An Empirical Evaluation of Ensemble Decision Trees to Improve Personalization on Advertisement

Eunyee Koh
Adobe Systems
321 Park Ave, San Jose, CA 95110
eunyee@adobe.com

Neha Gupta
Adobe Systems
601 Townsend St, San Francisco, CA 94103
nehagup@adobe.com

ABSTRACT

Ensemble decision tree algorithms are well known for good prediction accuracy in most cases, but not much research has been done on applying ensemble methods to improve personalization in the field of behavioral targeting in online advertisements. In behavioral targeting, the best ad is matched to the user based on his/her past activities and demographics. At present, most models used in the behavioral targeting are some form of linear models. Our goal in this paper is to analyze and understand the effect of ensemble techniques on large scale advertising data. Few of the main challenges of this kind of large scale data are sparse features and high dimensionality that make it hard for one single model to work the best. The form of ensemble method explored in this paper is the random forest based regression algorithm that combines the power of multiple decision trees to produce a more robust model which has a reduced variance as well bias.

Also, in the field of online advertising it is imperative to learn in an online fashion (while the advertising campaign is being run) as the customers want to get the most off their money at the earliest and the lifetime of such advertisements is short. So, some form of exploration vs. exploitation technique is also required to be used in the system. Our contributions in this paper are three fold. First, we develop a new technique to determine optimal parameters of the random forest algorithms. Second, we do a comparative analysis of random forest vs. logistic regression. Third, we combine ensemble decision tree algorithms with bandit algorithms to produce around 17% CTR improvement over random.

Categories and Subject Descriptors

H.4.M [**Information Systems Application**]: Miscellaneous; I.5.2 [**Pattern Recognition**]: Design Methodology-Feature evaluation and selection; I.2.6 [**Computing Methodologies**]: Learning

General Terms

Algorithms, Experimentation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
KDD'14, The Second Workshop on User Engagement Optimization, August 24–27, 2014, New York, NY, USA.
Copyright 2014 ACM X-XXXXX-XX-X/XX/XX.

Keywords

Personalization, Advertisement, Ensemble Decision Trees

1. INTRODUCTION

Many companies run internal marketing campaigns on their own websites to maximize their click-through-rate (CTR) and user experience. Each campaign usually consists of many different kinds of offers, and we would like our system to select the one that is most relevant to a given visitor. We can consider among many available information such as users' click history and geo location to make the decision of which offer to present. Human experts can develop the offer selection rule by their domain knowledge, but in order to make this scale, we want prediction algorithms to understand the users' behavior and make the decision automatically.

The problem of selecting a relevant or popular offer has been solved using bandit algorithms [7, 11]. There are two kinds of bandit algorithms; context free and contextual. Context free algorithms do not consider the users' context and only consider the popularity of offers to make decision. Contextual algorithms consider the users' context for selecting an offer by predicting the probability of click on an offer in a given context. So, better prediction accuracy of the click probability can improve the performance of contextual bandit algorithms.

Ensemble methods are well known for better prediction performance than other algorithms in most cases [9], but not much research has been done to apply ensemble methods to improve personalization on advertisements. In this paper, we propose our work-in-progress to leverage click predictions from ensemble decision trees in bandit algorithms and conduct empirical studies using real world datasets from Adobe Digital Marketing campaigns. Our empirical study results show that our approach performs around 17% better than random in term of CTR, and it is also better than other context free bandit algorithms that we choose to compare.

2. RELATED WORK

The general problem space of this work is that of targeted ad serving, which seeks to choose the optimal ad to serve a given user, based on features representing information about the user. An advertising campaign is costly to run, and therefore there is strong motivation to effectively infer the interests of a user and match them to the ad which is most likely to grab their attention. This has attracted wide interest from the research community, and many models have been constructed to address this problem [1, 5, 3, 10, 12]. These models generally work by learning from the behavior of past users targeted for a campaign in order to identify potential future conversions. While the aim of our model is the same, we take a more hybrid approach by combining ensemble decision trees with

Algorithm 1 Ensemble Decision Trees - EDT

Model Building: Given a training set *X, Y*

for *b*=1 through *B*:

1. Sample, with replacement, *n* training examples from *X, Y*; call X_b, Y_b
2. Train a regression tree T_b on X_b, Y_b with a random subset of available variables when splitting a node.

end for

Prediction: A testing sample x'

$\hat{f} = \frac{1}{B}\sum_{b=1}^{B}\hat{T}_b(x')$

exploration and exploitation strategies to get a better performance.

3. ALGORITHM

Ensemble methods use multiple learning algorithms to obtain better predictive performance than could be obtained from any of the constituent learning algorithms [6]. One example of the ensemble methods we choose to use for our empirical study is random forests [2]. Random forests are a classification or regression method that operates by constructing a multitude of decision trees at training time. Their output is a class or regression value that is the mode of the outputs by individual trees. The method combines bagging and the random selection of features at each split in order to construct a collection of decision trees with controlled variance. Algorithm 1 (referred to as EDT) outlines how ensemble decision trees are constructed and make predictions.

3.1 Algorithm Framework

We build the EDT regression model per offer in each marketing campaign data. For new testing data, we will get $P(click|offer_i)$ *for i=1,...,n* where *n* is the number of offers in campaign. Then, we can apply $\arg\max_i P(click|offer_i)$ for a greedy choice (exploitation).

Exploration is required to collect the click response from different offers and learn the regression model of the *offer*. So, we combine exploration and exploitation strategy with EDT regression model.

3.2 Training the Regression Model

Parameter tuning is one of the important tasks for prediction modeling. The number of trees parameter, *B* in Algorithm 1, can be a few hundreds to several thousands depending on the size and nature of the training data. Increasing the number of trees tends to decrease the variance of the model, without increasing the bias. As a result, the training and test error tend to level off after some number of trees have been fit. An optimal number of trees can be found using cross-validation, or by observing the out-of-bag error [2, 6].

To determine the optimal parameter on the fly, we have developed a new technique to decide the number of trees by considering the maximum coverage of available variables. At each node split decision, EDT chooses a subset of random variables to make the splitting decision. The goal of our technique is to choose the number of trees such that all the variables are covered, so the learning algorithm can consider and compare all the variables for the model.

The optimal number of trees to build is decided as:

$$(1-\frac{M}{N})^B < 10^{-4}$$

where *N* is the total number of variables, *M* is the number of variables selected at random for node splitting, and *B* is the number of trees to build. The probability of selecting *M* variables such that a variable is never selected is $\frac{C(N-1,M)}{C(N,M)} = (1-\frac{M}{N})$, and the probability of a variable is never selected for *B* trees is $(1-\frac{M}{N})^B$. We want the probability of a variable to be neglected in the model to be less than 0.01%.

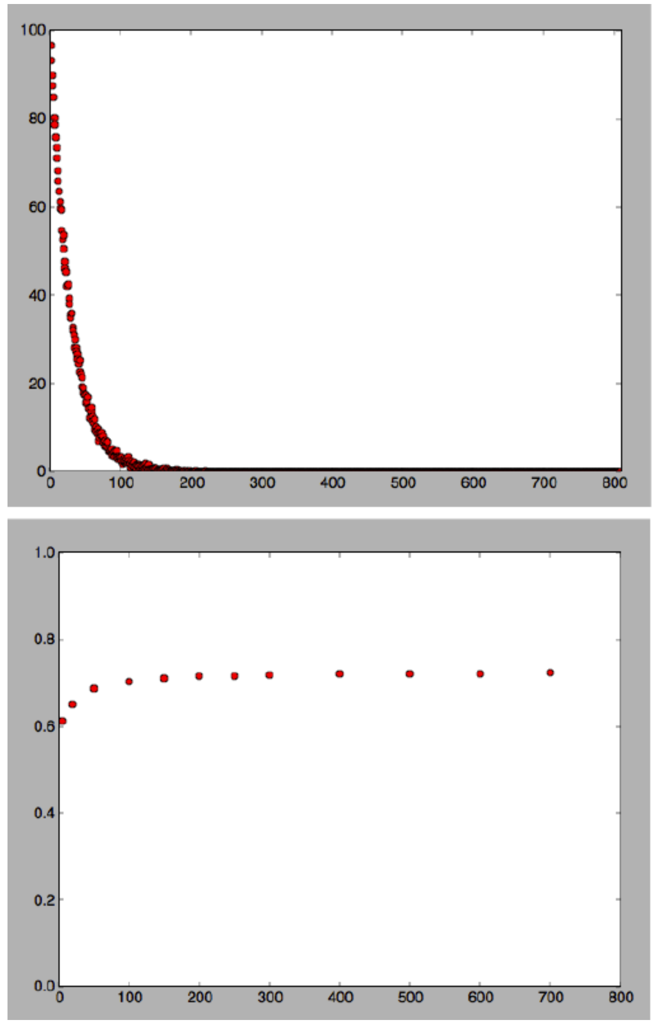

Figure 1: Parameter study results for selecting the optimal number of trees in ensemble decision trees with dataset 1; Upper graph: Y-axis for a probability of a variable not to be selected in EDT (scale in percentage), and X-axis for the number of trees in EDT; Bottom graph: Y-axis for AUC, and X-axis for the number of trees in EDT.

We conducted a parameter study with many different datasets to verify whether we can find the optimal one by using our proposed approach. Figure 1 shows the one of our parameter study results using dataset 1. Dataset 1 in Figure 1 has around 800 variables. Assuming we are selecting $sqrt(800)$ random variables each time [6], the optimal value for the number of trees parameter *B* computed by our approach is approximately 255. The bottom graph in Figure 1 shows the Area Under the Curve (AUC) change by the number of trees parameter variations. We can see that in about 200 number of trees, the change of AUC is stabilized. We can find a similar pattern in the results from the dataset 2 in Figure 2. The dataset 2 has around 340 variables, and the optimal value for the parameter *B* is approximately 172 computed by our technique with $sqrt(340)$ random variable selection. The parameter study results (bottom graph in Figure 2) confirms our proposal as the AUC change stabilizes around 172 number of trees.

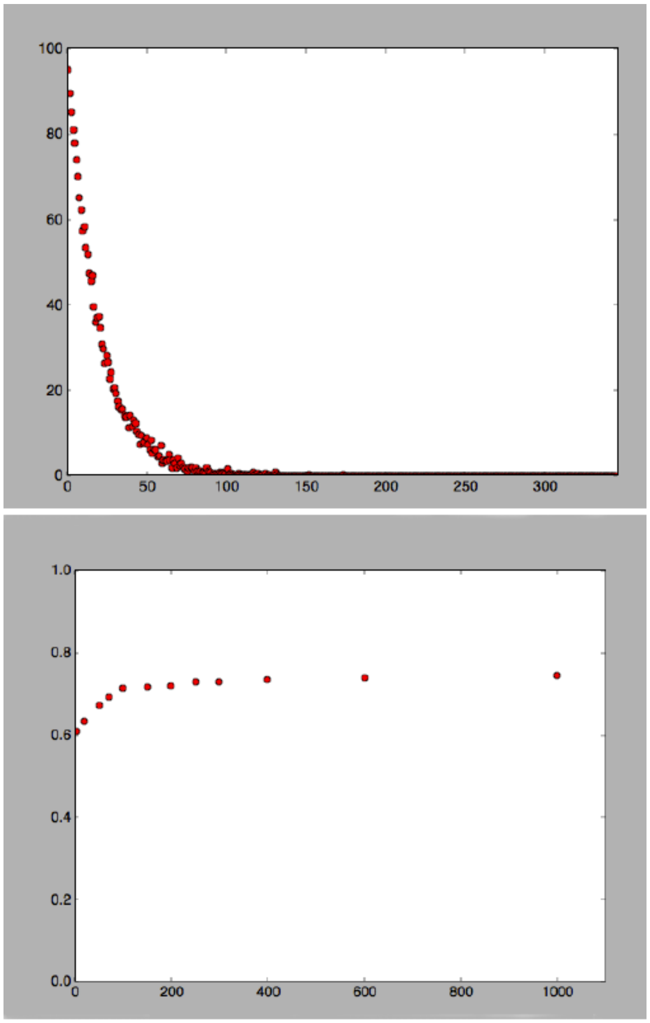

Figure 2: Parameter study results for selecting the optimal number of trees in ensemble decision trees with dataset 2; Upper graph: Y-axis for a probability of a variable not to be selected in EDT (scale in percentage), and X-axis for the number of trees in EDT; Bottom graph: Y-axis for AUC, and X-axis for the number of trees in EDT.

3.3 Exploration and Exploitation Strategy

A multitude of research has been done on improving exploration and exploitation strategies for online advertisements [11]. One of the well-known bandit algorithms, e-greedy, is to select the offer with highest current average reward with probability $(1 - e)$ and select the random offer with probability e. Softmax, the other bandit algorithm we choose to use, is using a similar strategy except that e is reduced on the basis of the learning progress instead of manual tuning. High fluctuations in the value estimates lead to a high e (exploration); low fluctuations to a low e (exploitation).

While e-greedy and softmax are context free algorithms, our approach uses the regression approach using EDT to predict the click probability for a given context. We combined EDT with e-greedy and softmax as exploration is essential in some cases (for example, when there is not enough data for the offer to learn good prediction model or when there is not enough variations in click probabilities among different offers).

There are many existing contextual bandit algorithms to compare against our approach, and we will leave it for our future work. The reason we did not try out the existing contextual bandit algorithms in the first place is that prediction accuracy for a given context is important and EDT performs better than linear models or bayesian probabilities in many cases [6].

Our approach learns in a batch mode. Recent research presents online EDT algorithms [4], and we can try this out in future. However, we do not expect a big performance change by switching the batch algorithm to online because we can learn quickly in mini-batches.

4. EXPERIMENTS

We used Adobe Digital Marketing campaign datasets to conduct our experiments. Each record in the dataset looks like $(x, offer, r)$ where x is the user context, *offer* is one of the available offers in the campaign that has been presented to the user, r contains click or non-click information. We used the dataset where offers are served randomly (random traffic) for our experiments.

4.1 Evaluation Methodology

The first evaluation we did is to compare our EDT model prediction accuracy with the other baseline, logistic regression. We used AUC metric because it is a good metric for measuring accuracy for the highly imbalanced dataset such as CTR.

For our second evaluation, we used the unbiased offline CTR estimation technique by Lihong et al. [8], and it works on the dataset from random traffic which we currently have.

4.2 Evaluation Results

For the contextual bandit algorithm with regression approach, we need to learn a regressor f to predict a click probability given $(x, offer)$. The prediction accuracy from a regressor can be an important factor for the contextual bandit to demonstrate the meaningful CTR lift, so we evaluated prediction accuracy with AUC. Then, we conducted our experiments to compute estimated CTR for our proposed approach to compare with existing bandit algorithms.

4.2.1 Prediction Accuracy with AUC

We used two datasets, campaign 1 and campaign 2, for AUC comparison experiments. The dataset 1 and the dataset 2 that we used for parameter study in section 3.2 are one of the offers from the campaign 1 and one of offers from the campaign 2.

The campaign 1 dataset has four different offers, 820311 records with the time period from July to August 2013. The offers are distributed almost evenly in the dataset due to it being from the random traffic. The campaign 2 dataset also has four different offers, 276905 records from random traffic, with the time period from Feb to March 2013.

We first divided the dataset per offer, and separated each subset of dataset into training (70%) and testing (30%) just like any supervised learning evaluation setting. We built LR models and EDT models in training data, and computed the AUC in testing data. Results show that EDT performs around 23% better than LR in prediction accuracy with AUC for the campaign 1 dataset, and EDT performs around 5% better than LR in the campaign 2 dataset (see Table 1).

4.2.2 CTR Estimation Results

We used the campaign 1 dataset to compare CTR estimation results using different algorithms; random, e-greedy, softmax, e-greedy+EDT, and softmax+EDT. Results with CTR improvements

Table 1: Performance Metrics AUC: Logistic Regression (LR), EDT

Campaign 1	LR	EDT
Offer 1	0.65	0.76
Offer 2	0.64	0.75
Offer 3	0.59	0.72
Offer 4	0.57	0.75
Average	0.61	**0.75**

Campaign 2	LR	EDT
Offer 1	0.84	0.88
Offer 2	0.83	0.88
Offer 3	0.83	0.88
Offer 4	0.85	0.89
Average	0.84	**0.88**

over random is in Figure 3. For each algorithm, we ran multiple experiments with parameter variations and selected the best one to compare with each other. We also used average results of each algorithm running 100 times because exploration in each algorithm can create a variance.

Results show our approaches improve over random a little more than 17% and our approaches perform better than other context free algorithms; e-greedy and softmax. A tuned e-greedy performs better than softmax for this dataset. This can mean that during the online learning offers were competing with each other quite often, and softmax did more exploration than e-greedy.

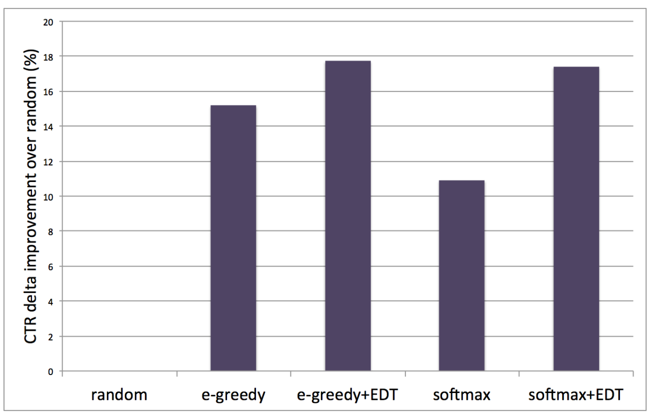

Figure 3: CTR improvements over random strategy

5. CONCLUSION

In this paper, we present our empirical study results of using ensemble decision trees (random forest) in the field of behavior targeting in online advertisements. Not many research has been done using ensemble decision trees to improve personalization on advertisements. As they are well known for good prediction accuracy compared to other algorithms in most cases, we propose to leverage prediction probabilities from ensemble decision trees in exploration and exploitation strategies.

We conducted three experiments. First, we developed a new technique to tune parameters for random forest algorithms and verified it with multiple parameter studies with our dataset. Second, we compared prediction accuracy of random forest vs. logistic regression as linear models have been explored many times in Ad personalization. Third, we combined ensemble decision trees with bandit algorithms to form a contextual bandit and compare it with context free bandit algorithms using CTR metrics.

Our experiment results show that ensemble decision trees perform better on CTR prediction than logistic regression, and the better prediction accuracy will lead to better exploitation strategy in exploration and exploitation algorithms. We compared our proposed approach only with context free algorithms, and we found that our approach performs better than tuned e-greed and softmax. Our future work will include other contextual bandit algorithms on our experiments.

6. REFERENCES

[1] M. Aly, A. Hatch, V. Josifovski, and V. K. Narayanan. Web-scale user modeling for targeting. In *Proceedings of the 21st international conference companion on World Wide Web*, WWW '12 Companion, pages 3–12, New York, NY, USA, 2012. ACM.

[2] L. Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, Oct. 2001.

[3] Y. Chen, D. Pavlov, and J. F. Canny. Large-scale behavioral targeting. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '09, pages 209–218, New York, NY, USA, 2009. ACM.

[4] M. Denil, D. Matheson, and N. de Freitas. Consistency of online random forests. In *ICML (3)*, volume 28 of *JMLR Proceedings*, pages 1256–1264. JMLR.org, 2013.

[5] N. Gupta, A. Das, S. Pandey, and V. K. Narayanan. Factoring past exposure in display advertising targeting. In *Proceedings of the 18th ACM SIGKDD*, pages 1204–1212. ACM, 2012.

[6] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer-Verlag, New York, 2001.

[7] L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 661–670, New York, NY, USA, 2010. ACM.

[8] L. Li, W. Chu, J. Langford, and X. Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*, WSDM '11, pages 297–306, New York, NY, USA, 2011. ACM.

[9] R. Maclin and D. W. Opitz. Popular ensemble methods: An empirical study. *CoRR*, abs/1106.0257, 2011.

[10] S. Pandey, M. Aly, A. Bagherjeiran, A. Hatch, P. Ciccolo, A. Ratnaparkhi, and M. Zinkevich. Learning to target: what works for behavioral targeting. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, CIKM '11, pages 1805–1814, New York, NY, USA, 2011. ACM.

[11] L. Tang, R. Rosales, A. Singh, and D. Agarwal. Automatic ad format selection via contextual bandits. In *Proceedings of the 22Nd ACM International Conference on Conference on Information and Knowledge Management*, CIKM '13, pages 1587–1594, New York, NY, USA, 2013. ACM.

[12] J. Yan, N. Liu, G. Wang, W. Zhang, Y. Jiang, and Z. Chen. How much can behavioral targeting help online advertising? In *Proceedings of the 18th international conference on World wide web*, WWW '09, pages 261–270, New York, NY, USA, 2009. ACM.