SMOILE: A Shopper Marketing Optimization and Inverse Learning Engine

Abhilash Reddy Chenreddy
University of Illinois at Chicago
Chicago, Illinois, USA
achenr2@uic.edu

Parshan Pakiman
University of Illinois at Chicago
Chicago, Illinois, USA
ppakim2@uic.edu

Selvaprabu Nadarajah*
selvan@uic.edu
University of Illinois at Chicago
Chicago, Illinois, USA

Ranganathan Chandrasekaran
University of Illinois at Chicago
Chicago, Illinois, USA
ranga@uic.edu

Rick Abens
Foresight ROI, Inc
Chicago, Illinois, USA
rickabens@foresightroi.com

ABSTRACT

Product brands employ shopper marketing (SM) strategies to convert shoppers along the path to purchase. Traditional marketing mix models (MMMs), which leverage regression techniques and historical data, can be used to predict the component of sales lift due to SM tactics. The resulting predictive model is a critical input to plan future SM strategies. The implementation of traditional MMMs, however, requires significant ad-hoc manual intervention due to their limited flexibility in (i) explicitly capturing the temporal link between decisions; (ii) accounting for the interaction between business rules and past (sales and decision) data during the attribution of lift to SM; and (iii) ensuring that future decisions adhere to business rules. These issues necessitate MMMs with tailored structures for specific products and retailers, each requiring significant hand-engineering to achieve satisfactory performance – a major implementation challenge.

We propose an SM Optimization and Inverse Learning Engine (SMOILE) that combines optimization and inverse reinforcement learning to streamline implementation. SMOILE learns a model of lift by viewing SM tactic choice as a sequential process, leverages inverse reinforcement learning to explicitly couple sales and decision data, and employs an optimization approach to handle a wide-array of business rules. Using a unique dataset containing sales and SM spend information across retailers and products, we illustrate how SMOILE standardizes the use of data to prescribe future SM decisions. We also track an industry benchmark to showcase the importance of encoding SM lift and decision structures to mitigate spurious results when uncovering the impact of SM decisions.

*contact author.

KDD '19, August 4–8, 2019, Anchorage, AK, USA

ACM ISBN 978-1-4503-6201-6/19/08…$15.00
https://doi.org/10.1145/3292500.3330788

CCS CONCEPTS

• **Computing methodologies → Planning and scheduling**; **Sequential decision making**; **Inverse reinforcement learning**; • **Applied computing → Marketing**; *Operations research.*

KEYWORDS

shopper marketing; inverse reinforcement learning; data-driven optimization

ACM Reference Format:
Abhilash Reddy Chenreddy, Parshan Pakiman, Selvaprabu Nadarajah, Ranganathan Chandrasekaran, and Rick Abens. 2019. SMOILE: A Shopper Marketing Optimization and Inverse Learning Engine. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19), August 4–8, 2019, Anchorage, AK, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3292500.3330788

1 SHOPPER MARKETING

The primary goal of any marketer is to reach shoppers at critical moments along their path to purchase. Shopper marketing (SM) involves understanding how shoppers behave and react to various marketing tactics and leveraging this intelligence to benefit marketers as well as the retailers [6, 7]. It is one of the fastest growing forms of marketing for consumer packaged goods. A stroll through the aisles of a retailer like Target or Walmart would reveal a variety of tactics that are used to influence consumer decisions. In-store tactics like paper signage, endcap displays, samples and live demos are combined with pre-store tactics such as coupons, advertisements, mobile and social media campaigns [16]. Given the multitude of tactics deployed, marketers face difficulties in estimating and isolating the revenue lifts associated with the SM tactics. Tactics are often used in combination, and it is not clear what tactics should be combined for specific product-market combinations.

A core goal of brands employing SM is to choose tactics to maximize the sales of a product at one or more retailers over a finite planning period (e.g., a quarter) while respecting budget constraints. Tackling this problem involves challenging *lift attribution* and *tactic planning* steps. During the lift attribution phase, the increase in sales (i.e., lift) due to SM tactics needs to be mined, that is, we require a model to predict the impact of using SM tactics over time. Historical weekly lift for a product and retailer can be calculated using sales analytics provided by external vendors such as Nielsen and IRI, where lift is often defined as weekly sales volume minus an estimate of base sales volume with no promotional activities. Since historical lift data accounts for all merchandising activities at a retailer in a given week, our predictive model must be able to attribute a component of this lift to individual SM tactics. This attribution step needs to also account for the impact of a second data source containing information on past SM campaigns, which includes tactic decisions, costs, and planning budget. Once such a predictive model is available, the tactic planning step uses it to prescribe future SM tactic decisions while respecting business rules.

Attributing lift to marketing strategies is a core task of marketing-mix models (MMMs) [19]. Traditional MMMs estimate lift by leveraging statistical techniques. Popular examples include Bayesian linear mixed models (BLMMs) [4, 17], which compute a posterior distribution of the regression coefficients supported by data and prior information. Although statistical techniques are effective at incorporating several exogenous sources of data, they have limited flexibility as methods for dealing with past SM decision data, as explained next. First, traditional MMMs do not capture the temporal nature of the SM planning process that generates the sales and tactic data. Second, they lack an explicit and interpretable coupling between the estimated SM lift model and past SM tactic decisions, even though past tactic decisions are clearly guided by an SM lift model. Finally, several business rules (e.g., spending constraints and limits on the number of active strategies) impact the chosen SM tactics but typically MMMs cannot capture rich constraints or become difficult to train when they are constrained.

The above challenges translate to an ad-hoc work flow that involves significant hand-engineering to apply MMMs for SM lift attribution and tactic planning. Specifically, MMMs need to be constructed and trained in the lift attribution phase and subsequently used in a tactic planning phase to validate its performance. If the performance is poor, parameters in the lift attribution phase (e.g., prior distribution in BLMMs) need to be modified largely by trial and error until acceptable results are observed in the planning phase. Further, these changes are data specific, which results in the need for maintaining specialized MMMs for different retailers and products – a major implementation hurdle.

1.1 Contributions

We introduce an SM Optimization and Inverse Learning Engine (SMOILE) to overcome the aforementioned issues with traditional MMMs. Our main contributions via SMOILE are summarized below.

- We design a data-driven modeling framework that is based on viewing SM tactic choice as a sequential decision making process in both lift attribution and tactic planning, thus making these steps consistent.
- We develop a method that trains an SM lift attribution model by coupling rich sources of data related to sales lift and SM decisions from external vendor and brand databases, respectively. To elaborate, the lift model parameters are trained on observed lift data while simultaneously ensuring that the trajectory of past SM decisions are "near-optimal" with respect to the model being trained and, in addition, respect business rules. Implementing SMOILE requires tuning two interpretable parameters that control for overfitting and the potential suboptimality of historical SM tactic decisions. The core ideas underpinning SMOILE extend approaches in empirical optimization (EO) [3] and inverse reinforcement learning (IRL) [12] to make them applicable to our SM setting.
- Using a case study involving a unique dataset from multiple retailers and products, we provide an illustration of how SMOILE can be configured for SM lift attribution and tactic planning. Our results highlight two useful insights: (i) making the training phase of the SM lift prediction model consistent with its downstream use for planning by leveraging data on past SM decisions leads to better predictive models; and (ii) spurious results on the impact of SM decisions can be mitigated by directly enforcing structure on the lift model capturing behavioral effects specific to SM. Example of such effects include the reduction in lift due to an SM tactic caused by waiting and satiation, where waiting is when a customer postpone using or being exposed to an SM tactic if this tactic is available for several future weeks, while satiation relates to fewer customers being engaged by an SM tactic if it has already been employed for a long time.

Through these contributions we standardize the use of SM decision data for lift attribution and planning. This standardization streamlines implementation and reduces the burden of maintaining multiple specialized MMMs. Moreover, the framework we develop is relevant for data-driven decision making beyond SM.

1.2 Related Work

Estimating the value of marketing strategies is a classical data mining problem tackled in marketing-mix or media-mix modeling [5, 8, 9, 14]. The review by [19] highlights revived interest in estimating lift due to marketing activities from multiple sources of data at different points in the path to purchase of a consumer. While advanced techniques based on causal inference and strategic consumer behavior represent the frontier of active academic research (e.g., [15]), such methods have not yet seen significant adoption by industry [19]. Rather, methods such as BLMMs, or special cases thereof, are commonly employed in practice [4, 17, 18]. The techniques in the aforementioned literature, however, do not account for how lift models are used for optimizing future marketing campaigns, nor do they focus on the implementation hurdles arising in an SM setting. In contrast, we propose a sequential data-driven modeling and solution framework that trains an SM lift model in a manner that is consistent with its use for tactic planning, and in addition, reduces the manual intervention needed during implementation.

EO provides a framework for model estimation in the context of sequential decision making processes [3]. EO has also been used to simultaneously account for regression and decision objectives during parameter estimation [10, 11]. A standard application of EO for SM would compute a parametric tactic scheduling rule that maximizes lift given a historical trajectory of realized lift due to each tactic. We extend the EO framework to the SM setting in two key ways. First, a historical lift due to individual SM tactics is unavailable. In fact, the main goal of the lift attribution phase in SM is to predict how much of the total lift is associated with the use of SM tactics. We thus incorporate this attribution problem within an EO approach. Second, within our EO framework we use IRL to handle an addition data source containing past SM campaign decisions, which is coupled with historical lift and the attribution problem discussed above.

IRL has been successfully employed to compute unknown reward functions, typically represented as a linear combination of features, from historical data trajectories of states and actions [12]. Early IRL algorithms assumed that historical decisions are optimal [12], while later methods accounted for the potential suboptimality of decisions [1, 20]. Suboptimality of past decisions is handled implicitly in these algorithms with goal of computing good future decisions. In contrast, our focus in SM is to obtain a good attribution of total lift across SM tactics that is consistent with past SM tactic decisions. Therefore, we explicitly allow for past decisions to be potentially suboptimal in an IRL model using an interpretable scalar parameter and embed this model as constraints in an EO formulation with a regularized objective function to handle any resulting degeneracy. We tune both the suboptimality and regularization parameters using cross validation.

2 LET'S SMOILE: DATA-DRIVEN LIFT ATTRIBUTION AND TACTIC PLANNING

Brands plan SM campaigns based on predictions of sales lift due to SM tactics. Thus, prediction and planning are closely connected. In §2.1, we present the model in SMOILE for planning SM tactics over a finite horizon assuming a model to predict SM lift is available. In §2.2, we introduce the data-driven lift attribution approach of SMOILE that trains a predictive SM lift model in a manner that is consistent with its use in the planning formulation of §2.1.

2.1 Tactic Planning

We consider a brand planning SM promotions for a product at R retailers indexed by $\mathcal{R} := \{1, 2, \dots, R\}$ over $P + 1$ planning periods belonging to set $\mathcal{T}^{\mathrm{P}} = \{0, 1, \dots, P\}$, where period 0 denotes the initial period. At each period, starting from period 0 and moving forward in time, the brand may not employ SM or choose from H distinct SM tactics, with indices in the set $\mathcal{H} = \{1, 2, \dots, H\}$. If SM tactic h is chosen, then the duration of its use must be specified. A planned tactic cannot be interrupted or modified while it is in progress.

Example 1. *In the first week of January, Mars is planning an SM campaign for one of its chocolate products at Walmart and Target for the next six weeks. It is considering the use of demos and in-store discounts. In our framework, this case corresponds to both the number of retailers (R) and number of SM tactics (H) being equal to* 2 *and a planning horizon length (P) of* 6*. Figure 1 displays a potential SM tactic plan using colored bars to indicate when a tactic is active.*

We begin by modeling the sequential process used to plan an SM campaign. At period $p \in \mathcal{T}^{\mathrm{P}}$ and for retailer $r \in \mathcal{R}$, we maintain information regarding active SM tactics. This information for an active SM tactic $h \in \mathcal{H}$ includes the number of periods remaining until the tactic ends and the number of periods since this tactic started, which we denote by $u^r_{p,h}$ and $v^r_{p,h}$, respectively. If tactic h is inactive at period p, then we set $u^r_{p,h} = v^r_{p,h} = 0$. Let

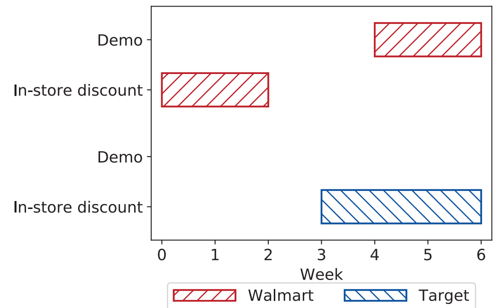

Figure 1: Illustration of SM tactic campaign for Example 1.

$\boldsymbol{u}^r_p := (u^r_{p,1}, \dots, u^r_{p,H})$ and $\boldsymbol{v}^r_p := (v^r_{p,1}, \dots, v^r_{p,H})$. In reinforcement learning parlance, the vector $(\boldsymbol{u}^1_p, \dots, \boldsymbol{u}^R_p, \boldsymbol{v}^1_p, \dots, \boldsymbol{v}^R_p)$ is the time p state of the process. At time zero, the state is known and denoted by $(\bar{\boldsymbol{u}}^1_0, \dots, \bar{\boldsymbol{u}}^R_0, \bar{\boldsymbol{v}}^1_0, \dots, \bar{\boldsymbol{v}}^R_0)$.

At each period p, the tactic choice involves specifying for each inactive tactic (i.e., all $h \in \mathcal{H}$ such that $u^r_{p,h} = 0$) the duration of promotion, which could be zero if the tactic is not chosen. To facilitate modeling, we define a binary auxiliary variable

$$x^r_{p,h} = \begin{cases} 1, & \text{if } u^r_{p,h} > 0, \\ 0, & \text{if } u^r_{p,h} = 0, \end{cases} \tag{1}$$

that specifies when a tactic h is active. Using this variable, the decision logic described above is captured using a quantity $a^r_{p,h}$ that belongs to set

$$\mathcal{A}_p(u^r_{p,h}) := \begin{cases} \{0, 1, \dots, P - p\}, & \text{if } x^r_{p,h} = 0, \\ \{-1\}, & \text{if } x^r_{p,h} = 1, \end{cases} \tag{2}$$

where we use −1 to indicate that promotion h is active and thus a decision regarding its future use cannot be currently made.

Tactic decisions made in each period result in transitions of the state. The transition of state variables satisfy for each $(r, p, h) \in \mathcal{R} \times \mathcal{T}^{\mathrm{P}} \setminus \{P\} \times \mathcal{H}$ the following conditions

$$u^r_{p+1,h} = u^r_{p,h} + a^r_{p,h}; \tag{3}$$

$$v^r_{p+1,h} = \begin{cases} v^r_{p,h} + x^r_{p,h}, & \text{if } x^r_{p,h} = 1, \\ 0, & \text{if } x^r_{p,h} = 0. \end{cases} \tag{4}$$

We define $\boldsymbol{x}^r_p := (x^r_{p,1}, \dots, x^r_{p,H})$ and refer to the triple $\boldsymbol{s}^r_p := (\boldsymbol{u}^r_p, \boldsymbol{v}^r_p, \boldsymbol{x}^r_p)$ as the SM state vector at time p. Table 1 illustrates the SM state trajectory corresponding to Figure 1.

We also replace $(r, h) \in \mathcal{R} \times \mathcal{H}$, $(r, h, p) \in \mathcal{R} \times \mathcal{H} \times \mathcal{T}^{\mathrm{P}}$, $(r, h, p) \in \mathcal{R} \times \mathcal{H} \times \mathcal{T}^{\mathrm{P}} \setminus \{P\}$ with the shorthand (r, h), (r, h, p), and $(r, h, p)_{(-P)}$, respectively. Then, a feasible state trajectory $(\boldsymbol{s}^r_p, (r, p) \in \mathcal{R} \times \mathcal{T}^{\mathrm{P}})$ belongs to the constraint set $\mathcal{F}(\mathcal{T}^{\mathrm{P}})$ defined as

$$u^r_{0,h} = \bar{u}^r_{0,h}, \qquad \forall (r, h) \tag{5}$$

$$v^r_{0,h} = \bar{v}^r_{0,h}, \qquad \forall (r, h) \tag{6}$$

$$x^r_{p,h} = \min(u^r_{p,h}, 1), \qquad \forall (r, h, p) \tag{7}$$

$$a^r_{p+1,h} \le (P - p)(1 - x^r_{p,h}) - x^r_{p,h}, \qquad \forall (r, h, p)_{(-P)} \tag{8}$$

Table 1: SM state trajectory and decisions for Figure 1.

			Week						
Retailer	Tactic	Variable	0	1	2	3	4	5	6
Walmart	Demo	u	0	0	0	0	3	2	1
		a	0	0	0	3	-1	-1	-1
		x	0	0	0	0	1	1	1
		v	0	0	0	0	0	1	2
	In-store discount	u	3	2	1	0	0	0	0
		a	-1	-1	-1	0	0	0	0
		x	1	1	1	0	0	0	0
		v	0	1	2	3	0	0	0
Target	Demo	u	0	0	0	0	0	0	0
		a	0	0	0	0	0	0	0
		x	0	0	0	0	0	0	0
		v	0	0	0	0	0	0	0
	In-store discount	u	0	0	0	4	3	2	1
		a	0	0	4	-1	-1	-1	-1
		x	0	0	0	1	1	1	1
		v	0	0	0	0	1	2	3

$$a^r_{p+1,h} \geq -x^r_{p,h}, \qquad \forall (r,h,p)_{(-P)} \tag{9}$$
$$u^r_{p,h} = a^r_{0,h} + \cdots + a^r_{p-1,h}, \qquad \forall (r,h,p) \tag{10}$$
$$v^r_{p+1,h} \leq x^r_{p,h}(p+1), \qquad \forall (r,h,p)_{(-P)} \tag{11}$$
$$v^r_{p+1,h} \leq v^r_{p,h} + x^r_{p,h}, \qquad \forall (r,h,p)_{(-P)} \tag{12}$$
$$v^r_{p+1,h} \geq v^r_{p,h} + x^r_{p,h} - (1 - x^r_{p,h})(p+1), \qquad \forall (r,h,p)_{(-P)} \tag{13}$$
$$x^r_{p,h} \text{ binary}; u^r_{p,h}, v^r_{p,h}, \text{ and } a^r_{p,h} \text{ integer}, \qquad \forall (r,h,p). \tag{14}$$

Constraints (5)-(6) capture the status of active promotions in period 0. Constraints (7), (8)-(9), (10), and (11)-(13) model conditions (1), (2), (3), and (4), respectively. Variable domains are specified by (14).

An optimal state trajectory belongs to $\mathcal{F}(\mathcal{T}^{\mathrm{P}})$ and, in addition, maximizes the lift in sales volume over the planning horizon. Lift is defined as the percentage increase in weekly sales volume over the sales volume when there are no promotions (referred to as base sales), where the percentage is computed with respect to base sales. The lift is the output of a predictive model and is a function of the SM state vector, as well as, other factors such as non-shopper marketing promotions, weather, etc. We thus define an exogenous state $\boldsymbol{es}^r_p$, which is a vector containing all exogenous features (possibly forecasts) related to retailer r at planning period p. Specifically, the period p lift corresponding to retailer r at state $(\boldsymbol{s}^r_p, \boldsymbol{es}^r_p)$ is denoted by the function $L^r_p(\boldsymbol{s}^r_p, \boldsymbol{es}^r_p)$.

The choice of SM tactics can be formulated as the optimization problem shown below.

$$\max_{s} \sum_{r \in \mathcal{R}} \sum_{p \in \mathcal{T}^{\mathrm{P}} \setminus \{0\}} L^r_p(\boldsymbol{s}^r_p, \boldsymbol{es}^r_p) \tag{TPO}$$
$$\text{s.t.} \quad (\boldsymbol{s}^r_p, (r,p) \in \mathcal{R} \times \mathcal{T}^{\mathrm{P}}) \in \mathcal{F}(\mathcal{T}^{\mathrm{P}}).$$

The math program (TPO) is a tactic planning model that determines the SM state trajectory. We will discuss in §3 several business rules and intuitive specifications of the lift model that can be used in this framework so that (TPO) can be solved as an integer program using an off-the-shelf commercial solver such as GUROBI or CPLEX.

2.2 Lift Attribution

The (TPO) formulation discussed in §2.1 employs predictive lift models $L^r_p(\boldsymbol{s}^r_p, \boldsymbol{es}^r_p)$ for each $(r,p) \in \mathcal{R} \times \mathcal{T}^{\mathrm{P}}$ to inform sequential SM tactic choice. Below we propose a data-driven approach to construct such a predictive model that is consistent with (TPO). To simplify exposition, we will assume that $\mathcal{T}^{\mathrm{P}}$ contains distinct time periods and that we want to train a different model for each period. For example, $\mathcal{T}^{\mathrm{P}}$ could contain all the weeks in a year.

The periods in our training set are contained in $\mathcal{T}^{\mathrm{T}}$. Historical lift is computed from point-of-sale and sales analytics data. Let S^r_t and BS^r_t denote the total sales and base sales volumes for training period t. We reiterate that base sales volume is an estimate of the sales volume if there were no promotions and is provided by an external vendor such as Nielsen. The lift estimate for training period t is then defined as $\tilde{L}^r_p := (S^r_t - BS^r_t)/BS^r_t$, where we use $\tilde{\cdot}$ to highlight that it is a historical value. Next, the decision data includes all the information regarding the SM tactics used during the training periods. This information can be encoded as a historical SM state trajectory $(\tilde{\boldsymbol{s}}^r_t, t \in \mathcal{T}^{\mathrm{T}})$. Similarly, historical data regarding non-SM features are captured in an exogenous state trajectory $(\tilde{\boldsymbol{es}}^r_t, t \in \mathcal{T}^{\mathrm{T}})$. For a given $t \in \mathcal{T}^{\mathrm{T}}$, we use $p(t)$ to reference the corresponding period in $\mathcal{T}^{\mathrm{P}}$. For example, if $\mathcal{T}^{\mathrm{P}}$ are indices of the weeks of a year, then the indices t_1 and t_2 for the third weeks of 2015 and 2016, respectively, satisfy $p(t_1) = p(t_2) = 3$.

We specify $L^r_p(\boldsymbol{s}^r_p, \boldsymbol{es}^r_p)$ to be a parametric model at a coarser time scale than $\mathcal{T}^{\mathrm{P}}$ (and hence $\mathcal{T}^{\mathrm{T}}$) so that the underlying model can learn the behavior of lift due to SM tactics at the finer time scale of the training set. For example, the planning and training periods could refer to weeks while the model parameters change for each month of the year. In more detail, we assume that parameters of the lift model are indexed by the time periods in set $\mathcal{T}^{\mathrm{M}}$ ($\subseteq \mathcal{T}^{\mathrm{P}}$) such that any $p \in \mathcal{T}^{\mathrm{P}}$ and $t \in \mathcal{T}^{\mathrm{T}}$ map to $m(p) \in \mathcal{T}^{\mathrm{M}}$ and $m(t) \in \mathcal{T}^{\mathrm{M}}$, respectively. We decompose lift into components attributable to external factors and SM tactics. Each component is represented as a linear combination of pre-defined "basis" functions. The k-th basis functions modeling the exogenous and SM lift components at period $m \in \mathcal{T}^{\mathrm{M}}$ are denoted by $\phi_k(\boldsymbol{s}^r_p)$ and $\psi_k(\boldsymbol{es}^r_p)$, respectively.

The predicted lift is

$$L^r_p(\boldsymbol{\theta}, \boldsymbol{\beta}; \boldsymbol{s}^r_p, \boldsymbol{es}^r_p) := \mathrm{EL}^r_p(\boldsymbol{\theta}; \boldsymbol{es}^r_p) + \mathrm{SML}^r_p(\boldsymbol{\beta}; \boldsymbol{s}^r_p);$$
$$\mathrm{EL}^r_p(\boldsymbol{\theta}; \boldsymbol{es}^r_p) := \left\langle \Psi(\boldsymbol{es}^r_p), \boldsymbol{\theta}^r_{m(p)} \right\rangle;$$
$$\mathrm{SML}^r_p(\boldsymbol{\beta}; \boldsymbol{s}^r_p) := \left\langle \Phi(\boldsymbol{s}^r_p), \boldsymbol{\beta}^r_{m(p)} \right\rangle;$$

where $\Phi(\cdot) = (\phi_1(\cdot), \ldots, \phi_K(\cdot))^\intercal$, $\Psi(\cdot) = (\psi_1(\cdot), \ldots, \psi_K(\cdot))^\intercal$, $\boldsymbol{\beta}^r_m := (\beta^r_{m,1}, \ldots, \beta^r_{m,K})$, $\boldsymbol{\theta}^r_m := (\theta^r_{m,1}, \ldots, \theta^r_{m,K})$, $\boldsymbol{\beta} := (\boldsymbol{\beta}^r_m, (r,m) \in \mathcal{R} \times \mathcal{T}^{\mathrm{M}})$, and $\boldsymbol{\theta} := (\boldsymbol{\theta}^r_m, (r,m) \in \mathcal{R} \times \mathcal{T}^{\mathrm{M}})$. The notation for the lift function $L^r_p(\boldsymbol{\theta}, \boldsymbol{\beta}; \boldsymbol{s}^r_p, \boldsymbol{es}^r_p)$ is an extended version of $L^r_p(\boldsymbol{s}^r_p, \boldsymbol{es}^r_p)$ explicitly showing model parameters. Computing the coefficient vectors $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ specifies the predictive model, and in particular, the vector $\boldsymbol{\beta}$ provides an attribution of lift to the use of SM tactics.

The parameters of the model are computed by solving the following lift attribution inverse optimization (LAIO) problem:

$$\max_{\boldsymbol{\beta}, \boldsymbol{\theta}} \sum_{r \in \mathcal{R}} \sum_{t \in \mathcal{T}^{\mathrm{T}}} L^r_{p(t)}\left(\boldsymbol{\theta}, \boldsymbol{\beta}; \tilde{\boldsymbol{s}}^r_t, \tilde{\boldsymbol{es}}^r_t\right) - \lambda \|(\boldsymbol{\theta}, \boldsymbol{\beta})\|_1$$

$$\text{s.t.} \quad L^r_{p(t)}\left(\boldsymbol{\theta}, \boldsymbol{\beta}; \tilde{s}^r_t, \tilde{es}^r_t\right) \le \tilde{L}^r_t, \forall (r,t) \in \mathcal{R} \times \mathcal{T}^{\mathrm{T}} \quad (15)$$

$$\sum_{t \in \mathcal{T}^{\mathrm{T}}} \left[L^r_{p(t)}\left(\boldsymbol{\theta}, \boldsymbol{\beta}; \boldsymbol{s}^r_t, \tilde{es}^r_t\right) - L^r_{p(t)}\left(\boldsymbol{\theta}, \boldsymbol{\beta}; \tilde{s}^r_t, \tilde{es}^r_t\right) \right] \le \mu$$

$$\forall \{\boldsymbol{s}^r_t, (r,t) \in \mathcal{R} \times \mathcal{T}^{\mathrm{T}}\} \in \mathcal{F}(\mathcal{T}^{\mathrm{T}}) \quad (16)$$

$$\boldsymbol{\beta}^r_p \in \mathcal{B}^r_p, \forall (r,p) \in \mathcal{R} \times \mathcal{T}^{\mathrm{P}}. \quad (17)$$

The objective of (LAIO) is to maximize the regularized cumulative lift across all retailers on the observed trajectory $(\tilde{s}^r_t, (r,t) \in \mathcal{R} \times \mathcal{T}^{\mathrm{T}})$. Regularization is used to avoid over-fitting and handle degenerate optimal solutions. Constraints (15) ensure that the predicted lift is a lower bound on the observed lift. Constraints (16) allow the cumulative lift on the observed SM state trajectory to be μ worse than the analogous lift values on other feasible SM state trajectories. Finally, conditions (17) specify the domain of the model parameters. Indeed, any additional side information, including business knowledge on the impact of promotions, can be added as explicit constraints in (LAIO).

Since regularization using a one-norm in the objective function has a well-known linear representation, (LAIO) is a linear program. It embeds useful information to help train a predictive lift model. First, the lower bound condition (15) makes explicit that the predictive lift model being built will, for all practical purposes, be unable to explain all the lift in the training set (or the validation set for that matter) using the features in $\boldsymbol{s}$ and $\boldsymbol{es}$. Thus, a conservative predictive model of lift is sought. Next, (LAIO) is consistent with (TPO) due to constraints (16). These constraints are a relaxation of the optimality condition originally used in IRL, where the observed SM state trajectory is required to provide lift that is no worse than all other feasible SM state trajectories. In particular, the business rules used for planning are implicitly incorporated in the set $\mathcal{F}(\mathcal{T}^{\mathrm{T}})$ indexing these constraints. This consistency with (TPO) is not free, however, as it leads to (LAIO) becoming a linear program with a potentially large number of constraints due to (16). Nevertheless, this issue can be handled as discussed in §3.3.

3 SMOILE CONFIGURATION AND IMPLEMENTATION

The SMOILE framework described in §2 can be configured in various ways. In this section, we illustrate some possibilities, which we also use in our numerical study. In §3.1 and §3.2 we describe business rules and lift models, respectively, that can be used with SMOILE. In §3.3, we provide possible steps for implementation.

3.1 Business Rules

The definition of a feasible SM state trajectory in $\mathcal{F}(\mathcal{T}^{\mathrm{P}})$ using constraints (5)-(14) has an integer linear representation. This definition can be extended to include various business rules common in SM tactic planning, while maintaining the integer linear structure of the feasible set. We outline a few examples below.

Budget constraint. SM campaigns are often planned around budget restrictions since using an SM tactic of type $h \in \mathcal{H}$ at retailer $r \in \mathcal{R}$ in period $p \in \mathcal{T}^{\mathrm{P}}$ incurs a cost, which we denote by $c^r_{p,h}$. Denoting the available budget for the campaign by B, the budget constraint can be modeled as

$$\sum_{p \in \mathcal{T}^{\mathrm{P}}} \sum_{r \in \mathcal{R}} \sum_{h \in \mathcal{H}} c^r_{p,h} x^r_{p,h} \le B. \quad (18)$$

Maximum active tactics. To avoid SM overexposure, a brand may want to impose a maximum N^r_p on the number of active tactics in a period p shown to customers at retailer r. This requirement can be enforced as

$$\sum_{h \in \mathcal{H}} x^r_{p,h} \le N^r_p, \forall (r,p) \in \mathcal{R} \times \mathcal{T}^{\mathrm{P}}. \quad (19)$$

Scheduling constraints. There may be period within the planning horizon where it may not be possible to use a particular tactic. For example, space may not be available to use a demo during a certain month at the retailer. The condition that a tactic h cannot be planned at retailer r during periods $\mathcal{T}^0$ can be enforced by the constraints $x^r_{p,h} = 0$, $\forall p \in \mathcal{T}^0$. A maximum on the duration of tactics can be enforced via constraints $u^r_{p,h} + v^r_{p,h} \le UB, \forall (r,p,h) \in \mathcal{R} \times \mathcal{T}^{\mathrm{P}} \times \mathcal{H}$. Other scheduling rules, such as a minimum number of periods between two active tactics, can also be modeled by appending additional auxiliary variables.

3.2 Lift Models

Constructing the lift model in §2.2 requires defining $\mathrm{EL}^r_p(\boldsymbol{\theta}; \boldsymbol{es}^r_p)$ and $\mathrm{SML}^r_p(\boldsymbol{\beta}; \boldsymbol{s}^r_p)$. Specifically, we need to specify the basis function sets $\Phi(\cdot)$ and $\Psi(\cdot)$ corresponding to the SM and exogenous components of lift, respectively. There is extant literature on exogenous feature selection, including the marketing literature, that deals with the choice of $\Psi(\cdot)$ and its impact. In contrast, work on the choice of $\Phi(\cdot)$, that is features related to decision making and in particular SM, are scant to the best of our knowledge. Therefore, we present below possible definitions of $\mathrm{SML}^r_p(\boldsymbol{\beta}; \boldsymbol{s}^r_p)$, which are driven by specific choices of $\Phi(\cdot)$. To keep our discussion concrete, we will assume $\mathcal{T}^{\mathrm{P}}$ and $\mathcal{T}^{\mathrm{T}}$ contain weekly periods while $\mathcal{T}^{\mathrm{M}}$ is at the monthly scale.

We begin with a simple SM lift model that depends only on the variable x in the state vector.

$$\text{Model 1: } \mathrm{SML}^r_p(\boldsymbol{\alpha}, \boldsymbol{\gamma}; \boldsymbol{s}^r_p) := \sum_{h \in \mathcal{H}} x^r_{p,h} \alpha^r_{m(p),h} + \sum_{h \in \mathcal{H}} \sum_{h' \in \mathcal{H}, h' > h} \mathbb{I}(x^r_{p,h} \cdot x^r_{p,h'}) \gamma^r_{m(p),h,h'}.$$

The coefficients $\boldsymbol{\alpha}$ attempt to learn the average effect of each SM tactic in a month, which we call the *base effect* of a tactic, while the parameters $\boldsymbol{\gamma}$ capture the *interaction effect* between tactics. Here, we have $\boldsymbol{\beta} = (\boldsymbol{\alpha}, \boldsymbol{\gamma})$, $\boldsymbol{\alpha} \ge 0$ to enforce that the base SM tactic effect does not hurt sales (although it could have no impact), and $\boldsymbol{\gamma}$ is unrestricted to allow for a pair of tactics to reinforce and cannibalize each other, respectively, if this coefficient is positive and negative.

We now extend $\mathrm{SML}^r_p(\boldsymbol{\alpha}, \boldsymbol{\gamma}; \boldsymbol{s}^r_p)$ to include structural features of lift due to SM tactics known in practice. Our first extension is

$$\text{Model 2: } \mathrm{SML}^r_p(\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\mu}; \boldsymbol{s}^r_p) := \sum_{h \in \mathcal{H}} x^r_{p,h} \alpha^r_{m(p),h} + \sum_{h \in \mathcal{H}} \sum_{h' \in \mathcal{H}, h' > h} \mathbb{I}(x^r_{p,h} \cdot x^r_{p,h'}) \gamma^r_{m(p),h,h'}$$

$$- \sum_{h \in \mathcal{H}} u^r_{p,h} \mu^r_{m(p),h}.$$

We have $\boldsymbol{\beta} = (\boldsymbol{\alpha} \geq 0, \boldsymbol{\gamma}, \boldsymbol{\mu} \geq 0)$. The parameter vector $\boldsymbol{\mu}$ accounts for a *waiting effect* as it is associated with the state variable $\boldsymbol{u}$, which is the number of weeks remaining of an active tactic.

Example 2. *Consider Figure 1, the retailer Target, the tactic in-store discount, and assume week 0 corresponds to the start of a month. The base effect due to an in-store discount is the same for weeks 0, 1, 2, and 3 because $\boldsymbol{\alpha}$ is defined over $\mathcal{T}^{\mathrm{M}}$, which corresponds to months. Therefore, the active in-store discount at Target from weeks 0 to 2 (see Figure 1) and a longer in-store discount at this retailer from weeks 0 to 3 would have the same base effect for the common weeks 0, 1, and 2, which is unrealistic because a customer can wait one extra week to benefit from the tactic in the latter case. Model 2 corrects for this issue, if the corresponding μ coefficient is strictly positive. For instance, in week 0, the number of remaining weeks in the 3- and 4-week long in-store discounts would be* 2 *and* 3*, respectively, which causes the week 0 SM lift for the longer tactic to be lower. The same applies for weeks 1 and 2. Note that the 4-week tactic can still have larger total lift than the 3-week tactic since its lift in week 3 may be positive while the latter tactic would be inactive in this week.*

Finally, we consider an SM lift specification that also accounts for a *satiation effect*. For a given tactic, this effect requires the lift for week p to be a decreasing function of the number of weeks this tactic has already run since initiation. Using the state variable v to account for tactic use, we have our third SM lift model:

$$\text{Model 3: } \mathrm{SML}^r_p(\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\mu}, \boldsymbol{\eta}; s^r_p) := \sum_{h \in \mathcal{H}} x^r_{p,h} \alpha^r_{m(p),h} + \sum_{h \in \mathcal{H}} \sum_{h' \in \mathcal{H}, h' > h} \mathbb{I}(x^r_{p,h} \cdot x^r_{p,h'}) \gamma^r_{m(p),h,h'} - \sum_{h \in \mathcal{H}} u^r_{p,h} \mu^r_{m(p),h} - \sum_{h \in \mathcal{H}} v^r_{p,h} \eta^r_{m(p),h},$$

where $\boldsymbol{\beta} = (\boldsymbol{\alpha} \geq 0, \boldsymbol{\gamma}, \boldsymbol{\mu} \geq 0, \boldsymbol{\eta} \geq 0)$.

3.3 SMOILE Implementation

Using SMOILE involves (i) choosing a lift model; (ii) determining the business rules needed for defining the planning process and set $\mathcal{F}(\mathcal{T}^{\mathrm{P}})$ (see §3.1); (iii) solving (LAIO) to train the lift model; and (iv) solving (TPO) to plan future SM tactics. We summarize below how we implement these steps in our numerical study.

Step (i). For the SM component of lift, we implement models 1, 2, and 3 shown in §3.2. We consider a very simple specification for the exogenous component of lift involving only an intercept in the numerical study, that is, $\mathrm{EL}^r_p(\boldsymbol{\theta}; \boldsymbol{es}^r_p) = \theta^r_{m(p),1}$. The reasoning for this admittedly simple choice is two fold. First, the focus of the numerical study and our paper is to highlight the importance of incorporating structural information in (LAIO) about SM decisions in order to obtain reasonable lift predictions for tactic planning. A simple specification for EL suffices for this purpose. Second, even an attempt to construct an elaborate model for EL will not be precise in reality, that is, there will very likely be exogenous factors missing [13, chapter 6.2]. Therefore, incorporating SM decisions remains relevant under a complex exogenous lift model.

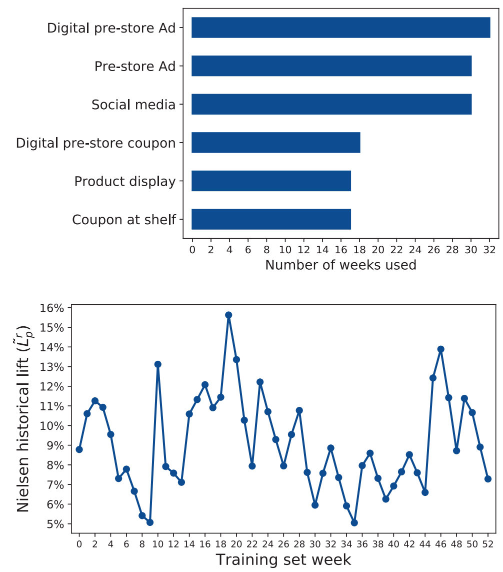

Figure 2: R1 historical lift and top SM tactic usage.

Step (ii). We add business rules (18) and (19) to our definition of the numerical study.

Step (iii). Although (LAIO) is a linear program, it has a large number of constraints (16), which makes its direct solution challenging. We thus employ constraint generation (CG) to solve this model (see, e.g, [2]), as we summarize below. The CG procedure starts with a relaxation of (LAIO) without constraints (16). This relaxation can be solved using an off-the-shelf solver such as GUROBI to obtain an optimal solution. At this solution, we find the most violated constraint in (16), add this to the strengthen the existing relaxation, and repeat. We stop when the solution to a one of these relaxations satisfies all the the constraints in (16). The calibration of (LAIO) depends only on two parameters μ and λ with intuitive interpretations (see §2.2). We choose them using cross validation.

Step (iv). When using the lift models and business rules discussed in §3.2 and §3.1, respectively, (TPO) is an integer program that can be solved using an off-the-shelf solver such as GUROBI.

4 NUMERICAL STUDY

We evaluate SMOILE for SM lift prediction and planning using real data across two retailers and two product categories. Our implementation follows §3.3. We describe the data and evaluation procedures in §4.1 and §4.2, respectively. We present results to highlight the computational performance of SMOILE in §4.3 and then discuss insights in §4.4.

4.1 Data and Computational Setup

We use data from brands B1 and B2 interested in planning SM tactics for the product categories "wings" and "frozen breakfast", respectively, at two US retailers, which we refer to as retailer R1 and retailer R2 (actual names anonymized for confidentiality). We have access to Nielson point-of-sales data and sales analytics for each retailer and product, from which we can compute lift as explained in §2.2. We also have data from the brands on the set of SM tactics used, the average cost per week associated with using an SM tactic, and the budget used for promotions. This information is available for all weeks in 2015 for both retailers and in 2016 for the first 12 and 8 weeks, respectively, for R1 and R2. The number of distinct SM tactics used at R1 and R2 were 11 and 9, respectively. Figure 2 provides some statistics from our dataset regarding sales lift and the usage of the most popular SM tactics at R1.

We constructed single retailer and multiple retailer SM planning instances with the training set containing the 52-weeks of data from 2015 (i.e., $P = 52$). For the single retailer case, we discuss results for only R1 for brevity and use the 12 weeks in 2016 for validation. In the multiple-retailer instance, the validation set contains the common 8 weeks in the 2016 data for both retailers.

4.2 Evaluation Procedure

SMOILE has both prediction and planning components. We thus require procedures to assess its performance in each of these tasks. The objectives of (TPO) and (LAIO) are the cumulative lift over the planning and training periods, respectively. To validate predictive performance, we benchmark against the lift estimate of Nielsen, which is a trusted reference in the retail industry. Specifically, we compute the mean lift per week predicted by the trained model when employing the historical SM tactic decisions in the validation set and compare this mean against the Nielsen mean weekly lift over this period. To avoid a particular week having an overbearing effect, we compute the mean and maximum of the absolute deviations between our model and Nielsen mean weekly lift values over subsets of the validation set obtained by moving the starting period forward sequentially. We use the maximum absolute deviation as our metric when performing cross-validation to understand the impact of μ and λ.

While the predictive performance of mean lift is important, the goal of (LAIO) is to attribute lift to SM tactics, that is, the component SML in the notation of §3.2. Understanding if the prediction of the SM component of lift is reasonable is difficult because there are no benchmarks available for comparison. For instance, Nielsen does not report SM specific lift. However, from the experience of our industry partner, average weekly SM-specific lift is typically close to 3% and less than 5%. We use this guideline as a sanity check.

Next, we ask if the IRL constraints (16), which explicitly leverage past SM decision data, add any value in the prediction process. To understand this, we compare against the predictive performance of models trained by a relaxation of (LAIO) without these constraints, which we refer to as (R-LAIO).

Finally, we evaluate the planning performance of (TPO) when using the best lift model that we identify from cross-validation. We compute the mean weekly lift due to the optimal actions determined by (TPO) and compare it against the Nielsen average lift. We also compare the mean SM lift per week when using optimal SM decisions of (TPO) and historical SM decisions in the validation set.

4.3 Computational Performance

We discuss the performance of the predictive and planning procedures in SMOILE on the R1 single retailer instance.

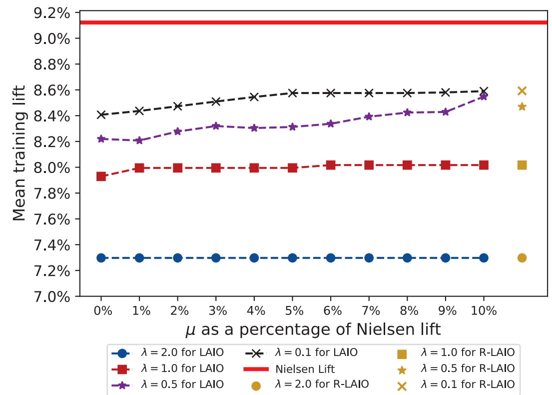

Figure 3: Weekly lift on the training set with Model 3.

Predictive performance. Figure 3 displays the mean weekly lift associated with Model 3 on the training set for $\lambda \in \{2, 1, 0.5, 0.1\}$ and $\mu \in \{0\%, 1\%, \ldots, 10\%\}$, where μ is expressed as a percentage of the Nielsen cumulative lift. This figure also shows the mean weekly lift associated with (R-LAIO). For a fixed λ, the cumulative lift of (LAIO) increases with μ, showing that the IRL constraints are not redundant. Thus, incorporating the SM decision data into the training process affects the values of the computed model parameters. Indeed, as μ becomes large, the IRL constraints become less constraining and the mean weekly lift tends to the mean lift of (R-LAIO), which has no IRL constraints. For λ equals 2, the optimal solution of (R-LAIO) already satisfies the IRL constraints and thus varying μ has no effect. The behavior for models 1 and 2 is similar. Compared to the Nielsen mean lift of 9.12% on the training set, the largest mean lift of Model 3 shown in Figure 3 is 8.58%.

On the validation set, Figure 4 shows the maximum absolute deviation between our weekly lift and the Nielsen weekly lift for Model 3. We find that the minimum of this metric is 1.95% and occurs for λ equals to 0.1 and μ equal to 8%. In contrast, the maximum absolute deviation for (R-LAIO) is 2.5%, that is, a larger value and occurs for λ equals to 0.1. This difference indicates that the presence of IRL constraints in (LAIO) results in a more robust calibration of Model 3 with improved predictive performance on our data.

Next, we summarize in Figure 5 the performance of models 1, 2, and 3 for their best (λ, μ) values of $(0.5, 3\%)$, $(0.5, 6\%)$, and $(0.1, 8\%)$, respectively, chosen via cross validation. This figure also includes the 95% confidence intervals of the lift estimates over the subsets of the validation set that we average over (see related discussion at the beginning of §4.2). At first sight, the performance of Model 1 appears to be the best in terms of predicted total weekly lift. However, a closer look suggests that this performance is potentially spurious. To elaborate, recall from §3.3 that the exogenous lift component we use is an intercept term that does not change with the week within a month. Therefore, if Model 1 explains all the lift in the validation set, it also implies that the SM component explains all the intra-month lift variation, which is unlikely since a host of other factors affect such weekly changes. Models 2 and 3, which incorporate behavioral structure, fall short of explaining all the lift

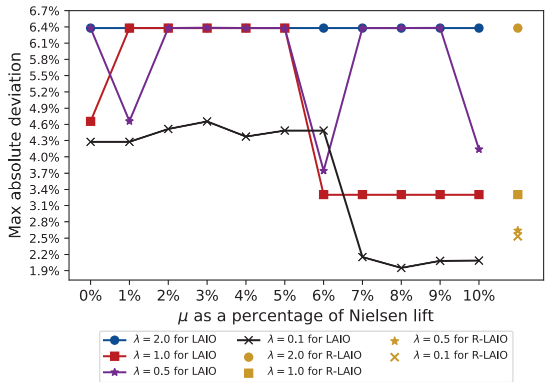

Figure 4: Maximum absolute deviation between weekly lifts of Model 3 and Nielsen on the validation set.

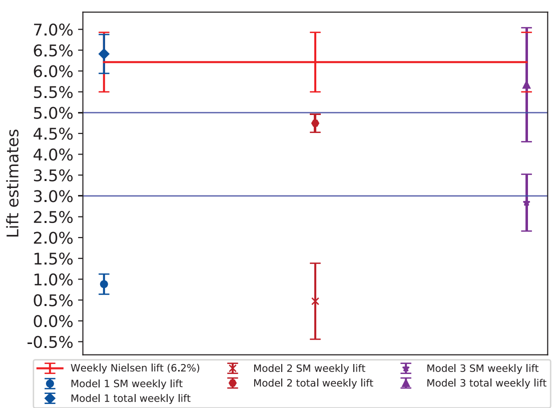

Figure 5: Performance of models 1, 2, and 3 for the optimal choices of (λ, μ) on the validation set.

and do not suffer from this behavior. It is also instructive to look at the SM component of models 2 and 3. The former has an SM lift percentage of 0.36% while this percentage in the latter model is roughly 2.8%. It is encouraging that the SM lift predicted by Model 3 is close to the practitioner guideline of 3%.

Planning performance. Table 2 reports the weekly lift due to historical and optimal SM tactics. For model 2, these two figures coincide but for models 1 and 3 they differ by more than 1.0-1.5%, which is significant when one considers the small net margin of retailers. The difference in the SM component of lift between the optimal and historical decisions is also striking. Based on the validation test analysis above, we identified Model 3 as being most reasonable, with its optimal SM tactics resulting in an increase in historical SM lift by over 1.5%, which is again significant.

4.4 Insights

The encouraging performance of Model 3 warrants a closer look into the actual SM tactics that it prescribes in the validation set. Figure

Table 2: Comparison of lift on validation set due to optimal and historical SM tactics.

	Weekly total lift		Weekly SM lift	
Model	Historical tactics	Optimal tactics	Historical tactics	Optimal tactics
1	7.66	9.25	0.83	4.71
2	8.75	8.75	0.61	2.91
3	7.11	8.84	3.02	4.75

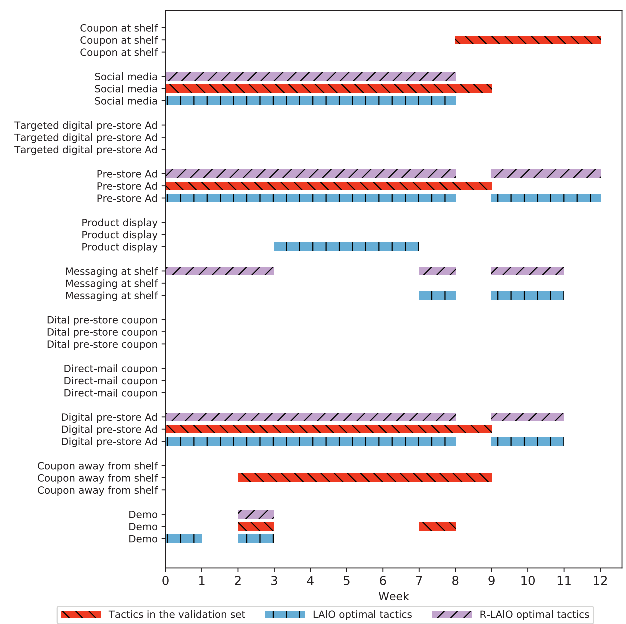

Figure 6: SM tactic decisions used in the validation set for the single-retailer (R1) instance.

6 displays this information for the single retailer instance. The optimal tactics differ from the historical tactics with and without the IRL constraints (i.e., LAIO and R-LAIO). In particular, it appears Model 3 deems it unfavorable to use the "coupon at shelf" and "coupon away from shelf" tactics that were both present in the validation set and instead employs more pre-store ads. The optimal tactics with IRL constraints also use "product display" which is not used by the model from (R-LAIO) on the validation set.

Figure 7 shows similar information on the multi-retailer instance. In this setting, it is clear that the SM tactics vary significantly by retailer. Specifically, several of the tactics used in Retailer 2 are stopped (e.g., "pre-store coupon", "social media", and "product display") even though the same categories appear favorable tactic investments at Retailer 1. The budget saved from not using these promotions is instead being redirected to tactics such as digital pre-store coupons. Finally, we investigated the percentage of the predicted lift of Model 3 attributable to the waiting and satiation effects on the two retailer instance. We found that these percentages are 1.21% and 0%, which suggests that waiting decreases lift prominently on our instance while the satiation effect appears negligible.

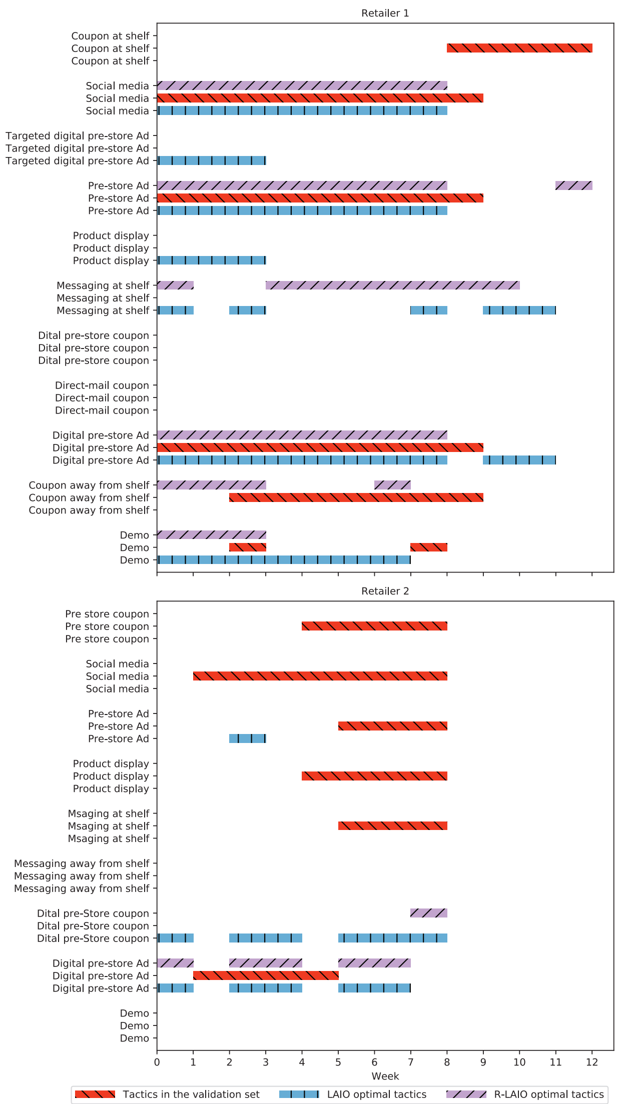

Figure 7: SM tactic decisions used in the validation set for the multi-retailer instance.

5 CONCLUSIONS

Shopper marketing (SM) represents an important form of marketing for retailers to increase sales. Uncovering the SM tactics, or combinations thereof, that drive sales lift remains a challenging task. Marketers typically rely on historical sales data at retail outlets as well as third-party supplied point of sales (POS) data to attribute lift to SM tactics and then subsequently use these estimates to plan new SM campaigns. Traditional statistical approaches used in this process find it difficult to directly incorporate salient aspects of SM planning: (i) the temporal link between SM decisions made over time; (ii) the impact of business rules on lift associated with SM tactics; and (iii) the need to adhere to these rules when planning future SM tactics. The result is an ad-hoc work flow with significant hand-engineering to obtain implementable SM campaigns.

We propose Shopper Marketing Optimization and Inverse Learning Engine (SMOILE), which is a data-driven approach that combines empirical optimization and inverse reinforcement learning. SMOILE streamlines the use of data, directly couples decision data with the training of the lift model, and results in a training process that is consistent with the planning process. Moreover, training a lift model via SMOILE requires tuning only two interpretable parameters, which can be done via cross validation, thus side-stepping ad-hoc manual intervention often needed when using traditional marketing mix models.

We use a dataset based on two US retailers and brands to illustrate the use of SMOILE and its benefits. We find that training the SM lift model in a manner consistent with the SM planning phase results in better predictive performance and mitigates spurious lift predictions. Moreover, the SMOILE SM tactic decisions outperform the historical SM tactics when evaluated using the trained lift model.

REFERENCES

[1] Pieter Abbeel and Andrew Y Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*. ACM, 1.

[2] Daniel Adelman and Diego Klabjan. 2012. Computing near-optimal policies in generalized joint replenishment. *INFORMS Journal on Computing* 24, 1 (2012), 148–164.

[3] Peter L Bartlett and Shahar Mendelson. 2006. Empirical minimization. *Probability Theory and Related Fields* 135, 3 (2006), 311–334.

[4] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67 (2015). Issue 1.

[5] David Chan and Mike Perry. 2017. *Challenges And Opportunities In Media Mix Modeling*. Technical Report. Google Inc.

[6] Tom Compernolle, Mark Baum, Pat Walsh, Paul Christman, and Jean Fally. 2015. *From category management to shopper-centric retailing: It can be done – here's how.* Technical Report. Delotte Consulting, Food Marketing Institute, and Winston, Weber, and Assocates.

[7] Mathew Egol and Edward Landry. 2009. *Shopper marketing 3.0: Unleashing the Next Wave of Value.* Technical Report. Booz and Company.

[8] Yuxue Jin, Yueqing Wang, Yunting Sun, David Chan, and Jim Koehler. 2017. *Bayesian Methods for Media Mix Modeling with Carryover and Shape Effects.* Technical Report. Google Inc.

[9] Anil Kamath. 2015. Optimizing Marketing Impact through Data Driven Decisioning. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1631–1631.

[10] Yi-hao Kao, Benjamin V Roy, and Xiang Yan. 2009. Directed regression. In *Advances in Neural Information Processing Systems*. 889–897.

[11] Yi-Hao Kao and Benjamin Van Roy. 2012. Directed time series regression for control. *arXiv preprint arXiv:1206.6141* (2012).

[12] Andrew Y Ng and Stuart J Russell. 2000. Algorithms for inverse reinforcement learning.. In *International connference on machine learning*. 663–670.

[13] Judea Pearl. 2009. *Causality*. Cambridge university press.

[14] Gregory Piatetsky-Shapiro and Brij Masand. 1999. Estimating campaign benefits and modeling lift. In *Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 185–193.

[15] Peter E Rossi. 2014. Even the rich can make themselves poor: A critical examination of IV methods in marketing applications. *Marketing Science* 33, 5 (2014), 655–672.

[16] Venkatesh Shankar, J Jeffrey Inman, Murali Mantrala, Eileen Kelley, and Ross Rizley. 2011. Innovations in shopper marketing: current insights and future research issues. *Journal of Retailing* 87 (2011), S29–S42.

[17] Tanner Sorensen and Shravan Vasishth. 2015. Bayesian linear mixed models using Stan: A tutorial for psychologists, linguists, and cognitive scientists. *arXiv preprint arXiv:1506.06201* (2015).

[18] Harald J Van Heerde and Scott A Neslin. 2008. Sales promotion models. In *Handbook of marketing decision models*. Springer, 107–162.

[19] Michel Wedel and PK Kannan. 2016. Marketing analytics for data-rich environments. *Journal of Marketing* 80, 6 (2016), 97–121.

[20] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. 2008. Maximum Entropy Inverse Reinforcement Learning.. In *AAAI*, Vol. 8. Chicago, IL, USA, 1433–1438.