

# Learning a Deep ConvNet for Multi-label Classification with Partial Labels

Thibaut Durand      Nazanin Mehrasa      Greg Mori  
 Borealis AI      Simon Fraser University  
 {tdurand, nmehrasa}@sfu.ca      mori@cs.sfu.ca

## Abstract

Deep ConvNets have shown great performance for single-label image classification (e.g. ImageNet), but it is necessary to move beyond the single-label classification task because pictures of everyday life are inherently multi-label. Multi-label classification is a more difficult task than single-label classification because both the input images and output label spaces are more complex. Furthermore, collecting clean multi-label annotations is more difficult to scale-up than single-label annotations. To reduce the annotation cost, we propose to train a model with partial labels i.e. only some labels are known per image. We first empirically compare different labeling strategies to show the potential for using partial labels on multi-label datasets. Then to learn with partial labels, we introduce a new classification loss that exploits the proportion of known labels per example. Our approach allows the use of the same training settings as when learning with all the annotations. We further explore several curriculum learning based strategies to predict missing labels. Experiments are performed on three large-scale multi-label datasets: MS COCO, NUS-WIDE and Open Images.

## 1. Introduction

Recently, Stock and Cisse [46] presented empirical evidence that the performance of state-of-the-art classifiers on ImageNet [44] is largely underestimated – much of the remaining error is due to the fact that ImageNet’s single-label annotation ignores the intrinsic multi-label nature of the images. Unlike ImageNet, multi-label datasets (e.g. MS COCO [35], Open Images [32]) contain more complex images that represent scenes with several objects (Figure 1). However, collecting multi-label annotations is more difficult to scale-up than single-label annotations [13]. As an alternative strategy, one can make use of partial labels; collecting partial labels is easy and scalable with crowdsourcing platforms. In this work, we study the problem of learning a multi-label classifier with partial labels per image.

The two main (and complementary) strategies to im-



	[a]	[b]	[c]
car	✓	✓	✓
person	✓		✗
boat	✗		✗
bear	✗	✗	✗
apple	✗		✗

Figure 1. Example of image with all annotations [a], partial labels [b] and noisy/webly labels [c]. In the partially labeled setting some annotations are missing (person, boat and apple) whereas in the webly labeled setting one annotation is wrong (person).

prove image classification performance are: (i) designing / learning better model architectures [41, 21, 47, 63, 15, 57, 50, 14, 43, 64, 36, 16] and (ii) learning with more labeled data [48, 37]. However, collecting a multi-label dataset is more difficult and less scalable than collecting a single label dataset [13], because collecting a consistent and exhaustive list of labels for every image requires significant effort. To overcome this challenge, [48, 34, 37] automatically generated the labels using web supervision. But the drawback of these approaches is that the annotations are noisy and not exhaustive, and [62] showed that learning with corrupted labels can lead to very poor generalization performance. To be more robust to label noise, some methods have been proposed to learn with noisy labels [53].

An orthogonal strategy is to use partial annotations. This direction is actively being pursued by the research community: the largest publicly available multi-label dataset is annotated with partial clean labels [32]. For each image, the labels for some categories are known but the remaining labels are unknown (Figure 1). For instance, we know there is a *car* and there is not a *bear* in the image, but we do not know if there is a *person*, a *boat* or an *apple*. Relaxing the learning requirement for exhaustive labels opens better opportunities for creating large-scale datasets. Crowdsourcing platforms like Amazon Mechanical Turk<sup>1</sup> and Google Image Labeler<sup>2</sup> or web services like reCAPTCHA<sup>3</sup> can scal-

<sup>1</sup><https://www.mturk.com/>

<sup>2</sup><https://crowdsource.google.com/imagelabeler/category>

<sup>3</sup><https://www.google.com/recaptcha/>

ably collect partial labels for a large number of images.

To our knowledge, this is the first work to examine the challenging task of learning a multi-label image classifier with partial labels on large-scale datasets. Learning with partial labels on large-scale datasets presents novel challenges because existing methods [52, 58, 56, 59] are not scalable and cannot be used to fine-tune a ConvNet. We address these key technical challenges by introducing a new loss function and a method to fix missing labels.

Our first contribution is to empirically compare several labeling strategies for multi-label datasets to highlight the potential for learning with partial labels. Given a fixed label budget, our experiments show that partially annotating all images is better than fully annotating a small subset.

As a second contribution, we propose a scalable method to learn a ConvNet with partial labels. We introduce a loss function that generalizes the standard binary cross-entropy loss by exploiting label proportion information. This loss automatically adapts to the proportion of known labels per image and allows to use the same training settings as when learning with all the labels.

Our last contribution is a method to predict missing labels. We show that the learned model is accurate and can be used to predict missing labels. Because ConvNets are sensitive to noise [62], we propose a curriculum learning based model [2] that progressively predicts some missing labels and adds them to the training set. To improve label predictions, we develop an approach based on Graph Neural Networks (GNNs) to explicitly model the correlation between categories. In multi-label settings, not all labels are independent, hence reasoning about label correlation between observed and unobserved partial labels is important.

## 2. Related Work

**Learning with partial / missing labels.** Multi-label tasks often involve incomplete training data, hence several methods have been proposed to solve the problem of multi-label learning with missing labels (MLML). The first and simple approach is to treat the missing labels as negative labels [49, 3, 38, 55, 48, 37]. The MLML problem then becomes a fully labeled learning problem. This solution is used in most webly supervised approaches [48, 37]. The standard assumption is that only the category of the query is present (*e.g.* *car* in Figure 1) and all the other categories are absent. However, performance drops because a lot of ground-truth positive labels are initialized as negative labels [26]. A second solution is Binary Relevance (BR) [52], which treats each label as an independent binary classification. But this approach is not scalable when the number of categories grows and it ignores correlations between labels and between instances, which can be helpful for recognition. Unlike BR, our proposed approach allows to learn a single model using partial labels.

To overcome the second problem, several works proposed to exploit label correlations from the training data to propagate label information from the provided labels to missing labels. [4, 58] used a matrix completion algorithm to fill in missing labels. These methods exploit label-label correlations and instance-instance correlations with low-rank regularization on the label matrix to complete the instance-label matrix. Similarly, [61] introduced a low rank empirical risk minimization, [56] used a mixed graph to encode a network of label dependencies and [38, 13] learned correlation between the categories to predict some missing labels. Unlike most of the existing models that assume that the correlations are linear and unstructured, [59] proposed to learn structured semantic correlations. Another strategy is to treat missing labels as latent variables in probabilistic models. Missing labels are predicted by posterior inference. [27, 54] used models based on Bayesian networks [23] whereas [10] proposed a deep sequential generative model based on a Variational Auto-Encoder framework [29] that also allows to deal with unlabeled data.

However, most of these works cannot be used to learn a deep ConvNet. They require solving an optimization problem with the training set in memory, so it is not possible to use a mini-batch strategy to fine-tune the model. This is limiting because it is well-known that fine-tuning is important to transfer a pre-trained architecture [30]. Some methods are also not scalable because they require to solve convex quadratic optimization problems [56, 59] that are intractable for large-scale datasets. Unlike these methods, we propose a model that is scalable and end-to-end learnable. To train our model, we introduce a new loss function that adapts itself to the proportion of known labels per example. Similar to some MLML methods, we also explore several strategies to fill-in missing labels by using the learned classifier.

Learning with partial labels is different from semi-supervised learning [6] because in the semi-supervised learning setting, only a subset of the examples is labeled with all the labels and the other examples are unlabeled whereas in the partial labels setting, all the images are labeled but only with a subset of labels. Note that [12] also introduced a partially labeled learning problem (also called ambiguously labeled learning) but this problem is different: in [12], each example is annotated with multiple labels but only one is correct.

**Curriculum Learning / Never-Ending Learning.** To predict missing labels, we propose an iterative strategy based on Curriculum Learning [2]. The idea of Curriculum Learning is inspired by the way humans learn: start to learn with easy samples/subtasks, and then gradually increase the difficulty level of the samples/subtasks. But, the main problem in using curriculum learning is to measure the difficulty of an example. To solve this problem, [31] used the defini-

tion that easy samples are ones whose correct output can be predicted easily. They introduced an iterative self-paced learning (SPL) algorithm where each iteration simultaneously selects easy samples and updates the model parameters. [24] generalizes the SPL to different learning schemes by introducing different self-paced functions. Instead of using human-designed heuristics, [25] proposed MentorNet, a method to learn the curriculum from noisy data. Similar to our work, [20] recently introduced the CurriculumNet that is a model to learn from large-scale noisy web images with a curriculum learning approach. However this strategy is designed for multi-class image classification and cannot be used for multi-label image classification because it uses a clustering-based model to measure the difficulty of the examples.

Our approach is also related to the Never-Ending Learning (NEL) paradigm [39]. The key idea of NEL is to use previously learned knowledge to improve the learning of the model. [33] proposed a framework that alternatively learns object class models and collects object class datasets. [5, 39] introduced the Never-Ending Language Learning to extract knowledge from hundreds of millions of web pages. Similarly, [7, 8] proposed the Never-Ending Image Learner to discover structured visual knowledge. Unlike these approaches that use a previously learned model to extract knowledge from web data, we use the learned model to predict missing labels.

### 3. Learning with Partial Labels

Our goal in this paper is to train ConvNets given partial labels. We first introduce a loss function to learn with partial labels that generalizes the binary cross-entropy. We then extend the model with a Graph Neural Network to reason about label correlations between observed and unobserved partial labels. Finally, we use these contributions to learn an accurate model that it is used to predict missing labels with a curriculum-based approach.

**Notation.** We denote by  $C$  the number of categories and  $N$  the number of training examples. We denote the training data by  $\mathcal{D} = \{(\mathcal{I}^{(1)}, \mathbf{y}^{(1)}), \dots, (\mathcal{I}^{(N)}, \mathbf{y}^{(N)})\}$ , where  $\mathcal{I}^{(i)}$  is the  $i^{\text{th}}$  image and  $\mathbf{y}^{(i)} = [y_1^{(i)}, \dots, y_C^{(i)}] \in \mathcal{Y} \subseteq \{-1, 0, 1\}^C$  the label vector. For a given example  $i$  and category  $c$ ,  $y_c^{(i)} = 1$  (resp.  $-1$  and  $0$ ) means the category is present (resp. absent and unknown).  $\mathbf{y} = [\mathbf{y}^{(1)}; \dots; \mathbf{y}^{(N)}] \in \{-1, 0, 1\}^{N \times C}$  is the matrix of training set labels.  $f_{\mathbf{w}}$  denotes a deep ConvNet with parameters  $\mathbf{w}$ .  $\mathbf{x}^{(i)} = [x_1^{(i)}, \dots, x_C^{(i)}] = f_{\mathbf{w}}(\mathcal{I}^{(i)}) \in \mathbb{R}^C$  is the output (before sigmoid) of the deep ConvNet  $f_{\mathbf{w}}$  on image  $\mathcal{I}^{(i)}$ .

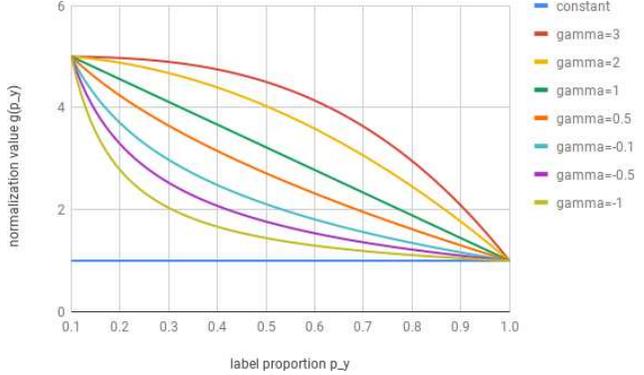


Figure 2. Examples of the weight function  $g$  (Equation 2) for different values of hyperparameter  $\gamma$  with the constraint  $g(0.1) = 5$ .  $\gamma$  controls the behavior of the normalization with respect to the label proportion  $p_{\mathbf{y}}$ .

#### 3.1. Binary cross-entropy for partial labels

The most popular loss function to train a model for multi-label classification is binary cross-entropy (BCE). To be independent of the number of categories, the BCE loss is normalized by the number of classes. This becomes a drawback for partially labeled data because the back-propagated gradient becomes small. To overcome this problem, we propose the partial-BCE loss that normalizes the loss by the proportion of known labels:

$$\ell(\mathbf{x}, \mathbf{y}) = \frac{g(p_{\mathbf{y}})}{C} \sum_{c=1}^C \left[ \mathbb{1}_{[y_c=1]} \log \left( \frac{1}{1 + \exp(-x_c)} \right) + \mathbb{1}_{[y_c=-1]} \log \left( \frac{\exp(-x_c)}{1 + \exp(-x_c)} \right) \right] \quad (1)$$

where  $p_{\mathbf{y}} \in [0, 1]$  is the proportion of known labels in  $\mathbf{y}$  and  $g$  is a normalization function with respect to the label proportion. Note that the partial-BCE loss ignores the categories for unknown labels ( $y_c = 0$ ). In the standard BCE loss, the normalization function is  $g(p_{\mathbf{y}}) = 1$ . Unlike the standard BCE, the partial-BCE gives the same importance to each example independent of the number of known labels, which is useful when the proportion of labels per image is not fixed. This loss adapts itself to the proportion of known labels. We now explain how we design the normalization function  $g$ .

**Normalization function  $g$ .** The function  $g$  normalizes the loss function with respect to the label proportion. We want the partial-BCE loss to have the same behavior as the BCE loss when all the labels are present *i.e.*  $g(1) = 1$ . We propose to use the following normalization function:

$$g(p_{\mathbf{y}}) = \alpha p_{\mathbf{y}}^{\gamma} + \beta \quad (2)$$

where  $\alpha$ ,  $\beta$  and  $\gamma$  are the hyperparameters that allow to generalize several standard functions. For instance with  $\alpha = 1$ ,

$\beta = 0$  and  $\gamma = -1$ , this function weights each example inversely proportional to the proportion of labels. This is equivalent to normalizing by the number of known classes instead of the number of classes. Given a  $\gamma$  value and the weight for a given proportion (e.g.  $g(0.1) = 5$ ), we can find the hyperparameters  $\alpha$  and  $\beta$  that satisfy these constraints. The hyperparameter  $\gamma$  controls the behavior of the normalization with respect to the label proportion. In Figure 2 we show this function for different values of  $\gamma$  given the constraint  $g(0.1) = 5$ . For  $\gamma = 1$  the normalization is linearly proportional to the label proportion, whereas for  $\gamma = -1$  the normalization value is inversely proportional to the label proportion. We analyse the importance of each hyperparameter in Sec.4. This normalization has a similar goal to batch normalization [22] which normalizes distributions of layer inputs for each mini-batch.

### 3.2. Multi-label classification with GNN

To model the interactions between the categories, we use a Graph Neural Network (GNN) [19, 45] on top of a ConvNet. We first introduce the GNN and then detail how we use GNN for multi-label classification.

**GNN.** For GNNs, the input data is a graph  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$  where  $\mathcal{V}$  (resp.  $\mathcal{E}$ ) is the set of nodes (resp. edges) of the graph. For each node  $v \in \mathcal{V}$ , we denote the input feature vector  $\mathbf{x}_v$  and its hidden representation describing the node’s state at time step  $t$  by  $\mathbf{h}_v^t$ . We use  $\Omega_v$  to denote the set of neighboring nodes of  $v$ . A node uses information from its neighbors to update its hidden state. The update is decomposed into two steps: message update and hidden state update. The message update step combines messages sent to node  $v$  into a single message vector  $\mathbf{m}_v^t$  according to:

$$\mathbf{m}_v^t = \mathcal{M}(\{\mathbf{h}_u^t | u \in \Omega_v\}) \quad (3)$$

where  $\mathcal{M}$  is the function to update the message. In the hidden state update step, the hidden states  $\mathbf{h}_v^t$  at each node in the graph are updated based on messages  $\mathbf{m}_v^t$  according to:

$$\mathbf{h}_v^{t+1} = \mathcal{F}(\mathbf{h}_v^t, \mathbf{m}_v^t) \quad (4)$$

where  $\mathcal{F}$  is the function to update the hidden state.  $\mathcal{M}$  and  $\mathcal{F}$  are feedforward neural networks that are shared among different time steps. Note that these update functions specify a propagation model of information inside the graph.

**GNN for multi-label classification.** For multi-label classification, each node represents one category ( $\mathcal{V} = \{1, \dots, C\}$ ) and the edges represent the connections between the categories. We use a fully-connected graph to model correlation between all categories. The node hidden states are initialized with the ConvNet output. We now detail the GNN functions used in our model. The algorithm

and additional information are given in the supplementary material.

**Message update function  $\mathcal{M}$ .** We use the following message update function:

$$\mathbf{m}_v^t = \frac{1}{|\Omega_v|} \sum_{u \in \Omega_v} f_{\mathcal{M}}(\mathbf{h}_u^t) \quad (5)$$

where  $f_{\mathcal{M}}$  is a multi-layer perceptron (MLP). The message is computed by first feeding hidden states to the MLP  $f_{\mathcal{M}}$  and then taking the average over the neighborhood.

**Hidden state update function  $\mathcal{F}$ .** We use the following hidden state update function:

$$\mathbf{h}_v^{t+1} = GRU(\mathbf{h}_v^t, \mathbf{m}_v^t) \quad (6)$$

which uses a Gated Recurrent Unit (GRU) [9]. The hidden state is updated based on the incoming messages and the previous hidden state.

### 3.3. Prediction of unknown labels

In this section, we propose a method to predict some missing labels with a curriculum learning strategy [2]. We formulate our problem based on the self-paced model [31, 24] and the goal is to optimize the following objective function:

$$\min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{v} \in \{0,1\}^{N \times C}} J(\mathbf{w}, \mathbf{v}) = \beta \|\mathbf{w}\|^2 + G(\mathbf{v}; \theta) \quad (7)$$

$$+ \frac{1}{N} \sum_{i=1}^N \frac{1}{C} \sum_{c=1}^C v_{ic} \ell_c(f_{\mathbf{w}}(\mathcal{I}^{(i)}), y_c^{(i)})$$

where  $\ell_c$  is the loss for category  $c$  and  $v_i \in \{0,1\}^C$  is a vector to represent the selected labels for the  $i$ -th sample.  $v_{ic} = 1$  (resp.  $v_{ic} = 0$ ) means that the  $c$ -th label of the  $i$ -th example is selected (resp. unselected). The function  $G$  defines a curriculum, parameterized by  $\theta$ , which defines the learning scheme. Following [31], we use an alternating algorithm where  $\mathbf{w}$  and  $\mathbf{v}$  are alternatively minimized, one at a time while the other is held fixed. The algorithm is shown in Algorithm 1. Initially, the model is learned with only clean partial labels. Then, the algorithm uses the learned model to add progressively new “easy” weak (*i.e.* noisy) labels in the training set, and then uses the clean and weak labels to continue the training of the model. We analyze different strategies to add new labels:

**[a] Score threshold strategy.** This strategy uses the classification score (*i.e.* ConvNet) to estimate the difficulty of a pair category-example. An easy example has a high absolute score whereas a hard example has a score close to 0. We use the learned model on partial labels to predict the

missing labels only if the classification score is larger than a threshold  $\theta > 0$ . When  $\mathbf{w}$  is fixed, the optimal  $\mathbf{v}$  can be derived by:

$$v_{ic} = \mathbb{1}[x_c^{(i)} \geq \theta] + \mathbb{1}[x_c^{(i)} < -\theta] \quad (8)$$

The predicted label is  $y_c^{(i)} = \text{sign}(x_c^{(i)})$ .

**[b] Score proportion strategy.** This strategy is similar to the strategy [a] but instead of labeling the pair category-example higher than a threshold, we label a fixed proportion  $\theta$  of pairs per mini-batch. To find the optimal  $\mathbf{v}$ , we sort the examples by decreasing order of absolute score and label only the top- $\theta\%$  of the missing labels.

**[c] Predict only positive labels.** Because of the imbalanced annotations, we only predict positive labels with strategy [a]. When  $\mathbf{w}$  is fixed, the optimal  $\mathbf{v}$  can be derived by:

$$v_{ic} = \mathbb{1}[x_c^{(i)} \geq \theta] \quad (9)$$

**[d] Ensemble score threshold strategy.** This strategy is similar to the strategy [a] but it uses an ensemble of models to estimate the confidence score. We average the classification score of each model to estimate the final confidence score. This strategy allows to be more robust than the strategy [a]. When  $\mathbf{w}$  is fixed, the optimal  $\mathbf{v}$  can be derived by:

$$v_{ic} = \mathbb{1}[E(\mathcal{I}^{(i)})_c \geq \theta] + \mathbb{1}[E(\mathcal{I}^{(i)})_c < -\theta] \quad (10)$$

where  $E(\mathcal{I}^{(i)}) \in \mathbb{R}^C$  is the vector score of an ensemble of models. The predicted label is  $y_c^{(i)} = \text{sign}(E(\mathcal{I}^{(i)})_c)$ .

**[e] Bayesian uncertainty strategy.** Instead of using the classification score as in [a] or [d], we estimate the bayesian uncertainty [28] of each pair category-example. An easy pair category-example has a small uncertainty. When  $\mathbf{w}$  is fixed, the optimal  $\mathbf{v}$  can be derived by:

$$v_{ic} = \mathbb{1}[U(\mathcal{I}^{(i)})_c \leq \theta] \quad (11)$$

where  $U(\mathcal{I}^{(i)})$  is the bayesian uncertainty of category  $c$  of the  $i$ -th example. This strategy is similar to strategy [d] except that it uses the variance of the classification scores instead of the average to estimate the difficulty.

## 4. Experiments

**Datasets.** We perform experiments on several standard multi-label datasets: Pascal VOC 2007 [17], MS COCO [35] and NUS-WIDE [11]. For each dataset, we use the standard train/test sets introduced respectively in [17], [40], and [18] (see subsection A.2 of supplementary for more details). From these datasets that are fully labeled, we create partially labeled datasets by randomly dropping some labels per image. The proportion of known labels is between 10% (90% of labels missing) and 100% (all labels present). We also perform experiments on the large-scale Open Images dataset [32] that is partially annotated: 0.9% of the labels are available during training.

---

### Algorithm 1 Curriculum labeling

---

**Input:** Training data  $\mathcal{D}$

- 1: Initialize  $\mathbf{v}$  with known labels
  - 2: Initialize  $\mathbf{w}$ : learn the ConvNet with the partial labels
  - 3: **repeat**
  - 4:   Update  $\mathbf{v}$  (fixed  $\mathbf{w}$ ): find easy missing labels
  - 5:   Update  $\mathbf{y}$ : predict the label of easy missing labels
  - 6:   Update  $\mathbf{w}$  (fixed  $\mathbf{v}$ ): improve classification model with the clean and easy weak annotations
  - 7: **until** stopping criteria
- 

**Metrics.** To evaluate the performances, we use several metrics: mean Average Precision (MAP) [1], 0-1 exact match, Macro-F1 [60], Micro-F1 [51], per-class precision, per-class recall, overall precision, overall recall. These metrics are standard multi-label classification metrics and are presented in subsection A.3 of supplementary. We mainly show the results for the MAP metric but results for other metrics are shown in supplementary.

**Implementation details.** We employ ResNet-WELDON [16] as our classification network. We use a ResNet-101 [21] pretrained on ImageNet as the backbone architecture, but we show results for other architectures in supplementary. The models are implemented with PyTorch [42]. The hyperparameters of the partial-BCE loss function are  $\alpha = -4.45$ ,  $\beta = 5.45$  (i.e.  $g(0.1) = 5$ ) and  $\gamma = 1$ . To predict missing labels, we use the bayesian uncertainty strategy with  $\theta = 0.3$ .

### 4.1. What is the best strategy to annotate a dataset?

In the first set of experiments, we study three strategies to annotate a multi-label dataset. The goal is to answer the question: what is the best strategy to annotate a dataset with a fixed budget of clean labels? We explore the three following scenarios:

- **Partial labels.** This is the strategy used in this paper. In this setting, all the images are used but only a subset of the labels per image are known. The known categories are different for each image.
- **Complete image labels or dense labels.** In this scenario, only a subset of the images are labeled, but the labeled images have the annotations for all the categories. This is the standard setting for semi-supervised learning [6] except that we do not use a semi-supervised model.
- **Noisy labels.** All the categories of all images are labeled but some labels are wrong. This scenario is similar to the webly-supervised learning scenario [37] where some labels are wrong.

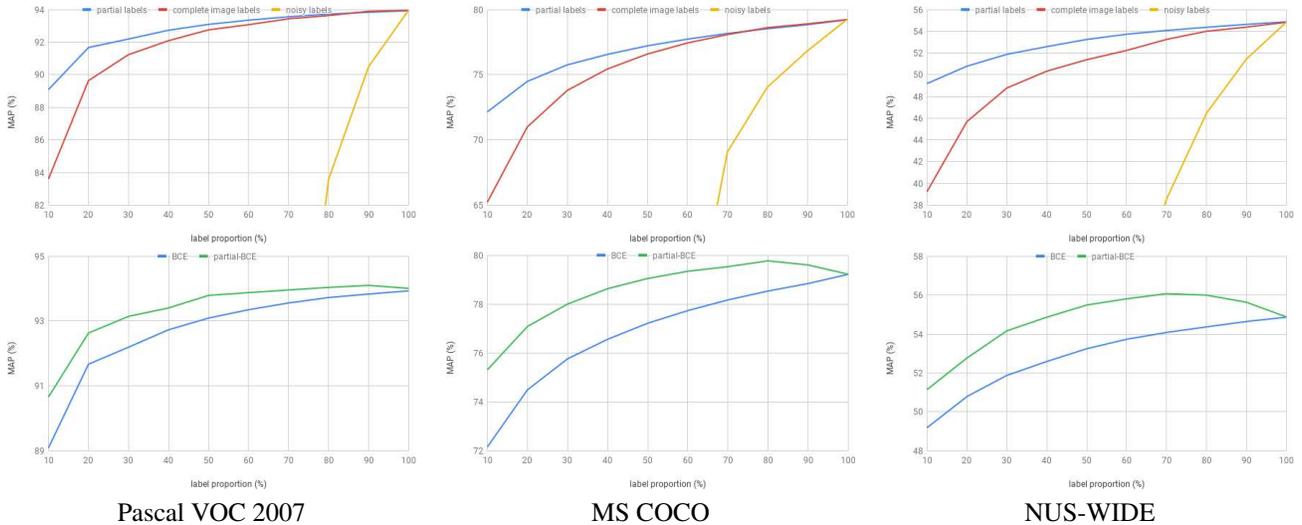


Figure 3. The first row shows MAP results for the different labeling strategies. On the second row, we shows the comparison of the BCE and the partial-BCE. The x-axis shows the proportion of clean labels.

To have fair comparison between the approaches, we use a BCE loss function for these experiments. The results are shown in Figure 3 for different proportion of clean labels. For each experiment, we use the same number of clean labels. 100% means that all the labels are known during training (standard classification setting) and 10% means that only 10% of the labels are known during training. The 90% of other labels are unknown labels for the partial labels and the complete image labels scenarios and are wrong labels for the noisy labels scenario. Similar to [48], we observe that the performance increases logarithmically based on proportion of labels. From this first experiment, we can draw the following conclusions: (1) Given a fixed number of clean labels, we observe that learning with partial labels is better than learning with a subset of dense annotations. The improvement increases when the label proportion decreases. A reason is that the model trained in the partial labels strategy “sees” more images during training and therefore has a better generalization performance. (2) It is better to learn with a small subset of clean labels than a lot of labels with some incorrect labels. Both partial labels and complete image labels scenarios are better than the noisy label scenario. For instance on MS COCO, we observe that learning with only 20% of clean partial labels is better than learning with 80% of clean labels and 20% of wrong labels.

**Noisy web labels.** Another strategy to generate a noisy dataset from a multi-label dataset is to use only one positive label for each image. This is a standard assumption made when collecting data from the web [34] *i.e.* the only category present in the image is the category of the query. From the clean MS COCO dataset, we generate a noisy dataset (named noisy+) by keeping only one positive la-

model	clean	partial 10%	noisy+
clean / noisy labels	100 / 0	10 / 0	97.6 / 2.4
MAP (%)	79.22	72.15	71.60

Table 1. Comparison with a webly-supervised strategy (noisy+) on MS COCO. Clean (resp. noisy) means the percentage of clean (resp. noisy) labels in the training set.

bel per image. If the image has more than one positive label, we randomly select one positive label among the positive labels and switch the other positive labels to negative labels. The results are reported in Table 1 for three scenarios: clean (all the training labels are known and clean), 10% of partial labels and noisy+ scenario. We also show the percentage of clean and noisy labels for each experiment. The noisy+ approach generates a small proportion of noisy labels (2.4%) that drops the performance by about 7pt with respect to the clean baseline. We observe that a model trained with only 10% of clean labels is slightly better than the model trained with the noisy labels. This experiment shows that the standard assumption made in most of the webly-supervised datasets is not good for complex scenes / multi-label images because it generates noisy labels that significantly decrease generalization.

## 4.2. Learning with partial labels

In this section, we compare the standard BCE and the partial-BCE and analyze the importance of the GNN.

**BCE vs partial-BCE.** The Figure 3 shows the MAP results for different proportion of known labels on three datasets. For all the datasets, we observe that using the partial-BCE significantly improves the performance: the

Relabeling	MAP	0-1	Macro-F1	Micro-F1	label prop.	TP	TN	GNN
2 steps (no curriculum)	-1.49	6.42	2.32	1.99	100	82.78	96.40	✓
[a] Score threshold $\theta = 2$	0.34	11.15	4.33	4.26	95.29	85.00	98.50	✓
[b] Score proportion $\theta = 80\%$	0.17	8.40	3.70	3.25	96.24	84.40	98.10	✓
[c] Postitive only - score $\theta = 5$	0.31	-4.58	-1.92	-2.23	12.01	79.07	-	✓
[d] Ensemble score $\theta = 2$	0.23	11.31	4.16	4.33	95.33	84.80	98.53	✓
[e] Bayesian uncertainty $\theta = 0.3$	0.34	10.15	4.37	3.72	77.91	61.15	99.24	
[e] Bayesian uncertainty $\theta = 0.1$	0.36	2.71	1.91	1.22	19.45	38.15	99.97	✓
[e] Bayesian uncertainty $\theta = 0.2$	0.30	10.76	4.87	4.66	57.03	62.03	99.65	✓
[e] Bayesian uncertainty $\theta = 0.3$	0.59	12.07	5.11	4.95	79.74	68.96	99.23	✓
[e] Bayesian uncertainty $\theta = 0.4$	0.43	10.99	4.88	4.46	90.51	70.77	98.57	✓
[e] Bayesian uncertainty $\theta = 0.5$	0.45	10.08	3.93	3.78	94.79	74.73	98.00	✓

Table 2. Analysis of the labeling strategy of missing labels on Pascal VOC 2007 val set. For each metric, we report the relative scores with respect to a model that does not label missing labels. TP (resp. TN) means true positive (resp. true negative) rate. For the strategy [c], we report the label accuracy instead of the TP rate.

	BCE	partial-BCE	GNN + partial-BCE
MAP (%)	79.01	83.05	83.36

Table 3. MAP results on Open Images.

lower the label proportion, the better the improvement. We observe the same behavior for the other metrics (subsection A.6 of supplementary). In Table 3, we show results on the Open Images dataset and we observe that the partial-BCE is 4 pt better than the standard BCE. These experiments show that our loss learns better than the BCE because it exploits the label proportion information during training. It allows to learn efficiently while keeping the same training setting as with all annotations.

**GNN.** We now analyze the improvements of the GNN to learn relationships between the categories. We show the results on MS COCO in Figure 4. We observe that for each label proportion, using the GNN improves the performance. Open Images experiments (Table 3) show that GNN improves the performance even when the label proportion is small. This experiment shows that modeling the correlation between categories is important even in case of partial labels. However, we also note that a ConvNet implicitly learns some correlation between the categories because some learned representations are shared by all categories.

### 4.3. What is the best strategy to predict missing labels?

In this section, we analyze the labeling strategies introduced in subsection 3.3 to predict missing labels. Before training epochs 10 and 15, we use the learned classifier to predict some missing labels. We report the results for different metrics on Pascal VOC 2007 validation set with 10% of labels in Table 2. We also report the final proportion of

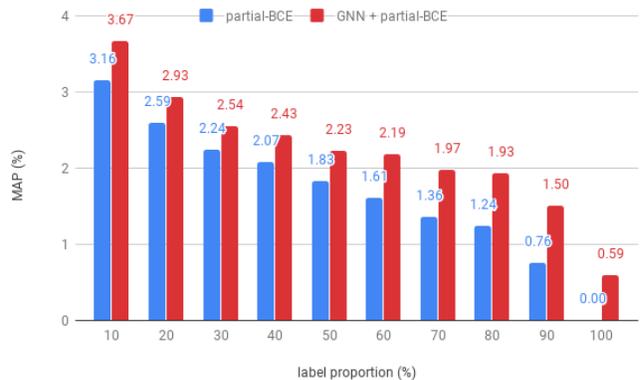


Figure 4. MAP (%) improvement with respect to the proportion of known labels on MS COCO for the partial-BCE and the GNN + partial-BCE. 0 means the result for a model trained with the standard BCE.

labels, the true positive (TP) and true negative (TN) rates for predicted labels. Additional results are shown in subsection A.9 of supplementary.

First, we show the results of a 2 steps strategy that predicts all missing labels in one time. Overall, we observe that this strategy is worse than curriculum-based strategies ([a-e]). In particular, the 2 steps strategy decreases the MAP score. These results show that predicting all missing labels at once introduced too much label noise, decreasing generalization performance. Among the curriculum-based strategies, we observe that the threshold strategy [a] is better than the proportion strategy [b]. We also note that using a model ensemble [d] does not significantly improve the performance with respect to a single model [a]. Predicting only positive labels [c] is a poor strategy. The bayesian uncertainty strategy [e] is the best strategy. In particular, we observe that the GNN is important for this strategy because it decreases the label uncertainty and allows the model to be

BCE	fine-tuning	partial-BCE	GNN	relabeling	MAP	0-1 exact match	Macro-F1	Micro-F1
✓					66.21	17.53	62.74	67.33
✓	✓				72.15	22.04	65.82	70.09
	✓	✓			75.31	24.51	67.94	71.18
	✓	✓	✓		75.82	25.14	68.40	71.37
	✓	✓		✓	75.71	30.52	70.13	73.87
	✓	✓	✓	✓	76.40	32.12	70.73	74.37

Table 4. Ablation study on MS COCO with 10% of known labels.

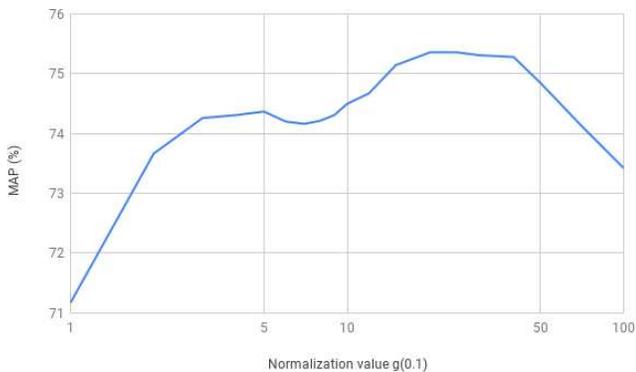


Figure 5. Analysis of the normalization value for a label proportion of 10% (i.e.  $g(0.1)$ ). (x-axis log-scale)

robust to the hyperparameter  $\theta$ .

#### 4.4. Method analysis

In this section, we analyze the hyperparameters of the partial-BCE and perform an ablation study on MS COCO.

**Partial-BCE analysis.** To analyze the partial-BCE, we use only the training set. The model is trained on about 78k images and evaluated on the remaining 5k images. We first analyse how to choose the value of the normalization function given a label proportion of 10% i.e.  $g(0.1)$  (it is possible to choose another label proportion). The results are shown in Figure 5. Note that for  $g(0.1) = 1$ , the partial-BCE is equivalent to the BCE and the loss is normalized by the number of categories. We observe that the normalization value  $g(0.1) = 1$  gives the worst results. The best score is obtained for a normalization value around 20 but the performance is similar for  $g(0.1) \in [3, 50]$ . Using a large value drops the performance. This experiment shows that the proposed normalization function is important and robust. These results are independent of the network architectures (subsection A.7 of supplementary).

Given the constraints  $g(0.1) = 5$  and  $g(1) = 1$ , we analyze the impact of the hyperparameter  $\gamma$ . This hyperparameter controls the behavior of the normalization with respect to the label proportion. Using a high value ( $\gamma = 3$ ) is better than a low value ( $\gamma = -1$ ) for large label proportions but is

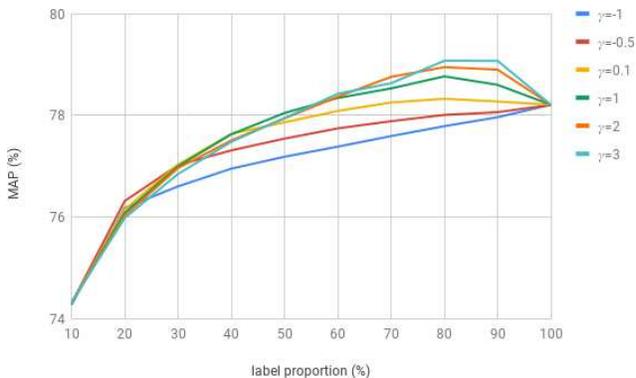


Figure 6. Analysis of hyperparameter  $\gamma$  on MS COCO.

slightly worse for small label proportions. We observe that using a normalization that is proportional to the number of known labels ( $\gamma = 1$ ) works better than using a normalization that is inversely proportional to the number of known labels ( $\gamma = -1$ ).

**Ablation study.** Finally to analyze the importance of each contribution, we perform an ablation study on MS COCO for a label proportion of 10% in Table 4. We first observe that fine-tuning is important. It validates the importance of building end-to-end trainable models to learn with missing labels. The partial-BCE loss function increases the performance against each metric because it exploits the label proportion information during training. We show that using GNN or relabeling improves performance. In particular, the relabeling stage significantly increases the 0-1 exact match score (+5pt) and the Micro-F1 score (+2.5pt). Finally, we observe that our contributions are complementary.

## 5. Conclusion

In this paper, we present a scalable approach to end-to-end learn a multi-label classifier with partial labels. Our experiments show that our loss function significantly improves performance. We show that our curriculum learning model using bayesian uncertainty is an accurate strategy to label missing labels. In the future work, one could combine several datasets with shared categories to learn with more training data.

## References

- [1] Ricardo A. Baeza-Yates and Berthier Ribeiro-Neto. *Modern Information Retrieval*. 1999. 5
- [2] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *International Conference on Machine Learning (ICML)*, 2009. 2, 4
- [3] S. S. Bucak, R. Jin, and A. K. Jain. Multi-label learning with incomplete class assignments. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011. 2
- [4] Ricardo S. Cabral, Fernando Torre, Joao P. Costeira, and Alexandre Bernardino. Matrix Completion for Multi-label Image Classification. In *Advances in Neural Information Processing Systems (NIPS)*, 2011. 2
- [5] Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr., and Tom M. Mitchell. Toward an Architecture for Never-Ending Language Learning. In *Conference on Artificial Intelligence (AAAI)*, 2010. 3
- [6] Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. *Semi-Supervised Learning*. 2010. 2, 5
- [7] X. Chen, A. Shrivastava, and A. Gupta. Neil: Extracting visual knowledge from web data. In *IEEE International Conference on Computer Vision (ICCV)*, 2013. 3
- [8] Xinlei Chen, Abhinav Shrivastava, and Abhinav Gupta. Enriching visual knowledge bases via object discovery and segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 3
- [9] Kyunghyun Cho, B van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. In *Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation (SSST-8)*, 2014. 4
- [10] Hong-Min Chu, Chih-Kuan Yeh, and Yu-Chiang Frank Wang. Deep Generative Models for Weakly-Supervised Multi-Label Classification. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [11] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng. NUS-WIDE: A Real-world Web Image Database from National University of Singapore. In *ACM International Conference on Image and Video Retrieval (CIVR)*, 2009. 5
- [12] Timothee Cour, Ben Sapp, and Ben Taskar. Learning from Partial Labels. *Journal of Machine Learning Research (JMLR)*, 2011. 2
- [13] Jia Deng, Olga Russakovsky, Jonathan Krause, Michael S. Bernstein, Alex Berg, and Li Fei-Fei. Scalable Multi-label Annotation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2014. 1, 2
- [14] Thibaut Durand, Taylor Mordan, Nicolas Thome, and Matthieu Cord. WILDCAT: Weakly Supervised Learning of Deep ConvNets for Image Classification, Pointwise Localization and Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [15] Thibaut Durand, Nicolas Thome, and Matthieu Cord. WELDON: Weakly Supervised Learning of Deep Convolutional Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1
- [16] Thibaut Durand, Nicolas Thome, and Matthieu Cord. Exploiting Negative Evidence for Deep Latent Structured Models. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2018. 1, 5
- [17] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The Pascal Visual Object Classes Challenge: A Retrospective. *International Journal of Computer Vision (IJCV)*, 2015. 5
- [18] Yunchao Gong, Yangqing Jia, Thomas Leung, Alexander Toshev, and Sergey Ioffe. Deep Convolutional Ranking for Multilabel Image Annotation. In *International Conference on Learning Representations (ICLR)*, 2014. 5
- [19] Marco Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. In *IEEE International Joint Conference on Neural Networks*, 2005. 4
- [20] Sheng Guo, Weilin Huang, Haozhi Zhang, Chenfan Zhuang, Dengke Dong, Matthew R. Scott, and Dinglong Huang. CurriculumNet: Weakly Supervised Learning from Large-Scale Web Images. In *European Conference on Computer Vision (ECCV)*, 2018. 3
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 5
- [22] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *International Conference on Machine Learning (ICML)*, 2015. 4
- [23] Finn V. Jensen and Thomas D. Nielsen. *Bayesian Networks and Decision Graphs*. 2007. 2
- [24] Lu Jiang, Deyu Meng, Qian Zhao, Shiguang Shan, and Alexander G Hauptmann. Self-Paced Curriculum Learning. In *Conference on Artificial Intelligence (AAAI)*, 2015. 3, 4
- [25] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. MentorNet: Learning Data-Driven Curriculum for Very Deep Neural Networks on Corrupted Labels. In *International Conference on Machine Learning (ICML)*, 2018. 3
- [26] Armand Joulin, Laurens van der Maaten, Allan Jabri, and Nicolas Vasilache. Learning visual features from large weakly supervised data. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [27] Ashish Kapoor, Raajay Viswanathan, and Prateek Jain. Multilabel classification using bayesian compressed sensing. In *Advances in Neural Information Processing Systems (NIPS)*, 2012. 2
- [28] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems (NIPS)*, 2017. 5
- [29] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*, 2014. 2
- [30] Simon Kornblith, Jonathon Shlens, and Quoc V. Le. Do Better ImageNet Models Transfer Better? 2018. 2
- [31] M. P. Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *Advances in Neural Information Processing Systems (NIPS)*, 2010. 2, 4

- [32] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Ujlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, and Vittorio Ferrari. The Open Images Dataset V4: Unified image classification, object detection, and visual relationship detection at scale. 2018. [1](#), [5](#)
- [33] L. J. Li, G. Wang, and Li Fei-Fei. OPTIMOL: automatic Online Picture collecTion via Incremental MOdel Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007. [3](#)
- [34] Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, Jesse Berent, Abhinav Gupta, Rahul Sukthankar, and Luc Van Gool. WebVision Challenge: Visual Learning and Understanding With Web Data. In *arXiv 1705.05640*, 2017. [1](#), [6](#)
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollr. Microsoft COCO: Common Objects in Context. 2014. [1](#), [5](#)
- [36] Chenxi Liu, Barret Zoph, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive Neural Architecture Search. In *European Conference on Computer Vision (ECCV)*, 2018. [1](#)
- [37] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the Limits of Weakly Supervised Pretraining. In *European Conference on Computer Vision (ECCV)*, 2018. [1](#), [2](#), [5](#)
- [38] Minmin Chen and Alice Zheng and Kilian Weinberger. Fast image tagging. In *International Conference on Machine Learning (ICML)*, 2013. [2](#)
- [39] T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. Never-Ending Learning. In *Conference on Artificial Intelligence (AAAI)*, 2015. [3](#)
- [40] Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic. Learning and Transferring Mid-Level Image Representations using Convolutional Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. [5](#)
- [41] Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. Is Object Localization for Free? - Weakly-Supervised Learning With Convolutional Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. [1](#)
- [42] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. [5](#)
- [43] Hieu Pham, Melody Y. Guan, Barret Zoph, Quoc V. Le, and Jeff Dean. Faster Discovery of Neural Architectures by Searching for Paths in a Large Model. In *International Conference on Learning Representations (ICLR)*, 2018. [1](#)
- [44] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 2015. [1](#)
- [45] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 2009. [4](#)
- [46] Pierre Stock and Moustapha Cisse. ConvNets and ImageNet Beyond Accuracy: Understanding Mistakes and Uncovering Biases. In *European Conference on Computer Vision (ECCV)*, 2018. [1](#)
- [47] Chen Sun, Manohar Paluri, Ronan Collobert, Ram Nevatia, and Lubomir Bourdev. ProNet: Learning to Propose Object-Specific Boxes for Cascaded Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [1](#)
- [48] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting Unreasonable Effectiveness of Data in Deep Learning Era. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. [1](#), [2](#), [6](#)
- [49] Yu-Yin Sun, Yin Zhang, and Zhi-Hua Zhou. Multi-label Learning with Weak Label. In *Conference on Artificial Intelligence (AAAI)*, 2010. [2](#)
- [50] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alex Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Conference on Artificial Intelligence (AAAI)*, 2017. [1](#)
- [51] Lei Tang, Suju Rajan, and Vijay K. Narayanan. Large scale multi-label classification via metalabeler. In *WWW*, 2009. [5](#)
- [52] Grigorios Tsoumakas and Ioannis Katakis. Multi-label classification: An overview. *International Journal of Data Warehousing and Mining (IJDWM)*, 2007. [2](#)
- [53] Arash Vahdat. Toward robustness against label noise in training deep discriminative neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. [1](#)
- [54] Deepak Vasisht, Andreas Damianou, Manik Varma, and Ashish Kapoor. Active Learning for Sparse Bayesian Multi-label Classification. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2014. [2](#)
- [55] Qifan Wang, Bin Shen, Shumiao Wang, Liang Li, and Luo Si. Binary Codes Embedding for Fast Image Tagging with Incomplete Labels. In *European Conference on Computer Vision (ECCV)*, 2014. [2](#)
- [56] Baoyuan Wu, Siwei Lyu, and Bernard Ghanem. ML-MG: Multi-Label Learning With Missing Labels Using a Mixed Graph. In *IEEE International Conference on Computer Vision (ICCV)*, 2015. [2](#)
- [57] Saining Xie, Ross Girshick, Piotr Dollr, Zhuowen Tu, and Kaiming He. Aggregated Residual Transformations for Deep Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [1](#)
- [58] Miao Xu, Rong Jin, and Zhi-Hua Zhou. Speedup Matrix Completion with Side Information: Application to Multi-Label Learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2013. [2](#)

- [59] Hao Yang, Joey Tianyi Zhou, and Jianfei Cai. Improving Multi-label Learning with Missing Labels by Structured Semantic Correlations. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [60] Yiming Yang. An evaluation of statistical approaches to text categorization. 1999. 5
- [61] Hsiang-Fu Yu, Prateek Jain, Purushottam Kar, and Inderjit S. Dhillon. Large-scale Multi-label Learning with Missing Labels. In *International Conference on Machine Learning (ICML)*, 2014. 2
- [62] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations (ICLR)*, 2017. 1, 2
- [63] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning Deep Features for Discriminative Localization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1
- [64] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1