Pattern Recognition Letters 69 (2016) 49–55

Contents lists available at ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

Weak supervision and other non-standard classification problems: A taxonomy☆

Jerónimo Hernández-González*, Iñaki Inza, Jose A. Lozano

Intelligent Systems Group, University of the Basque Country UPV/EHU, P. Manuel Lardizabal 1, 20018 Donostia, Spain

ARTICLE INFO

Article history:
Received 10 May 2015
Available online 24 October 2015

Keywords:
Weakly supervised classification
Partially supervised classification
Degrees of supervision

ABSTRACT

In recent years, different researchers in the machine learning community have presented new classification frameworks which go beyond the standard supervised classification in different aspects. Specifically, a wide spectrum of novel frameworks that use partially labeled data in the construction of classifiers has been studied. With the objective of drawing up a description of the state-of-the-art, three identifying characteristics of these novel frameworks have been considered: (1) the relationship between instances and labels of a problem, which may be beyond the *one-instance one-label* standard, (2) the possible provision of partial class information for the training examples, and (3) the possible provision of partial class information also for the examples in the prediction stage. These three ideas have been formulated as axes of a comprehensive taxonomy that organizes the state-of-the-art. The proposed organization allows us both to understand similarities/differences among the different classification problems already presented in the literature as well as to discover unexplored frameworks that might be seen as further challenges and research opportunities. A representative set of state-of-the-art problems has been used to illustrate the novel taxonomy and support the discussion.

1. Introduction

Supervised classification [24] is one of the most popular fields of machine learning. Its objective is to learn a classifier that reliably approximates a classification task which is inferred from a set of categorized examples of a problem of interest. The learnt classifier is posteriorly used in the prediction stage to anticipate the class label of new *unlabeled* examples. In this context, the term “supervised” indicates that, in the learning stage, the examples are *always* provided with their real class label (category).

On different real problems, obtaining a fully supervised dataset for training, as required in the standard framework, is costly, difficult or even impossible. Solutions proposed for learning from different kinds of partially labeled data have led to the foundation of a new subfield of machine learning called *weakly supervised classification* (a.k.a. partially supervised learning). Weak supervision refers to the lack of a full supervision for the provided data and, from this point of view, the popular semi-supervised learning [5] can be considered as the first and basic framework of the field. In the last decade, the field has experimented a rapid growth due to the determination of the machine learning community to solve new challenging classification problems with different constraints in the access to the class information. Different factors have been reported as being responsible for this restriction: the impossibility of observing examples individually [15], unaffordable or non-exhaustive labeling process [7,22], categorizations provided by labelers of arguable reliability [30], etc. Similarly, although the examples for prediction are traditionally provided to the classifier completely unlabeled, there exist situations where partial class information is available during the prediction stage [7,20]. For these problems, learning techniques have been proposed which, taking into account the available partial class information at prediction time, build classifiers that take advantage of that information, efficiently enhancing their performance. It is interesting to note that the kind of partial class information available in the prediction stage has to be known before the classifier is built in order to be able to exploit it.

In this paper, a taxonomy of weakly supervised classification problems is proposed. Apart from (a) the type of supervision in the data provided for learning and (b) the type of supervision in the data provided for prediction, the taxonomy considers another axis: (c) the instance-label relationship defined by the problem (e.g., in the multiple-instance learning framework [9] a group of instances is globally categorized by a single label). Although it is not an exclusive characteristic of weakly supervised classification problems, the inclusion in the taxonomy of the latter axis aims to avoid confusing

☆ This paper has been recommended for acceptance by Dr. D. Dembele

* Corresponding author. Tel.: +34 943 018 070; fax +34 943 015 590.
E-mail addresses: jeronimo.hernandez@ehu.eus (J. Hernández-González), inaki.inza@ehu.eus (I. Inza), ja.lozano@ehu.eus (J.A. Lozano).

http://dx.doi.org/10.1016/j.patrec.2015.10.008
0167-8655/

non-standard instance-label relationships with certain types of weak supervision which are observed in the provided data. This is the case of the first paper where a description of the field was presented [11]. In their enumeration of weakly supervised classification problems, they include the multiple-instance learning problem [9], where fully supervised training data is provided in a framework characterized by a non-standard instance-label relationship. Moreover, the possible availability of class information in the prediction stage is not considered in their collection. Our proposal tries to provide a complete organization of the currently large —and in part confusing— state-of-the-art, shedding light on the discussion about similarities and differences between frameworks. Its use should prevent a recurrence of parallel research, as is the case for instance of Musicant et al. [25] and Quadrianto et al. [28], who independently presented their proposals for the same weakly supervised problem (using different names) without being aware of each other. Additionally, the taxonomy reveals the existence of unexplored areas that could potentially lead to new challenging frameworks. Note that this is neither an exhaustive review of weakly supervised frameworks nor a discussion about the techniques used to solve them.

In the next section, the three axes of the taxonomy are formally explained and contextualized. Then, the derived taxonomy is presented and discussed, using a representative set of state-of-the-art problems to illustrate it. Along with the discussion, some ideas for further extension of the taxonomy and open challenges in the literature are pointed out.

2. Weakly supervised classification

Formally, a supervised classification problem [24] is described by a set of n predictive variables $(X_1, \ldots, X_n)$ and a class variable C. Each predictive variable X_i can take a value from its own set of possible values $\mathcal{X}_i$ and an instance is a tuple $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathcal{X} = (\mathcal{X}_1, \ldots, \mathcal{X}_n)$, where $\mathcal{X}$ is the set of all possible instances. Specifically, the set of values that the class variable can take, a.k.a. class labels, forms the label space $\mathcal{C}$. Assuming the existence of an unknown target function $H : \mathcal{X} \to \mathcal{C}$ that (i) *individually categorizes each instance with a single label*, supervised classification techniques learn (ii) *from a set of fully labeled examples* $\{(\boldsymbol{x}^1, c^1), \ldots, (\boldsymbol{x}^N, c^N)\}$ of the problem a mapping function or classifier $\hat{H}$ that approximates the real function H. The objective is to build a classifier $\hat{H}$ that accurately predicts the class label c of (iii) *new unlabeled examples*$(\boldsymbol{x}, ?)$.

A quick look at recent literature is enough to realize that the increasing number and variety of non-standard supervised classification problems cannot be described by means of this standard definition. In the previous paragraph, three well-established components of the definition have been emphasized. At least one of the indicated components is not fulfilled by the non-standard classification frameworks collected for this work. First of all, not all the problems involve samples which can be described by means of an instance-label pair: e.g., the multi-label framework [34], where the examples are categorized with one or more class labels. Secondly, some frameworks cannot provide a fully labeled dataset for training: e.g., the semi-supervised framework [5], where not all the training examples are labeled. Thirdly, certain class information can be known for the examples at prediction time: e.g., someone could be interested in categorizing a group of examples and it is known that they belong to different categories [20]. Each of these ideas, the three axes on which the proposed taxonomy is based, will be discussed in-depth in the following subsections.

2.1. Instance-label relationship

In standard supervised classification, each instance represents an example of the problem and is categorized with a single class label (single-instance single-label, SISL). There exist other popular state-of-the-art frameworks that do not follow this standard instance-label (IL) relationship: in the multi-label classification framework [34], each instance is categorized with multiple (one or more) class labels (SIML); in the multiple-instance learning problem [9], a set of instances (which represents an example) is categorized with a single class label (MISL); and the multi-instance multi-label framework [43] involves both examples of multiple instances and categorizations of multiple labels (MIML).

Table 1
Four possible definitions of the target function H. An example is composed of a single (SI) or multiple (MI) instances. The categorization is composed of a single (SL) or multiple (ML) class labels.

		Categorization	
		SL	ML
Example	SI	$H : \mathcal{X} \to \mathcal{C}$	$H : \mathcal{X} \to 2^{\mathcal{C}}$
	MI	$H : 2^{\mathcal{X}} \to \mathcal{C}$	$H : 2^{\mathcal{X}} \to 2^{\mathcal{C}}$

Note that the instance-label relationship can be used for characterizing both weakly and standard supervised classification problems, i.e., it is not an exclusive feature of weakly supervised classification problems. In the related literature, the interested reader can find classification problems with a non-standard IL relationship which provide standard fully supervised data (e.g., all the illustrative frameworks mentioned so far in the current subsection [9,34,43]), weakly supervised classification problems with the standard IL relationship [7,15] or problems that combine an alternative IL relationship with weak supervision [33,41]. However, the inclusion of this characteristic as an axis of our taxonomy allows us to leave this feature out of the discussion over weak supervision —the IL relationship has been confused several times with weak supervision [11].

In general, it may be agreed that a classifier $\hat{H}$ is built as an approximation of the real unknown target function H. The definition of the domain and image of the target function H determines the instance-label relationship of a problem. On the one hand, the domain of H comprises all the possible examples of the problem. There are two possible configurations: each example is represented (a) by a single instance, as in the standard framework [24], where the domain of H matches the instance space $\mathcal{X}$, or (b) by multiple instances [9], where the domain of H is the power set $2^{\mathcal{X}}$. On the other hand, the image of the target function H comprises all the possible categorizations. There are also two possible configurations: a categorization is represented (a) by a single class label, as in the standard framework, with the image of H matching the label space $\mathcal{C}$, or (b) by multiple class labels [34], where the image is the power set $2^{\mathcal{C}}$. Thus, both the examples and the categorizations can show a single or multiple configuration. Globally, there are four possible definitions of the unknown target function H (Table 1), and each of them implies a different IL relationship. This leads to a first subdivision of classification problems. From this section on, *example* and *categorization* are used as two general terms that take a particular meaning according to the instance-label relationship defined by the target function of each specific problem.

2.2. Supervision in the learning stage

According to the standard definition of supervised classification, a set of fully supervised examples has to be provided in the learning stage in order to infer a classifier. Loss functions, performance evaluation, feature subset selection or discretization techniques are a few examples of the different procedures that take advantage of this requirement. However, collecting such a complete set of examples is not always possible.

Many authors have dealt with classification problems in which the class information provided for the training examples is partial.