Proceedings of the 2013 Winter Simulation Conference
R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, eds.

EFFICIENT LEARNING OF DONOR RETENTION STRATEGIES FOR THE AMERICAN RED CROSS

Bin Han

Applied Mathematics, Statistics and Scientific Computation
University of Maryland
College Park, MD 20742, USA

Ilya O. Ryzhov

Robert H. Smith School of Business
University of Maryland
College Park, MD 20742, USA

Boris Defourny

Operations Research & Financial Engineering
Princeton University
Princeton, NJ 08544, USA

ABSTRACT

We present a new sequential decision model for adaptively allocating a fundraising campaign budget for a non-profit organization such as the American Red Cross. The campaign outcome is related to a set of design features using linear regression. We derive the first simulation allocation procedure for simultaneously learning unknown regression parameters and unknown sampling noise. The large number of alternatives in this problem makes it difficult to evaluate the value of information. We apply convex approximation with a quantization procedure and derive a semidefinite programming relaxation to reduce the computational complexity. Simulation experiments based on historical data demonstrate the efficient performance of the approximation.

1 INTRODUCTION

The American Red Cross (ARC) and many other disaster relief organizations often encounter the problem of retaining one-time "disaster donors"—people who have given their first donation in response to a major disaster. Only 30% of these donors return to give a second time after their initial contribution. A critical problem for the ARC is to motivate disaster donors to contribute to long-term operations and "post-disaster" programs such as community disaster preparation or emergency response training.

In a typical disaster-donor retention campaign, the ARC sends an appeal for funds (usually a mailed letter) to a segment of disaster donors. The appeal is designed using a combination of features related to the letter format and contents, donor and campaign type, and the use of various gift items. The number of possible designs is exponential in the number of features. The outcome is evaluated using the success rate, defined as the proportion of mailings in a campaign that elicited a donation. The goal is to find the best set of campaign attributes that maximize the success rate as quickly as possible, subject to constraints on the number of campaigns that can be conducted.

Linear regression based on historical data has been used to estimate the effects of the campaign attributes on the success rate (Ryzhov et al. 2013). The results of the estimation, however, are subject to change after every new campaign output, which provides new information about the campaign designs. This potential for information should be used as a criterion when designing the next campaign. The new outcome then feeds back into the linear model and guides the next campaign design, thus forming an "action–feedback"

978-1-4799-2077-8/13/$31.00 ©2013 IEEE

loop. This learning problem possesses two important characteristics: 1) A single campaign outcome may provide information about many other alternatives since different campaign designs can share a common set of features; 2) The variance of the observed campaign success rates is unknown. Although the variance can be estimated from historical data, an inaccurate estimate may bias the evaluation of new information.

Within the fundraising community, there exist numerous statistical models for predicting charitable giving; one example is the Tobit model (Tobin 1958), which is widely used to estimate effects of donor income or taxes on donations (see e.g. Lankford and Wyckoff 1991, Auten and Joulfaian 1996). Other determinants that drive charitable donations have also been studied, including donor demographics (Jones and Posnett 1991), prestige of the organization (Arnett et al. 2003) or certain behavioural factors (Fennis et al. 2009). Generalized linear models have been used to model the effects of attributes of donors and mailings on success rates of campaigns; see Bult et al. (1997) for a small-data analysis and Ryzhov et al. (2013) for a large-scale analysis. Measurement models provide insights on decision and policy making; most of them, however, fail to provide a real-time sequential decision model to guide fundraisers' actions that has the ability to dynamically adapt to new information. While various techniques from operations research—dynamic programming, linear or non-linear programming—have been applied to *for-profit* marketing to support sequential or non-sequential decisions on variables such as price, advertising expenditure or media coverage (see e.g. Lilien et al. 1995, Leeflang and Wittink 2000 for introductions), less attention has been given to developing mathematical models for sequential decision making in non-profit marketing.

In the ranking and selection (R&S) community, however, the problem of sequentially allocating an information budget has been widely studied. The Bayesian approach to R&S, which assigns prior distributions to unknown values, is the most relevant part of this literature for the present study; see Powell and Ryzhov (2012) for a recent and comprehensive review. Such methods include the optimal computing budget allocation (Chen et al. 1996) and the value of information procedure or VIP (Gupta and Miescke 1996). The VIP approach, based on an expected improvement criterion, allows one to handle correlated beliefs, where a single new observation provides information about multiple alternatives (Qu et al. 2012). Negoescu et al. (2011) shows how expected improvement can be used to learn in Bayesian linear regression, but assumes a known sampling variance. We present a Bayesian learning model that integrates the ability to handle linear regression together with the unknown-variance model of Chick et al. (2010). As a result, we capture both the ability to model correlations between regression parameters and the ability to learn an unknown sampling variance simultaneously with those parameters.

The best available algorithm for calculating expected improvement for correlated beliefs (Frazier et al. 2009) has prohibitively high computational complexity in the regression setting. We address this issue in two ways. First, we derive an easier computation in the special case where two-way interaction attributes are absent from the feature space. For more general cases, we apply a quantization procedure and approximate the non-convex value of information by a convex function, and optimize this function using semidefinite programming relaxations (Defourny et al. 2013). Thus, the present paper contributes to the literature on non-profit marketing as well as optimal learning.

The paper is organized as follows. Section 2 discusses the the R&S problem and derives a KG policy to allocate campaign budget. Section 3 describes our computational improvements to the implementation. Section 4 presents numerical experiments on historical data. Section 5 concludes.

2 CAMPAIGN LEARNING MODEL

The campaign decision variables controlled by the ARC mainly fall into three categories: campaign types, mailing designs, and donor segmentation. The campaign types indicate the basic style of the campaign and the general type of donors being targeted, e.g., whether we are reaching out to recent disaster donors or "converted" donors who regularly support the program. The mailing designs indicate detailed information and various items contained in the mailing. The donor segmentation indicates more detailed categorizations of the donors. Table 1 shows a full list of campaign attributes. We use 0 or 1 to denote the decision to include or exclude a certain attribute. For example if the ARC decides to send a batch of mailings with

Table 1: Decision variables in donor retention campaign

Category	Name	Description
Campaign Type	Acquisition	for new acquired one-time disaster donors
	Renewal	for converted donors with direct appeals for a contribution
	Cultivation	for converted donors with newsletter-like mailings
	Follow-up	other intermittent mailings
	Lapsed	for donors who have not responded in the past 18 months
Mailing Design	Personalization	inclusion of the donor's name and address
	Gift item	supporter cards, mailing labels or a glowstick
	Checkboxes	suggested donation amount options
	Stories	disaster preparedness story or generic story
	Online option	option to donate online
Donor Segment	Recency	the lapsed months since last donation (0-6, 7-12, etc)
	Level	the amount of last donation (low, median, high, etc)

an option (or "checkbox") to donate $50, a gift card and a newsletter to donors who have given a small amount within the past six months, we would set the corresponding attributes to 1. Statistical screening analysis, based on over 8 million historical records of communications with donors, has identified a small set of 10-20 key factors driving campaign success from the full list of features (Ryzhov et al. 2013).

2.1 Bayesian Learning

We let $\Phi = \{\varphi \in \{0,1\}^r | A\varphi = h\}$ be the campaign decision space, where r is the number of features and φ is a decision alternative consisting of 0s or 1s, whose vector form can be written as $(\varphi_0, \varphi_1, \ldots, \varphi_{r-1})^\top$. The linear constraints $A\varphi = h$ may come from the fact that there are interactive effects between attributes. For example, the combined effect of a disaster preparedness story with the Renewal campaign type may be greater than the sum of the individual effects. Then, if φ_i and φ_j represent the respective decisions to include a preparedness story and target the Renewal type, we will include an additional feature $\varphi_k = \varphi_i \times \varphi_j$. The condition $\{\varphi_k = \varphi_i \times \varphi_j; \varphi_. \in \{0,1\}\}$ is equivalent to a set of linear constraints, given by

$$\begin{aligned} &\varphi_k \leqslant \varphi_i, \\ &\varphi_k \leqslant \varphi_j, \\ &\varphi_i + \varphi_j - 1 \leqslant \varphi_k, \\ &\varphi_. \in \{0,1\}. \end{aligned} \tag{1}$$

By adding slack variables, we can denote the constraints in (1) as $A\varphi = h$ (see Section 3.2.2 for more discussion).

We let $K = |\Phi|$, which is the total number of alternatives. Notice that K depends exponentially on r if all or most of the attributes are controllable by the decision maker. The budget limits us to conduct N campaigns; at the nth stage we choose an alternative $\varphi^n \in \Phi$. After the campaign a new success rate $y^{n+1} \in (0,1)$ is observed, and we model the outcome-attributes relationship using a linear regression model

$$\eta^{n+1} = \text{logit}(y^{n+1}) = (\varphi^n)^\top \beta + \varepsilon^{n+1},$$

where β is an r-vector specifying the unknown coefficients of the campaign attributes, ε^{n+1} is the measurement noise of y^{n+1}, and 'logit' stands for the link function $\text{logit}(p) = \log(\frac{p}{1-p})$. Notice that the first component of β denotes the intercept, so $\varphi_0^n = 1$, which can also be included in $A\varphi = h$. We use a logit function to map y^{n+1} from $(0,1)$ to $\mathbb{R}$, to allow the response to range anywhere on the real number line. We assume $\{\varepsilon^n\}_{n=1}^N$ are i.i.d. with normal distribution $N(0, \frac{1}{\rho})$, where ρ is the unknown precision parameter.

We use a Bayesian prior to represent the organization's beliefs about the effectiveness of different campaigns. The number of possible campaigns is combinatorially large, but their performance is driven by a small set of r parameters, as well as the unknown precision ρ. The ARC's beliefs can thus be compactly represented by a Bayesian prior on β and ρ. We use a normal-gamma prior, where the conditional distribution (at stage n) of β given ρ is multivariate normal with mean θ^n and covariance matrix $\frac{1}{\rho}\Sigma^n$, and the marginal distribution of ρ is gamma with parameters a^n and b^n:

$$\beta|\rho \sim N(\theta^n, \frac{1}{\rho}\Sigma^n), \quad \rho \sim \Gamma(a^n, b^n).$$

Lemma 1 The posterior distribution after sampling y^{n+1} is still normal-gamma with updated parameters

$$\theta^{n+1} = \theta^n + \frac{\eta^{n+1} - (\varphi^n)^\top \theta^n}{1 + (\varphi^n)^\top \Sigma^n \varphi^n} \Sigma^n \varphi^n; \tag{2}$$

$$\Sigma^{n+1} = \Sigma^n - \frac{\Sigma^n \varphi^n (\varphi^n)^\top \Sigma^n}{1 + (\varphi^n)^\top \Sigma^n \varphi^n}; \tag{3}$$

$$a^{n+1} = a^n + \frac{1}{2}; \tag{4}$$

$$b^{n+1} = b^n + \frac{(\eta^{n+1} - (\varphi^n)^\top \theta^n)^2}{2(1 + (\varphi^n)^\top \Sigma^n \varphi^n)}. \tag{5}$$

After N campaigns we choose the alternative with the largest posterior mean. Our goal is to learn efficiently by choosing a policy π for designing campaigns to maximize the expected value of the final implementation decision. That is,

$$\pi^* = \arg\sup_{\pi \in \Pi} \mathbb{E}^\pi [\max_{\varphi \in \Phi} (\varphi^\top \theta^N)]. \tag{6}$$

2.2 KGUP Algorithm

The knowledge gradient (KG) algorithm (or VIP procedure) is a one-step look-ahead policy that allocates the simulation budget to alternatives that are believed to have the highest potential for improving the expected reward. It solves a simpler form of problem (6), assuming that at each step of simulation we only have one more observation to collect. We let $\mathscr{F}^n$ be the σ-algebra generated by φ^0, y^1, φ^1, y^2, ..., φ^{n-1}, y^n, i.e., all decisions and observations before stage n, and $\{\mathscr{F}^n\}_{n=0}^N$ be the filtration consisting of the sequence of σ-algebras. The KG quantity for choosing $\psi \in \Phi$ at stage n is defined as

$$v_\psi^{KG,n} = \mathbb{E}^n \left[\max_{\varphi \in \Phi} (\varphi^\top \theta^{n+1}) | \varphi^n = \psi \right] - \max_{\varphi \in \Phi} (\varphi^\top \theta^n), \tag{7}$$

where $\mathbb{E}^n[\cdot] = \mathbb{E}[\cdot|\mathscr{F}^n]$, i.e., the conditional expectation taken with respect to $\mathscr{F}^n$.

To compute the KG quantity in (7), it is necessary to find the conditional distribution of θ^{n+1} given $\mathscr{F}^n$ and $\varphi^n = \psi$, which is also known as the *predictive distribution*. Notice from (2) that at stage n the randomness in θ^{n+1} only comes from η^{n+1}. All other quantities in (2)-(5) are known at time n.

Lemma 2 The predictive distribution of η^{n+1} given $\mathscr{F}^n$ and $\varphi^n = \psi$ is a univariate Student's t-distribution with mean $\psi^\top \theta^n$, variance $b^n(1 + \psi^\top \Sigma^n \psi)/a^n$ and $2a^n$ degrees of freedom:

$$\eta^{n+1} \sim t\left(2a^n, \psi^\top \theta^n, \frac{b^n(1 + \psi^\top \Sigma^n \psi)}{a^n}\right).$$

By Lemma 2 we can use a standard Student's t random variable T_s (with degrees of freedom $s = 2a^n$) to represent η^{n+1}. We then use (2) to re-write θ^{n+1} as:

$$\theta^{n+1} = \theta^n + \Sigma^n \psi \sqrt{\frac{b^n}{a^n\left(1+\psi^\top \Sigma^n \psi\right)}} T_s. \tag{8}$$

Combining (7) and (8) we can derive a new formulation of the KG quantity, given by

$$v_\psi^{KG,n} = \mathbb{E}^n(\max_{\varphi\in\Phi} p_\varphi + q_\varphi(\psi)T_s) - \max_{\varphi\in\Phi} p_\varphi, \tag{9}$$

where

$$p_\varphi = \varphi^\top \theta^n \text{ and } q_\varphi(\psi) = \varphi^\top \Sigma^n \psi \sqrt{\frac{b^n}{a^n\left(1+\psi^\top \Sigma^n \psi\right)}}. \tag{10}$$

We sort the pairs $\{p_\varphi, q_\varphi(\psi)\}_{\varphi\in\Phi}$ and relabel them as $\{p_i, q_i\}_{i=1}^K$, such that q_i is in ascending order (see Powell and Ryzhov 2012 section 5.3 for details). Notice that the quantity inside the expectation in (9) is a piecewise linear function; we define $c_i = \frac{p_i - p_{i+1}}{q_{i+1}-q_i}$ as the set of "break-points" where the piecewise linear function changes slopes, and then compute the KG quantity as

$$v_\psi^{KG,n} = \sum_{i=1}^{K-1} \left(q_{i+1}(\psi) - q_i(\psi)\right)\left(\frac{s+c_i^2}{s-1} g_s(|c_i|) - |c_i|\left(1 - G_s(|c_i|)\right)\right), \tag{11}$$

where $g_s(\cdot)$ and $G_s(\cdot)$ are the pdf and cdf, respectively, of the standard Student's t-distribution with s degrees of freedom.

We define a new policy, called the *Knowledge Gradient with Unknown Precision* (KGUP), which chooses its measurement decision by

$$\psi^{KGUP,n} = \arg\max_{\psi\in\Phi} v_\psi^{KG,n}.$$

The KGUP policy extends the value of information approach to learn linear regression coefficients and unknown measurement variance (or precision). The Bayesian prior distribution is more compact by placing beliefs on the attributes rather than the alternatives. We are allowed to have uncertainty in our knowledge about the variance and use new collected information to update our beliefs and reduce the uncertainty.

3 ALGORITHM IMPROVEMENT

The current algorithm for computing the KG quantity in (11) with fixed ψ has computational complexity $O(K\log K)$, mostly due to the difficulty of finding the break-points; to find the largest KG quantity, we need to loop over the whole decision space Φ for each ψ, which requires K iterations. Thus the overall computational complexity is $O(K^2\log K)$ (Frazier et al. 2009). As K grows exponentially in the number of attributes r, i.e., $K \sim 2^r$, the computational complexity in terms of r is $O(r4^r)$. In practice, the algorithm might have difficulties in computing a single maximum of KG even for $r = 15$, which is a reasonable size of the campaign decision space for the ARC.

Here we discuss two approaches for computing the KG quantity with lower computational complexity. In case 1, we assume independent and additive campaign attributes and derive a direct calculation of KG without requiring any break-point computation. In case 2, we exploit the more general situation where combined effects between attributes are present and use convex approximation and SDP relaxation to improve the computational efficiency.

3.1 The KGUP$_2$ Algorithm for Independent Attributes

We assume all the campaign attributes have independent, linear and additive effects on the transformed campaign outcome, and the linear inequalities $A\varphi = h$ confining the decision space Φ only consist of $\varphi_0 = 1$. In this case, all the attributes are directly controllable except the first term (the intercept), thus $K = 2^{r-1}$. We take advantage of the fact that all decision variables in our problem are binary. Recalling that $\theta^n = (\theta_0^n, \theta_1^n, \ldots, \theta_{r-1}^n)^\top$, and $\mathbb{E}^n \theta^{n+1} = \theta^n$, we can re-write (7) as

$$v_\psi^{KG,n} = \mathbb{E}^n \left[\sum_{j \geq 1} (\theta_j^{n+1})^+ | \varphi^n = \psi \right] - \sum_{j \geq 1} (\theta_j^n)^+. \tag{12}$$

We define $u_j^n(\psi) = (\Sigma_{j\cdot}^n)^\top \psi \sqrt{\frac{b^n}{a^n(1+\psi^\top \Sigma^n \psi)}}$, and use (8) and (12) to obtain

$$v_\psi^{KG,n} = \sum_{j \geq 1, u_j^n \neq 0} \theta_j^n G_s \left(\left| \frac{\theta_j^n}{u_j^n(\psi)} \right| \right) + \frac{s(u_j^n(\psi))^2 + (\theta_j^n)^2}{(s-1)|u_j^n(\psi)|} g_v \left(\left(\frac{\theta_j^n}{u_j^n(\psi)} \right)^2 \right) - (\theta_j^n)^+ \quad (s > 1).$$

This is a direct calculation of the KG quantity requiring no additional calculation of the break-points or sorting; to find the maximum of KG we need K iterations to search over the whole decision space, thus the overall computational complexity is $O(K)$. The KGUP$_2$ algorithm has a dramatic improvement in efficiency compared to the algorithm in Section 2.2., although K may still be quite large. The main advantage is obtained by considering the special binary structure of the problem.

3.2 The KGUP$_3$ Algorithm for General Cases

The derivation of the KGUP$_2$ algorithm assumes that every feature is directly controllable by the decision-maker, that is, any combination of 0s and 1s is allowed. However, this is usually not the case in non-profit fundraising, because the regression model includes interactions between attributes, as discussed in Section 2.1. We let r_1 be the number of directly controllable features, r_2 be the number of two-way interaction features (by our choice), thus $r = 1 + r_1 + r_2$, where the additional 1 indicates the intercept.

In this case, we first apply a quantization procedure to approximate the infinite-dimensional optimization problem by a finite problem, and then approximate the finite non-convex problem using a convex relaxation based on semidefinite programming. The new problem can be solved by interior point methods, which have polynomial time complexity in r (Boyd and Vandenberghe 2004).

3.2.1 Finite Approximation

Note that the second term of the KG quantity in (9) is independent of ψ, and the decision made by the KGUP policy $\psi^{KGUP,n}$ only depends on the first term. Thus for convenience we omit the second term and redefine

$$v_\psi^{KG,n} = \mathbb{E}^n(f(T_s)), \tag{13}$$

where $f(t) = \max_{\varphi \in \Phi} p_\varphi + q_\varphi(\psi) t$. For fixed p_φ and $q_\varphi(\psi)$, the function $p_\varphi + q_\varphi(\psi)t$ is linear in t, and thus convex in t. The maximum over a finite family of convex function is still convex, thus $f(t)$ is convex in t.

The *Voronoi quantizer* for T_s is the function q_{vor}: $\mathbb{R} \to \{t_1, \ldots, t_M\}$ defined as

$$q_{vor}(T_s) = \sum_k t_k 1_{C(t_k)}(T_s),$$

where $\{C(t_k)\}_{1 \leq k \leq M}$ is a Borel partition of $\mathbb{R}$ with $C(t_k) = \{t \in \mathbb{R} | |t - t_k| \leq |t - t_j|, j = 1, \ldots, M\}$. For one-dimensional and unimodal random variable T_s and fixed M, there exists a unique $q_{vor}(\cdot)$ and corresponding sequence $t^M = \{t_k\}_{1 \leq k \leq M}$ that minimize the *quadratic quantization error* (Graf and Luschgy 2000)

$$D_M^{T_s,2} = \mathbb{E}|T_s - q_{vor}(T_s)|^2 = \int_{\mathbb{R}} \min_k |t_k - t|^2 g_s(t) dt.$$

By finding the quantization sequence t^M that minimizes $D_M^{T_s,2}$, we can approximate the KG quantity in (13) by

$$\hat{v}_M = \sum_{k=1}^{M} w_k f(t_k), \tag{14}$$

where $w_k = G_s(\frac{t_k+t_{k+1}}{2}) - G_s(\frac{t_{k-1}+t_k}{2})$, with the convention $t_0 = -\infty$ and $t_{M+1} = \infty$. Newton's method has been used to compute the quantization sequence for several distributions such as standard normal distribution (Pages and Printems 2003). Here we use a similar approach to find the quantization sequence for T_s.

3.2.2 Semidefinite Programming

By (10), (14) and the definition of f we have

$$\max_{\psi \in \Phi} \hat{v}_M = \max_{\psi \in \Phi} \max_{\{\varphi^k \in \Phi\}_{1 \le k \le M}} \sum_{k=1}^{M} w_k (\varphi^k)^\top (\theta^n + \Sigma^n t_k d_\psi^n),$$

where $d_\psi^n = \frac{\psi}{\sqrt{\frac{a^n}{b^n}(1+\psi^\top \Sigma^n \psi)}}$.

Lemma 3 Define

$$C_k = \frac{1}{2} \begin{bmatrix} 0 & (\theta^n)^\top & 0^\top \\ \theta^n & 0 & t_k \Sigma^n \\ 0 & t_k \Sigma^n & 0 \end{bmatrix}, Z_k = \begin{bmatrix} 1 \\ \varphi^k \\ d_\psi^n \end{bmatrix} \begin{bmatrix} 1 \\ \varphi^k \\ d_\psi^n \end{bmatrix}^\top = \begin{bmatrix} Z_k^{11} & Z_k^{1\varphi} & Z_k^{1d} \\ Z_k^{\varphi 1} & Z_k^{\varphi\varphi} & Z_k^{\varphi d} \\ Z_k^{d1} & Z_k^{d\varphi} & Z_k^{dd} \end{bmatrix}.$$

Then,

$$\max_{\psi \in \Phi} \hat{v}_M = \max_{\psi \in \Phi} \max_{\{\varphi^k \in \Phi\}_{1 \le k \le M}} \sum_{k=1}^{M} w_k \text{trace}(C_k Z_k). \tag{15}$$

We observe that C_k is a constant matrix and Z_k is a positive semidefinite matrix with rank 1, for all $k \in \{1, \ldots, M\}$. This problem is similar to SDP but has non-linear constraints consisting of: 1) the rank 1 constraint on Z_k; 2) the binary constraints on ψ and φ^k,i.e., $\psi, \varphi^k \in \{0,1\}^r$; 3) the non-linear constraints on d_ψ^n transformed from ψ. To formulate (15) as an SDP problem, we need to relax the non-linear constraints using a set of linear constraints. We first drop the rank 1 constraint on Z_k, then relax the binary constraints on ψ and φ^k as $\psi, \varphi^k \in [0,1]^r$, and finally develop a set of linear constraints on d_ψ^n from $A\psi = h$.

The constraints on ψ include the conditions in (1) and $\psi_0 = 1$. We transform the inequalities in (1) to equalities by adding slack variables $\varsigma \in \{0,1\}^{r_3}$. For convenience we abuse the notation a little and still denote the extended vector $[\psi; \varsigma]$ as ψ and $[\varphi^k; \varsigma^k]$ as φ^k. Here $r = 1 + r_1 + r_2 + r_3$, which correspond to the intercept, the independent variables, the two-way interaction variables and the slack variables respectively.

By $A\psi = h$, we have $\psi^\top A^\top A \psi = h^\top h$. Thus

$$d_\psi^n = \frac{\psi}{\sqrt{\psi^\top (\frac{a^n}{b^n}(\frac{A^\top A}{h^\top h} + \Sigma^n))\psi}} = \frac{\psi}{\sqrt{\psi^\top P \psi}},$$

where $P = \frac{a^n}{b^n}(\frac{A^\top A}{h^\top h} + \Sigma^n)$. It follows that $(d_\psi^n)^\top P d_\psi^n = 1$. We define $Y = d_\psi^n (d_\psi^n)^\top$, whence the condition is equivalent to

$$\text{trace}(PY) = 1. \tag{16}$$

By definition, Y is symmetric and positive semidefinite, and we also require Y to be non-negative:

$$Y_{i,j} \geq 0, \forall 1 \leq i, j \leq r. \tag{17}$$

To bound from above the elements of the semidefinite matrix Y, we define $\delta = \min_{\psi \in [0,1]^r, A\psi = h} \psi^\top P \psi$ and obtain

$$\text{diag}(Y) \leq 1_r / \delta, \tag{18}$$

with the convention that there is no upper bound if $\delta = 0$. The quantity δ can be easily found by solving a small quadratic program with linear constraints, using a convex programming solver.

Following Defourny et al. (2013), we add another constraint on φ that can strengthen the relaxations. Given $\xi \in \mathbb{R}^r$ with $\xi_i > 0$ for each i, we define $\zeta = \sup_{\varphi \in \Phi} \xi^\top \varphi$. Then for any $\varphi \in \Phi$ we have

$$\varphi \varphi^\top \preceq \zeta \text{Diag}(\varphi) \text{Diag}(\xi)^{-1}, \tag{19}$$

where $\text{Diag}(z)$ denotes the diagonal matrix with elements z_i. The value of ζ can be found by solving a small IP problem. If the IP cannot be solved to optimality, the best found upper bound on ζ should be used.

Combining (16), (17), (18), (19) and the linear constraints on φ^k, we formulate (15) as the SDP problem:

$$\begin{aligned}
\max \quad & \sum_{k=1}^{M} w_k \text{trace}(C_k Z_k) \\
\text{s.t. } \forall k: \quad & Z_k \succeq 0, \\
& Z_k^{11} = 1, A Z_k^{\varphi 1} = h, 0 \leq Z_k^{\varphi 1} \leq 1, \\
& A Z_k^{\varphi\varphi} A^\top = h h^\top, [Z_k^{\varphi\varphi}]_{i,j} \geq 0, \forall 0 \leq i, j \leq r, \\
& Z_k^{\varphi\varphi} \preceq \zeta \text{Diag}(Z_k^{\varphi 1}) \text{Diag}(\xi)^{-1}, \\
& Z_k^{dd} = Y, Y_{i,j} \geq 0, \forall 1 \leq i, j \leq r, \\
& \text{trace}(PY) = 1, \text{diag}(Y) \leq 1_r / \delta.
\end{aligned}$$

After solving this SDP, we obtain the matrix Y. With $\text{rank}(Y) = 1$, $\frac{d_\psi^n}{||d_\psi^n||}$ is equivalent to the unique normalized eigenvector of Y. After relaxing the condition, we can approximate $\frac{d_\psi^n}{||d_\psi^n||}$ by the normalized eigenvector associated to the largest eigenvalue of Y. We define the eigenvector as v and let $\tilde{v} = \frac{v}{\max_j (v_j)}$, which satisfies $\tilde{v} \in [0,1]^r$. Since ψ and d_ψ^n only differ by a scaling factor, we can interpret $\tilde{v}$ as the fractional approximation for the binary $\psi \in \{0,1\}^r$. To recover a binary ψ we can perform a rounding procedure by solving the small IP given by

$$\begin{aligned}
\min \quad & 1_r^\top z \\
\text{s.t.} \quad & \psi - \tilde{v} \leq z, \\
& \tilde{v} - \psi \leq z, \\
& A\psi = h, \psi \in \{0,1\}^r,
\end{aligned}$$

where we project the solution $\tilde{v}$ on the set $\{0,1\}^r$, by minimizing the L_1 norm of the difference between ψ and $\tilde{v}$.

4 NUMERICAL EXPERIMENTS

We simulate the performance of the KGUP policy using campaign historical data obtained from the ARC, which covers 8.6 million interactions and 1.2 million unique donors. The data are aggregated into 952 unique campaign designs, from which the linear regression coefficients and covariance matrix have been estimated (Ryzhov et al. 2013). We use the estimations as the initial prior for β, i.e., θ^0 and Σ^0. We experiment with different values (a^0, b^0) to represent different beliefs about ρ. In each of L sample runs, the "true" value of ρ is generated from $\rho \sim \Gamma(a^0, b^0)$ and the "true" value of β is generated from $\beta \sim N(\theta^0, \frac{1}{\rho}\Sigma^0)$. The algorithms do not see the true values when making decisions. We conduct N measurements with each policy; in the $(n+1)$th measurement, the observation η_π^{n+1} for each policy π is simulated by $\eta_\pi^{n+1} \sim N((\psi_\pi^n)^\top \beta, \frac{1}{\rho})$.

The first experiment examines the performances of the KGUP policy, the correlated KG (CKG) policy (Frazier et al. 2009) and the greedy policy. The CKG policy assumes the variance is known and maintains the belief through the entire learning process. We let CKG start with the assumption that $\rho = \frac{a^0}{b^0}$, whereas KGUP assumes that $\rho \sim \Gamma(a^0, b^0)$. The greedy policy chooses the alternative $\psi = \arg\max_{\varphi\in\Phi} \varphi^\top \theta^n$ and updates the beliefs according to (2) and (3). We select 10 features from the set of key factors described in Table 1, where the first 9 are directly controllable features, and the last one is an interaction feature. Thus $r_1 = 9$, $r_2 = 1$ and $K = 2^{r_1} = 512$. We set $L = 100$, $N = 20$. The ability of learning the unknown mean is quantified by the normalized opportunity cost C_π^n for each measurement n and policy π, defined by

$$C_\pi^n = \frac{\max_{\varphi\in\Phi} \varphi^\top \beta - (\arg\max_{\psi\in\Phi} \psi^\top \theta_\pi^n)^\top \beta}{\max_{\varphi,\psi\in\Phi}(\varphi^\top \beta - \psi^\top \beta)}. \tag{20}$$

The denominator in (20) confines C_π^n in $[0,1]$. The ability of learning the unknown variance is quantified by the *precision estimation error*, defined as $|\rho - \frac{a^n}{b^n}|$.

Figures 1(a)-1(c) show how performance values change over the number of measurements with different settings for (a^0, b^0). The graph is plotted with mean values plus or minus 2 standard deviations. The KGUP policy outperforms CKG by a statistically significant amount when the mean of the true measurement precision, i.e., $\frac{a^0}{b^0}$ is small (Figure 1(b) and 1(c)), and performs competitively against CKG when the mean is large (Figure 1(a)). We conjecture that this behaviour arises because, when the true precision is small, the true β tends to be far from the prior, in which case the CKG policy is more sensitive to an inaccurate precision value. Figure 1(d) shows how the error of estimation on ρ decreases for the KGUP policy. We see that the KGUP policy can learn the precision rapidly in the first few measurements, while the CKG policy maintains the initial error invariably.

In the second experiment we test the performance of the KGUP_3 algorithm with 2 different problem sizes. The smaller problem has the same features as in experiment 1, with $K = 512$; the larger problem has 3 extra independent features and 3 extra interaction features, and consequently $r_1 = 12$, $r_2 = 4$ and $K = 4096$. The prior for ρ is set as $a^0 = 0.5$, $b^0 = 1$. We let $L = 50$ and $N = 10$. Figure 2 shows that the KGUP_3 policy outperforms the CKG policy and the greedy policy in both problems, indicating that the KGUP_3 policy provides a comparatively good approximation to the KGUP policy in learning the unknown coefficients.

To further test the accuracy of the approximation in more general settings, we randomly generated 100 priors for β using $\theta^0 \sim N(0_r, I_{r\times r})$ and $\Sigma^0 = (s+s^\top)(s+s^\top)$, where $s_{i,j} \sim N(0,1), \forall 1 \le i,j \le r$. Otherwise, the problem remains the same as in the earlier experiment. For each of these 100 priors, we compute the approximate value of information using KGUP_3 (or CKG) and find the alternative φ' that maximizes this quantity. We then rank the alternatives according to their true values of information (computed by KGUP) and see how highly φ' places in that ranking. Figure 3(a) shows that KGUP_3 produces a better approximation to the KGUP policy than CKG. Over 50% of the alternatives chosen by KGUP_3 have values of information ranked in the top 50 (out of 512).

To test the computational complexity of the KGUP_3 algorithm with different problem sizes, we use the same setting as above but each time increase the number of independent features by 1. We record the CPU

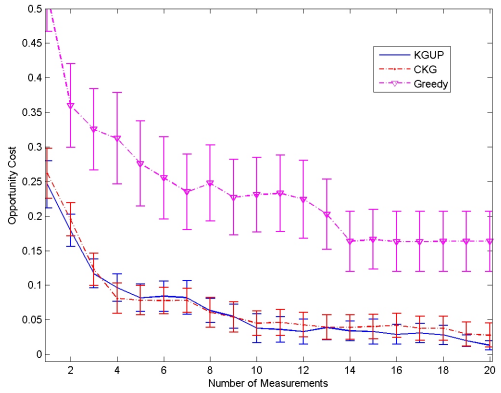

(a) Learning the regression coefficients, $a_0 = 10$, $b_0 = 1$

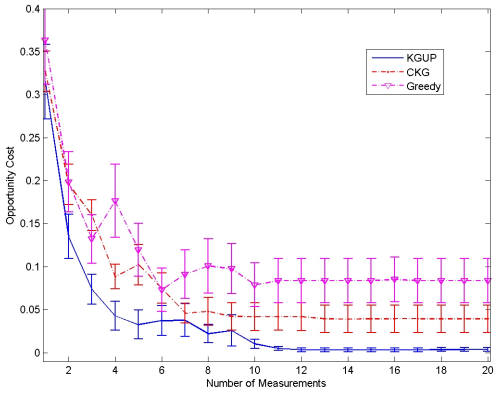

(b) Learning the regression coefficients, $a_0 = 1$, $b_0 = 10$

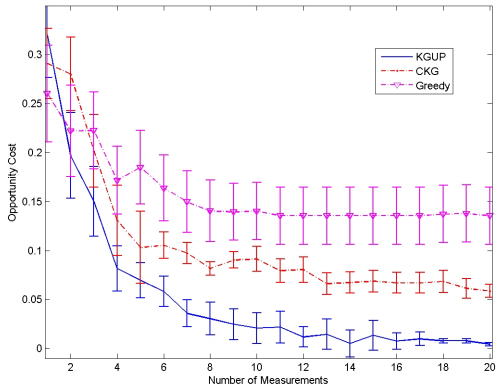

(c) Learning the regression coefficients, $a_0 = 0.5$, $b_0 = 1$

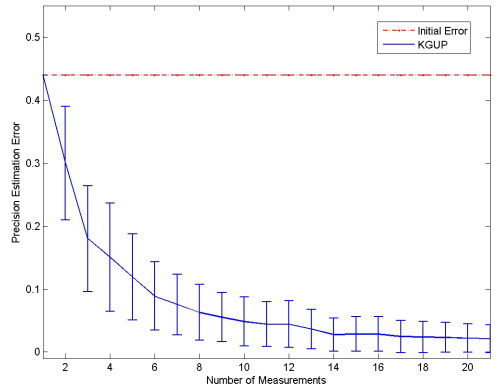

(d) Learning the unknown variance, $a_0 = 0.5$, $b_0 = 1$

Figure 1: Averaged opportunity cost and precision estimation error over time.

time lapsed for one run of the policy, averaged over 10 samples. Figure 3(b) shows that the computational complexity of the KGUP algorithm increases exponentially as r (linear increase in logarithm), while the computing time for KGUP_3 algorithm remains steady after 2^{14} alternatives. This indicates that the KGUP_3 algorithm is much more efficient than the KGUP algorithm when the number of alternatives is large.

5 CONCLUSION

We have proposed an optimal learning framework for non-profit campaign design and marketing, with the ability to simultaneously learn the regression coefficients and measurement precision. Experimental results suggest that this approach adds the most value when the true precision is small and the prior beliefs about the coefficients are far from the true values. We also provide two improved algorithms with higher computational efficiency for problems with large numbers of alternatives. We believe that our work offers a rigorous mathematical model for non-profit sequential decision making, and provides a new approach to optimal learning problems with linear regression structure, unknown variance and large decision space.

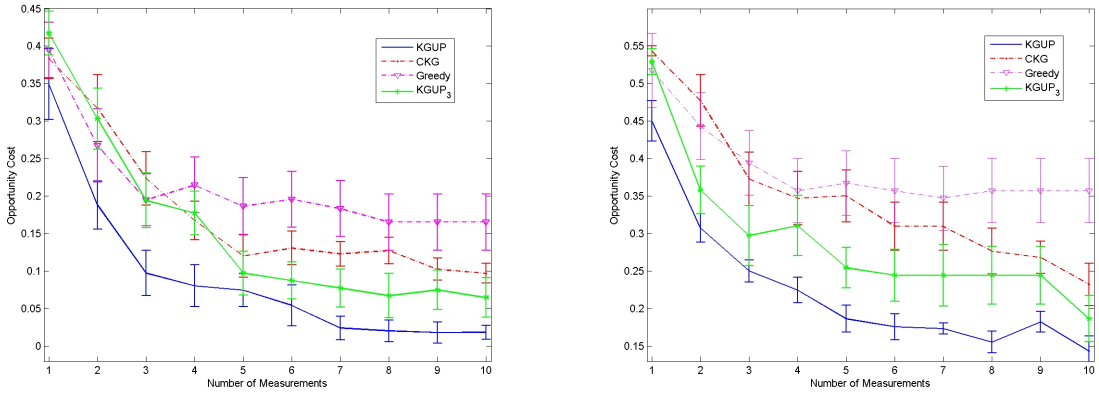

(a) Learn the regression coefficients, $N = 512$ (b) Learn the regression coefficients, $N = 4096$

Figure 2: Averaged opportunity cost over time.

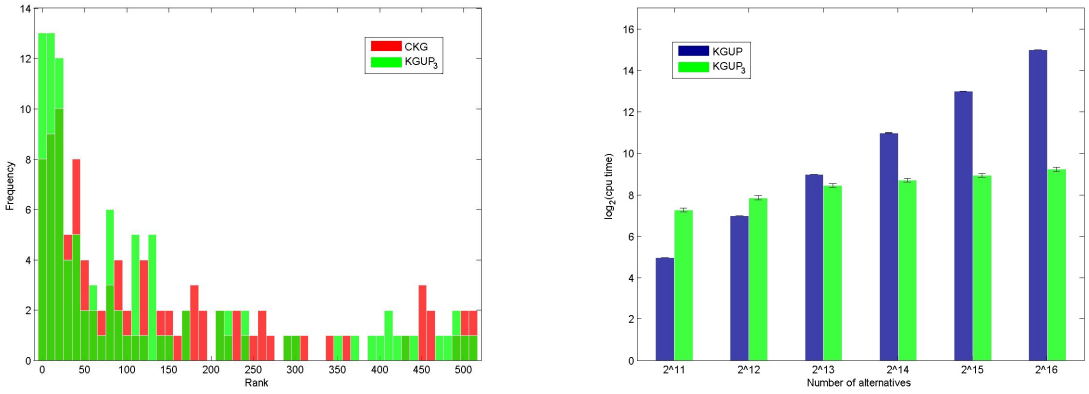

(a) Empirical distributions of value of information ranks (b) Averaged $\log_2$(CPU time) for 1 run of the policy

Figure 3: Accuracy and efficiency assessment of KGUP_3 using simulated prior for β.

REFERENCES

Arnett, D. B., S. D. German, and S. D. Hunt. 2003. "The Identity Salience Model of Relationship Marketing Success: The Case of Nonprofit Marketing". *Journal of Marketing* 67 (2): 89–105.

Auten, G., and D. Joulfaian. 1996. "Charitable Contributions and Intergenerational Transfers". *Journal of Public Economics* 59 (1): 55–68.

Boyd, S., and L. Vandenberghe. 2004. *Convex Optimization*. Cambridge, UK: Cambridge University Press.

Bult, J. R., H. V. der Scheer, and T. Wansbeek. 1997. "Interaction Between Target and Mailing Characteristics in Direct Marketing, with an Application to Health Care Fund Raising". *International Journal of Research in Marketing* 14 (4): 301–308.

Chen, C.-H., H.-C. Chen, and L. Dai. 1996. "A Gradient Approach for Smartly Allocating Computing Budget for Discrete Event Simulation". In *Proceedings of the 1996 Winter Sirnulation Conference*, edited by J. M. Charnes, D. J. Morrice, D. T. Brunner, and J. J. Swvain, 398–405. Washington, DC: Institute of Electrical and Electronics Engineers, Inc.

Chick, S. E., J. Branke, and C. Schmidt. 2010. "Sequential Sampling to Myopically Maximize the Expected Value of Information". *INFORMS Journal on Computing* 22 (1): 71–80.

Defourny, B., I. O. Ryzhov, and W. B. Powell. 2013. "Optimal Informational Blending with Measurement in the L2 Sphere". *Submitted for publication*.

Fennis, B. M., L. Janssen, and K. D. Vohs. 2009. "Acts of Benevolence: A Limited-Resource Account of Compliance with Charitable Requests". *Journal of Consumer Research* 35 (6): 906–924.

Frazier, P. I., W. B. Powell, and S. Dayanik. 2009. "The Knowledge-gradient Policy for Correlated Normal Rewards". *INFORMS Journal on Computing* 21 (4): 599–613.

Graf, S., and H. Luschgy. 2000. *Foundations of Quantization for Probability Distributions*. Berlin, Germany: Springer-Verlag.

Gupta, S. S., and K. J. Miescke. 1996. "Bayesian Look Ahead One-stage Sampling Allocations for Selection of the Best Population". *Journal of Statistical Planning and Inference* 47 (3): 229–244.

Jones, A., and J. Posnett. 1991. "Charitable Donations by UK Households: Evidence from the Family Expenditure Survey". *Applied Economics* 23 (2): 343–351.

Lankford, R. H., and J. H. Wyckoff. 1991. "Modeling Charitable Giving Using a Box-Cox Standard Tobit Model". *The Review of Economics and Statistics* 73 (3): 460–470.

Leeflang, P. S., and D. R. Wittink. 2000. "Building Models for Marketing Decisions: Past Present and Future". *International Journal of Research in Marketing* 18:105–126.

Lilien, G. L., P. Kotler, and K. S. Moorthy. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening and Multiple Comparisons*. New Delhi: Prentice Hall College Div.

Negoescu, D. M., P. I. Frazier, and W. B. Powell. 2011. "The Knowledge Gradient Algorithm for Sequencing Experiments in Drug Discovery". *INFORMS Journal on Computing* 23 (3): 346–363.

Pages, G., and J. Printems. 2003. "Optimal Quadratic Quantization for Numerics: the Gaussian Case". *Monte Carlo Methods and Applications* 9 (2): 135–166.

Powell, W. B., and I. O. Ryzhov. 2012. *Optimal Learning*. Hoboken, New Jersey: John Wiley & Sons, Inc.

Qu, H., I. O. Ryzhov, and M. C. Fu. 2012. "Ranking and Selection with Unknown Correlation Structures". In *Proceedings of the 2012 Winter Sirnulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher. Institute of Electrical and Electronics Engineers, Inc.

Ryzhov, I. O., B. Han, J. Bradic, and A. Bradic. 2013. "Cultivating Disaster Donors: a Case Application of Scalable Analytics on Massive Data". *Submitted for publication*.

Tobin, J. 1958. "Estimation of Relationships for Limited Dependent Variables". *Econometrica* 26:24–36.

AUTHOR BIOGRAPHIES

BIN HAN is a Ph.D. candidate in Applied Mathematics, Statistics, and Scientific Computation at the University of Maryland. His research interests are sequential decision making, stochastic optimization, statistics and machine learning. His email address is danielh@math.umd.edu.

ILYA O. RYZHOV is an Assistant Professor in the Robert H. Smith School of Business at the University of Maryland. He received a Ph.D. in Operations Research and Financial Engineering from Princeton University. His research deals with optimal learning and the broader area of stochastic optimization, with applications in disaster relief, energy, and revenue management. He was a recipient of WSC's Best Theoretical Paper Award in 2012. His work has appeared in *Operations Research*, and he is a co-author of the book *Optimal Learning*, published in 2012 by John Wiley & Sons. His email address is iryzhov@rhsmith.umd.edu.

BORIS DEFOURNY is an Assistant Professor in Industrial Systems Engineering at Lehigh University. He received an Electrical Engineering degree in 2005, and a Ph.D. in Applied Sciences in 2010, from the University of Liege, Belgium. His research is in sequential decision-making under uncertainty, using techniques from stochastic programming, robust optimization, and machine learning. He works on stochastic optimization problems for electric power systems. His email address is defourny@princeton.edu.