



# Semi-supervised robust deep neural networks for multi-label image classification

Hakan Cevikalp<sup>a</sup>, Burak Benligiray<sup>b,\*</sup>, Omer Nezh Gerek<sup>b</sup>

<sup>a</sup>Eskisehir Osmangazi University, Dept. of Electrical and Electronics Engineering, Turkey

<sup>b</sup>Eskisehir Technical University, Dept. of Electrical and Electronics Engineering, Turkey

## ARTICLE INFO

### Article history:

Received 6 April 2019

Revised 2 December 2019

Accepted 14 December 2019

Available online 27 December 2019

### Keywords:

Multi-label classification

Semi-supervised learning

Ramp loss

Image classification

Deep learning

## ABSTRACT

This paper introduces a robust method for semi-supervised training of deep neural networks for multi-label image classification. To this end, a ramp loss is utilized since it is more robust against noisy and incomplete image labels compared to the classic hinge loss. The proposed method allows for learning from both labeled and unlabeled data in a semi-supervised setting. This is achieved by propagating labels from the labeled images to their unlabeled neighbors in the feature space. Using a robust loss function becomes crucial here, as the initial label propagations may include many errors, which degrades the performance of non-robust loss functions. In contrast, the proposed robust ramp loss restricts extreme penalties from the samples with incorrect labels, and the label assignment improves in each iteration and contributes to the learning process. The proposed method achieves state-of-the-art results in semi-supervised learning experiments on the CIFAR-10 and STL-10 datasets, and comparable results to the state-of-the-art in supervised learning experiments on the NUS-WIDE and MS-COCO datasets. Experimental results also verify that our proposed method is more robust against noisy image labels as expected.

© 2019 Elsevier Ltd. All rights reserved.

## 1. Introduction

Multi-label image classification is a challenging task, and it has recently attracted great attention from the research community. The majority of the studies on image classification focus on single-label classification, and very successful methods have been proposed for this task [1,2]. However, using a single label for an image is generally not appropriate for real-world applications, as the majority of the images can be associated with multiple labels to describe its semantic contents, such as objects, scenes, actions and attributes. A successful multi-label classification framework should be able to learn the association between visual features and these complex tags.

Popularization of digital cameras, web-based services and social networks let users to upload, share and tag a tremendous number of images everyday. Clearly, there is an emerging need to correctly label images, learn from them and retrieve relevant ones when needed. Manually labeling all images is too costly and difficult in practice. As a consequence, image labels are mostly collected by (semi) automatic tools using image file names, the text surrounding the image or tags provided by the users. This type of labeling

causes several problems. One of these is the fact that the labels are going to be incomplete for most cases. Consider Fig. 1 as an example: A human annotator may label this image as *car* and *road*, but in addition to these, labels such as *tires*, *car lights*, *car mirrors*, *driving* are also correct. Similarly, in a hierarchical class model like ImageNet, labeling an object to be a *poodle* logically implies that it is a *dog* and a *mammal* as well. Those labels will often be absent because they are obvious. Besides, a poodle may still have additional attributes, such as *female* and *pink*. Even in a user-model-free pure attribute model, there are usually statistical image relationships. For example, *zebra* coincides with *striped*, and *car* typically implies *wheel*. ImageNet nominally has only one labeled class per image, but in reality, several others are often present, and even visually dominant. Therefore, a successful classification system must not consider all absent labels as negatives. Another property of image labels is that they are unlikely to be independent or uncorrelated. In fact, recent studies [3–7] show that there is a strong correlation among the labels of multi-labeled images. For example, *car* and *pedestrian* usually appear together, but *pet* and *whale* are rarely seen together. Therefore, there may be a need for modeling label dependencies and adopting a learning algorithm that utilizes this information.

The final critical problem with the labels is that some of them are incorrect, as they are mostly collected by (semi) automatic

\* Corresponding author.

E-mail address: [burakbenligiray@eskisehir.edu.tr](mailto:burakbenligiray@eskisehir.edu.tr) (B. Benligiray).



Fig. 1. Although a human annotator is going to use labels such as *car*, *road* for this image, additional labels such as *tires*, *car lights*, *driving* can also be used for labeling.

tools. For instance, Sun et al. [8] recently collected one of the largest visual object classification datasets, JFT-300M, which has more than 375M noisy labels for 300M images. The labels are collected automatically, which has resulted in many missing and noisy labels. Although the authors use sophisticated algorithms to clean the noisy labels, approximately 20% of the labels are still estimated to be erroneous. These kinds of examples show the urgent need for a robust multi-label learning algorithm.

**Our contributions:** In this paper, we introduce a semi-supervised multi-label image classification method that can learn from images with noisy and incomplete labels, and even unlabeled images. In order to achieve this, we interchange the hinge loss used in the weighted approximate ranking method [9] with the more robust ramp loss. To utilize unlabeled image data, we use label propagation by assigning labels to the unlabeled data based on their nearest labeled neighbors in the feature space, resulting in bootstrapping. Our proposed classifier can both be used with hand-crafted features, or jointly trained with the feature extractor. To achieve the latter, we integrated the proposed classifier within a deep convolutional neural network by using it as the loss function of the final classification stage.

**Related work:** Various methods have been proposed for multi-label image classification over the past few years. One of the most common approaches is based on label ranking, which has been applied to powerful deep convolutional neural networks successfully. Pair-wise ranking loss was first used in [10], and it was extended to weighted approximate ranking (WARP) in [9,11]. Wu et al. [12] used WARP for weakly labeled image classification problems. Kanehira and Harada [13] also used ranking and proposed multi-label PU (positive and unlabeled) method, which mostly focused on positive labels and took into consideration the fact that absent labels are not necessarily negative as we have described earlier. However, they assume that all assigned positive labels are definitely correct, thus their method is not very robust against noisy labels. Lapin et al. [14–16] introduced methods that optimize top-k error loss functions, such as smooth top-k hinge loss and top-k softmax loss, and they reported significant improvements on some multi-label image data sets.

Several methods treat the multi-label classification problem as a deep neural network-based object detection problem, and they follow a similar strategy to the R-CNN method of Girshick et al. [17], which returns promising candidate regions from an image and classifies them with a CNN. To achieve this, Wei et al. [18] use BING to return candidate regions, and a hypothesis extraction method to reduce and refine these regions using NCuts clustering. The resulting regions are then fed into a CNN classifier as in [17]. Differently from these, Zhang et al. [6] and Karpathy and Fei-Fei [3] both use a CNN to extract candidate regions and these regions are fed to an RNN to model label dependencies. However, both methods require bounding box annotations, which can be


seen as a major limitation, since bounding box annotation is even more costly than labeling itself. Similar to these, [4] and [5] also use CNN and RNN together, but they feed the entire image to CNN and use the resulting output in an RNN to produce multiple dependent labels. Therefore, these methods do not require bounding box information. More recently, Zhu et al. [19] have introduced a multi-label image classification method that exploits both semantic and spatial relations between labels with only image-level supervision. Yu et al. [20] proposed a deep dual-stream neural network method that utilizes local features and global image pairs for multi-label image classification. Zhang et al. [21] introduced an online deep neural network classifier that can detect novel classes.

Multi-labeled images are also studied in the context of large-scale image retrieval [22–24] and dimensionality reduction [25,26]. Among these, [22] and [24] propose deep neural network-based methods for jointly learning feature representations and compact hash codes. These methods work very well on simple single-label datasets such as CIFAR-10 and ImageNet, where the objects appear in the middle of the image. However, they are significantly outperformed on challenging multi-label datasets by the method in [23,27], where a hashing method that is robust to noisy labels and pre-computed CNN features are used. Gordo et al. [28] show that neural net-based hashing methods are very sensitive to noisy labels. All these results demonstrate the importance of using robust loss functions for multi-labeled data. In the context of dimensionality reduction, Kumar et al. [26] introduced a dimensionality reduction method that embeds both the feature vectors and label vectors onto a low-dimensional space. Mikalsen et al. [25] on the other hand proposed a semi-supervised dimension reduction method that can handle noisy multi-labeled images. In addition to these methods, there are more general multi-label classification methods using Support Vector Machines (SVM) or kernel methods [29], nearest neighbors [30], and decision trees [31]. A more comprehensive survey of multi-label classification methods can be found in [32].

Recently, there have been efforts [33–38] to train deep neural networks in a semi-supervised way to utilize both labeled and unlabeled data, as neural nets require a very large amount of data due to their size. Haeusser et al. [34] proposed a sophisticated semi-supervised deep neural network method that utilizes learning by association. Their proposed method encourages correct association cycles between the embeddings of labeled and unlabeled samples. In [35], a simple method is introduced, which assigns pseudo-labels to unlabeled data based on maximum predicted class membership probabilities, and trains a neural network in a supervised manner by using both original and predicted labels. In a recent work, Wang et al. [39] have proposed a semi-supervised deep neural network method that alternates iteratively between growing convolutional layers and selecting confidently pseudo-labeled data. There are also other semi-supervised learning methods that use auto-encoder structures [33], generative adversarial networks (GANs) [37,38,40], or label propagation [25]. In addition to these methods, there are also attempts to improve accuracy based on introduction of new regularization techniques or designing new neural network architectures. For example, Springenberg et al. [41] omit pooling layers and propose a network solely consisting of convolutional layers whereas De Vries and Taylor [42] introduce cutout technique, which masks out contiguous sections of the input images to improve the robustness of deep neural network models.

## 2. Method

In this work, we consider the case with  $l$  multi-labeled images  $\mathcal{L} = \{(I_1, \mathbf{y}_1), \dots, (I_l, \mathbf{y}_l)\}$  and  $u$  unlabeled images  $\mathcal{U} = \{I_{l+1}, \dots, I_{l+u}\}$ , where  $I_i$  represents the image, and  $\mathbf{y}_i \in \{-1, 1\}^m$  is



**Fig. 2.** The illustration of the Ramp loss function,  $R_s(t) = \min(1 - s, \max(0, 1 - t))$ . The Ramp loss can be written as the difference of two convex Hinge losses,  $R_s(t) = H_1(t) - H_s(t)$ , where  $H_a(t) = \max(0, a - t)$  is the classical Hinge loss. Here, we set  $s = -0.80$ .

the corresponding label vector with  $-1$  indicating a negative and  $1$  indicating a positive class. As mentioned earlier, labels of images are usually collected from image file names, nearby text or tags provided by users. Therefore, some labels may be incorrect or missing. Let us denote the visual feature vector of an image as  $\mathbf{x}_i \in \mathbb{R}^d$ . Supervised and semi-supervised learning process cases are as described below.

2.1. Supervised learning

Assuming that  $C_{\mathbf{x}_i}^+$  and  $C_{\mathbf{x}_i}^-$  denote the sets of indices with positive and negative labels for  $\mathbf{x}_i$ , our proposed classifier solves the following optimization problem

$$\min_{\mathbf{W}} \frac{\lambda}{2} \text{trace}(\mathbf{W}^T \mathbf{W}) + \sum_{i=1}^l \sum_{j=1}^{C_{\mathbf{x}_i}^+} \sum_{k=1}^{C_{\mathbf{x}_i}^-} L(r_j) R_s(\mathbf{w}_j^T \mathbf{x}_i - \mathbf{w}_k^T \mathbf{x}_i) + \kappa \sum_{i=1}^l \sum_{j=1}^m R_s(y_{ij}(\mathbf{w}_j^T \mathbf{x}_i)), \tag{1}$$

where  $L(\cdot)$  is a weighting function for different ranks,  $r_j$  is the rank of the  $i$ th image sample for the  $j$ th class,  $R_s(t) = \min(1 - s, \max(0, 1 - t))$  is the ramp loss function and  $-1 < s \leq 0$  is a parameter that must be set by the user. Note that  $\mathbf{w}_j$  corresponds to the  $j$ th column of the  $d \times m$  weight matrix  $\mathbf{W}$ ,  $\lambda$  is the regularization parameter that must be set by the user, and  $\kappa$  is also a user-defined parameter that controls the weight of the ramp loss. As illustrated in Fig. 2, the ramp loss can be written as the difference of two convex hinge losses,  $R_s(t) = H_1(t) - H_s(t)$ , where  $H_a(t) = \max(0, a - t)$  is the classic hinge loss. In our experiments, we set  $s = -0.80$ .

The first two terms in our optimization problem (Eq. (1)) are similar to WARP [9,11] in the sense that  $L(\cdot)$  is used to control the ranking. The first term in the optimization problem is a regularization term to ensure that the method can also be used with hand-crafted features. When the proposed method is integrated into a deep neural net, this regularization term is omitted, as described below. Differently from WARP, we use the robust ramp loss function instead of the classic hinge loss in the second term of the optimization problem. In addition, WARP does not include the last term in the optimization problem, which enforces the resulting classifier to return positive scores ( $> 1$ ) for positive class samples and negative scores ( $< -1$ ) for negative class samples.

We use the same weighting function as in [9,12], which is defined as

$$L(r) = \sum_{j=1}^r \alpha_j, \tag{2}$$

where  $\alpha_j$  is set to  $1/j$ . To compute rank  $r_j$ , we sample negative labels until we find a violation, i.e.  $(1 - \mathbf{w}_j^T \mathbf{x}_i + \mathbf{w}_k^T \mathbf{x}_i) \geq 0$ . Then, the rank is calculated according to the following formulation

$$r_j = \left\lfloor \frac{m - 1}{q} \right\rfloor, \tag{3}$$

where  $q$  is the number of sampling trials. It should be noted that the function  $L(\cdot)$  controls label rankings during optimization. More precisely,  $L(\cdot)$  assigns a small weight to the loss if a positive label is ranked at the top in the label list, whereas it assigns a large weight to push the positive label to the top if the positive label is not ranked at the top of the list.

As illustrated in Fig. 2, there is no bound on the hinge loss. Therefore, it produces very large losses for incorrect labels far from the margin. As a result, these mislabeled samples play a dominant role in determining the weight matrix, which deteriorates the overall classification performance. Using ramp loss instead of hinge loss is more advantageous: When ramp loss is used, a mislabeled image sample can only introduce a limited amount of loss, regardless of its position with respect to the margin. Therefore, the mislabeled samples cannot dominate the optimization. Similarly, in the case of missing labels, the score of a missing class can be higher than a positive labeled class. A good learning algorithm should not penalize such cases heavily as they are not actual mistakes. As in the previous case, using ramp loss is better suited for such cases, because it restricts extreme penalties.

Superiority of ramp loss over hinge loss is well-proven and demonstrated in the literature for both supervised and transductive learning [23,43], but hinge loss became more popular as it is convex, whereas ramp loss is not. However, Plessis et al. [44] demonstrate the necessity of using the non-convex ramp loss instead of the convex hinge loss while learning from positive and unlabeled data both theoretically and empirically. Therefore, the use of non-convex functions for multi-labeled semi-supervised classification tasks is essential. Lastly, it should be noted that the second term and the last term of the optimization problem (Eq. 1) work in a complementary way in the sense that they both aim to maximize the margin between the positive and negative classes. Besides, adding the last term especially improves performance on classification problems in which the images have a single positive class (see the Supplementary Material<sup>1</sup> for more discussion on this).

To solve the described optimization problem, we use the stochastic gradient (SG). The proposed method is fast and scales well with the amount of data, as it was proven that the run time of SG does not directly depend on the size of the training set for these types of quadratic loss functions [45].

<sup>1</sup> It is available at <http://mlcv.ogu.edu.tr/supp.pdf>.

## 2.2. Semi-supervised learning

Semi-supervised learning algorithms try to improve the generalization performance by using unlabeled data. Majority of the semi-supervised and transductive learning methods are built on the cluster assumption, which states that the decision boundary should lie in low-density regions to improve generalization performance. Until now, almost all binary transductive learning algorithms have used the hinge loss cost  $H_1(|\mathbf{w}^\top \mathbf{x}_i + b|)$  and its variants with unlabeled data. This cost function pushes the decision boundaries to less dense areas. It must be noted that to use this cost, the classifier must return positive scores for positive class samples and negative scores for negative class samples. WARP does not necessarily satisfy this condition, which is why we introduced the last term in our optimization function (Eq. (1)). In our multi-label classification setting, we can enforce this loss as

$$\begin{aligned} \min_{\mathbf{W}} \quad & \frac{\lambda}{2} \text{trace}(\mathbf{W}^\top \mathbf{W}) + \sum_{i=1}^l \sum_{j=1}^{C_i^+} \sum_{k=1}^{C_i^-} L(r_j) R_s(\mathbf{w}_j^\top \mathbf{x}_i - \mathbf{w}_k^\top \mathbf{x}_i) \\ & + \kappa \sum_{i=1}^l \sum_{j=1}^m R_s(y_{ij}(\mathbf{w}_j^\top \mathbf{x}_i)) + \eta \sum_{i=l+1}^{l+u} \sum_{j=1}^m H_1(|\mathbf{w}_j^\top \mathbf{x}_i|), \end{aligned} \quad (4)$$

where  $\eta$  is a user defined parameter that controls the weight of the loss associated to the unlabeled samples with respect to the labeled samples. The classifier trained by solving this optimization problem was shown to be more accurate experimentally. However, this improvement was not very significant even though the majority of the training data was unlabeled. Furthermore, we observed a slight decrease in the accuracy as the number of unlabeled data samples is increased. In contrast, the semi-supervised methods that use label propagation report significant improvements as the number of unlabeled data samples is increased [46,47]. These methods use the assumption that the nearest neighbors in the feature space tend to have the same labels by transitivity, thus we can pass the labels of labeled samples to their nearest unlabeled neighbors. By following this idea, we found the nearest labeled sample of each unlabeled data sample and computed a similarity score by using the heat kernel function

$$s_i = \exp\left(-\frac{d(\mathbf{x}_i, \tilde{\mathbf{x}}_i)}{t}\right), \quad i = l+1, \dots, l+u, \quad (5)$$

where  $d(\mathbf{x}_i, \tilde{\mathbf{x}}_i)$  is the Euclidean distance between the unlabeled data sample  $\mathbf{x}_i$  and its labeled nearest neighbor  $\tilde{\mathbf{x}}_i$ . Note that this function is widely used for measuring the similarity between samples in graph-based methods, including spectral clustering. The heat kernel function returns similarity scores between 0 and 1, and scores closer to 1 indicate that the samples are very close to each other, whereas scores closer to 0 mean that the samples are far from each other. Using the heat kernel function for computing similarity is important, as there may be image samples among unlabeled data that do not come from any of the classes used for training. The similarities returned for those irrelevant image samples will be low, thus their effect to the learning process will be diminished. In the proposed semi-supervised learning method, we pass the labels of labeled data samples to their nearest unlabeled samples, and the similarity score is used to control the reliability of this label assignment. By setting  $s_i = 1, i = 1, \dots, l$ , for labeled samples, our final semi-supervised learning objective function can be written as

$$\begin{aligned} \min_{\mathbf{W}} \quad & \frac{\lambda}{2} \text{trace}(\mathbf{W}^\top \mathbf{W}) + \sum_{i=1}^{l+u} \sum_{j=1}^{C_i^+} \sum_{k=1}^{C_i^-} s_i L(r_j) R_s(\mathbf{w}_j^\top \mathbf{x}_i - \mathbf{w}_k^\top \mathbf{x}_i) \\ & + \kappa \sum_{i=1}^{l+u} \sum_{j=1}^m s_i R_s(y_{ij}(\mathbf{w}_j^\top \mathbf{x}_i)). \end{aligned} \quad (6)$$

Here, using a robust ramp loss is very important, as the label propagation procedure may have generated many incorrect labels for unlabeled data. In contrast, the non-robust hinge loss can be adversely affected by these mislabeled image samples.

Our proposed semi-supervised multi-label classifier uses the optimization problem given in (Eq. (6)), and this optimization problem can be solved using the SG algorithm given in Algorithm 1 below. For hand-crafted features, we compute the similarities at the beginning of the algorithm and pass the labels to the unlabeled data. These assignments stay permanent until the end of the algorithm. However, while we are training deep neural nets, the visual features change over time, thus we update the similarities and label assignments as described in the next section.

---

**Algorithm 1** Stochastic Gradient Based Solver for Robust Semi-Supervised Multi-Label Classifier.

---

**Initialize**

$\mathbf{W}, T > 0, \lambda > 0, \kappa > 0, \alpha_0 > 0, \epsilon_w > 0, s_i, i = 1, \dots, l+u.$

**Description:**

**for**  $t \in 1, \dots, T$  **do**

$\alpha_t \leftarrow \alpha_0/t;$

$\mathbf{W}_{t-1} \leftarrow \mathbf{W}_t;$

**for**  $i \in \text{randperm}(l+u)$  **do**

– Compute  $L(r_j)$  for sample  $\mathbf{x}_i$  by using (2),

– Compute sub-gradients with respect to the first term of the optimization,

$\mathbf{g}_{w_j}^t = \frac{\lambda \mathbf{W}}{l+u};$

– Compute sub-gradients with respect to the second term of the optimization,

$$\mathbf{g}_{w_j}^t = \begin{cases} \frac{-s_i L(r_j) \mathbf{x}_i}{l+u}, & \text{if } (1 - \mathbf{w}_j^\top \mathbf{x}_i + \mathbf{w}_k^\top \mathbf{x}_i > 0) \\ & \& (s - \mathbf{w}_j^\top \mathbf{x}_i + \mathbf{w}_k^\top \mathbf{x}_i < 0) \\ 0, & \text{otherwise.} \end{cases}$$

$$\mathbf{g}_{w_k}^t = \begin{cases} \frac{s_i L(r_j) \mathbf{x}_i}{l+u}, & \text{if } (1 - \mathbf{w}_j^\top \mathbf{x}_i + \mathbf{w}_k^\top \mathbf{x}_i > 0) \\ & \& (s - \mathbf{w}_j^\top \mathbf{x}_i + \mathbf{w}_k^\top \mathbf{x}_i < 0) \\ 0, & \text{otherwise.} \end{cases}$$

– Compute sub-gradients with respect to the third term of the optimization,

$$\mathbf{g}_{w_j}^t = \begin{cases} \frac{-s_i L \kappa \mathbf{x}_i}{l+u}, & \text{if } (1 - \mathbf{w}_j^\top \mathbf{x}_i > 0) \& (s - \mathbf{w}_j^\top \mathbf{x}_i < 0) \\ 0, & \text{otherwise.} \end{cases}$$

– Update weights,

$\mathbf{W}_t \leftarrow \mathbf{W}_t - \alpha_t \mathbf{g}_{w_j}^t;$

**end for**

**if**  $\|\mathbf{W}_t - \mathbf{W}_{t-1}\| < \epsilon_w$ , **break**

**end for**

---

## 2.3. Robust semi-supervised multi-label deep neural networks

In contrast to applying a classifier to hand-crafted features, deep convolutional neural networks enable us to learn image features and classifiers simultaneously. However, how to handle multi-labeled images with deep convolutional neural networks still remains an open problem. CNN based systems perform well on synthetic-like single labeled data sets where the foreground objects appear in the middle of the image. In contrast, these networks are significantly outperformed by classifiers applied to hand-crafted features on more challenging multi-labeled image data sets where a variety of objects with different scales appear in images [23].



To train the proposed method for feature extraction and classification jointly, we integrate the proposed classifier to the deep convolutional neural networks and call the resulting method as Robust Multi-label Convolutional Neural Network (RML-CNN). In order to achieve this, we omit the first regularization term of the optimization problem (Eq. (6)), and optimize only for the remaining terms, as various regularization methods are already implemented in deep learning frameworks. For instance, Caffe [48] allows users to select L1 or L2 norm-based regularizers controlled by the weight decay parameter. In the proposed method, we take the gradients of the second and the last term separately, as the former focuses on maximizing the difference between positive and negative class scores, and requires updating two weight vectors at the same time. As indicated earlier, ramp loss can be written as the difference of two convex hinge losses, thus the gradient of the loss  $J$  for the second term with respect to  $\mathbf{w}_j$  and  $\mathbf{w}_k$  can be found by using

$$\psi \equiv (1 - \mathbf{w}_j^\top \mathbf{x}_i + \mathbf{w}_k^\top \mathbf{x}_i > 0) \ \& \ (s - \mathbf{w}_j^\top \mathbf{x}_i + \mathbf{w}_k^\top \mathbf{x}_i < 0)$$

$$\frac{\partial J}{\partial \mathbf{w}_j} = \begin{cases} \frac{-s_i L(r_j) \mathbf{x}_i}{l+u}, & \text{if } \psi \text{ is True} \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

$$\frac{\partial J}{\partial \mathbf{w}_k} = \begin{cases} \frac{s_i L(r_j) \mathbf{x}_i}{l+u}, & \text{if } \psi \text{ is True} \\ 0, & \text{otherwise} \end{cases}$$

In a similar manner, the gradient of the last term with respect to  $\mathbf{w}_j$  can be computed by using

$$\frac{\partial J}{\partial \mathbf{w}_j} = \begin{cases} \frac{-s_i \kappa \mathbf{x}_i}{l+u}, & \text{if } (1 - \mathbf{w}_j^\top \mathbf{x}_i > 0) \ \& \ (s - \mathbf{w}_j^\top \mathbf{x}_i < 0) \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

These gradients are then propagated back by using the chain rule. It is better to update similarity scores in each epoch, but this process will be too slow since it requires to compute all pair-wise distances between labeled and unlabeled data samples. Therefore, to alleviate the computational burden, we update the similarity scores and the labels of unlabeled data at every 20 epochs for the semi-supervised case, and each epoch sees all training data (both labeled and unlabeled). To initialize the scores and labels of the unlabeled data at the beginning of the algorithm, we use the CNN model trained only by using labeled data.

#### 2.4. Setting design parameters and implementation details

There are two parameters of the proposed method that must be set by the user: the regularization parameter  $\lambda$  and the weight  $\kappa$  of the ramp loss that enforces the classifier to have positive scores for positive class samples and negative scores for negative class samples, as seen in Eq. 6. For deep neural nets, the regularization parameter is commonly called weight decay. The common practice is to set weight decay to 0.00001 for most deep neural network architectures, thus we tried values closer to this number. We obtained the best results for 0.0005 on validation data for semi-supervised learning experiments, whereas the best validation accuracy is obtained by 0.00005 for supervised learning experiments. For  $\kappa$ , we set it to 5 for all experiments, and the results show that values between 2 and 8 produce similar and good results. Thus, we set  $\kappa$  to the mean of these two values.

We used the ResNet-101 architecture [49], and trained it with the SGD algorithm. Mini-batches of size 96 were used in the experiments. The momentum parameter was set to 0.9, and we initialized the weights both randomly and by using the weights of a pre-trained network. For semi-supervised learning experiments, the learning rate was set to 0.001 for randomly initialized models, and 0.0001 for fine-tuned models. Using a smaller value of learning rate for fine-tuned models yields better results since the network is already started with reasonable weights. On the other hand, we

need a higher value of learning rate for randomly initialized models to speed-up convergence. The learning rate was kept fixed for semi-supervised learning experiments. For supervised learning experiments, the learning rate was started from 0.00005, and divided by 10 at the 10Kth and 16Kth iterations. Semi-supervised learning experiments were conducted using 4 NVIDIA GTX 1080 (8 GB memory) GPUs, whereas supervised learning experiments were conducted using a Tesla K40 (12 GB memory) GPU. Lastly, we used data augmentation in all experiments.

### 3. Experiments

We tested the proposed methods<sup>2</sup> on four visual object classification datasets: CIFAR-10, STL-10, NUS-WIDE, and MS-COCO. CIFAR-10 and STL-10 are single-label datasets. We experimented on these to compare our semi-supervised learning method to the state-of-the-art, as the recent semi-supervised methods are generally tested on these datasets. NUS-WIDE and MS-COCO are multi-label image datasets that do not contain unlabeled data. As the network architecture, we used the recently proposed ResNet-101 [49], which performs very well for many classification tasks. To compare our cost functions, we also implemented WARP with the same network.

To assess the performance, we used the classical classification accuracy for single-label datasets, whereas we used precision, recall and F1 measures computed based on the top-3 highest ranked labels as in [9,19] for the multi-label datasets. In addition to these, we also report mean average precision (mAP) scores obtained from per-class average precisions for multi-label image classification experiments.


#### 3.1. Experiments on semi-supervised learning

##### 3.1.1. CIFAR-10 dataset

The CIFAR-10 dataset includes 60K  $32 \times 32$  images of 10 object classes: *airplane, automobile, bird, cat, deer, dog, frog, horse, ship* and *truck*. 50K samples are used as training data and they are split into 5 batches, whereas the remaining 10K samples are used for testing. We randomly selected 4000 labeled samples from each batch as in [33], and the remaining 46K samples are used as unlabeled data. This procedure is repeated 5 times for each batch, and the final accuracies are the averages over these trials. A mini-batch size of 96 was used. We conducted two sets of experiments. For the first set of experiments, we started the training from scratch by random initialization for all tested cost functions. In the second case, we applied fine-tuning by using the weights of a model trained on the ILSVRC 2013 dataset. To be able to use the images with the pretrained model, we resized all images to  $256 \times 256$ .

Results are given in Table 1. It should be noted that WARP, softmax, and the proposed methods RML-CNN and SS-RML-CNN are all integrated into the same ResNet-101 network, and therefore they are directly comparable. The other methods given in the table use different type of network structures, but they are trained and tested by using the same experimental settings we have used here. The best accuracy is achieved by the proposed semi-supervised method using ILSVRC initialization, which is followed by the networks using softmax and RML-CNN. The proposed method significantly outperforms all other recent state-of-the-art methods, as well as WARP. The accuracies obtained by random initialization are quite low for all tested losses since there are very few examples for such a wide and deep network structure. Using unlabeled data significantly improves the results over supervised RML-CNN. The improvement is especially very high for the deep networks using random initialization. More precisely, the proposed semi-supervised

<sup>2</sup> <https://github.com/bbenligiray/rml-cnn/>



**Fig. 3.** Example images from the STL-10 dataset. The images shown in the first five rows are from the 10 classes, whereas the last row shows unlabeled examples that do not belong to the 10 classes in the labeled data.

**Table 1**

Average classification accuracies on the CIFAR-10 dataset.

Method	Accuracy (%)
RML-CNN (ILSVRC init.)	89.24 ± 0.3
SS-RML-CNN (ILSVRC init.)	<b>91.46</b> ± 0.2
WARP (ILSVRC init.)	85.41 ± 0.4
Softmax (ILSVRC init.)	89.32 ± 0.4
RML-CNN (random init.)	52.37 ± 2.4
SS-RML-CNN (random init.)	58.86 ± 1.7
WARP (random init.)	50.52 ± 0.5
Softmax (random init.)	51.46 ± 1.1
Haeusser et al. [34]	--
DGL [39]	82.44 ± 0.3
Improved GAN [37]	81.37 ± 2.3
ALI [40]	82.01 ± 1.6
Ladder Network [33]	79.60 ± 0.5

**Table 2**

Average classification accuracies on the STL-10 dataset.

Method	Accuracy (%)
RML-CNN (ILSVRC init.)	91.11 ± 0.3
SS-RML-CNN (ILSVRC init.)	<b>94.1</b> ± 0.2
WARP (ILSVRC init.)	85.63 ± 0.6
Softmax (ILSVRC init.)	91.34 ± 0.3
RML-CNN (random init.)	40.00 ± 0.9
SS-RML-CNN (random init.)	44.25 ± 0.8
WARP (random init.)	39.15 ± 0.7
Softmax (random init.)	39.27 ± 1.7
Haeusser et al. [34]	81.00 ± --
Huang et al. [50]	76.80 ± 0.3

method introduces approximately 6.5% improvement over the supervised one. WARP does not perform well compared to softmax, as the dataset is composed of single-label examples. However, the proposed supervised method produces comparable or better results than softmax.

### 3.1.2. STL-10 dataset

The STL-10 dataset also includes images from 10 classes as in CIFAR-10. However, the images are of higher resolution ( $96 \times 96$ ), and the dataset is specifically developed for unsupervised and semi-supervised learning. This dataset includes 5K labeled samples and 100K unlabeled samples. Unlabeled data includes images from the 10 classes, as well as some additional classes that are not present in the labeled set. In contrast, the unlabeled data from the CIFAR-10 dataset only include images from the 10 classes.

Some images from the STL-10 dataset are illustrated in Fig. 3. This dataset is split into 10 pre-defined folds, where each fold contains 100 labeled images per class. The test set is composed of 8K images. We used 1000 labeled samples provided in each fold and all unlabeled data for learning and repeated this procedure for each fold. The final accuracy is the averages of accuracies obtained from 10 folds.

Accuracies are presented in Table 2. The results are similar to the results of the CIFAR-10 dataset in the sense that the proposed semi-supervised method again achieves the best accuracy, followed by softmax and RML-CNN. ILSVRC initialization produces much higher accuracies compared to random initialization for all tested methods. WARP performs worse than both softmax and RML-CNN for ILSVRC initialization, but its accuracy is similar to these methods for random initialization. Our reported results are based on 10 repetitions, but Haeusser et al. [34] performs only 1 experiment. It should be noted that the proposed methods significantly out-

**Table 3**  
Accuracies on the NUS-WIDE dataset on 81 classes for top-3 highest ranked labels.

Method	P-C	R-C	F1-C	P-O	R-O	F1-O	mAP
RML-CNN	44.3	54.8	45.3	55.7	69.0	61.7	57.4
WARP	40.2	51.8	42.2	53.9	66.7	59.7	48.5
Softmax	31.7	31.2	31.4	47.8	59.5	53.0	–
WARP [9]	31.7	35.6	33.5	48.6	60.5	53.9	–
CNN+RNN [4]	40.5	30.4	34.7	49.9	61.7	55.2	–
RLSD [6]	44.4	49.6	46.9	54.4	67.6	60.3	54.1
ResNet-107 (binary relevance) [19]	46.7	<b>56.8</b>	46.9	<b>55.9</b>	<b>69.2</b>	<b>61.8</b>	59.5
ResNet-101 (binary relevance) [19]	46.4	55.3	47.0	<b>55.9</b>	<b>69.2</b>	<b>61.8</b>	<b>60.1</b>
CNN+RMLC (ours)	<b>48.2</b>	56.0	<b>48.8</b>	<b>55.9</b>	69.0	<b>61.8</b>	58.8

**Table 4**  
Accuracies on the MS-COCO dataset on 80 classes for top-3 highest ranked labels.

Method	P-C	R-C	F1-C	P-O	R-O	F1-O	mAP
RML-CNN	62.4	61.1	59.8	62.8	65.4	64.1	71.5
WARP	60.7	59.8	58.5	61.5	64.1	62.8	63.2
WARP [9]	59.3	52.5	55.7	59.8	61.4	60.7	–
CNN+RNN [4]	66.0	55.6	60.4	<b>69.2</b>	66.4	<b>67.8</b>	–
RLSD [6]	<b>67.7</b>	56.4	<b>61.5</b>	70.5	59.9	64.8	67.4
ResNet-101 (binary relevance) [19]	65.3	<b>62.6</b>	61.3	64.1	<b>66.8</b>	65.4	<b>75.2</b>
CNN+RMLC (ours)	64.9	62.0	<b>61.5</b>	64.1	<b>66.8</b>	65.5	<b>75.2</b>

perform both the method of Haeusser et al. [34] and the previous state-of-the-art method of [50], which achieves 76.8% accuracy.

### 3.2. Experiments on multi-label learning

#### 3.2.1. NUS-WIDE dataset

This dataset contains 269, 648 images from Flickr, which are annotated manually. There are 81 categories, with 2.4 category labels per image on average. In addition to the objects classes, there are event/activity names among the labels such as *swimming* and *running*, as well as scene/location names such as *airport* or *ocean*. The officially provided train/test split is used in our experiments. More precisely we use 161, 789 images for training and validation, and 107, 859 images for testing. It should be noted that we learn all 81 categories, as opposed to some methods that use only the most frequent 21 categories.

Results are given in Table 3. It should be noted that RML-CNN and WARP given at the top of the table are integrated into the same network again, thus they are directly comparable. The proposed method, RML-CNN, yielded lower results than the ResNet-101 architecture with binary relevance cost for multi-label learning [19]. We trained the same network with the provided code and obtained slightly lower results compared to the proposed method RML-CNN.<sup>3</sup> Therefore, we also conducted new experiments by using the CNN features of the trained model provided by the authors of [19]. In this setting, we treated the CNN features as hand-crafted features and then trained our classifier using the cost function (Eq. (1)). The resulting classifier (CNN+RMLC) achieves the best P-C, F1-C, P-O, and F1-O scores. Especially, its F1-C score is 1.8% better compared to the results of Zhu et al. [19]. Both of the proposed methods, RML-CNN and CNN+RMLC, significantly outperform WARP. It should be noted that the RLSD method uses bounding box information for multi-label image classification, yet the proposed methods still outperform it by a small margin.

<sup>3</sup> We trained the network using the code and parameters provided by the authors of Zhu et al. [19]. The authors of Zhu et al. [19] use 4 GPUs with a batch size of 96 (i.e., 24 images per GPU, due to memory limitations). Since we had a single GPU, we collected gradients for 96 images in 4 mini-batches, and then made the updates. The difference may be due to this procedure.

#### 3.2.2. MS-COCO dataset

The Microsoft COCO (MS-COCO) dataset [51] is a large-scale multi-label benchmark dataset collected for several vision tasks such as recognition, segmentation and captioning. The training set includes 82, 783 images, whereas the test set includes 40, 504 images. The objects are categorized into 80 classes and there are approximately 2.95 object labels per image.

The accuracies based on the top-3 highest ranked labels are given in Table 4. As in the previous experiment, in addition to RML-CNN, we also conducted new experiments by using the CNN features of the ResNet-101 trained model provided by the authors of Zhu et al. [19]. We again treated the CNN features as hand-crafted features and trained our classifier using the cost function (Eq. (1)). The best F1-C scores are achieved by the proposed CNN+RMLC and RSLD methods, whereas the best mAP and R-O scores are obtained by the proposed CNN+RMLC and ResNet-101 method of Zhu et al. [19]. CNN+RNN achieves the best P-O and F1-O accuracies, and RSLD also obtains the best P-C accuracy. However, it should be kept in mind that RSLD needs bounding box annotations, which are not available for most visual recognition datasets.

#### 3.3. Experiments with noisy data

To demonstrate the robustness of the proposed method, we designed an experiment where image labels are deliberately corrupted. In this experiment, we used 5 batches of 4000 labeled examples from the CIFAR-10 dataset, similar to the supervised experiments in Section 3.1. In the same vein, we subsampled the MS-COCO and NUS-WIDE datasets while maintaining the class frequencies, which allowed us to run many experiments across different parameters.

To corrupt the dataset, a subset is chosen randomly, and the corresponding labels are altered stochastically, while maintaining the class frequencies of the original dataset. The same subset of labels are gradually corrupted between 10% and 50% for all parameters to reach consistent results. The models are initialized with ImageNet pre-trained models and the hyperparameters are optimized for the illustrated metrics.

The results are presented in Fig. 4. The most apparent result is that RML-CNN outperforms WARP across all datasets and parameters. The experiments with the MS-COCO and NUS-WIDE datasets are done with both 10% and 20% of the dataset, which gave consis-






Fig. 4. Classification performance on the CIFAR-10, MS-COCO and NUS-WIDE datasets with varying density of corrupted labels.

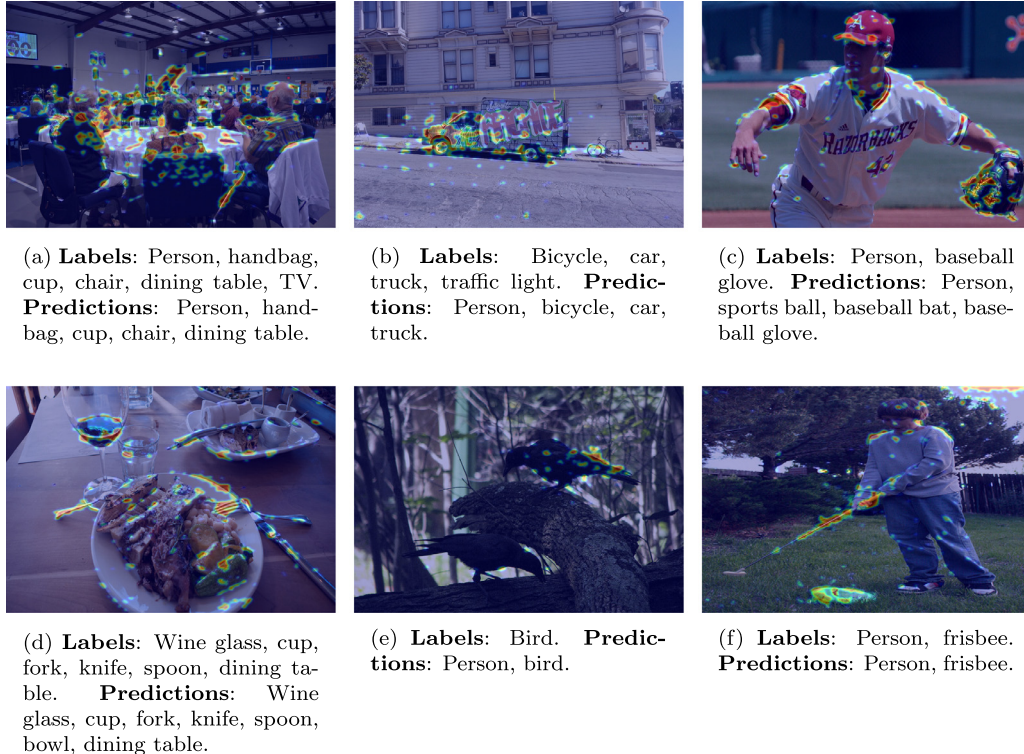


Fig. 5. Saliency maps superimposed on the images as heat maps for visualization. The images are from the MS-COCO dataset, and guided backpropagation [41] is used to produce the saliency maps. Predictions are the classes for which the produced scores were positive.

tent results. Both RML-CNN and WARP performances degrade linearly as the ratio of degraded labels increases. We can observe that using more data (i.e., 20%, instead of 10% of the dataset) provides a constant improvement across all parameters. Finally, it is surprising to see that both methods manage to converge even when half of the dataset is mislabeled.

### 3.4. Qualitative results

See Fig. 5 for qualitative results, for which we are going to provide commentary. The images are from the validation set of the MS-COCO dataset. A saliency map is produced for each of them using guided backpropagation [41], and these saliency maps are superimposed on the original images as heat maps using the method described in [52]. The labels in the subcaptions are gathered from the dataset, while the predictions are the classes for which our model has returned positive scores.

The first column (Fig. 5a–d) showcases performance in cluttered scenes. The proposed method recognizes the objects correctly, and also produces appropriate saliency maps. The second column (Fig. 5b–e) illustrates errors due to dataset imbalance. It

is much more difficult to gather a balanced multiclass dataset, and whether if this is desirable is up to debate. The model tends towards assigning higher scores for overly frequent classes (in this case, *person*). Although rarely, this results in the said object being recognized even in very irrelevant contexts (e.g., Fig. 5e). The final column (Fig. 5c–f) addresses underlying correlations between classes. For example, *baseball bat* is recognized in Fig. 5c, solely because there is a person with a baseball glove in the image. This implies that the CNN is able to model correlations between labels by using the underlying features (i.e., features extracted from the baseball glove increases the classification score for baseball bat). On the other hand, the model is also capable of performing correctly when the context is incoherent. For example, a person is hitting a frisbee with a golf club in Fig. 5f, yet this does not trigger a sports ball classification and the result is correct.

## 4. Conclusion

This study argues that image labels in large-scale datasets are typically noisy and incomplete, and traditional methods using non-robust loss functions may fail to classify them accurately. For such



cases, it is advantageous to use loss functions that are more robust against noisy and missing labels. In this paper, we proposed a robust semi-supervised method for multi-label image classification. To this end, we interchanged the classic non-robust hinge loss from the WARP method with the robust ramp loss and added another loss term that encourages margin maximization. To incorporate the unlabeled image data in the learning process, we also used label propagation by assigning labels to the unlabeled data based on their closest labeled neighbors. Our proposed classifiers achieved state-of-the-art results in semi-supervised classification on CIFAR-10 and STL-10 datasets. The proposed classifiers also significantly outperformed WARP on noisy datasets. Moreover, we obtained comparable results to the state-of-the-art on multi-labeled NUS-WIDE and MS-COCO datasets for multi-label classification. The proposed classifiers are robust against noisy and missing image labels, but they do not explicitly model label dependencies. As a future work, we are planning to add an RNN network to the current CNN network to build a structure that explicitly models such correlations between labels in an end-to-end manner.

### Declaration of Competing Interest

None.

### Acknowledgements

The authors would like to thank NVIDIA for the Tesla K40 GPU donation used in this study.

### References

- [1] A. Krizhevsky, I. Sutskever, G. Hinton, Imagenet classification with deep neural networks, NIPS, 2012.
- [2] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A.C. Berg, L. Fei-Fei, Imagenet large scale visual recognition challenge, *Int. J. Comput. Vision* 115 (2015) 211–252.
- [3] A. Karpathy, L. Fei-Fei, Deep visual-semantic alignments for generating image descriptions, CVPR, 2015.
- [4] J. Wang, Y. Yang, J. Mao, Z. Huang, C. Huang, W. Xu, Cnn-rnn: a unified framework for multi-label image classification, CVPR, 2016.
- [5] J. Jin, H. Nakayama, Annotation order matters: recurrent image annotator for arbitrary length image tagging, ICPR, 2016.
- [6] J. Zhang, Q. Wu, C. Shen, J. Zhang, J. Lu, Multi-label image classification with regional latent semantic dependencies, in: arXiv:1612.01082v1, 2016.
- [7] T.T. Nguyen, T.T.T. Nguyen, A.V. Luong, Q.V.H. Nguyen, A.W.-C. Liew, B. Stantic, Multi-label classification via label correlation and first order feature dependence in a data stream, *Pattern Recognit.* 90 (2019) 35–51.
- [8] C. Sun, A. Shrivastava, S. Singh, A. Gupta, Revisiting unreasonable effectiveness of data in deep learning era, ICCV, 2017.
- [9] Y. Gong, Y. Jia, T. Leung, A. Toshev, S. Ioffe, Deep convolutional ranking for multilabel image annotation, in: arXiv:1312.4894v1, 2014.
- [10] N. Usunier, D. Buffoni, P. Gallinari, Ranking with ordered weighted pairwise classification, ICML, 2009.
- [11] J. Weston, S. Bengio, N. Usunier, Wsabie: Scaling up to large vocabulary image annotation, IJCAI, 2011.
- [12] F. Wu, Z. Wang, Z. Zhang, Y. Yang, J. Luo, W. Zhu, Y. Zhuang, Weakly semi-supervised deep learning for multi-label image annotation, *IEEE Trans. Big Data* 1 (2015) 109–122.
- [13] A. Kanehira, T. Harada, Multi-label ranking from positive and unlabeled data, CVPR, 2016.
- [14] M. Lapin, M. Hein, B. Schiele, Top-k multiclass svm, NIPS, 2015.
- [15] M. Lapin, M. Hein, B. Schiele, Loss functions for top k-error: analysis and insights, CVPR, 2016.
- [16] M. Lapin, M. Hein, B. Schiele, Analysis and optimization of loss functions for multiclass, top-k, and multilabel classification, in: arXiv:1612.03663, 2016.
- [17] R. Girshick, J. Donahue, T. Darrell, J. Malik, Rich feature hierarchies for accurate object detection and semantic segmentation, CVPR, 2014.
- [18] Y. Wei, W. Xia, M. Lin, J. Huang, B. Ni, J. Dong, Y. Zhao, S. Yan, Hcp: a flexible cnn framework for multi-label image classification, *IEEE Trans. PAMI* 39 (2016) 1901–1907.
- [19] F. Zhu, H. Li, W. Ouyang, N. Yu, X. Wang, Learning spatial regularization with image-level supervisions for multi-label image classification, CVPR, 2017.
- [20] W.-J. Yu, Z.-D. Chen, X. Luo, W. Liu, X.-S. Xu, DELTA: a deep dual-stream network for multi-label image classification, *Pattern Recognit.* 91 (2019) 322–331.
- [21] Y. Zhang, Y. Wang, X.-Y. Liu, S. Mi, M.-L. Zhang, Large-scale multi-label classification using unknown streaming images, *Pattern Recognit.* 99 (2020) 1–10.
- [22] F. Zhao, Y. Huang, L. Wang, T. Tan, Deep semantic ranking based hashing for multi-label image retrieval, CVPR, 2015.
- [23] H. Cevikalp, M. Elmas, S. Ozkan, Towards category based large-scale image retrieval using transductive support vector machines, ECCV Workshops, 2016.
- [24] H. Lai, Y. Pan, S. Yank, Simultaneous feature learning and hash coding with deep neural networks, CVPR, 2015.
- [25] K.O. Mikalsen, C. Soguero-Ruiz, F.M. Bianchi, Noisy multi-label semi-supervised dimensionality reduction, *Pattern Recognit.* 90 (2019) 257–270.
- [26] V. Kumar, A.K. Pujari, V. Padmanabhan, V.R. Kagita, Group preserving label embedding for multi-label classification, *Pattern Recognit.* 90 (2019) 23–34.
- [27] H. Cevikalp, M. Elmas, S. Ozkan, Large-scale image retrieval using transductive support vector machines, *Comput. Vision Image Understanding* 173 (2018) 2–12.
- [28] A. Gordo, J. Almazan, J. Revaud, D. Larlus, Deep image retrieval: Learning global representations for image search, ECCV, 2016.
- [29] A. Elisseeff, J. Weston, A kernel method for multi-labeled classification, NIPS, 2001.
- [30] M.-L. Zhang, Z. Zhou, Ml-knn: a lazy learning approach to multi-label learning, *Pattern Recognit.* 40 (2007) 2038–2048.
- [31] A. Clare, R.D. King, Knowledge discovery in multi-label phenotype data, *Lecture Notes in Computer Science*, 2001.
- [32] M.-L. Zhang, Z.-H. Zhou, A review on multi-label learning algorithms, *IEEE Trans. Knowl. Data Eng.* 26 (2014) 1819–1837.
- [33] A. Rasmus, H. Valpoa, M. Honkala, M. Berglund, T. Raiko, Semi-supervised learning with ladder networks, NIPS, 2015.
- [34] P. Haeusser, A. Mordvintsev, D. Cremers, Learning by association a versatile semi-supervised training method for neural networks, CVPR, 2017.
- [35] D.-H. Lee, The simple and efficient semi-supervised learning method for deep neural networks, ICML Workshops, 2013.
- [36] J. Weston, F. Ratle, R. Collobert, Deep learning via semi-supervised embedding, ICML, 2008.
- [37] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, Improved techniques for training gans, in: arXiv:1606.03498v1, 2016.
- [38] D.P. Kingma, D.J. Rezende, S. Mohamed, M. Welling, Semi-supervised learning with deep generative models, in: arXiv:1406.5298v1, 2014.
- [39] G. Wang, X. Xie, J. Lai, J. Zhuo, Deep growing learning, ICCV, 2017.
- [40] V. Dumoulin, I. Belghazi, B. Poole, O. Mastropietro, A.L. abd M. Arjovsky, A. Courville, Adversarially learned inference, ICLR, 2017.
- [41] J.T. Springenberg, A. Dosovitskiy, T. Brox, M. Riedmiller, Striving for simplicity: the all convolutional net, ICLR Workshops, 2015.
- [42] T. Devries, G.W. Taylor, Improved regularization of convolutional neural networks with cutout, *CoRR* (2017) abs/1708.04552.
- [43] S. Ertekin, L. Bottou, C. Giles, Nonconvex online support vector machines, *IEEE Trans. PAMI* 33 (2011) 368–381.
- [44] M.C. du Plessis, G. Niu, M. Sugiyama, Analysis of learning from positive and unlabeled data, NIPS, 2014.
- [45] S. Shalev-Shwartz, Y. Singer, N. Srebro, Pegasos: primal estimated subgradient solver for svm, NIPS, 2007.
- [46] H. Cevikalp, J. Verbeek, F. Jurie, A. Klaser, Semi-supervised dimensionality reduction using pairwise equivalence constraints, in: *International Conference on Computer Vision Theory and Applications*, 2008.
- [47] B. Scholkopf, D. Zhou, J. Huang, Learning from labeled and unlabeled data on a directed graphs, ICML, 2005.
- [48] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, T. Darrell, Caffe: convolutional architecture for fast feature embedding, arXiv:1408.5093v1, (2014).
- [49] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, CVPR, 2016.
- [50] C. Huang, C.C. Loy, X. Tang, Unsupervised learning of discriminative attributes and visual representations, CVPR, 2016.
- [51] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollar, C.L. Zitnick, Microsoft coco: common objects in context, ECCV, 2014.
- [52] O.E. Yetgin, B. Benligiray, O.N. Gerek, Power line recognition from aerial images with deep learning, *IEEE Trans. Aerosp. Electron. Syst.* (2018).

**Hakan Cevikalp** received his Ph.D. degree from Vanderbilt University, Department of Electrical Engineering and Computer Science in 2005 and is currently a full professor in Eskisehir Osmangazi University, Department of Electrical and Electronics Engineering. His research interests include pattern recognition, neural networks, image and signal processing, optimization, and computer vision.

**Burak Benligiray** received his Ph.D. degree on deep learning-based vision at Eskişehir Technical University, Department of Electrical and Electronics Engineering and is a research assistant at the same department. His interests are deep learning, computer vision and smart contract oracles.

**Omer Nezh Gerek** is a full professor at Eskişehir Technical University, Department of Electrical and Electronics Engineering. Prof. Gerek is a member of TUBITAK (Turkish Scientific and Technological Research Council) EE-CS management committee. He is an IEEE senior member, and is serving in the editorial boards of Elsevier: DSP and TUBITAK: TJECS. His research interests include signal and image processing and analysis.