Theoretical Models for Learning from Multiple, Heterogenous and Strategic agents

(Doctoral Consortium)

Divya Padmanabhan
PhD Student
Indian Institute of Science
Bangalore, India

ABSTRACT

With the advent of internet enabled hand-held mobile devices, there is a proliferation of user generated data. Often there is a wealth of useful knowledge embedded within this data and machine learning techniques can be used to extract the information. However, as much of this data is user generated, it suffers from subjectivity. Any machine learning techniques used in this context should address the subjectivity in a principled way. We broadly study three problems in the context of learning from multiple agents, (1) Multi-label classification (2) Active Linear Regression (3) Sponsored Search Auctions.

1. INTRODUCTION

The problem of learning from multiple heterogenous sources is prevalent in the web today. With crowdsourcing platforms gaining popularity, it is possible to obtain labels (or any information required for supervised learning problems) cheaply. However the reliability of these multiple sources or agents is not known. Any learning algorithm may use the information provided by these sources, but must also evaluate the noise levels of these sources while leveraging the information provided by them. We study various learning paradigms in such scenarios including multi-label classification and linear regression and model the scenarios suitably.

A well known learning paradigm is active learning. In order to enhance an already learnt model, it would be useful to augment the training dataset with labels for some carefully chosen additional instances. In the context where multiple sources provide the labels, an important question that one must address is which of the agents must be selected to provide the label. We use techniques from the multi-armed bandits (MAB) literature for selecting the agent in active learning. The chosen agent could be strategic on the effort level he should put in, while providing the labels. Therefore, the centre would like the agent to put in her best efforts. We introduce incentives in such active learning scenarios to handle strategic agents and to induce the desired behaviour from them.

Multiple agents are also involved in sponsored search auctions (SSA). Multi-armed bandits (MAB) have been used extensively to provide solutions to this problem. To solve the strategic version of the problem, MAB mechanisms have been developed. However, the existing stochastic MAB mechanisms with a deterministic payment rule, proposed in the literature, necessarily suffer a regret of $\Omega(T^{2/3})$, where T is the number of time steps. This happens because the existing mechanisms consider the worst case scenario where the means of the agents' stochastic rewards are separated by a very small amount that depends on T. By making crucial observations about the separation between the rewards of the arms, we further introduce the notion of Δ-Regret and devise DSIC (dominant strategy incentive compatible) and IR (individually rational) mechanisms which suffer only logarithmic Δ-regret.

We now explain our contributions in detail in each of the problems.

2. MULTI-LABEL CLASSIFICATION FROM MULTIPLE AGENTS

Multi-label classification is an important problem in machine learning where an instance $\mathbf{d}$ is associated with multiple classes or labels. The task is to identify the classes for every instance. Multi-label classification finds applications in several areas, for example, text classification, image retrieval, etc. Topic models [2, 3, 4] have been used extensively in natural language processing tasks to model the process behind generating text documents. The model assumes that the driving force for generating documents arise from 'topics' which are latent. The topic models available in literature for multi-label setting include [10, 8]. However these models either involve too many parameters [8] or learn the parameters by heavily depending on iterative optimization techniques [10], thereby making it hard to adapt to the scenario where labels are provided by multiple agents. Moreover in all these models, the topics and hence words are assumed to be generated depending only on the classes that are present. They do not make use of the information provided by the *absence* of classes. The absence of a class often provides critical information about the words present. For example, a document labeled 'sports' is less likely to have words related to 'astronomy'. Similarly in the images domain, an image categorised as 'portrait' is less likely to have the characteristics of 'scenery'. Needless to say, such correlations are dataset dependent. However a principled analysis must account for such correlations. Motivated by this subtle observation, we introduce a novel topic model for multi-label classification.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

(1) Our model has the distinctive feature of exploiting any additional information provided by the absence of classes. Also, the use of topics enables our model to capture correlation between the classes. We refer to our topic model as ML-PA-LDA (Multi-Label Presence-Absence LDA).

(2) We enhance our model to account for the scenario where several heterogenous agents with unknown qualities provide the labels for the training set. We refer to this enhanced model as ML-PA-LDA-MNS (ML-PA-LDA with Multiple Noisy Sources).

(3) We test the performance of ML-PA-LDA and ML-PA-LDA-MNS on several datasets and establish its superior performance over state of the art.

For further details, the reader may refer to [6].

3. ACTIVE LINEAR REGRESSION FROM MULTIPLE STRATEGIC AGENTS

Active Learning [9] is an important problem in popular crowdsourcing platforms. These platforms are characterized by a large pool of diverse yet inexpensive annotators. To leverage these platforms for learning tasks actively, the following issues need to be addressed: (1) A learning model that encompasses parameter estimation and annotator quality estimation. (2) Identifying the best yet minimal set of instances from the pool of unlabeled data. (3) Determining an optimal subset of annotators to label the instances. (4) Providing suitable incentives to elicit best efforts from the chosen annotators under a budget constraint. We provide an end to end solution to address the above issues for a regression task.

Our contributions are:

(1) **Bayesian model for Regression**: We set up a Bayesian model for regression using labels from multiple annotators with varying noise levels.

(2) **Active learning for crowd regression and decoupling instance selection and annotator selection**: Next we focus on various active learning criteria as applicable to the proposed regression model. Interestingly, in our setting, we show that several criteria are equivalent.

(3) **Annotator selection with multi-armed bandits**: We look at the problem of selecting an annotator having least variance. We model this problem as a multi-armed bandit problem and establish regret bounds.

(4) **Handling strategic agents**: We consider the case of strategic annotators where the learner needs to induce them to put in their best efforts. For this, we propose the notion of '*quality compatibility*' and introduce a payment scheme that induces agents to put in their best efforts and is also individually rational.

Complete details of the above work can be found in [5].

4. SPONSORED SEARCH AUCTIONS AND STRATEGIC ADVERTISERS

In SSA, there are several advertisers (agents) who wish to display their ads along with the search results generated in response to a query from an internet user. There are two components that are of interest to the planner or the search engine, (1) *stochastic component*: click through rate (CTR) of the ads (2) *strategic component*: valuation of the agent for every click that the agent's ad receives. The search engine would seek to allocate a slot to an ad which has the maximum social welfare (product of click through rate and valuation). However neither the CTRs nor the valuations of the agents are known. This calls for a learning algorithm to learn the stochastic component (click through rate) as well as a mechanism to elicit the strategic component (valuation). This problem could become much harder as the agents may manipulate the learning process[1] to gain higher utilities.

We observe that the well-known characterization provided by Babaioff et al. [1] targets the worst case scenario. In particular, in the lower bound proof of $\Omega(T^{2/3})$, they consider an example scenario where the separation, Δ, between the expected rewards of the arms is a function of T. However, a dependence of Δ on T is severely restrictive for the case when the rewards are stochastic, even when the arms are non-strategic. Our contributions are the following (refer [7] for details):

(1) We make the crucial observation that in most MAB scenarios, the separation between the agents' rewards is rarely a function of T (the number of time steps). Moreover, in the case that the rewards of the arms are arbitrarily close, the regret contributed by such sub-optimal arms is negligible. We exploit this observation to allow the center to specify the resolution, Δ, with which the agents must be distinguished. We introduce the notion of Δ-Regret to formalize this regret.

(2) Using SSA as an example, we propose a dominant strategy incentive compatible (DSIC) and individually rational (IR) MAB mechanism with a deterministic allocation and payment rule, based on ideas from the UCB family of MAB algorithms. The proposed mechanism Δ-UCB achieves a Δ-regret of $O(\log T)$ for the case of single slot sponsored search auctions.

(3) We extend the above results to the case where multiple slots are to be allocated. Here again, our mechanism is DSIC, IR, and achieves a Δ-regret that is $O(\log T)$.

REFERENCES

[1] M. Babaioff, Y. Sharma, and A. Slivkins. Characterizing truthful multi-armed bandit mechanisms. *SIAM Journal on Computing*, 43(1):194–230, 2014.

[2] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *JMLR*, 3:993–1022, 2003.

[3] G. Heinrich. A generic approach to topic models. In *ECML PKDD*, pages 517–532. 2009.

[4] R. Krestel and P. Fankhauser. Tag recommendation using probabilistic topic models. *ECML PKDD Discovery Challenge*, page 131, 2009.

[5] D. Padmanabhan, S. Bhat, D. Garg, S. K. Shevade, and Y. Narahari. A robust UCB scheme for active learning in regression from strategic crowds. In *IEEE IJCNN*, pages 2212–2219, 2016.

[6] D. Padmanabhan, S. Bhat, S. K. Shevade, and Y. Narahari. Topic model based multi-label classification. In *IEEE ICTAI*, pages 996–1003, 2016.

[7] D. Padmanabhan, S. Bhat, S. K. Shevade, and Y. Narahari. A dominant strategy truthful, deterministic multi-armed bandit mechanism with logarithmic regret. In *AAMAS*, 2017 (accepted).

[8] T. N. Rubin, A. Chambers, P. Smyth, and M. Steyvers. Statistical topic models for multi-label document classification. *Machine Learning*, 88(1-2):157–208, 2012.

[9] B. Settles. Active learning literature survey. Technical report, University of Wisconsin–Madison, 2010.

[10] H. Wang, M. Huang, and X. Zhu. A generative probabilistic model for multi-label classification. In *ICDM*, pages 628–637, 2008.