Classifier Learning with Hidden Information

Ziheng Wang, Qiang Ji
ECSE Department, Rensselaer Polytechnic Institute, Troy, NY.

A wide variety of computer vision problems can be formulated as a classification problem. Traditional approach learns a classifier purely from a set of training instances $(x_1,y_1),\cdots,(x_n,y_n)$. Despite the substantial successes they have achieved for solving classification problems, data-driven approaches become very brittle and prone to overfitting when the training data is inadequate in either quantity or quality.

A natural solution to alleviate the limitations of data-driven approaches is incorporating additional prior information. In this paper we propose to incorporate a specific prior information, which is only available during training but not available during testing, to improve classifier learning. We denote such information as *hidden information* and the new learning problem as *learning with hidden information*.

Hidden information can be qualities, properties and context of the training instances, and can be found in a wide variety of applications. For example as shown in Figure 1, in object recognition, besides the image features and object labels, during training, the learner may also have access to object attributes which describe high-level properties of the objects in each image. In human action recognition, besides the RGB video features and human action labels, the learner may also obtain depth information and human joint positions about each human action instance. Both object attributes and joint positions could be obtained offline for training data yet are very expensive to obtain for all the testing instances.

Hidden information, also referred to as privileged information [6] or side information [1], has been exploited to enhance different learning tasks such as classifier learning [1, 4, 5, 6, 7], clustering [2], and metric learning [3]. However, research in this area remains limited. First, existing approaches are all designed to exploit certain type of hidden information for certain type of classifiers. Second, they are generally based on strong and even unrealistic assumptions. Third, existing methods typically assume hidden information is complete for each training sample. However, for many real world applications, hidden information may be incomplete. Finally, existing methods typically treat each piece of hidden information independent of each other, ignoring their relationships.

To address these limitations, we propose two approaches to capture hidden information. The proposed methods are general to capture different types of hidden information for different types of classifiers. We also extend the proposed methods to incorporate incomplete hidden information, where hidden information is complete for a fraction of training instances.

Our first method **loss inequality regularization (LIR)** treats hidden information as a secondary set of features. It is motivated by the assumption that secondary feature is more informative for classification than the primary feature. While seemingly strong, this assumption holds for many applications (e.g attributes are more discriminative than image features). In addition, it can be generally satisfied if we combine x and h together as secondary features. Mathematically, the assumption means that if we have a primary classifier $f(x,w)$ to classify y from x, and a secondary classifier $f(h,\tilde{w})$ to classify y from h, then the loss for classifying y with x should be higher than the loss for classifying y with h. It can can be encoded as a set of ε-insensitive loss inequality constraints shown below, where ε is used to account for uncertainties. The constraints are then encoded as regularization terms to influence the learning of both the primary and secondary classifiers.

$$\ell(y_i,x_i,w) \geq \ell(y_i,h_i,\tilde{w}) - \varepsilon_i, \;\; \varepsilon_i \geq 0, \;\; \forall 1 \leq i \leq n$$

Our second method **relationship preserving regularization (RPR)** treats hidden information as secondary targets. Our motivation is that we want predicted target labels by the learnt primary classifier to preserve their relationships with the secondary labels. For example, if an object is closely related to an attribute, then the predicted object to be closely related to this attribute (and vice-versa).

This is an extended abstract. The full paper is available at the Computer Vision Foundation webpage.

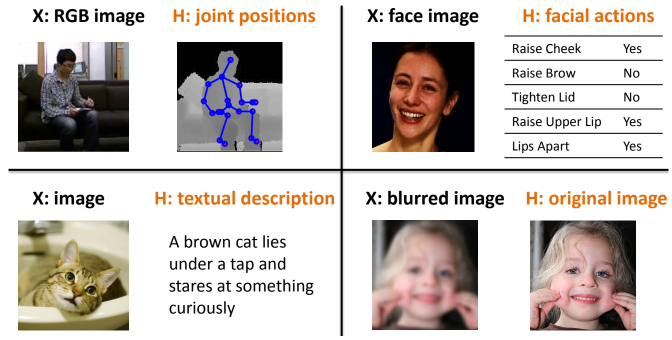

Figure 1: Examples of hidden information in different applications. X stands for primary measurement, H stands for hidden information.

Mathematically, denote $R(y,[h^1,\cdots,h^k])$ as a certain mathematical measure of the relationships between the primary label y and secondary labels $h^1,\cdots,h^k$, we propose the following relationship preserving constraint, where $\hat{y}$ is the predicted label of y. Similarly, the relationship preserving constraints are used to regularize the learning of the primary classifier.

$$|R(y,[h^1,\cdots,h^k]) - R(\hat{y},[h^1,\cdots,h^k])| < \varepsilon \tag{1}$$

Finally, the proposed methods are extended to deal with incomplete hidden information by imposing either loss inequality regularization or relationship preserving regularization only on the training instances with complete hidden information.

Compared to the existing methods, the proposed approaches are general (applicable to different types of hidden information and classifiers), without strong assumptions. We evaluate the proposed methods on two applications: facial expression recognition with facial action units as hidden information, and object recognition with human annotated attributes as hidden information. Experimental results demonstrate the effectiveness of the proposed method for exploiting hidden information, as well as its superior performance to the related methods.

Acknowledgement

The work described in this paper is supported in part by the grant IIS 1145152 from the National Science Foundation.

[1] Jixu Chen, Xiaoming Liu, and Siwei Lyu. Boosting with side information. In *ACCV*, 2012.

[2] Jan Feyereisl and Uwe Aickelin. Privileged information for data clustering. *Information Sciences*, 194:4–23, 2012.

[3] Shereen Fouad, Peter Tino, Somak Raychaudhury, and Petra Schneider. Learning using privileged information in prototype based models. In *ICANN*. 2012.

[4] Dmitry Pechyony and Vladimir Vapnik. On the Theory of Learning with Privileged Information. In *NIPS*, 2010.

[5] Viktoriia Sharmanska, Novi Quadrianto, and Christoph H Lampert. Learning to rank using privileged information. In *ICCV*, 2013.

[6] Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural Networks*, 22:544–557, July 2009. ISSN 0893-6080.

[7] Ziheng Wang, Xiaoyang Wang, and Qiang Ji. Learning with hidden information. In *ICPR*, pages 238–243. IEEE, 2014.