Proceedings of the 53rd Hawaii International Conference on System Sciences | 2020

Towards More Robust Uplift Modeling for Churn Prevention in the Presence of Negatively Correlated Estimation Errors

Frank Oechsle
esentri AG, Pforzheimer Str. 132
76275 Ettlingen, Germany
frank.oechsle@esentri.com

David W. Schönleber
esentri AG, Pforzheimer Str. 132
76275 Ettlingen, Germany
david.schoenleber@esentri.com

Abstract

The subscription economy is rapidly growing, boosting the importance of churn prevention. However, current true lift models often lead to poor outcomes in churn prevention campaigns. A vital problem seems to lie in instable estimations due to dynamic surrounding parameters such as price increases, product migrations, tariff launches of a competitor, or other events with uncertain consequences. The crucial challenge therefore is to make churn prevention measures more reliable in the presence of game-changing events. In this paper, we assume such events to be spatially finite in feature space, an assumption which leads to particularly bad churn prevention results if the selected customers lump in an affected region of the feature space. We then introduce novel methods which trade off uplift for reduced similarity in feature space when selecting customers for churn prevention campaigns and show that these methods can improve the robustness of uplift modeling.

1. Introduction

Referring to McKinsey's survey of US shoppers "Thinking inside the subscription box: New research on e-commerce consumers" [1], "the subscription e-commerce market has grown by more than 100 percent a year over the past five years. The largest such retailers generated more than \$2.6 billion in sales in 2016, up from a mere \$57.0 million in 2011". This survey was carried out in the end of the year 2017 and was published in the beginning of the year 2018. A similar development for the German subscription market is described by billwerk in their 2019 published white paper "subscription based services" [2]. They highlight that the revenues of German vendors of subscription-based services since 2015 are exponentially growing by more than 100 percent per year. At any rate, the subscription business is an economy gaining in importance, and is after the big successes in North America now conquering the European market. In consequence, churn management, and with it the subdomain churn prevention, will become of paramount prominence.

Yet, state-of-the-art uplift models often lead to poor outcomes in churn prevention campaigns, like any other common churn prevention approach as well. The crucial question thus is how to do churn prevention in a more reliable way, i.e., in a way that the benefit of a campaign is more probable. A churn managing company basically would like to know how each of their customers will react when being targeted within the scope of a churn prevention campaign such as a phone call with a specific contract renewal offer, in comparison to their behaviour when they are not targeted. At this context "reaction" or "behaviour" means in particular to announce churn or not to announce churn.

The underlying challenge is thus to predict the probability of a customer to announce churn, depending on the participation in a (specific) churn prevention campaign. This probability consists of two different probabilities, namely the probability of churning without being contacted and the probability of churning when being contacted. Even if only one of the probabilities is notably misestimated, the success of the whole churn prevention campaign is in danger.

The issue of estimating these probabilities is further aggravated by the rarity of churn *per se*, which implies that successful churn prevention cases are even rarer. Accordingly, the estimation of the probability of customers that can be successfully prevented from announcing churn when receiving an appropriate measure is both challenging and crucial, since failure provokes the opposite of the aimed target. Thus even partial failure in estimating churn probabilities can lead to increased churn rates, which is eminently adverse since it is much more expensive to acquire new customers than retaining the inventory customers [3]. Consequently, we need an approach that ensures the absence of failure as far as possible while still realizing existing chances of churn reduction.

URI: https://hdl.handle.net/10125/63931
978-0-9981331-3-3

In this paper, we explore the effect of game-changing events such as tariff launches of a competitor on uplift modeling. We assume these events to be spatially finite in feature space and evaluate different customer selection methods based on decision trees via Monte Carlo simulations, including novel selection methods which trade off uplift against more diversity in feature space and prove to be more robust in the presence of such events.

Note that when we use the term *churn* in this paper, we mean *churn announcement*. Only if a company's efforts in retaining the churn announcing customers are of no avail this results in churn. At this point, churn management is divided into prevention and retention. We clearly concentrate on churn prevention in this paper.

2. Related work

The basic theory underlying Lo's true lift model [4] is quite intuitive and well-defined, but surprisingly does not reliably succeed in the churn context. The definition according to Table 1 is neat and in essence considers the incremental effect a campaign has on the selected customers, whereas in the context of a traditional response model the focus is only on the response after the campaign (treatment) devoid of checking what it would have been without the campaign (control), that is maximizing the difference $A - C$. For instance A, B, C and D could denote the probability of purchase or churn for the corresponding customer segment.

Table 1. Definition of true lift, following Lo [4].

	Treatment	Control	Increment
Model-guided	A	B	A-B
Unguided	C	D	C-D
Difference	A-C	B-D	(A-B)-(C-D)

The true lift approach consequently results in a different selection method compared to the classic response model in that it selects by the delta of the customers churn probabilities when treated (= received the campaign treatment) or untreated (= not received the campaign treatment), that is the uplift. The uplift is calculated as

$$\Delta = p_0 - p_1, \quad (1)$$

where p_0 and p_1 are the churn probabilities without respectively with treatment. The selection method according to the uplift is illustrated in Figure 1. There, the lower-right corner represents the optimal point for selection (highest churn probability if untreated but lowest if treated, i.e. maximal Δ). The lower the value of the Δ selection threshold (e.g. due to higher available budget), the more customers are covered in the campaign treatment.

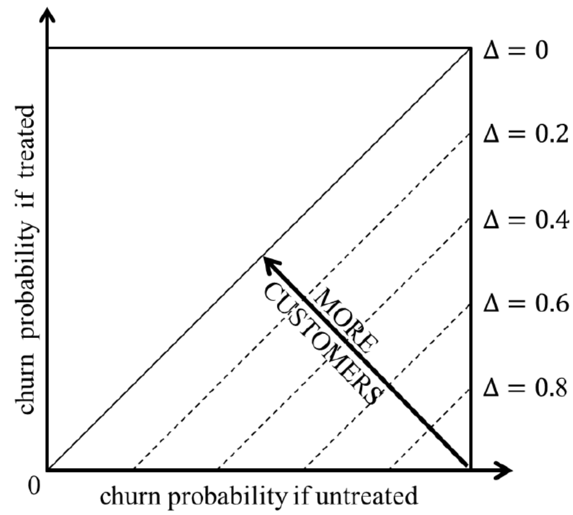

Figure 1. Graphical illustration of selection by delta from Oechsle et al. [5].

Even though the available research relating to uplift modeling is not inexhaustible, there is definitely adequate knowledge about predicting those increments used in Lo's true lift model. Kane et al. [6] depict this in a comprehensive way as well as Guelman et al. [7] do. Hence, the direct prediction of the difference between the probabilities in particular seems to be a mature approach. The leading alternative would be to predict the two probabilities separately, building the uplift afterwards via subtraction.

Nevertheless, empirical results are not satisfying. This picture emerges by for example comparing the corresponding work of Chickering and Heckerman [8], Hansotia and Rukstales [9], Manahan [10], Rzepakowski and Jaroszewicz [11], Radcliffe and Surry [12], Rzepakowski and Jaroszewicz [13] or Zaniewicz and Jaroszewicz [14]. The nonexistent established success stories in practice encourage this point of view.

Consequently, the focus of the most recent studies is uncertainty and estimation errors as a central root cause for the observed phenomenon of missing best practices in terms of churn prevention via uplift modeling. While Lo [15] accounts for the variability in estimates in a marketing context, Oechsle et al. [5] address estimation risks in the subscription business when it comes to churn. Athey and Imbens [16] analyse in general the increasing uncertainty in estimations arising from a smaller sample size. As they deal with decision trees, they propose to use different subsamples for splitting and estimating and try to control their results via confidence intervals. This approach appears not intuitive in so far as it multiplies the initial problem of small leaf size by further splits.

In the remainder of the paper, we pick up this latter stream and investigate the effect of suddenly upcoming estimation errors due to moving environments in the subscription business. We thereby contribute insight concerning the question why churn prevention all too often fails and how it can be done better.

3. Negatively correlated estimation errors due to uncontrollable events

The above mentioned moving environment accumulates dynamic surrounding parameters, which can be on the one hand company-intern changes such as mandatory price increases or product migrations owing to technical improvement, and on the other hand external factors like tariff launches of a competitor or other specific events influencing customer groups in undetermined ways. This fluctuation leads to instable estimations and suboptimal decisions at least. In the worst case, these dynamic parameters can lead to negatively correlated estimation errors within homogenous customer segments.

Consider, for instance, a company whose portfolio includes the cheapest tariff for the entire branch linked with a competitive common service package, and therefore successfully contracted bargain hunters. Consider further that it is suddenly confronted with a competitor launching an even cheaper tariff of adequate quality. In this case it is possible that the just very loyal (low p_0, high p_1) bargain hunters abruptly turn to unfaithful (high p_0, low p_1), change-oriented customers. Let now the prevention campaign process be at a point where this could not be recognized and incorporated any more, then it ends up in probability estimations that still pretend loyal bargain hunters but rather belong to potential churners. In other words underestimated p_0 and overestimated p_1, that is, negatively correlated estimation errors.

Such game-changing events can burst in on a prevention campaign at all times and cannot always be anticipated nor reliably excluded by, for example, smart definition of the target group of the prevention campaign. Their effect can only reasonably be assumed to be spatially finite, i.e. locally bounded with respect to their diffusion in feature space. The most detrimental impact of these events occurs when the customers that seem to be the most promising (and which are therefore selected for treatment) lump in a certain region of feature space which is affected by the upcoming game-changing event in such a way that the overall treatment effect is reversed. Consequently, targeting churn prevention measures at a customer group characterized by their similarity in feature space poses a potential threat to the success of these measures and should hence be penalized or avoided.

In this paper, we evaluate the effect of such game-changing events via Monte Carlo simulations and propose alternative strategies for customer selection, which are more robust to aforementioned events. To this end, we consider a two-dimensional feature space, sketched in Figure 2a). Note that the features are assumed to be normalized, such that the range of the feature values fall within the range $[0, 1]$. This feature space is then divided into rectangles, which is the general concept of decision trees, and afterwards, likewise typical for tree-based methods, each of the rectangles sustains a constant (often between 0 and 1 reflecting a relative frequency) representing the "model" according to the leaf. We further assume a game-changing event to have significant impact on customers at a specific, randomly selected point (E) in feature space as well as a distance-dependent impact on surrounding customers (or users U) that is spatially confined within a distance R from E.

More specifically, assume Δ to be the real and correctly estimated effect of the churn prevention campaign per customer in the absence of the upcoming game-changing event. For a customer with distance r to the center E of the game-changer, we calculate the modified effect Δ' of the churn prevention campaign in the presence of the game-changer according to

$$\Delta' = \begin{cases} \Delta & r > R, \\ \Delta\left[1 - 2\cos\left(\frac{\pi r}{2R}\right)\right], & r \leq R \end{cases} \qquad (2)$$

where R is the radius of the circle of influence of the game-changer and r specifies the Euclidean distance of the customer U to the center E of the game-changer. Hence, the effect of a game-changer is maximal for customers close-by while it decreases radially up to a distance R, where its effect vanishes, following the above cosine behaviour. At the center E of the game-changer, the true effect of the churn prevention campaign changes sign. For example, when the real and estimated effect of the prevention campaign would be 0.5 reduction of churn probability, it would switch to -0.5 reduction for a customer lying at the center of the game-changer, but only for the real effect. The estimated effect would still be 0.5.

This modeling is motivated by the fact that Δ is the difference between the churn probability without

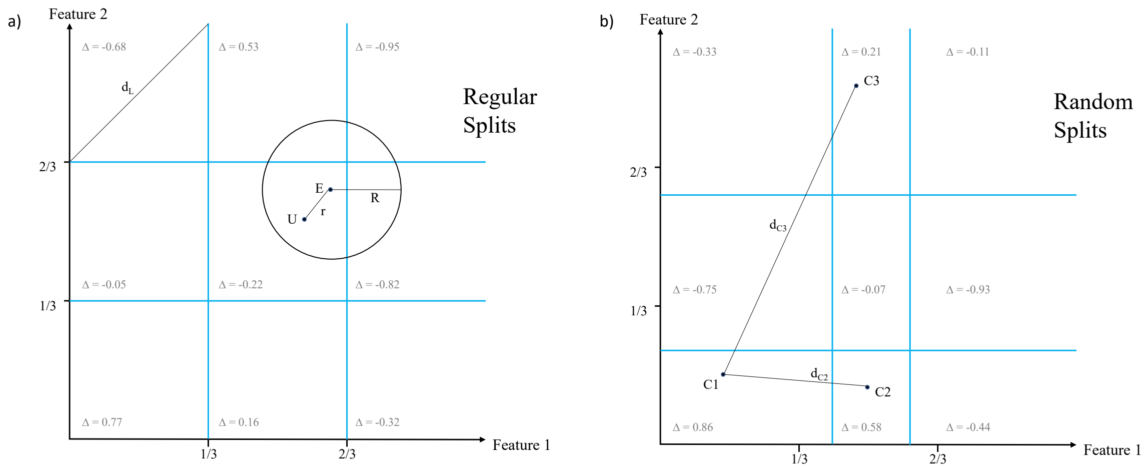

Figure 2. Radial estimation error and linkage via the centroid-method in regular or random decision trees.

treatment p_0 and the churn probability with treatment p_1 [cf. Eq. (1)]. Thus the modeled behaviour can be achieved by overestimating p_0 while underestimating p_1 at the same time and therefore would produce a negatively correlated estimation error. For instance, an initial $\Delta = 0.5$ which changes to $\Delta = -0.5$ in the presence of a game-changing event can be achieved by initially estimating $p_0 = 0.7$, $p_1 = 0.2$ whereas actually $p_0 = 0.2$ and $p_1 = 0.7$ in the presence of a game-changer.

4. Handling similarity

The spatial confinement and construction of the errors as described in the previous section suggests to prefer customers that are not similar to each other, in other words, that are distant in feature space. Since we cannot modify where the individual customer is located in feature space, we focus on handling the similarity of customers by reducing calculated benefits depending on the Euclidean distance between the customers and vary our selection approach using the centroid-method as linkage technique.

A linkage technique is relevant since we use the concept of decision trees in our analysis and thus do not select individual customers but customers aggregated in leaves. For this reason we split the two-dimensional feature space via two continuous splits, vertically and horizontally, resulting in nine rectangles, which cover the feature space and represent nine leaves of a decision tree. In our simulation, we then randomly assign three positive and six negative values to Δ with $|\Delta_i| \in [0, 1]$ for $i = 1, 2, ..., 9$, corresponding to the churn reduction estimated via the decision tree. By dint of the predominantly negative values of Δ we indicate a generally inauspicious base case for churn prevention measures. This is inspired by the absence of track records in churn prevention campaigns. The distances between the cluster centroids that are most relevant in our simulations are the ones between the rectangle (leaf) with the highest Δ and the two other rectangles (leaves) with positive value of Δ. Those two distances d_{C2} (best to second best) and d_{C3} (best to third best) are calculated via the Euclidean distance of the cluster centroids $C1$, $C2$ and $C3$. As $C1$ represents the centroid of the rectangle (leaf) with the highest Δ and $C2$, $C3$ analogical the second and third best leaf, d_{C2} is the Euclidean distance from $C1$ to $C2$ while d_{C3} denotes the Euclidean distance from $C1$ to $C3$. For the calculation of the cluster centroids themselves we use the average of both customer features in the corresponding segment. We furthermore assume the customers to be equally distributed in feature space, and hence these centroids typically do not deviate strongly from the centers of the rectangles (leaves).

The classic way of selection, independent of the business use case (sales, churn, etc.), is to select the best N customers the previously allocated budget is able to fund. It is not seldom that this budget is derived without having an idea about the probabilities and the chances or risks in the specific churn prevention campaign. In this approach, no similarity aspects are taken into account.

In contrast, in this paper, we introduce novel selection methods, which incorporate these similarity aspects by design. We thereby challenge the classic way of selecting customers with several distance-respecting methods. In particular, we introduce the following

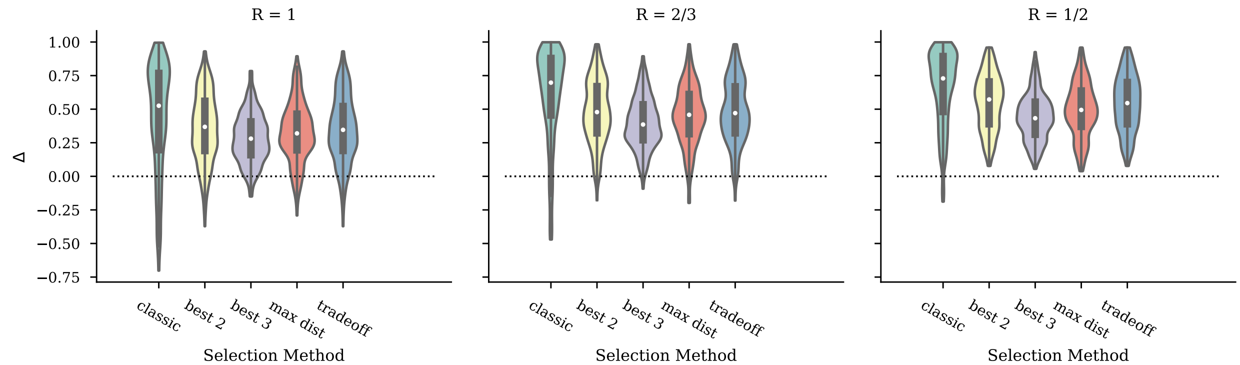

Figure 3. Comparison of selection methods for different radii R (in units of d_L) and regular decision tree splits.

selection methods: *best 2*, *best 3*, *max dist* and *tradeoff*, which are described below. Note that in the now following discussion the leaf represented by the centroid C_i is named leaf i and analogously all variables with index i correspond to leaf i. In addition *best* implies highest value of Δ among all leaves.

best 2 In this method we randomly select $N/2$ customers in the best leaf and $N/2$ customers in the second best leaf. Hence, we automatically trade off uplift against more diversity in feature space.

best 3 Here we randomly select $N/3$ customers in the best leaf, $N/3$ customers in the second best leaf, and $N/3$ customers in the third best leaf. Hence, within the scope of our simulation, we maximally diversify in feature space while trading in even more probability.

max dist In this method we randomly select $N/2$ customers in the best leaf, and $N/2$ customers in the leaf i where the distance to the best leaf is maximal, i.e., $d_i = \max(d_{C2}, d_{C3})$. Thus we prefer distance irrespective of the traded in probability.

tradeoff Here we randomly select $N/2$ customers in the best leaf, and $N/2$ customers in the leaf i which minimizes the quotient

$$q_i = \frac{\Delta_1 - \Delta_i}{d_{Ci}} \tag{3}$$

for $i = 2, 3$, i.e., $q_i = \min(\frac{\Delta_1 - \Delta_2}{d_{C2}}, \frac{\Delta_1 - \Delta_3}{d_{C3}})$. Therefore the focus is simultaneously on both of the parameters.

5. Simulation results

We now benchmark the selection methods introduced above in a Monte Carlo simulation. In this simulation, we consider two different decision tree splits: *regular* and *random*. In the *regular* case illustrated in Fig. 2a), the areas in feature space covered by the leaves are identical, which in practice could result from the constraint of a minimal leaf size. In the *random* case, illustrated in 2b), the decision tree splits are randomly chosen between zero and one, resulting in leaves of varying sizes.

We then randomly and uniformly distribute 225,000 customers in feature space, leading in the regular case to an average of about 25,000 customers per leaf. By adding a customer at the leaf center in each partition i we ensure that independent of the tree splitting there is no empty partition. To evaluate the effect of game-changing events (or negatively correlated estimation errors), we consider three different values for the impact radius R of the circle of influence that is depicted in 2a). At the maximal radius $R = d_L$, the radius coincides with the diagonal of a regular leaf size, and is consequently decreasing for the chosen values of $R = 2/3\, d_L$ and $R = d_L/2$.

For each of the two scenarios, *regular* and *random*, we simulate 250 decision trees and compare the expected Δ values achieved by the different selection methods in the presence of a game changing event. The statistical distributions of the results are depicted in Figs. 3 and 4.

Figure 3 visualizes the simulation results for three different impact radii R and regular decision tree splits. The violin plots depict the distributions of achieved uplift for the selection methods introduced in Section 4 according to the performed simulations. Most notably, while the classic selection method realizes more profitable outcomes overall (as can be seen from

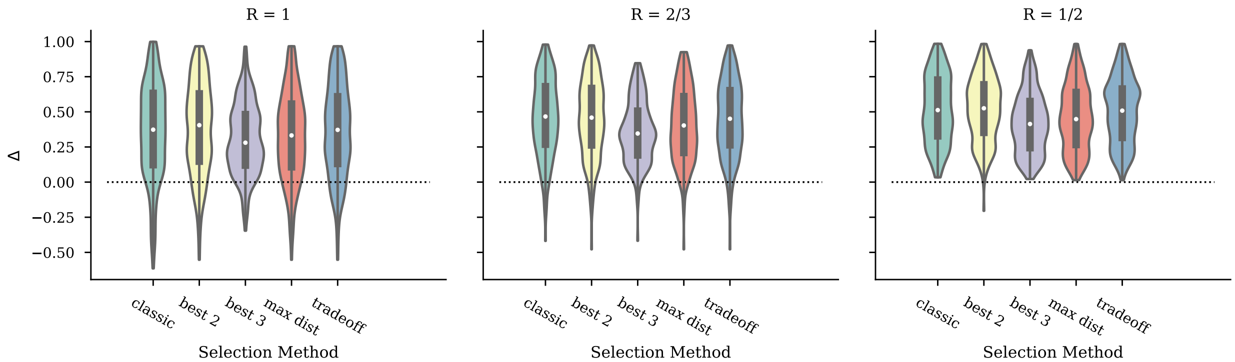

Figure 4. Comparison of selection methods for different radii R (in units of d_L) and random decision tree splits.

the high median value in the miniature boxplot inside the violin plot), this comes at the cost of seldom but significantly adverse negative outcomes. In contrast, more diversifying methods such as *best 2* or *best 3* achieve less favorable results on average, yet also the risk of negative outcomes with high costs are reduced as well. This behavior is most prominent for the largest impact radius $R = d_L$, but persists also for decreasing radii ($R = 2/3\, d_L$ and $R = d_L/2$). The same pattern is found for the *tradeoff* selection, yet this approach is able to reach a comparably low traded-in probability similarly to *best 2* and at the same time reduces failures in a superior way. The last one can be seen more clearly in Table 2 which gives a detailed summary of the simulation results.

For random splits, the benefit from trading off uplift for reduced similarity between customers in feature space disappears, as can be seen from Figure 4. This demonstrates that the architecture of the underlying decision tree in combination with the error rate (frequency and impact) is crucial. We conjecture that in the *random* case the repeatedly strong divergent sizes of leaves and radii of the error diffusion do not ensure a setting where either any churn prevention action is recommendable at all or it is necessary to desist from the classic selection approach.

6. Conclusion and discussion

The results presented in this paper clearly show that there is a notable uncertainty in the efficacy of churn prevention campaigns when spatially occurrent estimation errors are frequent and effectual enough and the general ecosystem is prevention unfriendly (6 out of 9 Δs < 0). Frequent and effectual enough here means that errors happen regularly and that the radius of the error impact has a minimum size relative to the leaf sizes of the decision tree. The latter can be observed in our results by comparing the regular splits scenario versus the random splits scenario. We will analyse the reasons for the blurring effect of random splits on the uplift in more detail in future work.

Despite this problematic setting with large and frequent game changing estimation errors we have illustrated that there are nevertheless methods of trading off uplift for reduced similarity between customers concerning their position in feature space that lead to more robust estimation results in terms of variance. In particular, it is possible in not too toxic ($R \leq d_L$) surroundings to reduce the number of disappointments (churn increasing churn prevention campaigns). To this end, we not only select the customers with the highest Δ ignoring whether they lump in a specific region, but also consider their similarity. Besides some straightforward methods (*best 2*, *best 3*, *max dist*) we also investigated a more elaborated *tradeoff*, namely a method which uses a linearly weighted combination of Δ-difference and Euclidean distance of the cluster centroids for customer selection. Especially this trading off method is promising, since it pays not the highest prices for the targeted risk reduction.

It is clear, however, that we will refine this tradeoff approach in our future work. It should be possible to derive an even superior tradeoff by for example taking a sigmoid relation as a basis, particularly with regard to the assumed cosine behaviour of the error diffusion. In either case, the appropriate rigor for tackling similarity depends on the industry's susceptibility to dynamic changes in customer churn probabilities. It is a parameter that could/should be optimized/learned over time if it could not already be observed in the past.

Another particularly exciting track we will focus on in further research is to directly influence the generation of the decision tree itself. This is to step in the splitting rules as well as to intervene in the pruning of the tree. In this regard it could be clever to perform additional random splits for artificially gaining more distance amongst the leaves, combined by an uplift and distance tradeoff selection method. A novel pruning approach we are currently considering is to aggregate partitions with low similarity respectively high distance in feature space and churn probabilities (or at least uplifts) that are as identical as possible.

We will witness the development of the subscription economy and how it influences the research on uplift modeling in the churn context. The curiosity about dependable solutions for the churn prevention challenge will certainly increase in the face of the current development.

References

[1] T. Chen, K. Fenyo, S. Yang, and J. Zhang, "Thinking inside the subscription box: New research on e-commerce consumers," 2018.

[2] billwerk GmbH, "Subscription based services," 2019.

[3] A. Gallo, "The value of keeping the right customers," 2014.

[4] V. S. Y. Lo, "The true lift model - a novel data mining approach to response modeling in database marketing," *SIGKDD Explorations*, vol. 4, pp. 78–86, 2002.

[5] F. Oechsle, T. Setzer, and S. M. Blanc, "On the assumptions of true lift models for churn prevention," in *Multikonferenz Wirtschaftsinformatik (MKWI) 2016, Technische Universität Ilmenau, 09.-11. März 2016, Band 2. Hrsg.: V. Nissen*, pp. 1233–1244, Universitätsverlag Ilmenau, 2016.

[6] K. Kane, V. Lo, and J. Zheng, "Mining for the truly responsive customers and prospects using true-lift modeling: Comparison of new and existing methods," *Journal of Marketing Analytics*, vol. 2, 12 2014.

[7] L. Guelman, M. Guillen, and A. M. Pérez-Marín, "Uplift random forests," *Cybernetics and Systems*, vol. 46, pp. 230–248, 2015.

[8] D. M. Chickering and D. Heckerman, "A decision theoretic approach to targeted advertising," in *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*, pp. 82–88, Morgan Kaufmann Publishers Inc., 2000.

[9] B. Hansotia and B. Rukstales, "Incremental value modeling," *Journal of Interactive Marketing*, vol. 16, no. 3, p. 35, 2002.

[10] C. Manahan, "A proportional hazards approach to campaign list selection," *SAS User Group International (SUGI) 30 Proceedings*, 2005.

[11] P. Rzepakowski and S. Jaroszewicz, "Decision trees for uplift modeling," in *2010 IEEE International Conference on Data Mining*, pp. 441–450, IEEE, 2010.

[12] N. J. Radcliffe and P. D. Surry, "Real-world uplift modelling with significance-based uplift trees," *White Paper TR-2011-1, Stochastic Solutions*, 2011.

[13] P. Rzepakowski and S. Jaroszewicz, "Uplift modeling in direct marketing," *Journal of Telecommunications and Information Technology*, pp. 43–50, 2012.

[14] L. Zaniewicz and S. Jaroszewicz, "Support vector machines for uplift modeling," in *2013 IEEE 13th International Conference on Data Mining Workshops*, pp. 131–138, IEEE, 2013.

[15] V. S. Lo and D. A. Pachamanova, "From predictive uplift modeling to prescriptive uplift analytics: A practical approach to treatment optimization while accounting for estimation risk," *Journal of Marketing Analytics*, vol. 3, no. 2, pp. 79–95, 2015.

[16] S. Athey and G. Imbens, "Recursive partitioning for heterogeneous causal effects," *Proceedings of the National Academy of Sciences*, vol. 113, no. 27, pp. 7353–7360, 2016.

Table 2. **Summary of simulation results. $\mathbb{E}[\Delta]$ denotes the expectation value, Successes and Failures correspond to simulation runs with $\Delta > 0$ and $\Delta \leq 0$, respectively.**

R/d_L	Splits	Selection Method	$\mathbb{E}[\Delta]$	# Successes	$\mathbb{E}[\Delta]$ Successes	# Failures	$\mathbb{E}[\Delta]$ Failures
1/2	regular	classic	0.66	244	0.67	6	-0.12
1/2	regular	best 2	0.55	250	0.55	0	–
1/2	regular	best 3	0.44	250	0.44	0	–
1/2	regular	max dist	0.49	250	0.49	0	–
1/2	regular	tradeoff	0.54	250	0.54	0	–
1/2	random	classic	0.52	250	0.52	0	–
1/2	random	best 2	0.52	249	0.52	1	-0.20
1/2	random	best 3	0.42	250	0.42	0	–
1/2	random	max dist	0.46	250	0.46	0	–
1/2	random	tradeoff	0.50	250	0.50	0	–
2/3	regular	classic	0.61	230	0.68	20	-0.24
2/3	regular	best 2	0.48	246	0.49	4	-0.07
2/3	regular	best 3	0.40	248	0.40	2	-0.06
2/3	regular	max dist	0.45	244	0.47	6	-0.11
2/3	regular	tradeoff	0.48	247	0.49	3	-0.09
2/3	random	classic	0.46	237	0.49	13	-0.11
2/3	random	best 2	0.45	237	0.48	13	-0.11
2/3	random	best 3	0.36	244	0.37	6	-0.13
2/3	random	max dist	0.41	241	0.43	9	-0.12
2/3	random	tradeoff	0.45	240	0.47	10	-0.13
1	regular	classic	0.44	212	0.57	38	-0.29
1	regular	best 2	0.36	231	0.40	19	-0.12
1	regular	best 3	0.28	239	0.30	11	-0.07
1	regular	max dist	0.33	231	0.36	19	-0.09
1	regular	tradeoff	0.35	235	0.38	15	-0.13
1	random	classic	0.36	220	0.44	30	-0.25
1	random	best 2	0.38	219	0.46	31	-0.16
1	random	best 3	0.31	229	0.35	21	-0.13
1	random	max dist	0.33	215	0.41	35	-0.16
1	random	tradeoff	0.37	216	0.45	34	-0.15