

A Batched Multi-Armed Bandit Approach to News Headline Testing

Yizhi Mao	Miao Chen	Abhinav Wagle	Junwei Pan	Michael Natkovich	Don Matheson
<i>Oath Inc</i>	<i>Oath Inc</i>	<i>Oath Inc</i>	<i>Oath Inc</i>	<i>Oath Inc</i>	<i>Oath Inc</i>
Sunnyvale, USA	Sunnyvale, USA	Sunnyvale, USA	Sunnyvale, USA	Sunnyvale, USA	Sunnyvale, USA
lolam@oath.com	miaoc@oath.com	awagle@oath.com	jwpan@oath.com	mln@oath.com	donm@oath.com

Abstract—Optimizing news headlines is important for publishers and media sites. A compelling headline will increase readership, user engagement and social shares. At Yahoo Front Page, headline testing is carried out using a test-rollout strategy: we first allocate equal proportion of the traffic to each headline variation for a defined testing period, and then shift all future traffic to the best-performing variation. In this paper, we introduce a multi-armed bandit (MAB) approach with batched Thompson Sampling (bTS) to dynamically test headlines for news articles. This method is able to gradually allocate traffic towards optimal headlines while testing. We evaluate the bTS method based on empirical impressions/clicks data and simulated user responses. The result shows that the bTS method is robust, converges accurately and quickly to the optimal headline, and outperforms the test-rollout strategy by 3.69% in terms of clicks.

Index Terms—Headline testing, Multi-armed bandit, Thompson Sampling

I. INTRODUCTION

Headline testing is important for publishers and media sites. Visitors to the homepage of online publishers are usually presented with lists or groups of headlines, sometimes along with snippets¹. A compelling headline would encourage visitors to click it and read the whole article, and thus help increase user engagement, social sharing and revenue. At Yahoo Front Page, we follow a *test-rollout* strategy for headline testing. Once an article is published, multiple title variations are displayed to randomized user buckets with equal size for a defined period to conduct bucket testing. When the test is complete, the headline variant that has most clicks during the bucket testing period is selected and displayed for the rest of the article life. This strategy has been adopted as a common practice to select the best title variants and ad versions for headline testing [1] and online advertising [2], respectively.

There are two limitations of this test-rollout strategy. First, during the initial testing period, we have to show the title variants with lower click-through rates (CTRs) to a sizable fraction of user population. For those users, we fail to optimize their engagement. Further, a large proportion of article traffic clusters in its early life, because freshness is a key factor of article popularity [3]. The test-rollout strategy conducts bucket

testing at the beginning of the article life, which may lead to significant click loss.

On the other hand, the performance of headline variations usually varies over time. The conclusion drawn from the initial testing period may not always hold throughout the whole article life. The test-rollout practice is not able to capture any changes subsequent to the bucket testing period.

To address the limitations mentioned above, we formulate headline testing as a multi-armed bandit (MAB) problem, and introduce a batched Thompson Sampling method to optimize user engagement while learning the performance of each headline variant. The MAB problem is defined as follows. There are K arms, each associated with an unknown reward distribution. The player iteratively plays one arm, observe the associated reward, and decides which arm to play in the next iteration [4]. There is a tension between selecting the current best-performing arm to harvest immediate gain (exploitation) and discovering the optimal arm, *i.e.*, the arm with the highest expected reward, but risking immediate loss (exploration). In the headline testing scenario, the headline variants of an article correspond to the arms, and the click count of each headline variant is the reward associated with each arm. Our goal is to maximize the sum of clicks across arms for each article.

There are many MAB algorithms such as ϵ -greedy [5], Upper Confidence Bound (UCB) algorithms [6] [7] [8], Thompson Sampling [9], and Gittins Index [10]. Among them we select Thompson Sampling due to its strong empirical results [11] [12], solid theoretical guarantees [13] [14] [4] [15], and wide industrial application [16] [17] [18] [19] [20].

In traditional Thompson Sampling, model parameters are updated for every single user response. It becomes a formidable computational burden, as our site has a high volume of incoming traffic at high velocity. This leads us to consider an algorithm that processes user responses after they arrive in batches over a certain time period. Thompson Sampling with batch updates had been studied in display advertising and news article recommendation to analyze how it performs in the case of delayed user response processing compared with other MAB algorithms [11]. Batch update is incorporated in recent industrial applications of Thompson Sampling [20] [21]. Although how they select the update frequency is not disclosed, [20] adopts a method to update traffic allocation once a day, and [21] updates twice a day.

¹For example, Yahoo Front Page (<https://www.yahoo.com/>), Google News (<https://news.google.com/>), The New York Times (<https://www.nytimes.com/>), and The Wall Street Journal (<https://www.wsj.com/>).

In this paper, we present a *batched Thompson Sampling (bTS)* method that is tuned for optimal performance when user feedback (*i.e.*, observed rewards) is processed in batches. The performance evaluation is based on empirical impressions/clicks of articles at Yahoo Front Page and user responses simulated from empirical CTRs, and shows that the bTS method is robust, converges quickly to the true optimal arms, and outperforms the test-rollout strategy. Our study is motivated by the headline testing problem, but one can apply this algorithm to other real-world problems where we need to accomplish optimization while testing with high volume and high velocity incoming data.

The rest of the paper is organized as follows. Section II introduces the current headline testing practice at Yahoo Front Page, as well as its limitations. Section III describes the batched Thompson Sampling (bTS) method that we propose to apply in headline testing to gain more reward. We show the evaluation of the bTS method in Section IV. Section V concludes and discusses the future work.

II. CURRENT HEADLINE TESTING PRACTICE

Currently at Yahoo Front Page, headline testing follows the test-rollout strategy which consists of two periods:

- *Testing period* is the first hour after an article is published. During the testing period, each viewing request is randomly assigned to one of the headline variants with equal probability. At the end of the testing period, we deem the headline variant with the highest CTR as the *winner headline*.
- *Post-testing period* refers to the remaining article lifespan after the testing period, during which we display the *winner headline* to all traffic.

The test-rollout strategy for headline testing is intuitive and easy to implement in the system. However, it has the following two limitations:

A. User Engagement Loss in the Testing Period

In the first hour testing period, impressions are equally allocated to each headline variant, which means that for an article with K arms, only $\frac{1}{K}$ of the traffic is assigned to the headline variant with the highest underlying CTR. Showing inferior headline variants to $\frac{K-1}{K}$ of the traffic will sacrifice user engagement (*e.g.*, article clicks) and user experience. The loss of user engagement and experience are both sizable because the traffic in the testing period covers as high as 24.36% of the total impressions across all articles. This result is based on empirical headline testing data from Yahoo Front Page, which will be described in Section III-B1.

B. Arm Performance Discrepancy Between Testing and Post-testing Period

The test-rollout practice is unable to capture any performance change of headline variations beyond the testing period. Empirical data from Yahoo Front Page, same as the data mentioned above, shows the headlines deemed best during the testing period changed their performance afterward: on

average there is a 12% discrepancy in their CTR between testing and post-testing periods. Thus, it is possible for the arm deemed optimal during the one-hour testing period to be actually sub-optimal over the article lifespan. Such a performance change cannot be observed and handled in the current practice.

To enhance this test-rollout practice, we formulate headline testing as an MAB problem and present a batched Thompson Sampling approach. It is able to explore for the optimal variant, and at the same time gradually shift traffic to the best-performing variant. The following section introduces the MAB headline testing methodology in detail.

III. METHODOLOGY

This section describes the *batched Thompson Sampling (bTS)* method that is able to gradually allocate traffic towards the well-performing arms, while leaving some traffic to other arms so as to explore for the possibly unobserved optimal arm. Section III-A formulates headline testing as a Bernoulli bandit problem, and introduces Thompson Sampling for Bernoulli bandits. In Section III-B, we explain the rationale of incorporating batch updates in Thompson Sampling, and introduce the factors associated with bTS. Section III-B also describes our empirical impressions/clicks data and how we simulate user responses based on them. Sections III-C to III-E present how we determine and tune the factors of bTS.

A. Preliminaries

1) *Headline Testing as a Bernoulli Bandit Problem*: Suppose an article has K headline variants written by editors. In the MAB framework, each headline variant is treated as an arm. Each headline, when displayed, yields either a click (success) or no click (failure) as the reward. The reward for headline $k \in \{1, \dots, K\}$ is Bernoulli distributed, with the *success probability* (*i.e.*, the probability of being clicked) as $\theta_k \in [0, 1]$. The reward distribution of arm k is fixed but unknown, in the sense that its parameter θ_k is unknown. At time step $t \in [1, T]$, we select an arm to display, and collect the reward observed from the selected arm. Here T is the total number of impressions we decide to run MAB experiment on. Our goal is to maximize the total clicks for this article.

2) *Thompson Sampling*: Thompson Sampling is a randomized Bayesian algorithm to solve the MAB problem of reward maximization [9]. The general idea of Thompson Sampling is to impose a prior distribution on the parameters of the reward distribution, update the posterior distribution using the observed reward, and play an arm according to its posterior probability at each time step.

In Bernoulli bandit problems, Thompson Sampling uses Beta distribution to model the success probability θ_k for arm k because the observed reward follows a Bernoulli distribution, and the Beta distribution is a conjugate prior for the Bernoulli distribution [4]. Initially, the Thompson Sampling algorithm imposes a Beta(1,1) prior on the success probabilities of all arms. It is a reasonable initial prior, because Beta(1,1) is the uniform distribution on the interval (0,1) [4] [11]. At time step

$t \in [1, T]$, Thompson Sampling algorithm draws a random sample from the Beta distribution of each arm, and displays the arm associated with the largest sampled value. Based on the observed feedback of the displayed arm, the Beta(α^t, β^t) distribution of the displayed arm is updated to Beta($\alpha^t + 1, \beta^t$) if the feedback is a click, or to Beta($\alpha^t, \beta^t + 1$) otherwise.

Many studies have demonstrated the strong performance of Thompson Sampling algorithm in the MAB problem, both theoretically and empirically. [13] investigates Thompson Sampling as Bayesian Learning Automaton, and shows that in the two-armed Bernoulli bandit problem, Thompson Sampling converges to only playing the optimal arm with probability one. [14] proposes an optimistic version of Thompson Sampling, and proves both Thompson sampling and the optimistic version result in optimal behaviour in the long term consistency sense described by [22]. Further, [4] and [15] provide regret bounds for Thompson Sampling that are asymptotically optimal in the sense defined by [23], so that it has the theoretical guarantee competitive to UCB algorithms. Empirically, [11] shows the performance of Thompson Sampling is competitive to or better than that of other alternative MAB algorithms, such as ϵ -greedy and UCB, on real-world problems like display advertising and news article recommendation. [11] also mentions Thompson Sampling can be implemented efficiently, in comparison with full Bayesian methods such as Gittins index. Recently, adaptations of Thompson Sampling have been applied in many domains, such as revenue management [24], recommendation system [25], online service experiments [19], website optimization [20], and online advertising [16] [17] [18].

B. Batch Updates for Real-world High-volume Traffic

In traditional Thompson Sampling for Bernoulli bandits, the Beta distribution of the selected arm is updated after every reward feedback is observed. In a real-world system, especially when both the volume and velocity of incoming traffic are high, the feedback is typically processed in batches over a certain period of time. This is the case for our current infrastructure of headline testing. Thus, to implement Thompson Sampling, it is necessary to apply the batched Thompson Sampling (bTS) that updates the posterior distribution after a fixed time period.

The general procedure of bTS is described as follows: within each fixed time interval (*i.e.*, batch), the Beta distribution of each arm remains unchanged. We allocate traffic across arms based on their random Beta distribution samples drawn for each incoming view event. At the end of a time interval, we aggregate the data collected within this batch, namely the numbers of clicks and impressions for each arm, and use the aggregated data to update the Beta distribution of each arm.

There are three factors to be tuned for bTS. The first one is how long the algorithm should run. We determine an algorithm stopping point after which the click gain is so little that the system cost is not worthwhile. The second is how we aggregate the feedback within each batch to update the posterior distributions. The last factor is the time interval

TABLE I
IMPRESSIONS AND BATCH SIZES FOR A SAMPLE ARTICLE

Timestamp	Impressions ²	Batch Index	Batch size
12:48:00	615	1	9,945
12:49:00	4,568	1	
12:50:00	4,762	1	
12:51:00	5,282	2	16,028
12:52:00	5,412	2	
12:53:00	5,334	2	

between each update. These three factors are determined based on the empirical as well as simulated data from our real-world headline testing platform, which will be described as follows.

1) *Empirical data*: The empirical data cover the articles with headline testing at Yahoo Front Page on two weekdays and one weekend day. In the testing period, the data consist of impressions and clicks of all headline variants for each article. While in the post-testing period, we only have data on the headline variant deemed best in the previous testing period for each article, because all traffic is shifted to this headline variant. The impressions and clicks are aggregated by every minute.

2) *Simulation practice*: Similar as how [11] evaluates Thompson Sampling in display advertising, we evaluate the performance of bTS under different factor values in a simulated environment, where impressions and CTRs are real, but the user responses are simulated based on the empirical CTR of each headline.

In the simulation of Thompson Sampling [11], the reward probability of each arm is modeled by a Beta distribution which is updated after an arm is selected. For batched Thompson Sampling, the Beta distribution of each arm is only updated at the end of each batch, *i.e.*, a fixed-length time interval. Accordingly, the number of events within a batch, which is referred to as *batch size*, is determined by the count of empirical impressions occurred in this fixed-length time interval. The batches of an article rarely have the same size, because the number of impressions usually varies a lot in different time intervals. Table I illustrates an example of how batch sizes are calculated based on the minute-level impressions when updates occur every 3 minutes.

We simulate user clicks following the practice of [2]: Upon a viewing request, suppose the algorithm selects to display arm k , then the user response is simulated from the Bernoulli($\hat{\theta}_k$) distribution, where we denote $\hat{\theta}_k$ as the estimated success probability of arm k , defined by its empirical CTR during the *testing period*. Note that for the arm deemed best during the testing period, we still use its testing period empirical CTR to calculate $\hat{\theta}_k$, although theoretically we can calculate its “overall” empirical CTR including both testing and post-testing period. This is because this “overall” empirical CTR is not comparable to the empirical CTRs of other arms,

²The numbers are for illustration purpose only. They are not actual impressions.

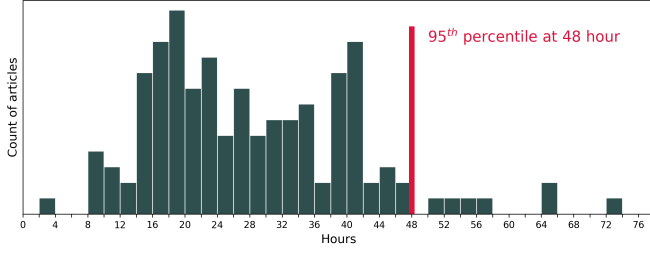


Fig. 1. Histogram of article active lifespan

which can only be calculated during the testing period. After the simulation is completed on all articles, we quantify the performance of a headline testing algorithm by total clicks summed across all articles over their lifespans.

C. Factor 1: Algorithm Stopping Point

Due to the observed CTR discrepancy between the testing and post-testing period illustrated in Section II-B, we would avoid stopping the algorithm too early, so as to cover any potential performance change among headline variations. One straightforward proposal is to run the algorithm throughout the whole article life. Although technically achievable, this proposal is not desired, given that the distribution of impressions over time usually has a very long right tail. When the impressions are very sparse, the click gain is so little that it is not worth the engineering overhead of running the algorithm. Thus, we would like to stop the algorithm when the majority of articles are no longer active.

We define the *active lifespan* of an article as the time it takes to reach 95% of its total impressions. Fig. 1 demonstrates that the active lifespans of 95% of the articles are under 48 hours. Thus, we take 48 hours as the algorithm stopping point.

D. Factor 2: Update Methods

In traditional Thompson Sampling with Bernoulli bandits, the Beta(α, β) distribution of the selected arm is updated to Beta($\alpha + 1, \beta$) if we observe a click, otherwise it is updated to Beta($\alpha, \beta + 1$). When observed responses come in batches, we need to aggregate the impressions and clicks data for each arm within each batch, before updating the corresponding Beta distribution. We consider two update methods to achieve this: *summation update* and *normalization update*.

To specifically describe the two update methods, we denote the Beta distribution of arm k in the t -th batch by Beta(α_k^t, β_k^t). S_k^t and F_k^t are the click and non-click counters for arm k in batch t . M^t denotes the number of impressions in the t -th batch.

Algorithm 1 explains the bTS method with *summation update*. It is a direct extension from the event-level update method of traditional Thompson Sampling, where α_k^t and β_k^t are updated by raw counts of clicks and non-clicks.

We also consider another update method named as *normalization update*. As illustrated in Algorithm 2, it increments

```

Initialize  $\alpha_k^1 = 1$  and  $\beta_k^1 = 1 \forall k$ ;
for batch index  $t \in \{1, \dots, T\}$  do
  Initialize  $S_k^t = 0$  and  $F_k^t = 0 \forall k$ ;
  for event  $i \in \{1, \dots, M^t\}$  do
    for arm  $k \in \{1, \dots, K\}$  do
      Draw random sample  $x_k$  from Beta( $\alpha_k^t, \beta_k^t$ )
    end
    Display arm  $k^* = \operatorname{argmax}_k x_k$  and observe  $r$ ;
    if  $r = 1$  then
       $S_{k^*}^t = S_{k^*}^t + 1$ 
    else
       $F_{k^*}^t = F_{k^*}^t + 1$ 
    end
  end
  for arm  $k \in \{1, \dots, K\}$  do
     $\alpha_k^{t+1} = \alpha_k^t + S_k^t$ ;
     $\beta_k^{t+1} = \beta_k^t + F_k^t$ ;
  end
end

```

Algorithm 1: Batched Thompson Sampling with summation update

```

Initialize  $\alpha_k^1 = 1$  and  $\beta_k^1 = 1 \forall k$ ;
for batch index  $t \in \{1, \dots, T\}$  do
  Initialize  $S_k^t = 0$  and  $F_k^t = 0 \forall k$ ;
  for event  $i \in \{1, \dots, M^t\}$  do
    for arm  $k \in \{1, \dots, K\}$  do
      Draw random sample  $x_k$  from Beta( $\alpha_k^t, \beta_k^t$ )
    end
    Display arm  $k^* = \operatorname{argmax}_k x_k$  and observe  $r$ ;
    if  $r = 1$  then
       $S_{k^*}^t = S_{k^*}^t + 1$ 
    else
       $F_{k^*}^t = F_{k^*}^t + 1$ 
    end
  end
  for arm  $k \in \{1, \dots, K\}$  do
     $\alpha_k^{t+1} = \alpha_k^t + \frac{M^t S_k^t}{S_k^t + F_k^t}$ ;
     $\beta_k^{t+1} = \beta_k^t + \frac{M^t}{K} \left(1 - \frac{S_k^t}{S_k^t + F_k^t}\right)$ ;
  end
end

```

Algorithm 2: Batched Thompson Sampling with normalization update

α and β by the number of normalized clicks and non-clicks respectively, assuming equal traffic allocation across arms.

Normalization update method addresses a side-effect of imbalanced traffic allocation across arms within each batch, which may fail to update in favor of the true optimal arm. More specifically, if arm k has few clicks in a batch, it may be a reflection of a low θ_k , or because the traffic allocated to this arm is not large enough to generate many clicks. Normalization update helps to mitigate the second possibility by assuming each arm has equal traffic allocation. The downside

of this method is that it lowers the noise tolerance of the algorithm, because it would give $\frac{1}{K}$ of the weight to potential noise in the data. Given that K is usually less than four, this method may magnify the potential noise.

We compare the performance of the two methods using the simulation practice described in Section III-B2. The algorithm runs for 48 hours as determined in Section III-C, and the performance of either update method is quantified by the total click counts for all articles over their lifespans.

According to the result shown in Table II, summation update consistently has better performance than normalization update across different update frequencies. One explanation is that, in the case of headline testing, data noise has more impact on the algorithm performance than the side-effect of imbalanced traffic allocation across arms. Thus, we use the summation update method described in Algorithm 1 for the bTS method.

E. Factor 3: Update Frequency

In Table III, we demonstrate the percentage gap of total clicks between the best-performing update frequency and the remaining frequencies. The table shows that for the bTS method with summation update, more frequent updates lead to more gain in clicks. The gain becomes marginal when the update frequency is lower than 5 minutes. Our infrastructure is capable of updating as frequent as every 5 minutes with almost no cost. However, going beyond 5-minute frequency creates a formidable challenge towards the infrastructure, due to the network transformation cost among multiple components of the system. In consideration of the marginal benefit associated with more granular updates, choosing 5 minutes as the update frequency is a reasonable decision for our use-case.

To conclude the factors we have selected for the bTS method in our real-world headline testing scenario to achieve the trade-off between gain in clicks and system cost: we run the algorithm for 48 hours, aggregate data in each batch using summation update method described in Algorithm 1, and update the Beta distribution of each arm every 5 minutes.

IV. METHODOLOGY EVALUATION

This section evaluates the bTS method in the setting of headline testing. In Sections IV-A, IV-B, and IV-C, we evaluate its false convergence rate, speed of optimization, and robustness, respectively. Then, we show bTS outperforms the test-rollout practice, in terms of click gain in Section IV-D, and less exposure of sub-optimal headlines in Section IV-E.

A. False Convergence Rate

We analyze the false convergence rate of the algorithm, defined as the proportion of articles that fail to allocate most traffic to the *optimal arm* – the arm with highest θ – when their traffic allocation across arms is stable. Our evaluation shows 99.25% of the articles converge correctly for the bTS method.

Fig. 2 demonstrates the traffic proportion over time of some sample articles that converge correctly: initially all arms have equal traffic allocation. When the algorithm begins running,

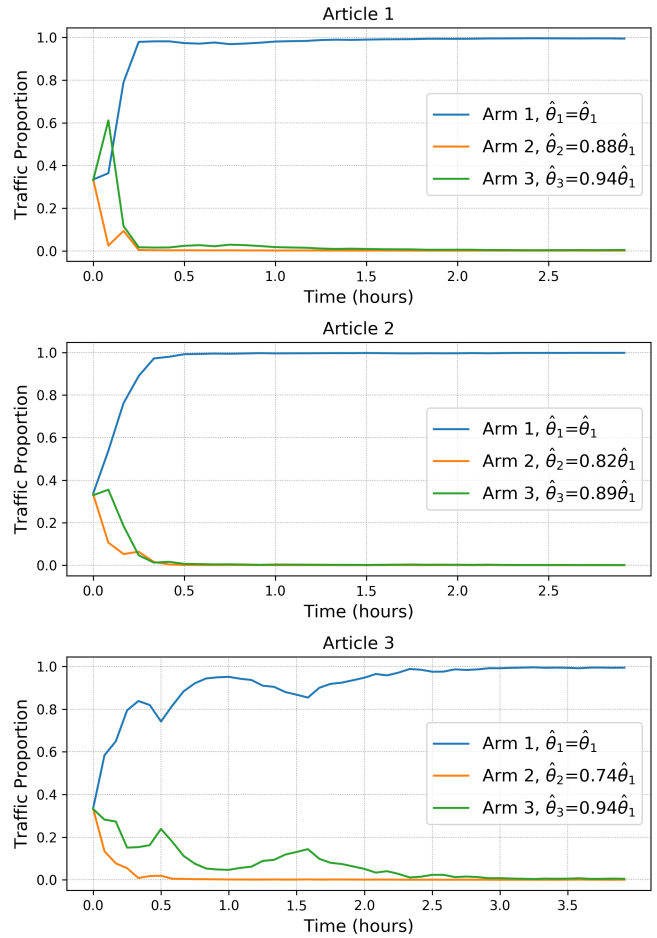


Fig. 2. Sample articles with correct convergence: traffic proportion allocated to each arm over time. For each article, the $\hat{\theta}$'s of arms are illustrated as the ratio to the highest $\hat{\theta}$.

it starts to explore for the optimal arm, while simultaneously allocating more traffic to the arm that performs well.

Fig. 2 also illustrates how Thompson Sampling gracefully handles exploration at the level of individual arms, as pointed out in [19]. bTS explores the clearly inferior arms (*e.g.*, Arm 2 in all sample articles) less frequently than arms that may be optimal, *i.e.*, good arms (*e.g.*, Arm 1 and Arm 3 in the sample articles). This increases the click gain by shifting traffic from a clearly inferior arm to arms with better performance. It also helps distinguish the optimal arm faster, because there are more samples to compare among the good arms.

B. Speed of optimization

We show the bTS method optimizes quickly, by analyzing the time it takes for the optimal arm to constantly have the largest traffic allocation. The histogram presented in Fig. 3 shows the distribution of the time to optimize among articles that converge correctly. The 80th percentile of the time to optimize is 30 minutes, which means after 30 minutes, 80% of the articles constantly have the largest traffic proportion allocated to their optimal arms.

TABLE II
COMPARISON BETWEEN SUMMATION UPDATE AND NORMALIZATION UPDATE

Total Clicks	Update frequency in minutes					
	1-min	3-min	5-min	10-min	30-min	60-min
Summation Update Method versus Normalization Update Method	+2.50%	+3.57%	+1.43%	+1.08%	+1.28%	+0.41%

TABLE III
COMPARISON BETWEEN DIFFERENT UPDATE FREQUENCIES FOR SUMMATION UPDATE METHOD

Total Clicks	Update frequency in minutes					
	1-min	3-min	5-min	10-min	30-min	60-min
% Gap from the Best-performing Update Frequency	0	-0.04%	-0.06%	-0.18%	-0.76%	-1.52%

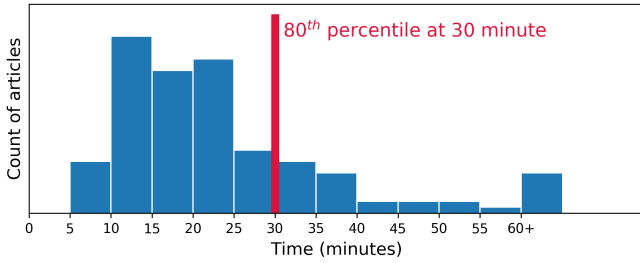


Fig. 3. Histogram of time to optimize. All articles that takes over 60 minutes to optimize are in the last bin.

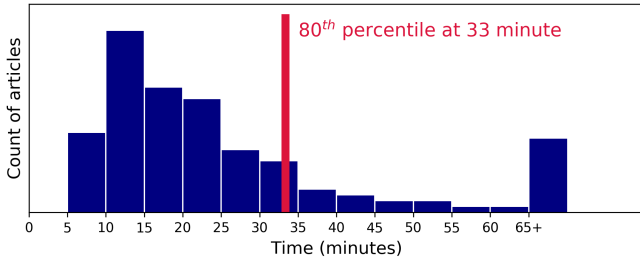


Fig. 4. Histogram of time needed for self-correction. All articles that takes over 65 minutes to self-correct are in the last bin.

C. Self-correction under Stress Test

We conduct a stress test via simulation to evaluate the robustness of our algorithm against unfavorable traffic allocation. In the stress test, we change the initialized beta distribution for each arm, so that around 90% of the traffic is allocated to the arm with the lowest θ , and the remaining traffic is equally allocated to the other arms. We analyze the time it takes for the optimal arm to have the highest likelihood to be displayed for five consecutive batches³, which we refer to as *self-correction*, and illustrate the distribution of the time needed for self-correction in Fig. 4. The vertical red line shows

³This is equivalent to the optimal arm having the highest mean of beta distribution among other arms for five consecutive batches.

that the 80th percentile is at 33 minutes – 80% of the articles are able to self-correct within 33 minutes.

D. Gain in Clicks

In this section, we illustrate that applying bTS to headline testing generates more clicks than the test-rollout strategy, especially during the testing period of an article when bTS dynamically allocate more traffic to the well-performing arms, while impressions are equally allocated to each arm in the test-rollout strategy. The *test-rollout baseline* is set up as follows: for a given article, denote the total baseline click number as $C^{baseline}$, the click number for arm k during the testing period as $C_k^{testing}$, and its click number during the post-testing period as C_k^{post} , then

$$C^{baseline} = \sum_{k=1}^K C_k^{testing} + C_{k^*}^{post} \quad (1)$$

$$\text{with } C_k^{testing} \sim \text{Binomial}\left(\frac{M^{test}}{K}, CTR_k^{testing}\right) \quad (2)$$

$$C_{k^*}^{post} \sim \text{Binomial}(M^{post}, CTR_{k^*}^{testing}) \quad (3)$$

$$k^* = \underset{k \in \{1, \dots, K\}}{\text{argmax}} C_k^{testing} \quad (4)$$

where M^{test} and M^{post} are the article impressions in the testing and post-testing period, respectively. $CTR_k^{testing}$ is the CTR for arm k in the testing period. Expressions (2) and (3) mean $C_k^{testing}$ and $C_{k^*}^{post}$ are simulated from the corresponding Binomial distributions.

We set up this test-rollout baseline instead of directly using the empirical click data, in order to achieve a fair comparison with the bTS method. Suppose an article has two arms, with their *testing period* empirical CTRs denoted as $CTR_1^{testing}$ and $CTR_2^{testing}$, and $CTR_2^{testing} > CTR_1^{testing}$. Thus, arm 2 is displayed to all traffic in the post-testing period, and we can calculate its “overall” empirical CTR, denoted as $CTR_2^{overall}$, by its empirical clicks divided by its empirical impressions during the overall article life. Note that $CTR_2^{overall}$ is not necessarily equal to $CTR_2^{testing}$ as illustrated in Section II-B.

TABLE IV
GAIN IN CLICKS: BTS VERSUS TEST-ROLLOUT BASELINE

Total clicks	First hour	Remaining hours	Total
% Increase in Clicks: bTS versus Test-rollout Baseline	13.54%	1.00%	3.69%

The simulation for bTS uses $CTR_1^{testing}$ and $CTR_2^{testing}$ as the estimated success probability of corresponding arms as explained in Section III-B2, and it is not comparable with the empirical click data, where we can consider $CTR_1^{testing}$ and $CTR_2^{overall}$ as the underlying success probabilities. The possible gap between $CTR_2^{testing}$ and $CTR_2^{overall}$ can also lead to different click counts between bTS and empirical click data, even if bTS performs the same as current practice. On the other hand, the test-rollout baseline we set up in (1) - (4) uses $CTR_1^{testing}$ and $CTR_2^{testing}$ as the underlying success probabilities, and thus is comparable with bTS.

The gain in clicks for bTS against the test-rollout baseline is summarized in Table IV. The gain in clicks is split into the “first-hour” period and the “remaining-hour” period, which correspond to the testing and post-testing periods in the test-rollout strategy. On average, bTS gives a 13.54% increase in total clicks during the first hour. It also has competitive accuracy in selecting the correct optimal arm to exploit, indicated by 1% increase in clicks during the post-testing period than test-rollout baseline.

It is worth noting that this baseline is conservative - the actual gain after implementation is likely to be larger. The baseline setup assumes the CTRs of arms are consistent between testing and post-testing period. This assumption leads to dilution in click gain during the post-testing period: the baseline almost always displays the true optimal arm in the post-testing period due to the long period of exploration⁴, and bTS cannot beat the baseline when the optimal arm is displayed in the baseline during the post-testing period. Thus, the 13.54% click gain harvested in testing period is diluted to 3.69% overall click gain. In practice, arm performance may change over time as mentioned in Section II-B. The current test-rollout practice is not able to capture the fluctuation, while the bTS algorithm would detect the change and adjust the traffic accordingly, which would potentially gain more clicks.

E. Fewer Impressions on Sub-optimal Headlines

Another benefit of bTS is that it improves user experience by exposing much fewer sub-optimal headlines. We analyze the overall percentage decrease in impressions on sub-optimal headlines of the bTS method versus the test-rollout practice. For the test-rollout strategy which we use as the baseline, we assume that no traffic is allocated to sub-optimal headlines during the post-testing period. Even in this best-case scenario for the baseline, the sub-optimal headline impressions of the

⁴There is a 97% accuracy rate for the baseline to display the optimal arm during the post-testing period.

bTS method decrease by 71.53% compared with the test-rollout practice.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we propose to apply an MAB approach to headline testing where data are processed in batches due to their high volume and high velocity. Although the bTS algorithm is developed under the news article headline testing scenario, the parameters used to tune our model can be generalized to other real-world MAB problems, including marketing campaigns, one-time event optimizations, and user purchase/registration funnels.

The stationary assumption beyond the first hour should be assessed after we have data on all arms beyond the first hour, which will be available after the bTS method is implemented. The current assumption is that θ_k , the success probability for arm k , is constant over time. It is possible for θ_k to change over time, and this formulates a non-stationary bandit problem. Note that for non-stationary environments, the proposed bTS method is already an improvement of the test-rollout strategy, because it continuously tests headline variants throughout the active lifespan of news articles. There is room for further enhancement of the methodology to achieve better performance, especially when θ_k changes so significantly and rapidly that the order among θ_k for $k \in \{1, \dots, K\}$ flips multiple times during the article lifespan.

Clickbait detection and prevention are known challenges to media sites [26]. We are also motivated to explore more sophisticated metrics than raw clicks to be used as the reward of MAB. As an example, we may discount the clicks with a short dwell time, and the reward of each headline variation becomes a number between 0 and 1. We can then optimize this “click-dwell” reward using an adaptation of Bernoulli Thompson Sampling algorithm introduced in [4], which supports the reward to be distributed in the interval [0,1]. When the discounting strategy of dwell time is properly selected, this should be helpful to mitigate the clickbait problem.

At Yahoo Front Page, the top module hosts multiple articles. When headline testing is conducted simultaneously on multiple articles within this module, they may interact with each other. That is, the success probability θ of a headline variant may be influenced by the titles of the remaining articles in the same module. This motivates us to extend our research and develop an MAB testing framework that optimizes the article titles for the entire module, e.g., with a multivariate MAB algorithm introduced in [20].

ACKNOWLEDGMENT

The authors would like to express their gratitude to Russell Chen, Shriram Kumar, Jigar Patel, and Rupert Wu for valuable discussions; to Virendra Pratap Singh for his help with empirical data collection; and to Sameer Raheja and Kelly Hirano for their support of this study.

REFERENCES

- [1] "Headline testing page at optimizely." [Online]. Available: <https://www.optimizely.com/optimization-glossary/headline-testing/>
- [2] E. M. Schwartz, E. T. Bradlow, and P. S. Fader, "Customer acquisition via display advertising using multi-armed bandit experiments," *Marketing Science*, vol. 36, no. 4, pp. 500–522, 2017.
- [3] Y. Keneshloo, S. Wang, E.-H. S. Han, and N. Ramakrishnan, "Predicting the shape and peak time of news article views," in *Big Data (Big Data), 2016 IEEE International Conference on*. IEEE, 2016, pp. 2400–2409.
- [4] S. Agrawal and N. Goyal, "Analysis of thompson sampling for the multi-armed bandit problem," in *Conference on Learning Theory*, 2012, pp. 39–1.
- [5] V. Kuleshov and D. Precup, "Algorithms for multi-armed bandit problems," *arXiv preprint arXiv:1402.6028*, 2014.
- [6] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2-3, pp. 235–256, 2002.
- [7] J.-Y. Audibert, R. Munos, and C. Szepesvári, "Exploration–exploitation tradeoff using variance estimates in multi-armed bandits," *Theoretical Computer Science*, vol. 410, no. 19, pp. 1876–1902, 2009.
- [8] J.-Y. Audibert and S. Bubeck, "Regret bounds and minimax policies under partial monitoring," *Journal of Machine Learning Research*, vol. 11, no. Oct, pp. 2785–2836, 2010.
- [9] W. R. Thompson, "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples," *Biometrika*, vol. 25, no. 3/4, pp. 285–294, 1933.
- [10] J. Gittins, K. Glazebrook, and R. Weber, *Multi-armed bandit allocation indices*. John Wiley & Sons, 2011.
- [11] O. Chapelle and L. Li, "An empirical evaluation of thompson sampling," in *Advances in neural information processing systems*, 2011, pp. 2249–2257.
- [12] S. L. Scott, "A modern bayesian look at the multi-armed bandit," *Applied Stochastic Models in Business and Industry*, vol. 26, no. 6, pp. 639–658, 2010.
- [13] O.-C. Granmo, "A bayesian learning automaton for solving two-armed bernoulli bandit problems," in *2008 Seventh International Conference on Machine Learning and Applications*. IEEE, 2008, pp. 23–30.
- [14] B. C. May, N. Korda, A. Lee, and D. S. Leslie, "Optimistic bayesian sampling in contextual-bandit problems," *Journal of Machine Learning Research*, vol. 13, no. Jun, pp. 2069–2106, 2012.
- [15] E. Kaufmann, N. Korda, and R. Munos, "Thompson sampling: An asymptotically optimal finite-time analysis," in *International Conference on Algorithmic Learning Theory*. Springer, 2012, pp. 199–213.
- [16] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich, "Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft's bing search engine." Omnipress, 2010.
- [17] D. Agarwal, "Computational advertising: the linkedin way," in *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*. ACM, 2013, pp. 1585–1586.
- [18] D. Agarwal, B. Long, J. Traupman, D. Xin, and L. Zhang, "Laser: A scalable response prediction platform for online advertising," in *Proceedings of the 7th ACM international conference on Web search and data mining*. ACM, 2014, pp. 173–182.
- [19] S. L. Scott, "Multi-armed bandit experiments in the online service economy," *Applied Stochastic Models in Business and Industry*, vol. 31, no. 1, pp. 37–45, 2015.
- [20] D. N. Hill, H. Nassif, Y. Liu, A. Iyer, and S. Vishwanathan, "An efficient bandit algorithm for realtime multivariate optimization," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2017, pp. 1813–1821.
- [21] S. L. Scott, "Overview of content experiments: Multi-armed bandit experiments," 2014. [Online]. Available: <https://support.google.com/analytics/answer/2844870>
- [22] Y. Yang, D. Zhu *et al.*, "Randomized allocation with nonparametric estimation for a multi-armed bandit problem with covariates," *The Annals of Statistics*, vol. 30, no. 1, pp. 100–121, 2002.
- [23] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.
- [24] K. Ferreira, D. Simchi-Levi, and H. Wang, "Online network revenue management using thompson sampling," 2017.
- [25] J. Kawale, H. H. Bui, B. Kveton, L. Tran-Thanh, and S. Chawla, "Efficient thompson sampling for online matrix-factorization recommendation," in *Advances in neural information processing systems*, 2015, pp. 1297–1305.
- [26] A. Chakraborty, B. Paranjape, S. Kakarla, and N. Ganguly, "Stop clickbait: Detecting and preventing clickbaits in online news media," in *Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. IEEE Press, 2016, pp. 9–16.