ADS THAT STICK: NEAR-OPTIMAL AD OPTIMIZATION THROUGH PSYCHOLOGICAL BEHAVIOR MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Optimizing the timing and frequency of advertisements (ads) is a central problem in digital advertising, with significant economic consequences. Existing scheduling policies rely on simple heuristics, such as uniform spacing and frequency caps, that overlook long-term user interest. However, it is well-known that users' long-term interest and engagement result from the interplay of several psychological effects (Curmei, Haupt, Recht, and Hadfield-Menell, ACM CRS, 2022).

In this work, we model change in user interest upon showing ads based on three key psychological principles: *mere exposure*, *hedonic adaptation*, and *operant conditioning*. The first two effects are modeled using a concave function of user interest with repeated exposure, while the third effect is modeled using a *temporal decay* function, which explains the decline in user interest due to overexposure. Under our psychological behavior model, we ask the following question: Given a continuous time interval T, how many ads should be shown, and at what times, to maximize the user interest towards the ads?

Towards answering this question, we first show that, if the number of displayed ads is fixed, then the optimal ad-schedule only depends on the operant conditioning function. Our main result is a quasi-linear time algorithm that outputs a *near-optimal* ad-schedule, i.e., the difference in the performance of our schedule and the optimal schedule is exponentially small. Our algorithm leads to significant insights about optimal ad placement and shows that simple heuristics such as uniform spacing are sub-optimal under many natural settings. The optimal number of ads to display, which also depends on the mere exposure and hedonistic adaptation functions, can be found through a simple linear search given the above algorithm. We further support our findings with experimental results, demonstrating that our strategy outperforms various baselines.

1 INTRODUCTION

Digital advertising forms the backbone of today's trillion-dollar internet economy, serving as a primary channel for both acquiring new customers and sustaining engagement with existing ones. Since user attention is scarce and capturing it carries significant economic value, *optimizing the timing and frequency of ads* becomes a critical challenge for advertisers. This issue arises across many settings, including placing ads in (live) video or audio streams, sending push notifications to app users, or embedding sponsored content within or across user sessions. In each case, the objective is to maximize user engagement and recall while mitigating long-term *fatigue or satiation*.

A long line of empirical work in behavioral psychology has shown that the temporal spacing and frequency of ads have a significant effect on the memory retention and fatigue of the customer (Singh et al., 1994; Sahni, 2015; Curmei et al., 2022). For instance, following an initial positive neural response to repeated stimuli, individuals tend to revert toward a baseline level of interest, causing the "thrill" of the same message to fade. This initial boost is known as *mere exposure*, while the subsequent tapering-off is termed *hedonic adaptation* in the behavioral psychology literature (Curmei et al., 2022). Moreover, insufficient spacing between two ads can drain attention and affect memory retention in the long run (Singh et al., 1994). This effect is referred to as *operant conditioning*.

Despite the importance of this problem, there has been relatively little work studying the long-term cognitive and psychological effects of ad impressions (Sahni, 2015; Aravindakshan & Naik, 2015;

Rafieian, 2023) (see Section 1.1 for more details). Most existing approaches for ad placement rely on simple heuristics such as uniform spacing, front-loading, or frequency caps (Aravindakshan & Naik, 2015; Rafieian, 2023; Despotakis et al., 2021), or on short-sighted policies that treat each ad impression as independent of the others (Theocharous et al., 2015).

Recent empirical studies have tried to move beyond these simple heuristics. For instance, (Freeman et al., 2022) ran an experiment with 327 participants to test where ads should be placed in order to reduce negative reactions such as *anger, irritation, and perceived intrusiveness*. Their main hypothesis was that *"Mid-roll ads will elicit more anger and be seen as more intrusive than preroll ads"*, and their experiments confirmed this. They also tested several other hypotheses, all arriving at the same conclusion that placing ads at the beginning is generally more effective than inserting them in the middle. Similarly, (Ritter & Cho, 2009) conducted an experiment with 129 participants on audio podcasts. Their experiments supported the hypothesis that *"Advertising at the beginning of podcasts will generate less intrusiveness, less irritation, more favorable attitudes toward an ad, and less ad avoidance than advertising in the middle of podcasts"*. Other studies also suggest similar trends. For example, (Goldstein et al., 2011) argues that a mix of ad placements, at the beginning, at the end, and a small number evenly spaced in the middle, can create a more positive overall experience for users.

These studies motivate us to ask whether there is an underlying psychological reward model that can provide an explanation for these empirical findings:

Question 1.1. *Can we design a theoretical reward model for ad scheduling that captures the users' psychological behavior?*

Ideally, the optimal policy under this reward model should be consistent with the empirical observations cited above. Moreover, this model should have tunable parameters so that it can adapt to different real-world settings. This further motivates an algorithmic question about computing the optimal ad schedule under various parameter settings:

Question 1.2. *Under such a reward model, can the optimal ad schedule be computed efficiently?*

To answer Question 1.1, we study a dynamic model of user interest under repeated ad exposures. Let $n+1$ ads be shown at time $\bar{t} = (t_0, t_1, \ldots, t_n)$, where each $t_i \in [0, T]$ and $t_i \leq t_{i+1}$ for all $i \in \{0, 1, \ldots, n\}$. Here, t_i represents the time at which the i^{th} ad is shown, and without loss of generality, we set $t_0 = 0$ and $t_n = T$. We refer to $\bar{t}$ interchangeably as a *strategy* or a *schedule*. Given a time horizon T and a strategy $\bar{t}$, the reward obtained from showing the i^{th} ad at time t_i is denoted by $R(\bar{t}, i)$. [1] The total reward associated with strategy $\bar{t}$ is then defined as $R(\bar{t}) = \sum_{i=0}^{n} R(\bar{t}, i)$.

To capture user psychology in our model, we study $R(\bar{t}, i)$ as a combination of two functions. The first function captures the (positive) mere exposure and the hedonic adaptation effects and depends only on the number of ads shown previously. We denote this function by $B : \mathbb{Z}_{\geq 0} \to \mathbb{R}_+$, a concave function with respect to the number of ads shown till now, and represents the reward when the i^{th} ad is shown. The concavity of the function B is justified by the 'diminishing returns' property implied by mere exposure and hedonic adaptation, and is common in recent literature that considers the dynamic effect of actions (Patil et al., 2023; Blum & Ravichandran, 2025). The second function is a temporal exponential decay function with parameter $\delta \in [0, 1]$ which captures the (negative) operant conditioning effect so that reward at time t decreases by $\delta^{t-t'}$ for any ad shown previously at time t'. This temporal decay function is motivated by the classical theory of (Ebbinghaus, 1913) on forgetting curves, which hypothesizes the exponential decline of memory retention over time. This model is similar to the influential *goodwill stock* model of (Nerlove & Arrow, 1962), and *exponential discounting* models used in control theory (Leqi et al., 2021).

The parameter δ plays a key role in capturing different types of user psychological behavior, such as anger, irritation, intrusiveness, or interest. A low value of δ indicates that past ad impressions have little effect on the user and therefore additional ads do not strongly reduce engagement. In contrast, a high value of δ implies that ads have more long-term effects on the user (Curmei et al., 2022). This makes a user highly sensitive when an ad is shown, which can quickly lead to irritation, loss of interest, or a perception of intrusiveness. Thus, the parameter δ provides a compact way to encode the psychological effects observed in experimental studies such as (Freeman et al., 2022; Ritter & Cho, 2009; Goldstein et al., 2011).

[1] Throughout, we use the term "reward" to describe the advertiser's optimization objective. Depending on the context, this could represent probability of purchase, user satisfaction with product, or engagement with ads.

Given these functions, the goal is to find the number of ads $n+1$ and the strategy $\bar{t} = (t_0, \ldots, t_n) \in [0, T]$, such that the total reward $R(\bar{t})$ is maximized. Given this model of user reward, we first show that the problem has a special structure: if the number of ads $n+1$ is fixed, then the optimal display timing only depends on the function capturing operant conditioning.

To answer Question 1.2, our main algorithmic result is a quasi-linear time algorithm that, given a fixed number of ads $n+1$, outputs a *near-optimal* ad-schedule. Let t_i^* denote the optimal time to place the i-th ad. We show that each t_i^* can be approximated by t_i (produced by our algorithm) with exponentially small error, namely, $t_i^* - \frac{1}{2^n} \leq t_i \leq t_i^* + \frac{1}{2^n}$. Our result is based on several key insights: (1) the operant conditioning function is strictly convex, leading to a single global optima, (2) the global optima has a special structure such that each t_i can be recursively found using $t_1, \ldots, t_{i-1}$, (3) t_1 can be found using binary search with *bounded error* and, (4) error propagation in the recursive computation can be controlled. Next, to determine the optimal number of ads to display, which depends on the mere exposure and hedonic adaptation functions, we perform a simple linear search, using the result discussed above.

We further support our theoretical findings using simulations where we compare the performance of different strategies under our reward model. We demonstrate that our ad placement strategy outperforms other baselines for various values of $\delta \in [0, 1]$.

Key insights from our model and relation to previous empirical findings. Recall that experimental studies (Freeman et al., 2022; Ritter & Cho, 2009; Goldstein et al., 2011) suggest that placing ads at the beginning and end is generally more effective than placing them in the middle. While these findings provide useful evidence regarding effective ad placement, they are limited in scope and cannot establish that such a strategy is universally optimal across all scenarios. This motivates the need for a theoretical framework that can explain and generalize these observations. In particular, our model offers deeper insights into ad placement strategies by explicitly accounting for all values of $\delta \in [0, 1]$. To systematically understand how δ influences the optimal policy and how ad placement strategies may vary across different values of δ, we analyze the optimal ad policy in our model for all values of $\delta \in [0, 1]$. Our analysis demonstrates that the structure of the optimal schedule changes as δ changes. Building on this, our near-optimal ad scheduling algorithm provides not only practical scheduling strategies but also valuable theoretical insights into how ad placement should adapt across the full spectrum of δ. To be precise, we observe that for small δ, the ads in the optimal strategy $\bar{t}$ are placed almost evenly in $[T]$. As δ increases, more and more ads start to concentrate at 0 and T such that the first $t_i > 0$ moves towards 0, and the last $t_j < T$ moves towards T. The remaining ads are evenly spaced between t_i and t_j. All these insights show that the optimality of heuristics such as uniform spacing or placing more ads at the beginning, depend on the value of δ, highlighting the need to adapt schedules based on user behavior (i.e., δ).

1.1 Related Work

Behavioral psychology. While the work on behavioral psychology is vast, in our work we follow (Curmei et al., 2022) and focus on three well-studied phenomena from psychology: mere exposure, operant conditioning, and hedonic adaptation. Several works in behavioral psychology has empirically shown the effect of mere exposure and hedonic adaptation under various settings (Cox & Cox, 2002; Fang et al., 2007; Chugani et al., 2015; Yang & Galak, 2015). The most relevant to our work is the study of (Hekkert et al., 2013) who found that attractiveness of a product increased with the number of times it was shown to a user. Moreover, in a similar context (Nelson & Meyvis, 2008) also found evidence of hedonic adaptation as a function of the number of exposures. Similarly, operant conditioning has also been well-studied (See (Cooper et al., 2007) and references therein). While operant conditioning might have additional connotations in the psychology literature, we mainly use it to model the ‘annoyance’ or ‘satiation’ effect of repeated exposures (Sahni, 2015). There has been some effort on psychology-aware recommendation systems (Curmei et al., 2022; Jesse & Jannach, 2021), however, incorporating psychological effects in advertising has received less attention.

Value of δ in real world. There exists some work related to determining the exact curve to analyze user retention of content. In the work of (Curmei et al., 2022), the authors note that in real-world scenarios, the value of δ is 0.98. Several other works (Murre & Dros, 2015; Goldstein et al., 2011) provide motivation for setting δ between 0.7 and 0.99. This is also the regime where uniform spacing is a bad strategy as compared to the optimal ad schedule (see Section 6).

Ad scheduling. The problem of optimizing ad schedules has been studied across various communities such as marketing, operations research, and machine learning. One of the earliest work is due to (Nerlove & Arrow, 1962), who proposed the goodwill-stock model that treats advertising as an investment that builds a "stock" of consumer goodwill (or awareness) which then depreciates over time. Specifically, they considered a dynamical model of change in consumer goodwill given an ad impression, and cast the problem of optimizing the ad schedule as an *optimal control* problem. However, they do not consider memory/satiation effects, and the optimal policy under their model is to show most of the ads at the start of the time-horizon. (Naik et al., 1998) extended this work to consider memory effects, however, the optimal policy under their dynamical system is not interpretable. (Sahni, 2015) conducted field experiments to demonstrate that *temporal spacing* of ads has a large effect on the memory retention and satiation of the users. More recently, the problem of optimizing ad schedules has been casted as a reinforcement learning problem (Rafieian, 2023; Theocharous et al., 2015). While these algorithms tend to be general-purpose, their solutions are hard to interpret and can be difficult to execute in practice, given the complexity of ad exchanges (Despotakis et al., 2021).

Dynamic rewards and restless bandits. (Leqi et al., 2021) introduced a bandit framework that models user satiation using linear dynamical systems and showed that the greedy strategy is optimal when all arms have the same base reward and decay profile. More broadly, several works have explored reward structures where the current payoff depends on historical actions and decays over time (Heidari et al., 2016; Levine et al., 2017; Seznec et al., 2019). In contrast, (Kleinberg & Immorlica, 2018) studies a setting where rewards increase with time since the last pull. Related ideas appear in models where rewards evolve based on the number of pulls or the time elapsed since the last interaction (Cella & Cesa-Bianchi, 2020; Basu et al., 2019; Warlop et al., 2018; Mintz et al., 2020).

Digital advertising. In (Schwartz et al., 2017), the authors investigated customer acquisition through advertisements on online platforms using multi-armed bandits, and designed a policy that achieves an 8% improvement in acquisition rate. (Adany et al., 2013) addressed the problem of allocating personalized ads to users, considering each user's profile and estimated viewing capacity. (Seshadri et al., 2015) explored advertisement scheduling to meet advertisers' campaign goals while maximizing ad-sales revenue. (Dobrița et al., 2025) proposed a framework that leverages k-nearest neighbors to predict ad positions, enhancing ad scheduling optimization. Aiming to maximize viewership under budget constraints, (Czerniachowska, 2019) presented a scheduling solution aligned with advertisers' budgets. (Sumita et al., 2017) developed a $(1-\epsilon)$-competitive algorithm for envy-free allocation of video ads where $\epsilon > 0$ is a constant. We also note that digital advertising is a vast field with many practical considerations – here we summarize work that is closely related to our theoretical modeling.

1.2 Organization

In Section 2, we formalize the problem, and in Section 3, we present a near-optimal algorithm for scheduling ads. Next, in Section 4, we analyze this algorithm, followed by a discussion of the broader implications of our work in Section 5. We support our theoretical findings with experiments in Section 6, and conclude by outlining the key takeaways, limitations, and possible extensions in Section 7. All remaining details and proofs are provided in the Appendix.

2 Problem Description

Let T denote the total time horizon, which is known to us in advance. We consider the setting where we have to display $n+1$ homogeneous ads in the continuous time interval $[0, T]$. We assume for now that displaying these ads is instantaneous, though in Appendix E, we will show how to extend this to the case when each ad needs the same amount of time. As discussed before, the function $R(\bar{t}, i)$ is composed of two parts. The first part $B(i)$ is the reward when the i^{th} ad is shown (note that i could also be 0 and corresponds to the ad shown at t_0), and is simply a function of the number of times we have shown the ad previously, and not of the times at which we show ads. In this work, we consider $B(i)$ to be a concave function, for example, the sigmoid function $B(i) = \frac{1}{1+e^{-ci}}$, but in general, any function would work. The second function depends on both the number of times we show ads and the times t_i at which ads were shown. Let $\bar{t} = (t_0, t_1, \ldots, t_{n-1}, t_n)$ denote the time and order in which ads were displayed. The second function is $\gamma. \sum_{i=0}^{l-1} \delta^{t_l - t_i}$, where $\delta \in [0, 1]$ and $\gamma \in \mathbb{R}_+$ is a constant that parameterizes the strength of this effect. The term $\sum_{i=0}^{l-1} \delta^{t_l - t_i}$, $\delta \in [0, 1]$ captures the temporal decay in the loss when the l^{th} ad is shown. Hence, the reward for the i^{th} ad under strategy $\bar{t}$ is given by $R(\bar{t}, i) = B(i) - \gamma \sum_{j=0}^{i-1} \delta^{t_i - t_j}$.

Let $\tilde{n} \in \mathbb{N}$ be an upper bound on the number of ads that can be shown in $[0, T]$. Our objective is to find the number of ads $n + 1 \leq \tilde{n}$, and the strategy $\bar{t} = (t_0, \ldots, t_n)$ according to which ads should be shown so that the reward $R(\bar{t})$ is maximized.

The reward $R(\bar{t})$ can be written as:

$$R(\bar{t}) = \sum_{i=0}^{n} R(\bar{t}, i) = \sum_{i=0}^{n} \left(B(i) - \gamma \sum_{j=0}^{i-1} \delta^{t_i - t_j} \right) = \left(\sum_{i=0}^{n} B(i) \right) - \left(\gamma \sum_{j<i} (\delta^{t_i - t_j}) \right).$$

Observe that the first term $\sum_i B(i)$ is independent of the times at which we show the ads, and only depends on the number of ads themselves. The coefficient of the second term, i.e., γ, only plays the role of a scaling factor. Hence, maximizing $R(\bar{t})$ for a given value of n is equivalent to minimizing the *loss* $L(\bar{t})$, where $L(\bar{t})$ is defined as $L(\bar{t}) = \sum_{j<i} (\delta^{t_i - t_j})$.

Once we can find the optimal value of $R(\bar{t})$ for strategies that involve showing $n + 1$ ads, we can iterate over $\tilde{n}$ possible values to find the optimal number of ads to be shown. Next, we provide an overview of the algorithm to find the optimal number of ads to be shown to maximize the reward, and the optimal strategy $\bar{t}$ to show these ads.

3 Near-Optimal Algorithm for Ad Scheduling

In this section, we provide an overview of our near-optimal algorithm for ad scheduling. We first consider the case when we know $n + 1$, the number of ads to show. Given the time horizon T and δ, we want to compute the schedule $\bar{t} = (t_0, t_1, \ldots, t_n)$, such that for all $0 \leq i \leq n - 1$, $t_i \leq t_{i+1}$.

The first step in our algorithm is to compute the number of ads to show at time 0 and T. Note that multiple ads can be shown at the same time, given the instantaneous nature of the ads (we relax this assumption in Appendix E). Let $t_a > 0$ be the first non-zero time at which an ad is shown, i.e., $a - 1$ ads are shown at time 0. We will define the quantity T_i as $T_i := \delta^{t_i}$, and throughout the paper we derive our results in terms of T_i for $i \in \{0, \ldots, n\}$. Given the parameters n, T, δ, our first algorithm (Algorithm 1) outputs the value of a and $T_a = \delta^{t_a}$. Later in our analysis, we show that all the subsequent ad timings can be found once we know a and T_a. We now describe the algorithm to find a and T_a.

Algorithm 1 Algorithm to obtain a and T_a

1: **Input:** n, δ, T.
2: Find the smallest value of $a \in [n/2]$ s.t. $\frac{a^{n-2a}}{(a+1)^{n-2a}} > \delta^T$ and $\frac{1}{a^2} \cdot \frac{(a\delta^T)^{n+2-2a}}{(a\delta^T+1)^{n-2a}} < \delta^T$ holds.
3: For the above value of a, define the function $h(T_a) = \frac{1}{a^2} \cdot \frac{(aT_a)^{n-2a+2}}{(1+aT_a)^{n-2a}}$
4: Compute the solution for the equation $h(T_a) = \delta^T$ via binary search. For binary search, initialize the search space to be $[\delta^T, 1]$. (see Appendix C for more details)
5: **return** (a, T_a)

As Algorithm 1 returns a, we now know the first non-zero time to show an ad. This means that we show $a - 1$ ads each at $t_0 = 0$ and at time $t_n = T$. Note that for the corner case when the value of a does not exist, and for a detailed algorithm, refer to Appendix D. Once a and T_a is known, Algorithm 2 shows how to find the near-optimal schedule.

Algorithm 2 Algorithm to obtain near-optimal schedule

1: **Input:** n, δ, T, a, T_a
2: Set $t_i = 0, \forall i \in \{0, \ldots, a-1\}$ and $t_j = T, \forall j \in \{n - 2a + 1, \ldots, n\}$
3: Set $t_a = \ln(T_a)/\ln(\delta)$ and $t_{n-a} = T - t_a$
4: Distribute the remaining $t_{i+a}, \forall i \in \{0, \ldots, n-2a\}$ such that $t_{i+a} = t_a + i \cdot (t_{n-a} - t_a)/(n-2a)$

Algorithm 2 provides a near-optimal ad schedule when the number of ads n is known in advance. If n is unknown, we use Algorithm 3 to determine it. Below, we show how to compute the optimal number of ads n, given an upper bound $\tilde{n}$ on the total ads that can be shown within the time horizon T.

Algorithm 3 Optimal number of ads

1: **Input:** $\delta, T, \tilde{n}$
2: For each $n \in [\tilde{n}]$, compute the near-optimal strategy $\tilde{t_n}$ using Algorithm 2
3: For each near-optimal strategy, compute $R(\tilde{t_n})$
4: **return** the value of n that maximizes $R(\tilde{t_n})$

Algorithm 3 returns the optimal number of ads to show and the near-optimal strategy to display them. For clarity, we focus our analysis on the case $a = 1$, i.e., the case where only one ad is shown at time 0 and time T, respectively. The general case $a \neq 1$ is discussed in Appendix D and follows a similar approach. When $a = 1$, our goal is to show that approximating T_1 helps us approximate $t_1, \ldots, t_n$, resulting in a nearly optimal schedule. In the next section, we outline our approach for the $a = 1$ case, with full proofs provided in Section 4.

4 Analysis of Our Algorithm for $a = 1$

In this section, we present our analysis of the near-optimal strategy for the case $a = 1$. The analysis is divided into three parts: Section 4.1 proves that L has a unique minimum. Section 4.2 derives closed-form expressions showing that each T_i can be written in terms of T_1. Section 4.3 first approximates T_1 and then t_1 to an additive error of $1/2^n$. Once T_1 and t_1 are approximated, the remaining T_i and t_i values can be computed easily. We now start our analysis by showing that L has a unique minima.

4.1 L has a Unique Minima

In this section, we want to argue about the number of minima for the loss function L. Specifically, we want to show that L is a strictly convex function, and hence, it has at most one minima. This would imply that local minima is also the global minimum (Boyd & Vandenberghe, 2004). This characterization helps us to find the minima for L.

For a strategy $\bar{t}$, to form a feasible solution, it must satisfy $t_i \leq t_{i+1}$ and $0 \leq t_i \leq T$. Then, the set of feasible solutions, which we denote by $\mathcal{D}$, is defined as follows: $t_i \in \mathbb{R}, 0 \leq t_i \leq T, t_i \leq t_{i+1}, t_0 = 0, t_n = T$. We first show why $t_0 = 0$ and $t_n = T$ are required to minimize L.

Lemma 4.1. *In any optimal solution $\bar{t}$ that minimizes the loss function L, we have $t_0 = 0$ and $t_n = T$.*

For $L = \sum_{i>j} \delta^{t_i - t_j}$, let $a_{ij} = t_i - t_j$, where $i > j$ and denote $L' = \sum_{i>j} \delta^{a_{ij}}$. The feasible solution $\mathcal{D}'$ for L' is given by: $a_{ij} \in \mathbb{R} \ \ \forall i > j, a_{ij} \geq 0, a_{n0} = T, a_{ij} + a_{jk} = a_{ik}$.

Using Lemma 4.1, we only consider those $\bar{t}$ where $t_0 = 0$ and $t_n = T$. Since t_0 and t_n are fixed, all references to ∇L hereafter are with respect to $t_1, t_2, \ldots, t_{n-1}$. To show that L has a unique global minima, we divide our analysis into four parts: (1) We establish that there exists exactly one global minima in $\mathcal{D}'$ for the function L', by showing that L' is strictly convex and the space $\mathcal{D}'$ is compact. (2) We show a bijection between $\mathcal{D}$ and $\mathcal{D}'$. (3) We show that, according to the previous bijection, L' indeed models L. (4) We finally show that the minima of L' corresponds exactly to the minima of L, and vice versa. Note that the missing proof of Lemma 4.1, the technical details to prove that L has a unique minimum, and the proof of the following theorem, can be found in Appendix A.

Theorem 4.2. *The function L admits a unique minima in $\mathcal{D}$.*

As we have shown that L has a unique minima. To find the minima, we first argue that all T_i's can be expressed in terms of T_1.

4.2 Dependency of T_i's on T_1

Let $T_i = \delta^{t_i}$ for $i \in \{0, 1, \ldots, n\}$. In this section, we show that $T_2, T_3 \ldots, T_n$ can be expressed in terms of T_1. As shown in Section 4.1, we know that $L = \sum_{j>i} \delta^{t_j - t_i}$ has a unique minima and therefore our eventual goal is to find this minima. Let $H_0 = \left(\frac{1}{\delta^{t_0}}\right), H_1 = \left(\frac{1}{\delta^{t_0}} + \frac{1}{\delta^{t_1}}\right)$, similarly $H_i = \left(\frac{1}{\delta^{t_0}} + \frac{1}{\delta^{t_1}} + \cdots + \frac{1}{\delta^{t_i}}\right) = \sum_{j=0}^{i} \frac{1}{\delta^{t_j}}$. In the following lemma, we show how T_i can be expressed in terms of T_{i-1}, H_{i-1}, H_{i-2}, which will later help us to express T_i in terms of T_1.

Lemma 4.3. *For $1 \leq i \leq n-1$, for any strategy $\bar{t} = (0, t_1, \cdots, t_{n-1}, T)$ corresponding to which $\nabla L = 0$, the following relations hold:*

$$T_i = \frac{-1 + \sqrt{1 + 4H_{i-2}H_{i-1}(T_{i-1}^2)}}{2H_{i-1}} \tag{1}$$

We now use Lemma 4.3 to show that every T_i can be expressed in terms of T_1, which is the main theorem of this section.

Theorem 4.4. *For any strategy $\bar{t} = (0, t_1, \cdots, t_{n-1}, T)$ corresponding to which $\nabla L = 0$, for all i in $\{1, \cdots, n-1\}$, T_i can be written in terms of T_1 as follows:*

$$T_i = \frac{T_1^i}{(1 + T_1)^{i-1}} \tag{2}$$

Note that in Theorem 4.4, we have shown that $T_2, \ldots, T_{n-1}$ can be expressed in terms of T_1. We now show how to express T_n in terms of T_1.

Lemma 4.5. *For $1 \leq i \leq n$, for any strategy $\bar{t} = (0, t_1, \ldots, t_{n-1}, T)$ corresponding to which $\nabla L = 0$, the following relation holds:*

$$\frac{T_1^n}{(1 + T_1)^{n-2}} = T_n \tag{3}$$

To conclude, we derived many interesting relations between $T_1, T_2, \ldots, T_n$ in this section. The missing details and proofs can be found in Appendix B. In the following section, we use the theorem and lemmas of this section to find an approximate value of T_i and t_i for all $i \in \{1, \ldots, n-1\}$.

4.3 Approximating t_i's

In Section 4.1, we showed that a unique minima exists for L, and in Section 4.2, we showed that $(T_2, \ldots, T_n)$ can be expressed in terms of T_1. In this section, we show how to find an approximate solution for $(T_2, \ldots, T_n)$ by finding an approximate solution for T_1. Note that as $T_1 = \delta^{t_1}$, once we obtain an approximate solution for T_1, we also get a solution for t_1. To this end, we describe the outline of this section before proving every detail in Appendix C. We first show that T_1 has a unique solution and then demonstrate how to approximate T_1 using binary search. Once we obtain an approximation for T_1, we use it to approximate T_i, $i \in \{2, \cdots, n\}$ using Theorem 4.4. Next, we approximate t_1 using T_1, which further helps in approximating $t_2, \ldots, t_n$.

For an optimal t_1^*, let $T_1^* = \delta^{t_1^*}$ be the optimal value of T_i that satisfies $\nabla L = 0$. Similarly, define $T_2^*, \ldots, T_n^*$. First, we show the following result for T_i.

Lemma 4.6. *Assuming $T_1^*(1-\epsilon) \leq T_1 \leq T_1^*(1+\epsilon)$, then T_i can be bounded by $T_i^* \cdot (1 - 4^n\epsilon) \leq T_i \leq T_i^* \cdot (1 + 4^n\epsilon)$.*

We now use the above lemma to approximate t_i, which is our main result.

Lemma 4.7. *For $\epsilon < \frac{1}{2}$, $t_i = \log_\delta T_i$ is bounded by $t_i^* - \frac{2\ln(2)\epsilon}{\log(1/\delta)} \leq t_i \leq t_i^* + \frac{2\ln(2)\epsilon}{\log(1/\delta)}$.*

Corollary 4.8. *Setting $\epsilon = \frac{1}{2^n} \cdot \frac{\log(1/\delta)}{2\ln(2)}$ we have, $t_i^* - \frac{1}{2^n} \leq t_i \leq t_i^* + \frac{1}{2^n}$.*

We have thus shown that each t_i can be approximated within an exponentially small error of the optimal t_i^*, completing the analysis of our near-optimal strategy for $a = 1$. We next discuss several important implications and behaviors of the solution.

5 Implications

Having already established the optimality for the objective function L, we now proceed to analyze its structural properties. These properties reveal how the placement of advertisements varies with δ, offering deeper insights into the behavior of our strategy under different values of δ. We summarize here the behavior of our near-optimal solution, which is mathematically proved in Appendix E.

Observation 5.1. *The near-optimal ad schedule exhibits the following patterns as δ varies: (a) As $\delta \to 0$, ads are placed at evenly spaced intervals, (b) as δ increases, more ads concentrate at times 0 and T, (c) as δ increases, the first $t_i > 0$ moves towards 0, and the last $t_j < T$ moves towards T. The remaining ads are evenly spaced between t_i and t_j.*

The above properties indicate a form of clustering behavior in our near-optimal solution. When $\delta \to 0$, it is optimal to display the advertisements at uniformly spaced intervals. As δ increases, a greater number of advertisements are placed at the endpoints, 0 and T, while the remaining ones are evenly distributed in the interior of the interval. In the limit as $\delta \to 1$, the majority of the ads are concentrated at 0 and T, with only a few ads placed uniformly in between.

6 Experiments

In this section, we evaluate the performance of our near-optimal strategy through four distinct experiments: (1) how does the near-optimal strategy vary as the value of δ changes? (2) comparison between our strategy and other baseline strategies? (3) how does the loss function change with a change in the number of ads? (4) how to find the optimal number of ads? We provide a brief overview of each experiment below. For detailed setup and additional results, please refer to Appendix F.

6.1 Variation in near optimal strategy with change in δ

In this experiment, we illustrate how our strategy evolves as the parameter δ increases from 0 to 1. Figure 1a depicts the outcome when the number of ads $n + 1$ is odd. Please refer to the Appendix F for the odd case. Initially, for $\delta \leq 0.4$, the ads are placed nearly equidistantly. As δ increases beyond 0.4, the ads gradually bifurcate—half of them shift towards $t = 0$, and the other half move towards $t = 20$. This experiment aligns precisely with the behavior we obtained in Observation 5.1.

6.2 Our strategy vs baseline strategies

We compare our near-optimal strategy against three baselines - Uniform, Corner, and Random, which are explained in Appendix F. Our experiment models a video streaming setting, where users engage for 1.5–2 hours and are shown approximately 15 ads. Therefore, we have $n + 1 = 15$ and $T = 100$. For this experiment, we focus on a high value of δ (greater than 0.9) as suggested in (Curmei et al., 2022). As observed in our experiment (Figure 1b) and in Observation E.13, when δ is small, the Uniform strategy performs well. Conversely, when δ approaches 1, the Corner strategy becomes more effective. Our results show that the near-optimal strategy adapts to δ and consistently outperforms all baseline strategies. Specifically, at δ around 0.98, our strategy outperforms all other baseline strategies at least by 10%.

6.3 Change in loss with number of ads

We now conduct another experiment to quantify the loss incurred by our near-optimal strategy due to the effect of operant conditioning, as the number of ads increases. Let $L^{\#}(n + 1)$ denote the loss associated with showing n ads under our strategy. A natural intuition is that if showing n ads results in a loss of $L^{\#}(n + 1)$, then doubling the number of ads would roughly double the loss, i.e., $L^{\#}(2(n + 1)) \approx 2 \cdot L^{\#}(n + 1)$. Our experiments actually support this intuition (see Figure 1c).

Interestingly, for smaller values of δ (around 0.7), the loss remains relatively stable even as the number of ads increases. However, we observe a sharp rise in loss as δ increases from 0.9 to 0.99, indicating increased sensitivity to operant conditioning in this regime (see Figure 1c).

6.4 Optimum number of ads

In our final experiment, we demonstrate that for a given user (fixed δ), the optimal number of ads $n+1$ changes under different mere exposure and hedonic adaptation functions. We use a sigmoid reward function $B(i) = k \cdot \frac{1}{1+e^{-ci}}$, where k captures the overall strength of these effects and c controls sensitivity to the number of ads. For $k, c > 0$, the function is concave and increasing. As shown in Figure 1d, the reward initially rises with more ads due to the dominance of mere exposure. Beyond a point, the negative impact of hedonic adaptation and operant conditioning becomes significant, causing the reward to decline. The peak of this curve corresponds to the optimal number of ads. For other related experiments, refer to Appendix F.

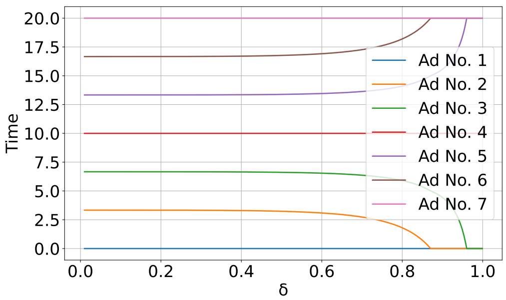

(a) Optimal strategy for 7 Ads, Time = 20

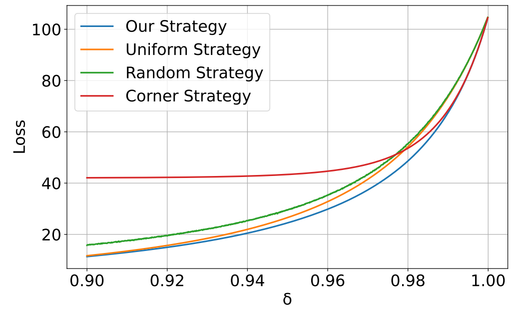

(b) Loss vs δ for $(n+1) = 15$, $T = 100$

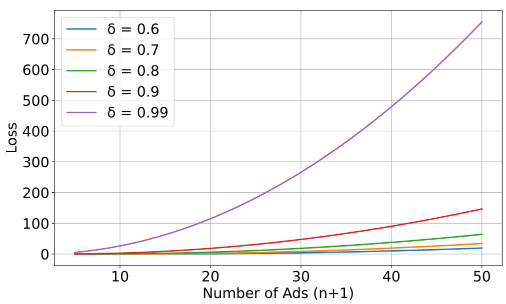

(c) Loss vs $(n+1)$ ads

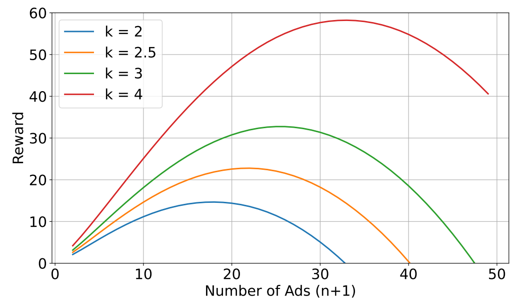

(d) Reward v/s $(n+1)$ for $B(i) = k \cdot \frac{1}{1+e^{-i/5}}$

Figure 1: (a) Change in near-optimal strategy with δ for 7 ads and $T = 20$. (b) Loss between near-optimal and baseline strategies for $n+1 = 15$, $T = 100$. (c) Loss increases as the number of ads increases. (d) Gain function used to find the optimal number of ads for a user with $\delta = 0.9$.

7 Conclusion

In this paper, we consider a model that incorporates dynamic psychological effects – mere exposure, hedonic adaptation, and operant conditioning – into the problem of ad scheduling. We present a near-optimal strategy to schedule ads based on our behavioral model. Our strategy leads to several insights into the problem of ad scheduling, for example, equal spacing of ads might not be optimal under many settings, and it might be better to show more ads in the beginning and at the end as compared to the middle of the time-horizon. We also support these theoretical results using simulations.

7.1 Limitations/Extensions

Seasonality and non-stationary rewards. While our model can work well for real-world settings where the rewards are approximately stationary, such as inserting ads into a (live) video stream, our model does not handle scenarios where the rewards are affected by seasonality or time-of-day effect. For example, it is unlikely that sending a push notification at night will result in user attention. It will be interesting to extend our model to account for non-stationary rewards.

Competition between ads. In our work, we consider the optimization of the ad schedule from the point of view of a single advertiser or an ad agency running multiple homogeneous ads. In the future it will be interesting to account for externalities in the form of competing ads across various channels.

Incorporating context or side-information. Our model does not incorporate the context or side-information of the user or the ad into the optimization problem. This is motivated by the fact that under many scenarios, advertisers do not have access to user information, such as advertising on streaming platforms. Another future direction is to incorporate additional context of the user and ad.

Learning the reward function. Our current setup assumes knowledge of the reward function (including the parameter δ). While our methodology is flexible enough to allow various types of reward functions, it will be interesting to study our problem as a joint learning and optimization problem in a multi-armed bandits setting.

REFERENCES

Ron Adany, Sarit Kraus, and Fernando Ordonez. Allocation algorithms for personal tv advertisements. *Multimedia systems*, 19:79–93, 2013.

Ashwin Aravindakshan and Prasad A Naik. Understanding the memory effects in pulsing advertising. *Operations Research*, 63(1):35–47, 2015.

Soumya Basu, Rajat Sen, Sujay Sanghavi, and Sanjay Shakkottai. Blocking bandits. *Advances in Neural Information Processing Systems*, 32, 2019.

Avrim Blum and Kavya Ravichandran. Nearly-tight approximation guarantees for the improving multi-armed bandits problem. In *Algorithmic Learning Theory*, pp. 228–245. PMLR, 2025.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

Leonardo Cella and Nicolo Cesa-Bianchi. Stochastic bandits with delay-dependent payoffs. In *International Conference on Artificial Intelligence and Statistics*, pp. 1168–1177. PMLR, 2020.

Sunaina K Chugani, Julie R Irwin, and Joseph P Redden. Happily ever after: The effect of identity-consistency on product satiation. *Journal of Consumer Research*, 42(4):564–577, 2015.

John O Cooper, Timothy E Heron, William L Heward, et al. Applied behavior analysis. 2007.

Dena Cox and Anthony D Cox. Beyond first impressions: The effects of repeated exposure on consumer liking of visually complex and simple product designs. *Journal of the Academy of Marketing Science*, 30(2):119–130, 2002.

Mihaela Curmei, Andreas A Haupt, Benjamin Recht, and Dylan Hadfield-Menell. Towards psychologically-grounded dynamic preference models. In *Proceedings of the 16th ACM Conference on Recommender Systems*, pp. 35–48, 2022.

Kateryna Czerniachowska. Scheduling tv advertisements via genetic algorithm. *European Journal of Industrial Engineering*, 13(1):81–116, 2019.

Stylianos Despotakis, Ramamoorthi Ravi, and Amin Sayedi. First-price auctions in online display advertising. *Journal of Marketing Research*, 58(5):888–907, 2021.

Gabriela Dobrița, Simona-Vasilica Oprea, and Adela Bâra. Federated learning's role in next-gen tv ad optimization. *Computers, Materials & Continua*, 82(1), 2025.

Hermann Ebbinghaus. A contribution to experimental psychology. *New York, NY: Teachers College, Columbia University*, 1913.

Xiang Fang, Surendra Singh, and Rohini Ahluwalia. An examination of different explanations for the mere exposure effect. *Journal of consumer research*, 34(1):97–103, 2007.

Jason Freeman, Lewen Wei, Hyun Yang, and Fuyuan Shen. Does in-stream video advertising work? effects of position and congruence on consumer responses. *Journal of Promotion Management*, 28 (5):515–536, 2022.

Daniel G Goldstein, R Preston McAfee, and Siddharth Suri. The effects of exposure time on memory of display advertisements. In *Proceedings of the 12th ACM conference on Electronic commerce*, pp. 49–58, 2011.

Hoda Heidari, Michael J Kearns, and Aaron Roth. Tight policy regret bounds for improving and decaying bandits. In *IJCAI*, pp. 1562–1570, 2016.

Paul Hekkert, Clementine Thurgood, and TW Allan Whitfield. The mere exposure effect for consumer products as a consequence of existing familiarity and controlled exposure. *Acta psychologica*, 144 (2):411–417, 2013.

Mathias Jesse and Dietmar Jannach. Digital nudging with recommender systems: Survey and future directions. *Computers in Human Behavior Reports*, 3:100052, 2021.

Robert Kleinberg and Nicole Immorlica. Recharging bandits. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 309–319. IEEE, 2018.

Liu Leqi, Fatma Kilinc Karzan, Zachary Lipton, and Alan Montgomery. Rebounding bandits for modeling satiation effects. *Advances in Neural Information Processing Systems*, 34:4003–4014, 2021.

Nir Levine, Koby Crammer, and Shie Mannor. Rotting bandits. *Advances in neural information processing systems*, 30, 2017.

Yonatan Mintz, Anil Aswani, Philip Kaminsky, Elena Flowers, and Yoshimi Fukuoka. Nonstationary bandits with habituation and recovery dynamics. *Operations Research*, 68(5):1493–1516, 2020.

Jaap MJ Murre and Joeri Dros. Replication and analysis of ebbinghaus' forgetting curve. *PloS one*, 10(7):e0120644, 2015.

Prasad A Naik, Murali K Mantrala, and Alan G Sawyer. Planning media schedules in the presence of dynamic advertising quality. *Marketing science*, 17(3):214–235, 1998.

Leif D Nelson and Tom Meyvis. Interrupted consumption: Disrupting adaptation to hedonic experiences. *Journal of Marketing Research*, 45(6):654–664, 2008.

Marc Nerlove and Kenneth J Arrow. Optimal advertising policy under dynamic conditions. *Economica*, pp. 129–142, 1962.

Sirius XM Pandora. Measuring consumer sensitivity to audio advertising: A long-run field experiment on pandora internet radio.

Vishakha Patil, Vineet Nair, Ganesh Ghalme, and Arindam Khan. Mitigating disparity while maximizing reward: tight anytime guarantee for improving bandits. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 4100–4108, 2023.

Omid Rafieian. Optimizing user engagement through adaptive ad sequencing. *Marketing Science*, 42 (5):910–933, 2023.

RedCircle. Less is more: How many ads should there be in an ad break?, 2023. URL `https://redcircle.com/blog/less-is-more-how-many-ads-should-there-be-in-an-ad-break/`. Accessed: 2025-09-02.

Eric A Ritter and Chang-Hoan Cho. Effects of ad placement and type on consumer responses to podcast ads. *CyberPsychology & Behavior*, 12(5):533–537, 2009.

Navdeep S Sahni. Effect of temporal spacing between advertising exposures: Evidence from online field experiments. *Quantitative Marketing and Economics*, 13:203–247, 2015.

Eric M Schwartz, Eric T Bradlow, and Peter S Fader. Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36(4):500–522, 2017.

Sridhar Seshadri, Sriram Subramanian, and Sebastian Souyris. Scheduling spots on television. 2015.

Julien Seznec, Andrea Locatelli, Alexandra Carpentier, Alessandro Lazaric, and Michal Valko. Rotting bandits are no harder than stochastic ones. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2564–2572. PMLR, 2019.

Surendra N Singh, Sanjay Mishra, Neeli Bendapudi, and Denise Linville. Enhancing memory of television commercials through message spacing. *Journal of Marketing Research*, 31(3):384–392, 1994.

Hanna Sumita, Yasushi Kawase, Sumio Fujita, Takuro Fukunaga, and RA Center. Online optimization of video-ad allocation. In *IJCAI*, pp. 423–429, 2017.

Georgios Theocharous, Philip S. Thomas, and Mohammad Ghavamzadeh. Personalized ad recommendation systems for life-time value optimization with guarantees. In Qiang Yang and Michael J. Wooldridge (eds.), *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pp. 1806–1812. AAAI Press, 2015.

Romain Warlop, Alessandro Lazaric, and Jérémie Mary. Fighting boredom in recommender systems with linear reinforcement learning. *Advances in Neural Information Processing Systems*, 31, 2018.

Yang Yang and Jeff Galak. Sentimental value and its influence on hedonic adaptation. *Journal of personality and social psychology*, 109(5):767, 2015.