Knowledge-Based Systems 000 (2017) 1–9

Contents lists available at ScienceDirect

Knowledge-Based Systems

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

Combining ontology and reinforcement learning for zero-shot classification

Bin Liu, Li Yao*, Zheyuan Ding, Junyi Xu, Junfeng Wu

Science and Technology on Information Systems and Engineering Laboratory, National University of Defense Technology, Changsha 410073, China

ARTICLE INFO

Article history:
Received 9 June 2017
Revised 19 December 2017
Accepted 21 December 2017
Available online xxx

Keywords:
Image classification
Zero-shot classification
Ontology
Reinforcement learning
Adaptive

ABSTRACT

Zero-Shot Classification (ZSC) has received much attention recently in computer vision research. Traditional classifiers are unable to handle ZSC because test data labels are significantly different from training data labels. Attribute-based methods have long dominated ZSC. However, classical attribute-based methods fail to distinguish between discriminative attributes and non-discriminative attributes and do not distinguish the different contributions each attribute makes to classification. We propose CORL (Combining Ontology and Reinforcement Learning) for ZSC. CORL first obtains hierarchical classification rules from attribute annotations of object classes based on ontology. These rules contain only discriminative attributes. Reinforcement learning is used to adaptively determine the discriminative degrees of the rules. The most discriminative rules are then selected for ZSC. Experiments on three benchmark datasets showed that CORL achieved higher accuracies than baseline classifiers. This suggests that CORL effectively discovers the most discriminative rules for ZSC.

1. Introduction

Zero-Shot Classification (ZSC) is a focus of computer vision research [1–6]. Different from traditional image classification in machine learning, ZSC considers cases in which training sample labels are completely different from test sample labels. Developing zero-shot classifiers could significantly extend machine learning abilities for handling practical problems. For example, it is difficult to obtain labeled training samples for some specific classes due to the large number of object classes. Some classes even have no training samples because new classes are defined ‘on the fly’ [3]. Zero-shot classifiers fit classifications under these circumstances.

Traditional classifiers such as support vector machine [7] and Bayesian classifiers [8] cannot handle ZSC because test stage classes (seen classes) are materially different from training stage classes (unseen classes). Zero-shot learning [9] is proposed as a solution for ZSC problems. Before long, attribute-based ZSC methods are proposed to solve ZSC problems specifically [3]. Generally speaking, semantic attributes are humanly understandable [10] and shared among all object classes [11]. Attribute knowledge used in describing unseen classes can be obtained by learning from examples of seen classes. Therefore, knowledge of unseen classes is obtained by transferring knowledge about the attributes [12]. Attribute-based methods have become popular for ZSC recently [4,13]. There are two realizations: (1) Direct Attribute Prediction (DAP); and, (2) Indirect Attribute Prediction (IAP). They realize ZSC by employing attribute annotations of object classes derived from human prior knowledge, or from data analyses, and by training attribute predictors using low-level features extracted from images. In attribute annotations used by DAP and IAP, all object classes are annotated using the same set of attributes and all attributes are treated equally.

Not all attributes have a role in ZSC [3]. Choosing the best attribute combination benefits ZSC [11]. Attributes which distinguish one set of classes from another in scope are called ‘discriminative attributes’ [14]. Non-discriminative attributes have no role in ZSC, and the best combination excludes them. Furthermore, not all discriminative attributes have the same ability to distinguish different pairs of classes [10] and should be treated differently. For example, CAAP (Class-Attribute Association Prediction) [6] predicts class-attribute associations and discriminates between any negative and positive associations of the discriminative attributes and the unseen classes. Enriched descriptions of unseen classes make CAAP more competitive for ZSC. This study indicates that seeking attributes associated with unseen classes and differentiating these associations improve performances of attribute-based zero-shot classifier.

List of Abbreviations: CORL, Combining Ontology and Reinforcement Learning for Zero-Shot Classification; ZSC, Zero-Shot Classification; HCR, Hierarchical Classification Rule; MC, Monte Carlo; DAP, Direct Attribute Prediction; SVM, support vector machine.

* Corresponding author.
E-mail address: 913381959@qq.com (L. Yao).

https://doi.org/10.1016/j.knosys.2017.12.022

Hierarchy, another kind of prior knowledge, also benefits ZSC. A zero-shot classifier is realized by employing a hierarchy in [12]. It trains classifiers for superclasses of unseen classes in the hierarchy and classifies objects based on these classifiers. It obtains higher accuracy than other methods that do not use hierarchical knowledge. In a different way to use hierarchy, AHLE (Attribute and Hierarchy Label Embedding) [15] and SJE (Structured Joint Embedding) [16] derive class embedding from hierarchy as an input of additional information. They also achieve higher ZSC accuracy than other methods that lack hierarchical knowledge. HAT (Hierarchical Attribute Transfer) [10] is another attribute-based classifier that yields higher ZSC accuracy by training attribute predictors at different abstraction levels hierarchically. At the test stage, it first predicts each attribute value using attribute predictors at different abstraction levels. Then it calculates a score for every unseen class using all attributes values. Finally, it returns a classification result according to the scores. As a result, attributes at different abstraction levels contribute to ZSC differently. However, it has not been solved to adjust different attribute contributions using adaptive weights in HAT yet.

In this paper, we propose a method called CORL (Combining Ontology and Reinforcement Learning) for ZSC. By introducing the concept '*hierarchical classification rule*' (HCR) and the value function Q to evaluate HCRs, CORL is able to apply reinforcement learning to ZSC. CORL has two learning stages. At the training stage, CORL learns HCRs from attribute annotations guided by ontology. A predictor is trained for each attribute if it appears in an HCR. This enables it to predict object attributes using low-level features extracted from images. At the test stage, object attributes are predicted and used as a precondition to select HCRs. A ε-greedy policy in reinforcement learning [17] is employed to select HCRs that satisfy the preconditions for classifying the object hierarchically. The value function, Q, is adaptively determined by a Monte Carlo (MC) reinforcement learning method using feedbacks from classifications. The ε-greedy policy for HCR selection is adaptively adjusted and optimized along with the update of Q.

This study is believed to provide two significant contributions. First, the *hierarchical classification rule* (HCR) and ontologically-guided HCR learning from attribute annotations, allow using only the attributes necessary for classification, i.e. discriminative attributes. HCR improves prior knowledge expressiveness and is different from previous representations of hierarchical knowledge as it favors flexible deployment such as reinforcement learning. Second, an MC reinforcement learning method is used to adaptively optimize the policy for HCR selection, with the result being that only the most discriminative HCRs are used which achieves higher ZSC accuracies.

Experiments on a benchmark ZSC dataset showed that O-ZSC (the method which classifies object using only HCRs without reinforcement learning) obtained accuracy comparable to DAP while using fewer attributes. CORL achieves higher accuracies via adaptive adjustment of HCRs using reinforcement learning. The experiments validated that CORL's reinforcement learning is convergent. CORL outperforms most baseline zero-shot classifiers on 3 popular ZSC benchmark datasets suggesting that it selects the most discriminative HCRs for classification.

2. Related work

Attribute-based methods, which categorize samples of unseen classes using shared attributes between seen classes and unseen classes, is an important method for solving ZSC. The attribute-based method for ZSC proposed by Lampert, et al. [3], first links attribute annotations which are human prior knowledge with low-level features extracted from images. Two approaches implement this method: (1) Direct Attribute Prediction (DAP); and, (2) Indirect Attribute Prediction (IAP). DAP trains a predictor for each attribute appearing in attribute annotations using a support vector machine (SVM), aiming at estimating object attributes using low-level features extracted from images. IAP trains a probabilistic multi-class classifier to predict seen classes using SVM and low-level features extracted from images. The predicted probabilities and attribute annotations of seen classes are then used to estimate object attribute values. Utilizing the attributes which are comprehensible features as a bridge, DAP and IAP first associate low-level features extracted from images to the human prior knowledge, and then realize ZSC based on attributes. The technique of predicting attributes used by DAP is simple and direct, thus CORL adapts it to attribute predictions.

By querying domain ontology and removing classes having only one subclass, a hierarchical model is built in HAT [10]. It is used for building attribute predictors at different abstraction levels. CORL obtains a hierarchy in the same way. The hierarchy, as part of the ontology, is devoted to generating HCRs other than training attribute predictors. A hierarchy called AwA-10 [18] is built for the dataset "Animals with Attributes" [3] by querying the WordNet ontology [19]. This hierarchy is useful in generating HCRs for the ZSC conducted on "Animals with Attributes".

Hierarchical classifications use pre-defined semantic taxonomies [20]. A taxonomy contains only one root class and defines "Is-A" relationships between classes. An "Is-A" relationship is transitive and asymmetrical. A hierarchical classification is the opposite of a flat classification in terms of "whether semantic taxonomy is involved". In CORL, multiple local classifiers built by learning HCRs, are constructed for each superclass in the taxonomy. All HCRs are top-down classifiers. The latter part of CORL, classifying objects using attributes, is a top-down classification. Top-down hierarchical classifiers such as Selective Classifier [21] and Selective Feature Representation Top-Down [22] improve predictive abilities by selecting the best local classifier for each superclass using validations. HCR learning differs significantly from this. First, no validation data can be provided in ZSC. Second, HCR learning aims at obtaining rules that contain only discriminative attributes by decomposing attribute annotations other than the most predictive rules.

Ontology-based classifiers have been used extensively. They are constructed based on ontologies and realized through ontological reasoning. They take advantage of ontology's ability to represent various kinds of knowledge and data [23,24]. Ontology organizes knowledge hierarchically [25] and integrates rules in various ways [23,24,26]. Rules are integrated into ontology-based classifiers. Most are flat classification rules which directly classify objects taxonomically from root to leaf. Some flat classification rules are obtained by learning from decision trees [24,27]. However, the ontology is not incorporated into learning nor does it play a role in learning the flat classification rules.

Extracting rules from decision trees is a popular method for obtaining classification rules [28]. C4.5 is an algorithm for inducing a decision tree [29]. It attempts to select the best attributes for classification by using an information measure, *gain ratio*, and obtains a smaller tree via pessimistic pruning. CORL relies on C4.5′s attribute selection and feature pruning to build local classifiers for each superclass in taxonomy by calling C4.5 in order to obtain shorter rules.

Reinforcement learning is one of three typical machine learning techniques [30]. It perceives states of environment and chooses the best actions to change the environment. Trials and rewards are two important properties of reinforcement learning [31]. The MC reinforcement learning method is a method used for episodic tasks [17]. It estimates value function by averaging rewards from multiple episodes (also called trials). MC obtains the best policy from online experience without having full knowledge of the com-

Please cite this article as: B. Liu et al., Combining ontology and reinforcement learning for zero-shot classification, Knowledge-Based Systems (2017), https://doi.org/10.1016/j.knosys.2017.12.022