



**HAL**  
open science

# Feature Selection Methods for Uplift Modeling and Heterogeneous Treatment Effect

Zhenyu Zhao, Yumin Zhang, Totte Harinen, Mike Yung

► **To cite this version:**

Zhenyu Zhao, Yumin Zhang, Totte Harinen, Mike Yung. Feature Selection Methods for Uplift Modeling and Heterogeneous Treatment Effect. 18th IFIP International Conference on Artificial Intelligence Applications and Innovations (AIAI), Jun 2022, Hersonissos, Greece. pp.217-230, 10.1007/978-3-031-08337-2\_19 . hal-04668649

**HAL Id: hal-04668649**

**<https://inria.hal.science/hal-04668649v1>**

Submitted on 7 Aug 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

# Feature Selection Methods for Uplift Modeling and Heterogeneous Treatment Effect

Zhenyu Zhao<sup>1</sup>, Yumin Zhang<sup>2</sup>, Totte Harinen<sup>3</sup>, and Mike Yung<sup>4</sup>

<sup>1</sup> Tencent [zzy287@gmail.com](mailto:zzy287@gmail.com)

<sup>2</sup> Purdue University [zhan2013@purdue.edu](mailto:zhan2013@purdue.edu)

<sup>3</sup> Toyota Research Institute [th.harinen@gmail.com](mailto:th.harinen@gmail.com)

<sup>4</sup> Spotify [yungmsh@gmail.com](mailto:yungmsh@gmail.com)

**Abstract.** Uplift modeling is a causal learning technique that estimates subgroup-level treatment effects. It is commonly used in industry and elsewhere for tasks such as targeting ads. In a typical setting, uplift models can take thousands of features as inputs, which is costly and results in problems such as overfitting and poor model interpretability. Consequently, there is a need to select a subset of the most important features for modeling. However, traditional methods for doing feature selection are not fit for the task because they are designed for standard machine learning models whose target is importantly different from uplift models. To address this, this paper introduces a set of feature selection methods explicitly designed for uplift modeling, drawing inspiration from statistics and information theory. Empirical evaluations are conducted on the proposed methods on publicly available datasets, demonstrating the advantages of the proposed methods compared to traditional feature selection. We make the proposed methods publicly available as a part of the *CausalML* open-source package.

**Keywords:** Feature Selection · Uplift Modeling · Causal Learning.

## 1 Introduction

Uplift modeling [8,10,12,13,18,21,24,30,32,34,35,36,37,39], also known as heterogeneous treatment effect estimation or incremental modeling, is a technique designed to estimate the individual treatment effect (ITE) or the conditional average treatment effect (CATE) of an intervention. It is often used for user targeting, budget allocation, and personalization applications.

In practice, there is often a rich set of features that can be used to build an uplift model. However, using all of the available features in the model can lead to computational inefficiency, over-fitting, high maintenance workload, and model interpretation challenges. A feature selection method screens the large feature space and picks the important features for the model, and then the uplift model can be built based on the most important features. As Radcliffe and Surry [27] noted, feature selection is actually of greater importance to improve model quality and stability in uplift modeling than in conventional modeling because

uplift models estimate the difference between two outcomes, which is often small relative to the direct outcomes (which conventional models estimate), increasing the risk of over-fitting markedly.

This work contributes to the feature selection of uplift modeling from both methodological and empirical perspectives, specifically:

- Five filter methods and two types of embedded methods are proposed as feature selection methods for uplift modeling.
- These feature selection methods are empirically evaluated and compared with ordinary feature selection methods on two data sets. One synthetic dataset is generated for this study and made available online [38], and the other dataset is from Megafon Uplift Competition training data [1].
- This study demonstrates that the ordinary feature selection methods for conventional machine learning models are sub-optimal in the uplift modeling context.
- To make these feature selection methods easily accessible for broad applications, the proposed methods are implemented in the *CausalML* [7] open-source package.

To simplify the discussion, the rest of this paper assumes the outcome variable to be binary, covering most common use cases such as advertisement click-through, new user conversion, and existing user retention.

The structure of the paper is as follows. Section 2 discusses existing work on feature selection methods for uplift modeling, and the difference from feature selection for conventional models and observational causal inference. Section 3 reviews the critical concepts of uplift modeling. Section 4 introduces five filter methods and three embedded methods for uplift modeling. Section 5 evaluates these methods with benchmark methods on two data sets. Finally, Section 7 summarizes the findings and makes recommendations for choosing and using the proper feature selection methods for uplift modeling applications.

## 2 Related Work

Feature selection methods for conventional machine learning problems have been well studied [4,5,6,33]. Nevertheless, as this paper will show, these methods are ineffective for solving feature selection problems for uplift modeling. The main reason is that such methods try to select predictive features of the outcome variable, which may not be the features related with treatment effect variability.

This paper assumes the data for uplift modeling is collected from a randomized experiment (a.k.a A/B testing [17]), which is often the case in practice. It is essential to differentiate the feature selection methods for uplift modeling based on data from randomized experiments and the feature selection methods developed for observational causal inference. Feature selection algorithms for observational causal inference, such as the lasso-based approach proposed by [31], are designed to help models whose goal is to reduce confounding. These methods are out of scope for discussion in this paper as the framework and goals are different.

Several feature selection methods for uplift modeling purposes have been discussed in the literature. A filter method named net information value (NIV) [19] is built based on the net weight of evidence (NWOE), where NWOE is the difference between the weight of evidence (WOE) of treatment and WOE of control, and NIV is a weighted average of NWOE. Radcliffe [27] introduced the pessimistic qini estimate that uses one feature to train an uplift model at a time, and the resulting model accuracy measure (qini) is taken as the feature importance score. Certain regularizations are added to this method to improve stability. The pessimistic qini method can be more computationally costly as it involves training and potentially tuning an uplift model.

### 3 Uplift Models

Uplift modeling can be viewed as a way to estimate ITE using machine learning models. Following the commonly used Neyman-Rubin causal model [15,22,28,29], the ITE for unit  $i \in \{1, 2, \dots, N\}$  can be expressed as:

$$Y_i(1) - Y_i(0) \tag{1}$$

where  $Y_i(1)$  and  $Y_i(0)$  denotes the outcome variable for individual  $i$  under treatment condition and control condition respectively.

A closely related concept is the conditional average treatment effect (CATE). Let  $X = (X_1, X_2, \dots, X_M)$  denote  $M$  covariates (also called features) and  $x_i = (x_{i1}, x_{i2}, \dots, x_{iM})$  denote the observed realisation of  $X$  for unit  $i$ . The conditional average treatment effect (CATE) is the expected treatment effect within groups of units that have a similar realisations of  $X$ :

$$\tau(x) := E[Y(1) - Y(0) | X = x] \tag{2}$$

In practice, many uplift modeling approaches train on CATEs. [13]

There are two main categories of uplift models. The first category is known as meta-learners ([18,23]). The models in this category are built based on conventional machine learning models as base learners in such a way that they can estimate CATEs, such as the “*TwoModel*” approach (a.k.a. *T-Learner* [14]), *X-Learner* [18], and *R-Learner* [23]. The other category of uplift models is known as uplift trees or causal trees. These models are based on modifying the loss functions within the classification or regression trees, such that the split is optimized for estimating the heterogeneous treatment effect [2,3,11,12,30,35]. In this paper, feature selection methods are evaluated with models from both uplift modeling categories to test the generality of the proposed feature selection methods.

### 4 Feature Selection Methods for Uplift Modeling

To identify the important features for uplift modeling, we propose both filter methods, which are easy and fast to use as a screening step for the data, and embedded methods, which are a by-product from training an uplift model.

#### 4.1 Filter Methods

As discussed above, a feature’s importance in an uplift modeling task depends on how well it predicts the CATE. A filter method calculates the importance score for each feature based on the marginal relationship between the treatment effect and the feature. It can be used for quickly screening the feature space and selecting important features for downstream modeling tasks.

**F Filter** The F filter method is named after the F statistic for testing the significance of the interaction between the treatment indicator and a feature in linear regression. To capture possible nonlinear associations between features and the treatment effect, we extend this method by including higher-order terms of the features in the regression. For a given feature  $X_j$  ( $j \in \{1, \dots, M\}$ ), the heterogeneous treatment effect of  $X_j$  can be studied in the following regression:

$$Y = \alpha + \delta I + \sum_{r=1}^R \beta_r X_j^r + \sum_{r=1}^R \theta_r I X_j^r + \epsilon \quad (3)$$

where  $Y$  stands for the response variable,  $I$  is the treatment indicator (1 for treatment, and 0 for control),  $R$  is a hyperparameter controlling the order to be studied ( $X_j^r$  is  $X_j$  to the power of  $r$ ),  $\epsilon$  represents the error term, and  $\alpha$ ,  $\delta$ ,  $\beta$ ,  $\theta$  are the coefficients, where the significance of  $\theta$  indicates the strength of the heterogeneous treatment effect in the dimension of  $X_j$ .

The significance of  $\theta$  can be studied by contrasting the model above with the following reduced model:

$$Y = \alpha' + \delta' I + \sum_{r=1}^R \beta'_r X_j^r + \epsilon' \quad (4)$$

The feature importance score by F filter method is defined as the F-statistic for the coefficient of the interaction term:

$$F = \frac{(RSS - RSS')/R}{RSS'/(N - R - 2)} \quad (5)$$

with

$$RSS = \sum_{i=1}^N (y_i - \hat{\alpha} - \hat{\delta} I_i - \sum_{r=1}^R \hat{\beta}_r x_{ij}^r - \sum_{r=1}^R \hat{\theta}_r I_i x_{ij}^r)$$

$$RSS' = \sum_{i=1}^N (y_i - \hat{\alpha}' - \hat{\delta}' I_i - \sum_{r=1}^R \hat{\beta}'_r x_{ij}^r)$$

where  $N$  is the total number of observations,  $RSS$  is the Residual Sum of Squares for fitted model 3 and  $RSS'$  is the Residual Sum of Squares for fitted model 4, that are calculated by plugging in the sample data and fitted coefficients.

This F statistic follows an F distribution with  $(R, N - R - 2)$  degrees of freedom assuming that the true value of  $\theta$  equals 0. A byproduct of the F filter method is a p-value for the correlation between the feature and the treatment effect, which can be used to deem whether heterogeneity in a given feature dimension counts as statistically significant.

Setting the hyperparameter  $R > 1$  will enable the F filter to capture the nonlinear relationship between feature and HTE. As empirical results show in Section 5, F filter with  $R = 2$  outperforms F filter with  $R = 1$  remarkably.

The F filter method works for both the regression uplift modeling problem where  $Y$  is continuous and the classification uplift modeling problem where  $Y$  is binary. Both continuous and binary features can be evaluated by the F filter. A multi-class categorical feature can fit in this method with one-hot-encoding approach (transforming one multi-class feature to multiple binary features).

**LR Filter** The LR (Likelihood Ratio) filter defines the feature importance score as the likelihood ratio test statistic for the interaction coefficient in a logistic regression model. Similar to the F filter, the LR filter for any feature  $X_j$  is constructed by contrasting two logistic regression models:

$$\begin{aligned} \text{logit}(p(X_j, I; \alpha, \delta, \beta_1, \dots, \beta_R, \theta_1, \dots, \theta_R)) &= \alpha + \delta I + \sum_{r=1}^R \beta_r X_j^r + \sum_{r=1}^R \theta_r I X_j^r \quad (6) \\ \text{logit}(p(X_j, I; \alpha', \delta', \beta'_1, \dots, \beta'_R)) &= \alpha' + \delta' I + \sum_{r=1}^R \beta'_r X_j^r \quad (7) \end{aligned}$$

where  $p(X_j, I; \alpha, \delta, \beta_1, \dots, \beta_R, \theta_1, \dots, \theta_R)$  and  $p(X_j, I; \alpha', \delta', \beta'_1, \dots, \beta'_R)$  are probability representations of  $Pr(Y = 1|X_j, I)$  under two functions, and  $R$  is a hyperparameter.

The significance of interaction coefficient  $\theta$  can be tested through a likelihood ratio test. Let  $\hat{\alpha}, \hat{\delta}, \hat{\beta}_1, \dots, \hat{\beta}_R, \hat{\theta}_1, \dots, \hat{\theta}_R, \hat{\alpha}', \hat{\delta}', \hat{\beta}'_1, \dots, \hat{\beta}'_R$  denote the fitted coefficient estimates for the logistic regression models. The likelihood ratio statistic can be calculated by plugging in the sample data and fitted coefficients:

$$LR = -2 \sum_{i=1}^N [y_i \log(p(x_{ij}, I_i; \hat{\alpha}, \hat{\delta}, \hat{\beta}_1, \dots, \hat{\beta}_R, \hat{\theta}_1, \dots, \hat{\theta}_R)) \quad (8)$$

$$+ (1 - y_i) \log(p(x_{ij}, I_i; \hat{\alpha}, \hat{\delta}, \hat{\beta}_1, \dots, \hat{\beta}_R, \hat{\theta}_1, \dots, \hat{\theta}_R)) \quad (9)$$

$$- y_i \log(p(x_{ij}, I_i; \hat{\alpha}', \hat{\delta}', \hat{\beta}'_1, \dots, \hat{\beta}'_R)) \quad (10)$$

$$- (1 - y_i) \log(p(x_{ij}, I_i; \hat{\alpha}', \hat{\delta}', \hat{\beta}'_1, \dots, \hat{\beta}'_R))] \quad (11)$$

The LR statistic is taken as the feature importance score by the LR filter. Assuming that the true value of  $\theta$  is 0, then the LR statistic follows a  $\chi^2_R$  distribution. This LR filter method can also produce a p-value for feature importance. Similar to the F filter, setting  $R > 1$  for the LR filter extends its capability for capturing the nonlinear importance of features.

The LR filter method can be used for categorical (binary and multi-class) outcome, as well as continuous and binary features. Similar to the F filter, the LR filter can be applied to a multi-class categorical feature with one-hot-encoding approach.

**Bin-based Divergence Filter (KL Filter, ED Filter, Chi Filter)** The section introduces three variants of the bin-based divergence filter method, which are direct applications from the split criteria for uplift trees proposed by [30]. Similar to the NIV method [19], for a given feature, the bin-based method first divides the samples into  $S$  (preferably equally sized) bins, where  $S$  is a hyperparameter. The importance score is defined as the divergence measure of the treatment effect over these  $S$  bins.

Formally, let  $p_k$  and  $q_k$  denote the sample proportion of class  $Y = 1$  in the  $k$ th ( $k = 1, \dots, K$ ) bin for the treatment group and control group respectively. The KL (Kullback-Leibler divergence) filter, ED (squared Euclidean Distance), and Chi (chi-squared divergence) filter feature importance scores are defined as follows:

$$\Delta_{KL} := \sum_{k=1}^K \frac{N_k}{N} KL(p_k, q_k) = \sum_{k=1}^K \frac{N_k}{N} \left[ p_k \log \frac{p_k}{q_k} + (1 - p_k) \log \frac{1 - p_k}{1 - q_k} \right] \quad (12)$$

$$\Delta_{ED} := \sum_{k=1}^K \frac{N_k}{N} ED(p_k, q_k) = 2 \sum_{k=1}^K \frac{N_k}{N} (p_k - q_k)^2 \quad (13)$$

$$\Delta_{Chi} := \sum_{k=1}^K \frac{N_k}{N} \chi^2(p_k, q_k) = \sum_{k=1}^K \frac{N_k}{N} \left[ \frac{(p_k - q_k)^2}{q_k} + \frac{(p_k - q_k)^2}{1 - q_k} \right] \quad (14)$$

where  $N_k$  is the sample size in the  $k$ th bin.

Even though these bin-based filter methods share the same divergence measures as an uplift tree, they are simpler to implement and compute than training an uplift model. For data with  $M$  features, the time complexity of the bin-based filter methods is linear with sample size,  $O(M \cdot N)$ , while the time complexity of building a single complete uplift tree is  $O(M \cdot N \log(N))$ ,

The bin-based divergence filter method can be directly applied to categorical outcome and continuous features. For discrete features, the divergence measure can be applied without the binning step.

## 4.2 Embedded Methods

The final category of methods that we propose, the embedded methods, obtain feature importance as a byproduct of training an uplift model. The way in which the importance scores are calculated differs depending on whether the target model is a meta-learner or an uplift tree. The embedded methods can be used for continuous and categorical outcomes, as well as continuous and discrete features.



**Embedded Methods by Meta-learners** For meta-learners, feature importance can be obtained from the base-learners, which are the composite models making up a meta-learner.

The simplest meta-learner, the *OneModel* or *S-Learner* [20], is based on just one base-learner and predicts treatment effects as the difference between  $E[Y | X = x, I = 1]$  and  $E[Y | X = x, I = 0]$  where  $I$  is the treatment indicator. Because the importance scores from the single base-learner in *S-Learner* tend to be similar to an outcome based model (the only difference between the two is that the base-learner in *S-Learner* has a treatment indicator as an additional feature), we will not differentiate these two embedded methods and let *Outcome* embedded represent these two methods in the following discussion.

The *TwoModel* embedded feature selection method is derived from the *TwoModel* uplift model, which feature importance score is defined as the sum of its embedded importance scores produced by the two base-learners. As the embedded methods derived from Meta-learners are based on ordinary feature selection methods in base learners (conventional model), the feature selection performance for uplift modeling is expected to be poor as it does not consider HTE during the process.

**Embedded Methods by Uplift Trees** For uplift trees, the importance score for a feature can be defined as the cumulative contribution to the loss function during the tree node splits in the trees. This is similar to the well-known embedded feature importance for standard classification trees, except the score is obtained from an uplift tree with a special splitting criterion. At each split, we calculate the gain in the distribution divergence:

$$\Delta = \sum_{k \in \{\text{left}, \text{right}\}} \frac{n_k}{n} D(p_k, q_k) - D(p, q), \quad (15)$$


where  $n$  is the sample size in the parent node,  $n_k$  is the sample size in the child node,  $D()$  is divergence measure defined as  $KL(), ED(), \chi^2()$  in Eq. (12),  $p$  and  $q$  denote the proportion of  $Y = 1$  for the treatment group and control group separately in the parent node, and  $p_k, q_k$  are corresponding proportions in the child nodes. The feature importance score is calculated by summing over all the  $\Delta$  from the tree node splits where the feature is used.

The time complexity for embedded methods depend on the learners used. For random forest algorithms, it is at order of  $O(t_{tree} \cdot m_s \cdot N \cdot \log(N))$  where  $t_{tree}$  is the number of trees and  $m_s$  is the maximum features considered in each split.

## 5 Empirical Evaluation

This section compares different feature selection methods empirically, and evaluate their performance consistency with different uplift models.

### 5.1 Experiment 1: Evaluation with Synthetic Data



**Fig. 1.** Feature Association Pattern with Outcome by Experiment Group in Experiment 1. The first two plots demonstrate a heterogeneous treatment effect associated with uplift features in a linear and sine pattern respectively. The last two plots illustrate classification features are correlated with outcome, but not treatment effect.

For evaluating the feature selection methods, a specific synthetic data generation process is designed, such that the data contains three types of features: (1) uplift features influencing the treatment effect on the conversion probability; (2) classification features affecting the conversion probability but independent of the treatment effect; and (3) irrelevant features that are independent of both conversion probability and the treatment effect.

The binary response variable is generated based on the probability  $Pr(Y = 1|X, I)$ , where  $X$  denotes the feature vector and  $I$  is the treatment indicator. Assuming there are  $m_1$  classification features and  $m_2$  uplift features, the probability  $Pr(Y = 1|X, I)$  itself is modeled as:

$$\text{logit}(Pr(Y = 1|X, I)) = a_1 + \sum_{j=1}^{m_1} f_j(X_j)\beta_j + I \cdot (a_2 + \sum_{j=m_1+1}^{m_1+m_2} f_j(X_j)\beta_j) + \epsilon$$

where  $a_1, a_2, \{\beta_j\}_1^{m_1+m_2}$  are coefficient parameters,  $\epsilon$  is a random noise added,  $\{f_j(\cdot)\}_1^{m_1+m_2}$  is transformation function for the features, such that features have six types of association with the outcome: linear, quadratic, cubic, ReLU (Rectified Linear Unit [9]), trigonometric function sine, and cosine. Example feature patterns are plotted in Figure 1.

In this study, there are 36 features in total, including  $m_1 = 10$  classification features,  $m_2 = 6$  uplift features, and  $m_3 = 20$  irrelevant features. The values for the coefficients are set such that the average control conversion probability is around 0.2 and the average treatment effect is around 0.1. The generated synthetic data is published online [38].

Eleven feature selection methods are evaluated, including six uplift filter methods (F filter, LR filter, KL filter, Chi filter, ED filter, NIV filter), two uplift embedded methods (*TwoModel* embedded and KL embedded), two ordinary filter methods (ordinary Mutual Information and ordinary F Score [25]), and one ordinary embedded method *Outcome embedded* as a benchmark [25]. For F filter and LR filter, variants (F, F2, F3, LR, LR2, LR3) are created based on  $R =$

1, 2, 3. For the bin-based filter methods, the number of bins is set at 10 as default. Four uplift models are used to evaluate the performance of the feature selection methods: *TwoModel*, *X-Learner*, *R-Learner*, and uplift random forest (with KL criterion), given the results are similar in the order of feature selection method performance, only *TwoModel* results are presented in the paper. A random forest classifier ( $n\_estimators = 10, max\_depth = 10, min\_child\_samples = 100$ ) [25] is used as the base learner.

Each simulation trial consists of four steps. First, the data generator simulates the data with a new random seed and randomly splits the data into training and testing (with 50% : 50% ratio). Second, each feature selection method is applied to the training data and ranks the features by importance score. Third, for each feature selection method, the top  $m^*$  (for  $m^* \in \{2, 4, 6, 8, 10\}$ ) features are used to build uplift models using training data. Fourth, the testing data is used to evaluate the accuracy of the uplift models based on the top features selected by each feature selection method. Each trial generates 20,000 samples (10,000 for control and 10,000 for treatment) and the simulation study consists of  $t = 50$  trials.

**Table 1.** RMSE (and Standard Deviation of RMSE) for estimating ITE using *TwoModel* as evaluating model based on Synthetic Data (RMSE lower the better).

Top Features	2	4	6	8	10
F filter	0.183(0.005)	0.174(0.005)	0.174(0.006)	0.174(0.006)	0.174(0.007)
F2 filter	0.184(0.004)	0.163(0.003)	0.165(0.005)	0.166(0.006)	0.168(0.005)
F3 filter	0.185(0.004)	0.163(0.003)	0.163(0.005)	0.165(0.006)	0.167(0.005)
LR filter	0.184(0.005)	0.175(0.006)	0.174(0.006)	0.176(0.006)	0.175(0.006)
LR2 filter	0.184(0.004)	0.163(0.003)	0.166(0.005)	0.168(0.006)	0.170(0.005)
LR3 filter	0.186(0.004)	0.163(0.003)	0.165(0.005)	0.167(0.005)	0.168(0.005)
KL filter	0.184(0.004)	0.163(0.005)	0.153(0.006)	0.162(0.006)	0.165(0.005)
ED filter	0.184(0.003)	0.162(0.004)	0.151(0.004)	0.157(0.006)	0.160(0.005)
Chi filter	0.184(0.004)	0.163(0.005)	0.154(0.007)	0.162(0.006)	0.165(0.005)
NIV filter	0.184(0.004)	0.163(0.005)	0.152(0.006)	0.158(0.005)	0.162(0.005)
KL embedded	0.187(0.005)	0.171(0.011)	0.165(0.009)	0.165(0.008)	0.167(0.007)
TwoModel embedded	0.210(0.007)	0.198(0.010)	0.191(0.008)	0.184(0.007)	0.181(0.006)
Outcome embedded	0.212(0.006)	0.207(0.009)	0.199(0.009)	0.190(0.009)	0.185(0.008)
Ordinary FScore	0.208(0.008)	0.195(0.011)	0.185(0.008)	0.181(0.003)	0.181(0.006)
Ordinary MutualInfo	0.211(0.006)	0.207(0.009)	0.202(0.010)	0.198(0.010)	0.196(0.010)

A preferred feature selection method is expected to provide important features that lead to accurate ITE estimation by the downstream uplift model based on the features selected. Table 1 summarizes the RMSE (Root Mean Square Error) of ITE estimation and its standard deviation over synthetic data trials. As a benchmark, if all 36 features are used in an uplift model, then the RMSE by *TwoModel* is 0.186 (with standard deviation 0.003).

**Table 2.** Proportion of Uplift Features Selected in Top 6 Positions by Method (Feature Recall). The table is ranked by the 'All Uplift' column, that indicates proportion of all uplift features (6 in total) being captured in the top 6 features ranked by each method. A breakdown of feature recall score by different uplift feature pattern is presented.

Method	All Uplift	Linear	Quadratic	Cubic	ReLU	Sin	Cos
F filter	55%	100%	12%	100%	100%	6%	12%
F2 filter	69.3%	100%	100%	100%	100%	4%	12%
F3 filter	70.3%	100%	100%	100%	100%	10%	12%
LR filter	55%	100%	10%	100%	100%	6%	14%
LR2 filter	69.3%	100%	100%	100%	100%	6%	10%
LR3 filter	69.3%	100%	100%	100%	100%	6%	10%
KL filter	97.7%	100%	100%	90%	100%	98%	98%
ED filter	99.7%	100%	100%	98%	100%	100%	100%
Chi filter	95.3%	100%	100%	80%	98%	98%	96%
NIV filter	98.3%	100%	100%	90%	100%	100%	100%
KL embedded	75.7%	94%	88%	82%	86%	42%	62%
TwoModel embedded	32.7%	66%	8%	22%	92%	6%	2%
Outcome embedded	23%	66%	20%	8%	42%	0%	2%
Ordinary FScore	38.7%	96%	0%	46%	88%	2%	0%
Ordinary MutualInfo	20%	24%	22%	14%	20%	12%	28%

The results show that 1) the ordinary feature selection methods (ordinary F Score and ordinary Mutual Information) have poor performance and are not suitable for selecting top features for uplift modeling; 2) the bin-based divergence filter methods (KL filter, ED filter, Chi filter, NIV filter) have consistent top performance in all scenarios, followed by KL embedded, and F2, F3, LR2, LR3 filters; 3) adding higher-order terms in F filter and LR filter ( $R > 1$ ) improved the performance; 4) the outcome embedded method and the *TwoModel* embedded method have essentially similar logic as ordinary feature selection methods, thus failing in the tasks.

To better understand what features are selected by each method, we report the proportion of uplift features selected in top 6 positions in Table 2. The proportion is averaged across the 100 trials. The detailed breakdown by uplift feature type explains why some methods are not performing well. F filter and LR filter fail to capture quadratic features, Sin features, and Cos features. Setting  $R > 1$  for these two methods improved the capability for capturing quadratic features. The ordinary filter methods and *Outcome* embedded and *TwoModel* embedded methods fail to capture uplift features in general.

## 5.2 Experiment 2: Evaluation with MegaFon Uplift Competition Data

This example uses a data set that is publicly available for an online competition [1], that contains 600,000 observations and 50 features.

AUUC based on Top 20 Features

Feature Selection Method	KL-RF	TwoModel-LR	TwoModel-RF	TwoModel-LGBM	TwoModel-NN	X-Learner-LR	X-Learner-RF	X-Learner-LGBM	X-Learner-NN	R-Learner-LR	R-Learner-RF	R-Learner-LGBM	R-Learner-NN
All Features	103.1%	57.8%	115.8%	138.9%	142.9%	61.7%	131.2%	137.4%	144.5%	64.0%	123.4%	137.0%	123.9%
F1 filter	77.2%	59.5%	83.6%	92.9%	109.3%	59.0%	87.9%	94.4%	109.8%	60.7%	78.4%	91.5%	67.4%
F2 filter	107.5%	65.4%	114.5%	134.7%	141.0%	64.2%	128.6%	133.7%	138.5%	64.1%	121.8%	133.4%	136.8%
F3 filter	104.5%	65.4%	113.7%	134.7%	140.9%	64.2%	128.3%	133.7%	138.7%	64.1%	122.4%	133.6%	136.9%
LR1 filter	77.9%	57.0%	86.4%	95.2%	112.1%	57.8%	90.1%	96.1%	111.7%	61.0%	82.1%	93.6%	96.5%
LR2 filter	104.8%	58.3%	112.7%	132.4%	135.5%	58.0%	126.5%	131.6%	135.4%	62.9%	119.5%	131.1%	131.0%
LR3 filter	104.5%	58.3%	112.4%	132.4%	134.9%	58.0%	126.6%	131.6%	134.3%	63.0%	119.0%	131.2%	133.4%
KL filter	104.6%	56.1%	112.6%	132.8%	137.3%	57.6%	126.7%	131.6%	134.1%	62.2%	119.7%	131.3%	132.9%
ED filter	106.5%	63.4%	112.9%	134.0%	139.5%	62.4%	127.3%	132.8%	136.7%	62.4%	121.6%	132.2%	138.7%
Chi filter	102.3%	56.1%	112.1%	132.8%	136.8%	57.6%	126.6%	131.6%	138.1%	62.2%	120.0%	131.2%	125.6%
NIV filter	106.0%	56.1%	112.7%	132.8%	136.2%	57.7%	126.7%	131.6%	136.4%	62.2%	119.6%	131.4%	125.3%
KL embedded	107.3%	55.1%	115.3%	136.2%	143.1%	59.5%	129.9%	134.6%	142.9%	63.7%	123.5%	134.3%	142.1%
TwoModel embedded	100.2%	60.7%	109.1%	125.3%	128.3%	59.0%	119.7%	124.7%	129.9%	60.6%	112.1%	124.0%	126.0%
Outcome embedded	84.4%	56.5%	90.0%	94.3%	107.5%	54.6%	92.8%	95.9%	103.4%	56.7%	83.8%	94.5%	96.8%
Ordinary FScore	78.7%	59.2%	88.0%	99.1%	118.5%	57.3%	94.1%	99.8%	119.0%	61.5%	85.7%	97.1%	90.5%
Ordinary MutualInfo	91.0%	57.1%	97.6%	110.2%	121.1%	57.3%	103.4%	109.8%	122.4%	58.2%	94.6%	108.6%	96.5%

**Fig. 2.** AUUC of Uplift Models Using the Top 20 Selected Features by Different Feature Selection Methods in the MegaFon Data Experiment with 300,000 training data and 300,000 testing data (AUUC larger the better). The Y-axis shows the feature selection method and the X-axis shows the uplift model (in a {uplift model - base learner} format) used for producing the AUUC score. The first row represents the AUUC using all 50 features as reference.

To test the generality of the feature selection methods for different uplift models, the uplift models considered include: (1) *TwoModel-LR*, *X-Learner-LR*, *R-Learner-LR* using { Logistic Regression Classifier & Linear Regression Regressor } as base learners; (2) *TwoModel-LGBM*, *X-Learner-LGBM*, *R-Learner-LGBM* using { Gradient Boosting Classifier & Gradient Boosting Regressor } from *LightGBM* implementation [16] as base learners, with hyperparameter values ( $n\_estimators = 100, max\_depth = 10, min\_child\_samples = 100$ ); (3) *TwoModel-RF*, *X-Learner-RF*, *R-Learner-RF* using { Random Forest Classifier & Random Forest Regressor } as base learners, with hyper-parameter values ( $n\_estimators = 100, max\_depth = 10, min\_samples\_leaf = 100$ ); (4) *TwoModel-NN*, *X-Learner-NN*, *R-Learner-NN* using { Neural Network (Multi-layer Perceptron) Classifier & Neural Network Regressor } from scikit-learn [26] implementation as base learners: each contains two layers with 100 neurons on the first layer and 10 neurons on the second layer; (5) *KL-RF* as the uplift random forest using KL divergence criterion with hyper-parameter values ( $n\_estimators = 20, max\_depth = 10, min\_samples\_leaf = 100$ ).

The evaluation with this data is conducted by equally splitting the data into training data and testing data, each with 300,000 observations, where training data is used for feature selection and uplift modeling training, and testing data is used for evaluating the accuracy of uplift modeling. The results are summarized in Figure 2, reporting the AUUC (area under the uplift curve) scores [32,37,30,13] from the uplift models using the top 20 features selected by each feature selection

method. The relative performance of different feature selection methods can be compared within each column given the same uplift model. The results are roughly aligned with the synthetic data results, except the performance of F2, F3, LR2, LR3 filters become comparable with other bin-based divergence methods. A possible reason is that this dataset does not have extreme nonlinear feature patterns like the sine or cosine forms in the previous synthetic data. Despite the differences in uplift models, the relative order of feature selection method performance is quite consistent across different uplift models. Models using all 50 features still perform better than models using 20 features. One reason is that the sample size is large compared with the feature space, in which case over-fitting is unlikely to occur.

## 6 Discussion

Our experiments demonstrate that the proposed methods can select important features based on their association with heterogeneous treatment effects and outperform ordinary feature selection methods for uplift modeling tasks. The F filter and LR filter can enhance performance by adding nonlinear terms, and such methods can perform well in noisy data settings. The bin-based divergence filter methods also show good performance, although their performance is expected to suffer if the feature space makes it hard to create equal-sized bins. The embedded method with uplift random forest also perform competitively because the split criterion for such trees directly models HTE in child nodes. This approach, however, is more time-consuming compared to the other variable selection methods because it involves uplift model training. Additionally, its performance depends on the uplift model’s hyperparameters.

## 7 Conclusion

We started by hypothesizing that traditional variable selection methods are not suitable for uplift models. We developed a number of alternative approaches that took into account the unique character of uplift models, namely the fact that the target in such models is the treatment effect. We then tested our proposed methods in a number of experiments, establishing that they conclusively outperform traditional variable selection approaches. We further highlighted different strengths and weaknesses of the proposed methods.

Given the increasing use of uplift model in industry applications and elsewhere, we believe the proposed methods to prove practically relevant for several purposes, including accuracy, efficiency and interpretability. For this reason, we have made the proposed methods publicly available in the *Causal ML* open source library. Promising areas for future development include covering more of the ever-increasing group of uplift models, applying the methods to observational in addition to experimental data, and integrating the methods with emerging model interpretability approaches.

The future work may include a detailed empirical study of feature selection methods for continuous outcome, as well as discussing feature selection methods for selecting the best set of features (collaboratively instead of individually).

## References

1. Megafon Uplift Competition (2021), <https://ods.ai/tracks/df21-megafon/competitions/megafon-df21-comp/data>
2. Athey, S., Imbens, G.: Recursive partitioning for heterogeneous causal effects (Apr 2015)
3. Athey, S., Tibshirani, J., Wager, S.: Generalized random forests (Oct 2016)
4. Bolón-Canedo, V., Sánchez-Marroño, N., Alonso-Betanzos, A.: A review of feature selection methods on synthetic data. *Knowledge and information systems* **34**(3), 483–519 (2013)
5. Bommert, A., Sun, X., Bischl, B., Rahnenführer, J., Lang, M.: Benchmark for filter methods for feature selection in high-dimensional classification data. *Computational Statistics and Data Analysis* **143**, 106839 (2020)
6. Chandrashekar, G., Sahin, F.: A survey on feature selection methods. *Computers & Electrical Engineering* **40**(1), 16–28 (2014)
7. Chen, H., Harinen, T., Lee, J.Y., Yung, M., Zhao, Z.: Causalml: Python package for causal machine learning. *arXiv preprint arXiv:2002.11631* (2020)
8. Chen, X., Liu, Z., Yu, L., Yao, L., Zhang, W., Dong, Y., Gu, L., Zeng, X., Tan, Y., Gu, J.: Imbalance-aware uplift modeling for observational data (2022)
9. Dahl, G.E., Sainath, T.N., Hinton, G.E.: Improving deep neural networks for lvcsr using rectified linear units and dropout. In: 2013 IEEE international conference on acoustics, speech and signal processing. pp. 8609–8613. IEEE (2013)
10. Grimmer, J., Messing, S., Westwood, S.J.: Estimating heterogeneous treatment effects and the effects of heterogeneous treatments with ensemble methods. *Polit. Anal.* **25**(4), 413–434 (Oct 2017)
11. Guelman, L., Guillén, M., Pérez-Marín, A.M.: Random forests for uplift modeling: An insurance customer retention case. In: *Modeling and Simulation in Engineering, Economics and Management*. pp. 123–133. Springer Berlin Heidelberg (2012)
12. Guelman, L., Guillén, M., Pérez-Marín, A.M.: Uplift random forests. *Cybernetics and Systems* **46**(3-4), 230–248 (2015)
13. Gutierrez, P., Gerardy, J.Y.: Causal inference and uplift modeling a review of the literature. *JMLR: Workshop and Conference Proceedings* **67** (2016)
14. Hansotia, B., Rukstales, B.: Incremental value modeling. *Research Council Journal* (2001)
15. Holland, P.W.: Statistics and causal inference. *J. Am. Stat. Assoc.* **81**(396), 945–960 (1986)
16. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.Y.: Lightgbm: A highly efficient gradient boosting decision tree. In: *Advances in neural information processing systems*. pp. 3146–3154 (2017)
17. Kohavi, R., Tang, D., Xu, Y.: *Trustworthy online controlled experiments: A practical guide to a/b testing*. Cambridge University Press (2020)
18. Künzel, S.R., Sekhon, J.S., Bickel, P.J., Yu, B.: Meta-learners for estimating heterogeneous treatment effects using machine learning (Jun 2017)
19. Larsen, K.: Data exploration with weight of evidence and information value in r (2015)

20. Lo, V.S.: The true lift model: a novel data mining approach to response modeling in database marketing. *ACM SIGKDD Explorations Newsletter* **4**(2), 78–86 (2002)
21. Mouloud, B., Olivier, G., Ghaith, K.: Adapting neural networks for uplift models. arXiv preprint arXiv:2011.00041 (2020)
22. Neyman, J.: Sur les applications de la théorie des probabilités aux expériences agricoles: Essai des principes. *Roczniki Nauk Rolniczych* **10**, 1–51 (1923)
23. Nie, X., Wager, S.: Quasi-Oracle estimation of heterogeneous treatment effects (Dec 2017)
24. Olaya, D., Coussement, K., Verbeke, W.: A survey and benchmarking study of multitreatment uplift modeling. *Data Mining and Knowledge Discovery* **34**(2), 273–308 (2020)
25. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
26. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al.: Scikit-learn: Machine learning in python. *Journal of machine learning research* **12**(Oct), 2825–2830 (2011)
27. Radcliffe, N.J., Surry, P.D.: Real-world uplift modelling with significance-based uplift trees. White Paper TR-2011-1, Stochastic Solutions pp. 1–33 (2011)
28. Rubin, D.B.: Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* **66**(5), 688–701 (1974)
29. Rubin, D.B.: Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association* **100**(469), 322–331 (2005)
30. Rzepakowski, P., Jaroszewicz, S.: Decision trees for uplift modeling with single and multiple treatments. *Knowl. Inf. Syst.* **32**(2), 303–327 (Aug 2012)
31. Shortreed, S.M., Ertefaie, A.: Outcome-adaptive lasso: variable selection for causal inference. *Biometrics* **73**(4), 1111–1122 (2017)
32. Sołtys, M., Jaroszewicz, S., Rzepakowski, P.: Ensemble methods for uplift modeling. *Data Min. Knowl. Discov.* **29**(6), 1531–1559 (Nov 2015)
33. Tang, J., Alelyani, S., Liu, H.: Feature selection for classification: A review. *Data classification: algorithms and applications* p. 37 (2014)
34. Teinmaa, I., Albert, J., Goldenberg, D.: Uplift modeling: From causal inference to personalization (2021)
35. Wager, S., Athey, S.: Estimation and inference of heterogeneous treatment effects using random forests (Oct 2015)
36. Zaniewicz, L., Jaroszewicz, S.: Support vector machines for uplift modeling. In: 2013 IEEE 13th International Conference on Data Mining Workshops. pp. 131–138 (Dec 2013)
37. Zhao, Y., Fang, X., Simchi-Levi, D.: Uplift modeling with multiple treatments and general response types (May 2017)
38. Zhao, Z.: Synthetic Data for Uplift Modeling and Heterogenous Treatment Effect with Known Counterfactuals and ITE (Mar 2022), <https://doi.org/10.5281/zenodo.6342553>
39. Zhao, Z., Harinen, T.: Uplift modeling for multiple treatments with cost optimization. arXiv preprint arXiv:1908.05372 (2019)