standard co-training [9]. Since standard co-training employs two views, we compare with all the six possible combinations of features that we can generate by assigning one feature to each view or the concatenation of two features to one view and a single feature to the other view.

The experimental results are reported in Fig. 5, for both the Scene-15 and Caltech-101 data sets. Results are reported only for the classifier obtained by combining the scores on the two views, since its performance are always better than those of each single-view classifier. We report the results in the transductive scenario only.

It can be seen that the results obtained by CURL are better than those obtained by any combination of features with standard co-training. In particular we can observe that on the Scene-15 data set, the co-training variants using only PHOG features as a view obtain the worst performance. On the Caltech-101 data set instead, the co-training variants using only GIST features as a view obtain the worst performance. This is not surprising since GIST were specifically designed to provide a rough description (i.e. the gist) of the scene.

### 5.3. Leveraging CNN features in CURL

In this further experiment we want to test if the proposed classification strategy works when more powerful features are used. Recent results indicate that the generic descriptors extracted from pre-trained Convolutional Neural Networks (CNN) are able to obtain consistently superior results compared to the highly tuned state of the art systems in all the visual classification tasks on various datasets [67]. We extract a 4096-dimensional feature vector from each image using the Caffe [70] implementation of the deep CNN described by Krizhevsky et al. [34]. The CNN was discriminatively trained on a large dataset (ILSVRC 2012) with image-level annotations to

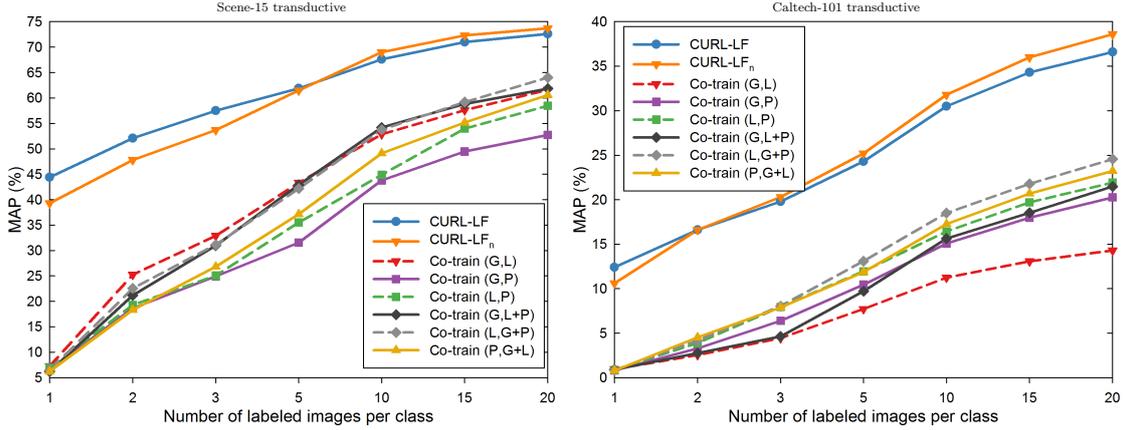


Figure 5: Mean Average Precision (MAP) varying the number of labeled images per class, obtained on the Scene-15 data set (left), and on the Caltech-101 data set (right). Results are compared with standard co-training on the possible six subdivision of the GIST (G), PHOG (P), and LBP (L) features in the two views.

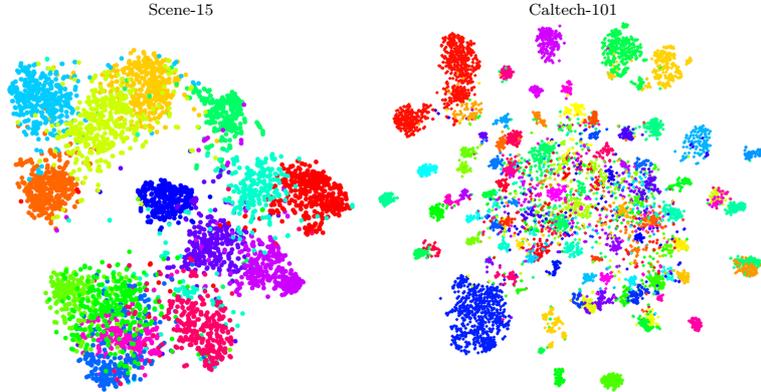


Figure 6: 2D projections for the CNN features on the two data sets used: Scene-15 (left) and Caltech-101 (right). Different classes are represented in different colors.

classify images into 1000 different classes. Briefly, a mean-subtracted  $227 \times 227$  RGB image is forward propagated through five convolutional layers and two fully connected layers. Features are obtained by extracting activation values of the last hidden layer. More details about the network architecture can be found in [34, 70].

We leverage the CNN features in CURL using them as a fourth feature in addition to the three used in Section 4. The discriminative power of these CNN features alone can be seen in Fig. 6, where their 2D projections obtained applying the t-SNE [69] method are reported.

The experimental results using the four features, are reported in Fig. 7, for both the Scene-15 and Caltech-101 data sets. We report the results in the transductive scenario only. It can be seen that the results using the four features are significantly better than those using only three features mainly due to the discriminative power of the CNN features. Furthermore, the CURL variants achieve better results than the baselines. This suggests that CURL is able to effectively leverage both low/mid level features as LBP, PHOG and GIST, and more powerful features as CNN.

#### 5.4. Second embodiment of CURL using LapSVM

In this Section we want to evaluate the CURL performance in a different embodiment. Specifically, we substitute the EP and LR components with LapSVM-based ones. In the LapSVM, first, an unsupervised geometrical deformation of the feature kernel is performed. This deformed kernel is then used for classification by a standard SVM thus by-passing an explicit definition of a new feature representation. In this CURL embodiment we exploit the unsupervised step as surrogate of the URL component, and SVM as C component. The confidence scores are computed by converting the signed distance from the separating hyperplanes in the kernel space into probabilities [71].

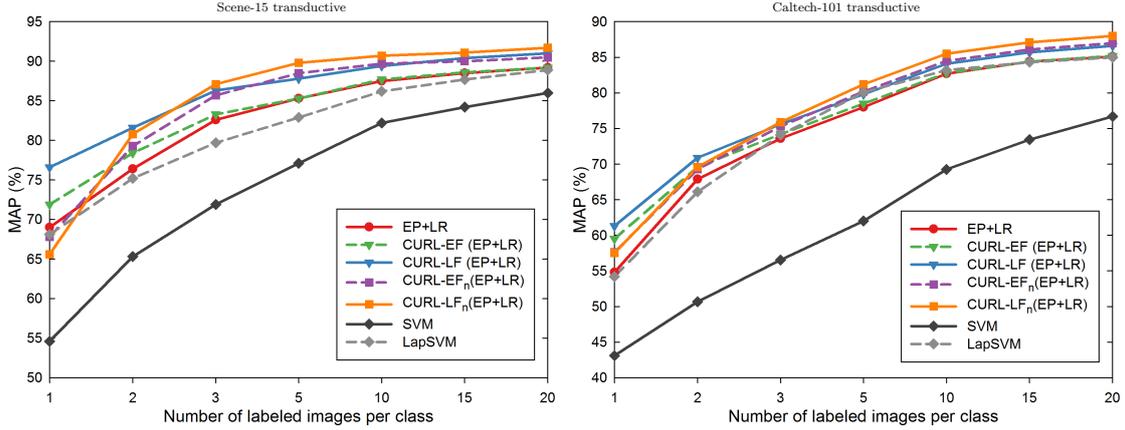


Figure 7: Mean Average Precision (MAP) varying the number of labeled images per class, obtained on the Scene-15 data set (left), and on the Caltech-101 data set (right). Results are obtained using GIST, PHOG, LBP and CNN features.

The EF view is obtained concatenating the GIST, PHOG, LBP and CNN features and generating the corresponding kernel, while the LF one is obtained by a linear combination of the four kernels computed on each feature. This is similar to what is done in multiple kernel learning [59]. Due to its performance in the previous experiments, the  $\chi^2$  kernel is used for both views. The experimental results on the Scene-15 and Caltech-101 data sets in the transductive scenario, are reported in Fig. 8. We named the variants of this CURL embodiment by adding the suffix (LapSVM). It can be seen that the behavior of the different methods is the same of the previous plots, with the LapSVM-based CURL outperforming the standard LapSVM. The plots confirm that CURL-LF(LapSVM) is better suited for small sets of labeled images, while CURL-LF<sub>n</sub>(LapSVM) is to be preferred when more labeled examples are available.

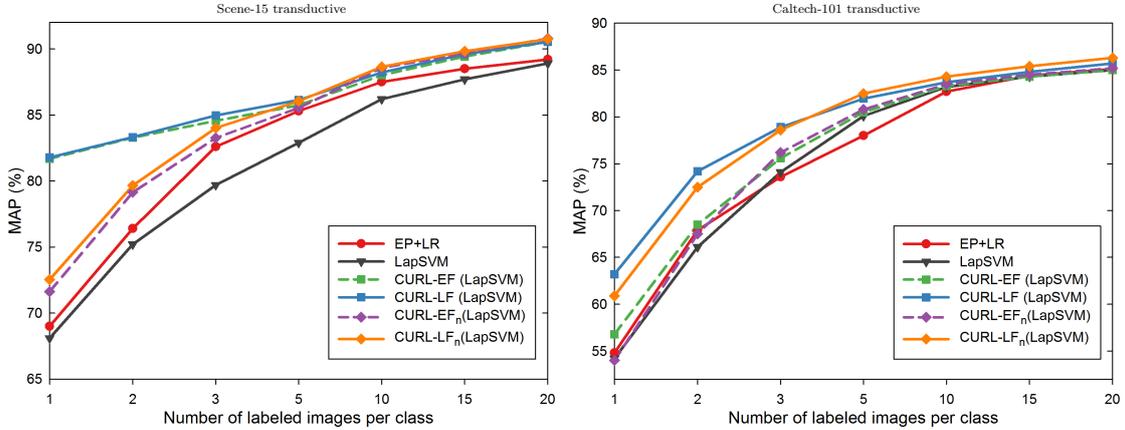


Figure 8: Mean Average Precision (MAP) varying the number of labeled images per class, obtained on the Scene-15 data set (left), and on the Caltech-101 data set (right). Results are obtained using GIST, PHOG, LBP and CNN features.

In Fig. 10 and 11 qualitative results for the ‘Panda’ class of the Caltech-101 data set are reported: the results are relative to the case in which a single instance is available for training and one single example is added at each co-training round (i.e. each pair of rows correspond to CURL-EF(LapSVM) and CURL-LF(LapSVM) respectively). The left part of Fig. 10 contains the training examples that are added by the CURL-EF(LapSVM) and CURL-LF(LapSVM) at each co-training round, while the right part and Fig. 11 contain the first 40 test images ordered by decreasing classification confidence. Samples belonging to the current class are surrounded by a green bounding box, while a red one is used for samples belonging to other classes.

In the sets of training images, it is possible to see that after the first co-training round, CURL-

LF(LapSVM) selects new examples to add to the training set, while CURL-EF(LapSVM) adds examples selected by CURL-LF(LapSVM) in the previous round. This is a pattern that we found to occur also in other categories when very small training sets are used.

In the sets of test images, it is possible to see that more and more positive images are recovered. Moreover, we can see how the images belonging to the correct class tends to be classified with increasing confidence and move to the left, while the confidences of images belonging to other classes decrease and are pushed to the right.

### 5.5. Large scale experiment

In this experiment we want to test the proposed classification strategy on a large scale data set, namely the ILSVRC 2012 which contains a total of 1000 different classes. The experiment is run on the ILSVRC 2012 validation set since the training set was used to learn the CNN features. The ILSVRC 2012 validation set, which contains a total of 50 images for each class, has been randomly divided into a training and a test set containing each 25 images per class. Again, different numbers of training images per class were tested (i.e. 1, 2, 3, 5, 10, and 20). The second embodiment of CURL is used in this experiment.

The experimental results are reported in Fig. 9 and represent the average performance over ten runs with random labeled-unlabeled feature splits.

Given the large range of MAP values, the plot of MAP improvements with respect to LapSVM baseline is also reported. It can be seen that the behavior is similar to that of the previous plots, with the LapSVM-based CURL variants outperforming the LapSVM. As for the previous data sets, the plots show that CURL-EF(LapSVM) and CURL-LF(LapSVM) are better suited for small sets of labeled images, while CURL-EF<sub>n</sub>(LapSVM) and CURL-LF<sub>n</sub>(LapSVM) are to be preferred when more labeled examples are available. It is remarkable that the proposed classification strategy is able to improve the results of the LapSVM, since the CNN features were specifically learned for the ILSVRC 2012.

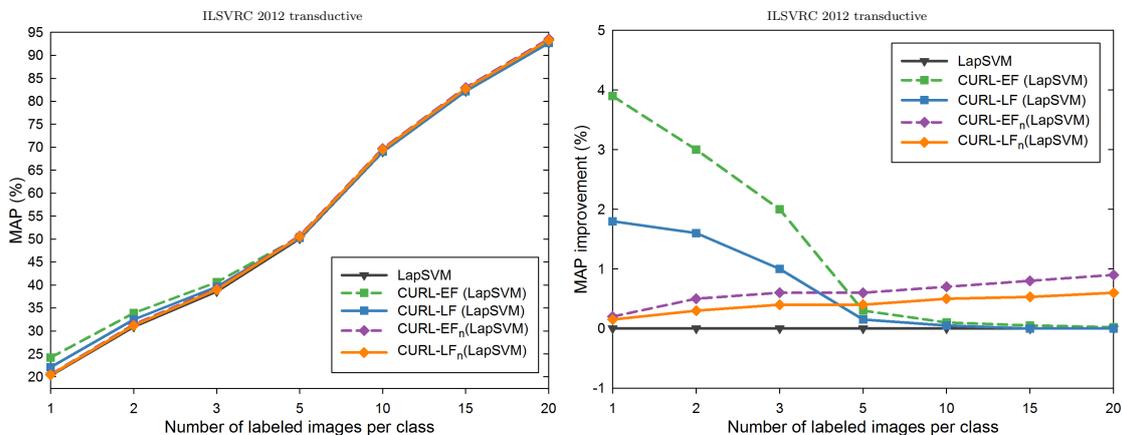


Figure 9: Mean Average Precision (MAP) varying the number of labeled images per class, obtained on the ILSVRC 2012 data set: MAP values (left) and MAP improvements over LapSVM baseline (right). Results are obtained using GIST, PHOG, LBP and CNN features.

## 6. Conclusions

In this work we have proposed CURL, a semi-supervised image classification strategy which exploits unlabeled data in two different ways: first two image representations are obtained by unsupervised learning; then co-training is used to enlarge the labeled training set of the corresponding classifiers. The two image representations are built using two different fusion schemes: early fusion and late fusion.

The proposed strategy has been tested on the Scene-15, Caltech-101, and ILSVRC 2012 data sets, and compared with other supervised and semi-supervised methods in three different experimental scenarios: inductive learning, transductive learning, and self-taught learning. We tested

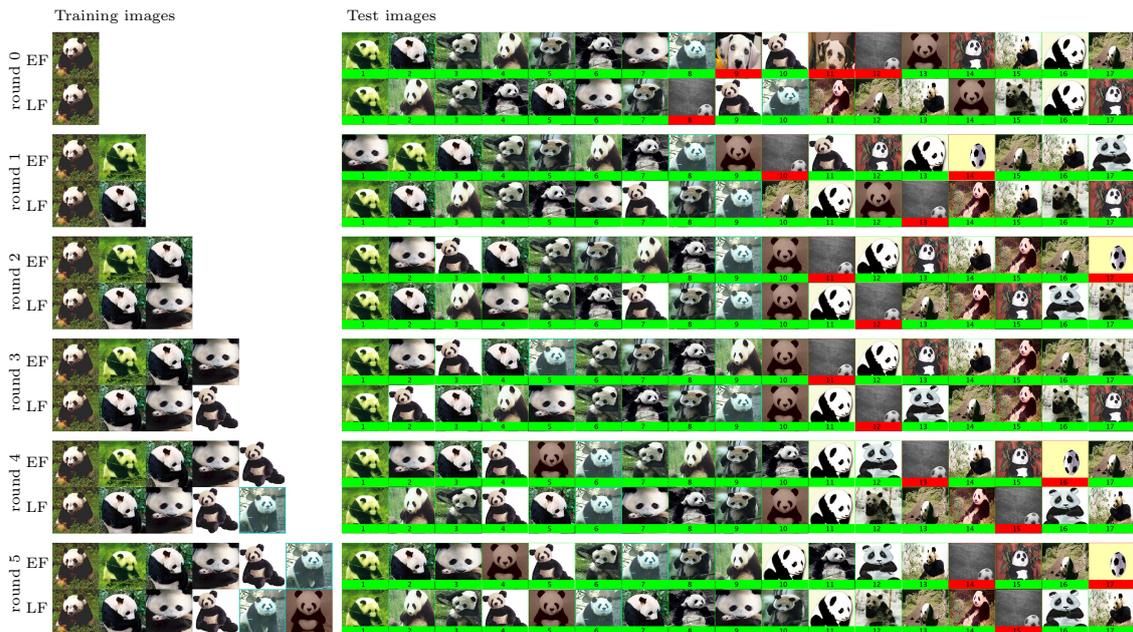


Figure 10: Qualitative results of the proposed strategy for the ‘Panda’ class of the Caltech-101 data set over five co-training rounds. Train images are on the left, the first 17 test images, ordered by decreasing classification confidence are on the right. Test images from 18 to 40 are reported in Fig. 11.

two embodiments of CURL and several variants differing in the co-trained classifier used and in the number of pseudo-labeled examples that are added at each co-training round. The experimental results showed that the CURL embodiments outperformed the other methods in the state of the art included in the comparisons. In particular, the variants that add a single pseudo-labeled example per class at each co-training round, resulted to perform best in the case of a small number of labeled images, while the variants adding more examples at each round obtained the best results when more labeled data are available.

Moreover, the results of CURL using a combination of low/mid and high level features (i.e. LBP, PHOG, GIST, and CNN features) outperform those obtained on the same features by state of the art methods. This means that CURL is able to effectively leverage less discriminative features (i.e. LBP, PHOG, GIST) to boost the performance of more discriminative ones (i.e. CNN features).

## References

- [1] O. Chapelle, B. Schölkopf, A. Zien, et al., *Semi-supervised learning*, MIT press, 2006.
- [2] J. Cheng, J. Wang, S. Jiang, Z.-H. Zhou, E. Hancock, Special edition on semi-supervised learning for visual content analysis and understanding, *Pattern Recognition* 44 (1011) (2011) 2242–2243, *semi-Supervised Learning for Visual Content Analysis and Understanding*.
- [3] K. Nigam, A. McCallum, S. Thrun, T. Mitchell, Text classification from labeled and unlabeled documents using em, *Machine learning* 39 (2-3) (2000) 103–134.
- [4] A. Fujino, N. Ueda, K. Saito, A hybrid generative/discriminative approach to semi-supervised classifier design, in: *Proc. of the National Conf. on Artificial Intelligence*, 2005, pp. 764–769.
- [5] A. Blum, S. Chawla, Learning from labeled and unlabeled data using graph mincuts, in: *Proc. 18th Int’l Conf. on Machine Learning*, 2001, pp. 19–26.
- [6] O. Chapelle, J. Weston, B. Schölkopf, Cluster kernels for semi-supervised learning, in: *Advances in neural information processing systems*, 2002, pp. 585–592.

Test images



Figure 11: Qualitative results of the proposed strategy for the ‘Panda’ class of the Caltech-101 data set over five co-training rounds. The images are ordered by decreasing classification confidence. Training image, and test images from 1 to 17 are reported in Fig. 10.

- [7] T. Joachims, Transductive inference for text classification using support vector machines, in: Proc. 16th Int’l Conf. on Machine Learning, Vol. 99, 1999, pp. 200–209.
- [8] M. Belkin, P. Niyogi, V. Sindhwani, Manifold regularization: A geometric framework for learning from labeled and unlabeled examples, *The Journal of Machine Learning Research* 7 (2006) 2399–2434.
- [9] A. Blum, T. Mitchell, Combining labeled and unlabeled data with co-training, in: Proc. of the 11th annual Conf. on Computational learning theory, 1998, pp. 92–100.
- [10] Z.-H. Zhou, M. Li, Semi-supervised learning by disagreement, *Knowledge and Information Systems* 24 (3) (2010) 415–439.
- [11] G. Iyengar, H. J. Nock, Discriminative model fusion for semantic concept detection and annotation in video, in: Proceedings of the eleventh ACM international conference on Multimedia, 2003, pp. 255–258.
- [12] P. Gehler, S. Nowozin, On feature combination for multiclass object classification, in: Computer Vision, 2009 IEEE 12th International Conference on, 2009, pp. 221–228.
- [13] P. Natarajan, S. Wu, S. Vitaladevuni, X. Zhuang, S. Tsakalidis, U. Park, R. Prasad, Multimodal feature fusion for robust event detection in web videos, in: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, 2012, pp. 1298–1305.
- [14] Z.-H. Zhou, M. Li, Tri-training: Exploiting unlabeled data using three classifiers, *Knowledge and Data Engineering, IEEE Transactions on* 17 (11) (2005) 1529–1541.
- [15] M. Li, Z.-H. Zhou, Improve computer-aided diagnosis with machine learning techniques using undiagnosed samples, *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on* 37 (6) (2007) 1088–1098.
- [16] Z.-H. Zhou, When semi-supervised learning meets ensemble learning, *Frontiers of Electrical and Electronic Engineering in China* 6 (1) (2011) 6–16.

- [17] M. Wang, X.-S. Hua, Y. Dai, L.-R. and Song, Enhanced semi-supervised learning for automatic video annotation, in: IEEE Int'l Conf. on Multimedia and Expo, 2006, pp. 1485–1488.
- [18] S. Gupta, J. Kim, K. Grauman, R. Mooney, Watch, listen & learn: Co-training on captioned images and videos, in: Machine Learning and Knowledge Discovery in Databases, 2008, pp. 457–472.
- [19] A. Levin, P. Viola, Y. Freund, Unsupervised improvement of visual detectors using cotraining, in: Proc. of IEEE Int'l Conf. on Computer Vision, 2003, pp. 626–633.
- [20] C. Christoudias, K. Saenko, L. Morency, T. Darrell, Co-adaptation of audio-visual speech and gesture classifiers, in: Proc. of the Int'l Conf. on Multimodal interfaces, 2006, pp. 84–91.
- [21] H. Feng, T.-S. Chua, A bootstrapping approach to annotating large image collection, in: Proc. of the ACM SIGMM Int'l Workshop on Multimedia Information Retrieval, 2003, pp. 55–62.
- [22] H. Bhatt, S. Bharadwaj, R. Singh, M. Vatsa, A. Noore, A. Ross, On co-training online biometric classifiers, in: Int'l Joint Conf. on Biometrics, 2011, pp. 1–7.
- [23] S. Tong, E. Chang, Support vector machine active learning for image retrieval, in: Proc. of ACM Int'l Conf. on Multimedia, 2001, pp. 107–118.
- [24] M. Guillaumin, J. Verbeek, C. Schmid, Multimodal semi-supervised learning for image classification, in: IEEE Conf. on Computer Vision and Pattern Recognition, 2010, pp. 902–909.
- [25] Y. Hong, W. Zhu, Spatial co-training for semi-supervised image classification, Pattern Recognition Letters 63 (2015) 59–65.
- [26] O. Javed, S. Ali, M. Shah, Online detection and classification of moving objects using progressively improving detectors, in: Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on, Vol. 1, IEEE, 2005, pp. 696–701.
- [27] F. Tang, S. Brennan, Q. Zhao, H. Tao, Co-tracking using semi-supervised support vector machines, in: Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on, IEEE, 2007, pp. 1–8.
- [28] W. Wang, Z. Zhou, Analyzing co-training style algorithms, in: Proc. of the European Conf. on Machine Learning, 2007, pp. 454–465.
- [29] S. Goldman, Y. Zhou, Enhancing supervised learning with unlabeled data, in: Proc. of the Int'l Conf on Machine Learning, 2000, pp. 327–334.
- [30] M. Chen, Y. Chen, K. Q. Weinberger, Automatic feature decomposition for single view co-training, in: Proceedings of the 28th International Conference on Machine Learning (ICML-11), 2011, pp. 953–960.
- [31] W. Wang, Z.-H. Zhou, Co-training with insufficient views, in: Asian Conference on Machine Learning, 2013, pp. 467–482.
- [32] W. Wang, Z. Zhou, A new analysis of co-training, in: Proc. of the Int'l Conf on Machine Learning, 2010, pp. 1135–1142.
- [33] J. Mairal, J. Ponce, G. Sapiro, A. Zisserman, F. R. Bach, Supervised dictionary learning, in: D. Koller, D. Schuurmans, Y. Bengio, L. Bottou (Eds.), Advances in Neural Information Processing Systems 21, 2009, pp. 1033–1040.
- [34] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, in: Advances in neural information processing systems, 2012, pp. 1097–1105.
- [35] Z. Zuo, G. Wang, B. Shuai, L. Zhao, Q. Yang, X. Jiang, Learning discriminative and shareable features for scene classification, in: D. Fleet, T. Pajdla, B. Schiele, T. Tuytelaars (Eds.), Computer Vision ECCV 2014, Vol. 8689 of Lecture Notes in Computer Science, Springer International Publishing, 2014, pp. 552–568.

- [36] G. Hinton, S. Osindero, Y. Teh, A fast learning algorithm for deep belief nets, *Neural computation* 18 (7) (2006) 1527–1554.
- [37] K. Jarrett, K. Kavukcuoglu, M. Ranzato, Y. LeCun, What is the best multi-stage architecture for object recognition?, in: *IEEE Int'l Conf. on Computer Vision*, 2009, pp. 2146–2153.
- [38] G. Hinton, Training products of experts by minimizing contrastive divergence, *Neural computation* 14 (8) (2002) 1771–1800.
- [39] H. Bourlard, Y. Kamp, Auto-association by multilayer perceptrons and singular value decomposition, *Biological cybernetics* 59 (4-5) (1988) 291–294.
- [40] Y. Bengio, Learning deep architectures for ai, *Foundations and trends in Machine Learning* 2 (1) (2009) 1–127.
- [41] A. Coates, A. Y. Ng, Learning feature representations with k-means, in: *Neural Networks: Tricks of the Trade*, 2012, pp. 561–580.
- [42] G. Csurka, C. Dance, L. Fan, J. Willamowski, C. Bray, Visual categorization with bags of keypoints, in: *ECCV Workshop on statistical learning in computer vision*, 2004, pp. 1–2.
- [43] S. Lazebnik, C. Schmid, J. Ponce, Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories, in: *IEEE Conf. on Computer Vision and Pattern Recognition*, Vol. 2, 2006, pp. 2169–2178.
- [44] H. Zhang, X. Bai, J. Zhou, J. Cheng, H. Zhao, Object detection via structural feature selection and shape model, *Image Processing, IEEE Transactions on* 22 (12) (2013) 4984–4995.
- [45] X. Bai, H. Zhang, J. Zhou, Vhr object detection based on structural feature extraction and query expansion, *Geoscience and Remote Sensing, IEEE Transactions on* 52 (10) (2014) 6508–6520.
- [46] B. A. Olshausen, D. Field, Sparse coding with an overcomplete basis set: A strategy employed by v1?, *Vision research* 37 (23) (1997) 3311–3325.
- [47] J. Mairal, F. Bach, J. Ponce, G. Sapiro, Online learning for matrix factorization and sparse coding, *The J. of Machine Learning Research* 11 (2010) 19–60.
- [48] M. Lewicki, T. Sejnowski, Learning overcomplete representations, *Neural computation* 12 (2) (2000) 337–365.
- [49] A. Y. Ng, M. I. Jordan, Y. Weiss, et al., On spectral clustering: Analysis and an algorithm, *Advances in neural information processing systems* 2 (2002) 849–856.
- [50] S. Lafon, Y. Keller, R. R. Coifman, Data fusion and multicue data matching by diffusion maps, *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 28 (11) (2006) 1784–1797.
- [51] B. Xiao, E. R. Hancock, R. C. Wilson, Graph characteristics from the heat kernel trace, *Pattern Recognition* 42 (11) (2009) 2589–2606.
- [52] D. Dai, L. V. Gool, Ensemble projection for semi-supervised image classification, in: *Computer Vision (ICCV), 2013 IEEE International Conference on*, IEEE, 2013, pp. 2072–2079.
- [53] P. K. Atrey, M. A. Hossain, A. El Saddik, M. S. Kankanhalli, Multimodal fusion for multimedia analysis: a survey, *Multimedia systems* 16 (6) (2010) 345–379.
- [54] C. G. Snoek, M. Worring, A. W. Smeulders, Early versus late fusion in semantic video analysis, in: *Proceedings of the 13th annual ACM international conference on Multimedia*, ACM, 2005, pp. 399–402.
- [55] W. S. Noble, et al., Support vector machine applications in computational biology, *Kernel methods in computational biology* (2004) 71–92.

- [56] S. Ayache, G. Quénot, J. Gensel, Classifier fusion for SVM-based multimedia semantic indexing, Springer, 2007.
- [57] Z. Wu, L. Cai, H. Meng, Multi-level fusion of audio and visual features for speaker identification, in: *Advances in Biometrics*, Springer, 2005, pp. 493–499.
- [58] G. R. Lanckriet, N. Cristianini, P. Bartlett, L. E. Ghaoui, M. I. Jordan, Learning the kernel matrix with semidefinite programming, *The Journal of Machine Learning Research* 5 (2004) 27–72.
- [59] M. Gönen, E. Alpaydm, Multiple kernel learning algorithms, *The Journal of Machine Learning Research* 12 (2011) 2211–2268.
- [60] D.-H. Lee, Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks, in: *Workshop on Challenges in Representation Learning, ICML*, 2013.
- [61] M.-F. Balcan, A. Blum, K. Yang, Co-training and expansion: Towards bridging theory and practice, in: *Advances in neural information processing systems*, 2004, pp. 89–96.
- [62] L. Fei-Fei, R. Fergus, P. Perona, Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories, *Computer Vision and Image Understanding* 106 (1) (2007) 59–70.
- [63] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, IEEE, 2009, pp. 248–255.
- [64] A. Oliva, A. Torralba, Modeling the shape of the scene: A holistic representation of the spatial envelope, *International journal of computer vision* 42 (3) (2001) 145–175.
- [65] A. Bosch, A. Zisserman, X. Muoz, Image classification using random forests and ferns, in: *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, 2007, pp. 1–8.
- [66] T. Ojala, M. Pietikainen, T. Maenpaa, Multiresolution gray-scale and rotation invariant texture classification with local binary patterns, *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 24 (7) (2002) 971–987.
- [67] A. S. Razavian, H. Azizpour, J. Sullivan, S. Carlsson, Cnn features off-the-shelf: an astounding baseline for recognition, in: *Computer Vision and Pattern Recognition Workshops (CVPRW), 2014 IEEE Conference on*, IEEE, 2014, pp. 512–519.
- [68] V. Sindhwani, P. Niyogi, M. Belkin, Beyond the point cloud: from transductive to semi-supervised learning, in: *Proceedings of the 22nd international conference on Machine learning*, ACM, 2005, pp. 824–831.
- [69] L. Van der Maaten, G. Hinton, Visualizing data using t-sne, *Journal of Machine Learning Research* 9 (2579-2605) (2008) 85.
- [70] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, T. Darrell, Caffe: Convolutional architecture for fast feature embedding, in: *Proceedings of the ACM International Conference on Multimedia*, ACM, 2014, pp. 675–678.
- [71] T.-F. Wu, C.-J. Lin, R. C. Weng, Probability estimates for multi-class classification by pairwise coupling, *The Journal of Machine Learning Research* 5 (2004) 975–1005.