



You must have clicked on this ad by mistake! Data-driven identification of accidental clicks on mobile ads with applications to advertiser cost discounting and click-through rate prediction

Gabriele Tolomei¹ · Mounia Lalmas² · Ayman Farahat³ · Andrew Haines⁴

Received: 9 January 2018 / Accepted: 1 April 2018
© Springer International Publishing AG, part of Springer Nature 2018

Abstract

In the *cost per click* pricing model, an advertiser pays an ad network *only* when a user clicks on an ad; in turn, the ad network gives a share of that revenue to the publisher where the ad was impressed. Still, advertisers may be unsatisfied with ad networks charging them for “valueless” clicks, or so-called *accidental clicks*. These happen when users click on an ad, are redirected to the advertiser website and bounce back without spending *any* time on the *ad landing page*. Charging advertisers for such clicks is detrimental in the long term as the advertiser may decide to run their campaigns on other ad networks. In addition, machine-learned click models trained to predict which ad will bring the highest revenue may overestimate an ad click-through rate, and as a consequence negatively impacting revenue for both the ad network and the publisher. In this work, we propose a data-driven method to detect accidental clicks from the perspective of the ad network. We collect observations of time spent by users on a large set of ad landing pages—i.e., *dwelt time*. We notice that the majority of per-ad distributions of dwell time fit to a *mixture of distributions*, where each component may correspond to a particular type of clicks, the first one being accidental. We then estimate dwell time *thresholds* of accidental clicks from that component. Using our method to identify accidental clicks, we then propose a technique that *smoothly* discounts the advertiser’s cost of accidental clicks at billing time. Experiments conducted on a large dataset of ads served on Yahoo mobile apps confirm that our thresholds are stable over time, and revenue loss in the short term is marginal. We also compare the performance of an existing machine-learned click model trained on all ad clicks with that of the same model trained only on non-accidental clicks. There, we observe an increase in both ad click-through rate (+3.9%) and revenue (+0.2%) on ads served by the Yahoo Gemini network when using the latter. These two applications validate the need to consider accidental clicks for both billing advertisers and training ad click models.

Keywords Accidental ad clicks · Online mobile advertising · Dwell time · Mixture of distributions · Ad cost discounting · Click-through rate prediction

All the authors contributed to this work while employed at Yahoo Research.

✉ Gabriele Tolomei
gtolomei@math.unipd.it
Mounia Lalmas
mounia@acm.org
Ayman Farahat
afarahat@amazon.com
Andrew Haines
haines@oath.com

¹ University of Padua, Padua, Italy

² Spotify, London, UK

1 Introduction

The time a user spends on a web page, referred as *dwell time*, varies considerably by user and type of content. Still, for a given web page, dwell time can be used as an effective *proxy* of engagement with its content [4,10,16,29]. If the content is of no interest to the user or is presented poorly, the dwell time on the web page will often be short. By contrast, for a user engaged with the content the dwell time will be longer. A third case, which does not depend on the content or its

³ Amazon, Palo Alto, CA, USA

⁴ Yahoo Research @ Oath, London, UK




Fig. 1 Empirical distribution of the logarithm of dwell time observed on an ad landing page

presentation, corresponds to *extremely* short dwell time. This paper focuses on these extremely short dwell times.

In online advertising, when a user clicks on an advertisement, or ad for short, she is redirected to the advertiser web page, i.e., the *ad landing page*. The dwell time on the ad landing page is measured as the time between the ad click and the user returning to the publisher site where the ad was impressed. There are three *types* of clicks:

- *Accidental click* The user clicks on the ad, likely by mistake, reaches the ad landing page, and immediately bounces back; the user spends no time on the landing page.
- *Short click* The user intended to click on the ad but once on the landing page decides to bounce back; the post-click experience is not satisfactory, due to the low quality of the landing page or its low relevance.
- *Long click* Once landing on it, the user engages with the ad landing page and spends time on the advertiser site.

We plot in Fig. 1 the distribution of the natural logarithm of dwell time values observed on an ad landing page. The distribution is not unimodal; a small component can be identified with extremely low dwell time values, around $e^{7.5} \approx 1.8$ s, as representative of accidental clicks, whereas the other two components capture the short and long clicks, respectively, demonstrating the existence of the above mentioned three types of clicks. In this paper, we propose a data-driven approach for estimating the dwell time *thresholds* to identify whether a click is accidental or not.

Properly accounting for accidental clicks is business-critical to online advertising. Consider the widely adopted *cost per click* (CPC) pricing model, when an advertiser is charged by the ad network only when a user clicks on an ad. Users may accidentally click on an ad, get redirected to the ad landing page and then bounce back without looking at it. This behaviour is even more severe on smartphones, as their limited screen size makes it more prone for users to click on ads by mistake [12,27]. It is therefore not unusual for advertisers running CPC campaigns to complain when charged for such clicks, as these are “valueless” to them. Ignoring these complaints can be detrimental in the long term, as it may affect the relationship between the ad network and the advertiser, who might at worst switch to other ad networks to run their campaigns.

A main challenge is for the ad network to accurately identify accidental clicks. Current solutions use hand-coded thresholds of dwell time on landing page to determine whether an ad click is accidental or not; for example, all visits to ad landing pages shorter than 5 s may be considered as accidental. With this approach, the threshold is fixed and set arbitrarily, and does not take into account empirically observed dwell time of ad clicks. Plus, while 5 s may be a reasonable threshold for ad clicks on a tablet, it may not be the right one for a smartphone. As a first and major contribution of this paper, we propose an unsupervised learning method that estimates thresholds of accidental clicks by fitting observed dwell time data to mixture models, which capture the three types of clicks defined above.

Further, we deploy our method for identifying accidental clicks on two applications. The first one is concerned with a controlled approach to discount accidental clicks when charging advertisers. In principle, accidental clicks could be simply discarded once detected, and advertisers not charged for them. However, this strategy will negatively impact revenue for both the ad network and, consequently, the publisher. Therefore, as a second contribution of this paper, we propose a *smooth* discounting method based on the proportion of detected accidental clicks. The third contribution of this paper presents how accidental clicks identified by our approach are discarded from an existing machine-learned clicks model used to predict an ad *click-through rate* (CTR for short). Our intuition is that by removing valueless clicks we can feed machine-learned models with “cleaner” training data, mitigating any possible overestimation of CTR.

The rest of the paper is organised as follows. In Sect. 2, we motivate why we use dwell time as a proxy of an ad click value, and discuss why it is appropriate to adopt a data-driven solution for discovering accidental clicks. Sections 3 and 4 describe our data-driven approach to detect accidental clicks and experiments carried out with real-world data, respectively. Sections 5 and 6 present two applica-

Table 1 Statistics on the natural logarithm of dwell time

Conversion	Mean	SE	SD
yes	5.729	0.064	2.582
no	3.264	0.006	2.569

tions where our proposed approach to identify accidental click was deployed. In Sect. 7, we discuss related work and position our contributions. Finally, Sect. 8 concludes our work.

2 Accidental clicks

We show first how we estimate the value of an ad click using dwell time, and then motivate using a data-driven approach to identify accident clicks.

2.1 Using dwell time as the value of an ad click

Advertisers may wish to be fully charged only for valuable clicks, i.e., those that lead to *conversions*,¹ whereas all remaining clicks should be proportionally discounted. Implementing this strategy requires an accurate estimate of the *conversion rate*, i.e., the probability of conversion conditioned on a click. However, conversion rate suffers from three major problems. First, it is hard to estimate since conversion data are often not available for a large number of advertisers. Second, those data are not missing at random, as advertisers sharing their conversion data would be a biased sample. Finally, using conversion data to identify valuable clicks may lead to high false positive rates, since clicks not followed by conversions are not necessarily “valueless”. In fact, those clicks represent profitable feedbacks for advertisers.

It is therefore more appropriate to focus on identifying valuable clicks as clicks that are *not* valueless, thus using so-called accidental clicks as complementary of conversions. We propose to use dwell time as our proxy measure of the value of an ad click. Using dwell time can help alleviate the aforementioned problems provided that dwell time is indeed a good proxy for conversion in our data. This was shown to be indeed the case in [10], and we further validate this in our context.

Table 1 shows statistics on the natural logarithm of dwell time from two samples, one containing observations of dwell time leading to conversions (*yes*) and the other made up of those that do not (*no*) for 40 ads managed by a large ad

¹ There is no consensus around what a *conversion* is; it is up to the advertiser to specify it.

Table 2 Coefficients of linear and logit regression model

	Linear (Eq. 1)	Logit (Eq. 2)
β_0 (intercept)	− 0.00384	− 6.424
β_1 (dwell time)	0.00413	0.301
AIC	14,399	14,297

network, Yahoo *Gemini*.² We test against the null hypothesis that the mean of dwell time computed from the *yes*-sample ($e^{5.729} \approx 307.7$ s) is the same as that calculated from the *no*-sample ($e^{3.264} \approx 26.2$ s). We run a two-tailed two-sample *t-test* and are able to reject the null hypothesis stating that the difference of those two means is 0 (level of significance $\alpha = 0.01$, p value $\ll 0.01$). Therefore, the average dwell time of the two samples are indeed statistically significantly different.

In addition, we test two regression models to verify if dwell time is a good predictor of conversion. Let $Y_{i,j}$ be the binary random variable representing the conversion event occurring after the j th click on the i th ad. We also denote by $X_{i,j}$ the logarithm of dwell time on the i th ad landing page after the j th click. The simplest model we test is a linear regression of $Y_{i,j}$ on $X_{i,j}$, that is:

$$Y_{i,j} = \beta_0 + \beta_1 X_{i,j} + \epsilon_{i,j} \quad (1)$$

where $\epsilon_{i,j} \sim \mathcal{N}(0, \sigma_{i,j}^2)$ is a zero-mean Gaussian error term. A variant of this model simply applies the logit operator to the right-hand side of Eq. 1, namely:

$$Y_{i,j} = \text{logit}(\beta_0 + \beta_1 X_{i,j}) + \epsilon_{i,j} \quad (2)$$

Table 2 shows the regression coefficients obtained with the two models above. The best performing model is the logit, which is selected as the one with the smallest *Akaike Information Criterion* (AIC) [1]. We observe that the coefficient associated with the natural logarithm of dwell time (β_1) is both positive and significant, hence is a good predictor of conversion. It is worth noting that a *unit* increase in natural logarithm of dwell time (i.e., ≈ 2.72 s, since 1 unit = $\ln(x)$, where x is dwell time in seconds) will increase the odds of conversion by 30%. This provides further justification in using dwell time as proxy of conversion, and therefore of ad click value.

2.2 A data-driven approach to identifying accidental clicks

Using dwell time as proxy of the “value” of an ad click, this paper puts forward a data-driven approach—based on

² <https://gemini.yahoo.com/>.

actual dwell time observations—to identify whether an ad click is accidental or not. On the other hand, using fixed thresholds on dwell time to identify accidental clicks—say 1 s—is instead a common practice. The reason being is that this is straightforward, as no data processing nor computation is required, and can be easily integrated with any existing production systems. This simple approach however prevents capturing subtleties arising from hidden latent factors, such as the device (desktop vs. mobile), the application (e.g., mail vs. news stream) or the network type (e.g., 3G/4G vs. Wi-Fi). A data-driven approach can account for all these factors, thus providing more reliable *estimates* of dwell time thresholds of accidental clicks from observed data. Moreover, the analysis of observed dwell times may characterise other phenomena; e.g., not only the lowest values (accidental clicks) but also large and extremely large ones (short and long clicks, respectively). We leave this for future work, as we focus on accidental clicks in this paper.

3 Discover accidental clicks

Our data-driven approach for detecting accidental clicks on ads consists of two steps: *data modelling* and *dwell time thresholding*. The former fits observations of dwell time of a large set of ad landing pages to a probabilistic model, whereas the latter estimates threshold of dwell time to identify accidental clicks. We employ an unsupervised learning approach, as no supervised learning approaches to classify ad clicks as accidental can be designed as this would require building a ground truth, which is not achievable.

3.1 Data modelling

We assume that observations of dwell time of ad clicks are generated by an underlying probabilistic model. Ideally, such a model has to simultaneously represent the three types of clicks shown in Fig. 1, *accidental*, *short*, and *long*.

Generally speaking, a *mixture of distributions* is a probabilistic model that captures the presence of “sub-populations” within an overall population [17,19]. As such, it is a good candidate to describe our observations of dwell time. More formally, a continuous random variable X (e.g., dwell time) is distributed according to a mixture of K (discrete) component distributions if its probability density function (pdf) f_X is a convex combination of K pdfs f_1, \dots, f_K :

$$\begin{aligned} f_X(x; \Theta) &= w_1 f_1(x; \theta_1) + \dots + w_K f_K(x; \theta_K) \\ &= \sum_{i=1}^K w_i f_i(x; \theta_i) \end{aligned} \quad (3)$$

where:

- each f_i belongs to the same (parametric) family of distributions (e.g., Normal,³ log-normal, gamma, Weibull, etc.);
- w_i is the *mixture weight* (or prior probability) associated with the i th component;
- $w_i \geq 0$ and $\mathbf{w} = (w_1, \dots, w_K)^T$ is the K -dimensional vector of mixture weights, so that $\sum_{i=1}^K w_i = 1$;
- θ_i is the vector of *parameters* associated with the i th component, e.g., if f_i is the pdf of a Normal distribution $\mathcal{N}(\mu_i, \sigma_i^2)$ then $\theta_i = (\mu_i, \sigma_i^2)$;
- $\Theta = (\mathbf{w}, \theta_1, \dots, \theta_K)$ is the *overall* vector of parameters of the mixture model;
- there exists a *latent random variable* denoted by I governing which component each observation of X is drawn from. This random variable is distributed according to a categorical distribution whose parameter is the vector of mixture weights \mathbf{w} , so that:
 1. $I \sim \text{Categorical}(\mathbf{w}) \mapsto$ pick the component distribution f_i with probability w_i ;
 2. $X = x' \mid I = i \mapsto$ generate a value for X from the component distribution f_i .

We describe next how to model dwell time on ad landing pages using a mixture of distributions.

3.1.1 Mixture of distributions for dwell time

Representing the three types of clicks can be done by having the observed dwell times on ad landing pages generated by a mixture of $K = 3$ distributions. Let M be the number of ads. For each ad $j \in \{1, \dots, M\}$ we consider a sample \mathcal{D}_j of n_j i.i.d. *positive* random variables $X_{j,1}, \dots, X_{j,n_j}$, with each $X_{j,k}$ representing an observation of the dwell time associated with the k th click on the ad j : $\mathcal{D}_j = \{X_{j,1} = x_{j,1}, \dots, X_{j,n_j} = x_{j,n_j}\}$. Each $X_{j,k}$ is drawn from a mixture of up to $K = 3$ components, so that the pdf of $X_{j,k}$ is:

$$f_{X_j}(x; \Theta) = \sum_{i=1}^3 w_i f_i(x; \theta_i) \quad (4)$$

In addition, each f_i is the pdf of the same parametric distribution, although we only consider probability distributions with positive domain (e.g., log-normal, gamma, Weibull) as dwell time cannot be negative.

Next, we discuss how the parameters Θ of the mixture model can be estimated from the observed data.

³ We refer to *Normal* and *Gaussian* distribution, interchangeably.

3.1.2 Parameter estimation

For each ad, we estimate the overall vector of parameters $\Theta = (\mathbf{w}, \theta_1, \theta_2, \theta_3)$ from the observed dwell times using *maximum likelihood estimation* (MLE). For each ad j we know that the pdf of each of its observations is a mixture of three distributions, defined as in Eq. 4. Since all observations in \mathcal{D}_j are independent and identically distributed, we compute their joint probability density f_{X_j} as:

$$f_{X_j}(\mathcal{D}_j; \Theta) = f_{X_j}(x_{j,1}, \dots, x_{j,n_j}; \Theta) = \prod_{k=1}^{n_j} f_{X_j}(x_{j,k}; \Theta) \tag{5}$$

From the joint probability density, we derive the *likelihood function* $L(\Theta; \mathcal{D}_j)$, as $L(\Theta; \mathcal{D}_j) = f_{X_j}(\mathcal{D}_j; \Theta)$. Although the two functions are the same, the likelihood function emphasises that the dataset is fixed and the parameters Θ are variable. The aim of MLE is thus to find a value of Θ —i.e., an estimate $\hat{\Theta}_{MLE}$ —maximising the likelihood function:⁴

$$\begin{aligned} \hat{\Theta}_{MLE} &= \operatorname{argmax}_{\Theta} [L(\Theta; \mathcal{D}_j)] \\ &= \operatorname{argmax}_{\Theta} \left[\prod_{k=1}^{n_j} f_{X_j}(x_{j,k}; \Theta) \right] \\ &= \operatorname{argmax}_{\Theta} \left[\prod_{k=1}^{n_j} \sum_{i=1}^3 w_i f_i(x_{j,k}; \theta_i) \right] \end{aligned} \tag{6}$$

Different likelihood functions to be maximised can be obtained depending on which pdf we fill in Eq. 6 with. However, if the resulting $L(\Theta; \mathcal{D}_j)$ is differentiable in Θ , we can find $\hat{\Theta}_{MLE}$ as a solution of the system of equations:

$$\nabla L(\Theta; \mathcal{D}_j) = \frac{\partial L(\Theta; \mathcal{D}_j)}{\partial \theta} = 0, \theta \in \Theta \tag{7}$$

where ∇ is the *gradient* of the likelihood function, i.e., the vector of partial derivatives of the likelihood function with respect to each parameter in Θ .

Unfortunately, a maximum likelihood estimation of the parameters Θ is not straightforward, since often there are no closed-form solutions to Eq. 7 available; as such, we cannot solve directly $\nabla L(\Theta; \mathcal{D}_j) = 0$ [23]. A typical solution is to use the *expectation-maximization* (EM) algorithm [6]. EM is an *iterative*, numerical approximation procedure that starts with an initial random guess for the values of the parameters and converges to a local maximum (or to a saddle point) of

⁴ In practice, we often seek for $\hat{\Theta}_{MLE}$ so as to maximise the *log-likelihood function* $\ln(L(\Theta; \mathcal{D}_j))$, since this is equivalent (the natural logarithm is monotonically increasing) but simpler because products change into summations.

the observed-data likelihood. Although EM does not guarantee convergence to a global maximum, in practice there are a variety of heuristic approaches for escaping a local maximum: multiple restarts, clever initialization, and modifications to the EM algorithm itself [9]. Finally, for each ad, we compute the set of model parameters $\hat{\Theta}_{MLE}$ which maximise the likelihood function.

3.1.3 Model selection

To choose the “best” model for each ad, we cannot just select the one with parameters $\hat{\Theta}_{MLE}$, i.e., the one that best fits to the observed data. In general the more complex (flexible) is the model the better will be its goodness-of-fit to the observed data; in other words, the higher will be its likelihood function computed with respect to the observed data. At the same time, the more complex (flexible) is the model the less it generalises to unseen data; in other words, the higher is the chance of the model to *overfit* the observed data [13]. Therefore, if we choose the model having the highest likelihood we always end up selecting the one having the maximum degree of freedom, i.e., the maximum number of components $K = 3$.

Therefore, to avoid overfitting and find a trade-off between complexity and interpretability⁵, we use tools such as the *Akaike Information Criterion* (AIC) [1] or *Bayesian Information Criterion* (BIC) [24]. The former is computed as $AIC = -2 \ln(L) + 2K$, whereas the latter as $BIC = -2 \ln(L) + K \ln(n)$, where K is the number of components of the model, L is the likelihood function as maximised by the parameters of the model estimated from the observed data, and n is the dataset size. Both criteria try to penalise models that are unnecessarily too complex, and finally select the one with the *smallest* AIC or BIC.

So far, we have estimated and selected the mixture model that best describes the observed dwell times on each ad landing page. Next, we present how we can use this model to compute dwell time thresholds to identify accidental clicks.

3.2 Dwell time threshold of accidental clicks

For each ad, we fit the observed dwell times on its landing page to a mixture of distributions using MLE and one of the model selection criteria described above. We then focus on the subset of ads exhibiting exactly *all* three components, namely ads with dwell times fitting a mixture of three distributions. Intuitively, these are the ads showing all the three categories of clicks we have conjectured the existence of, namely *accidental*, *short* and *long*.

Given an ad and the set of parameters of all its components, we compute statistics such as the expected value or the

⁵ This is also often referred to as the *bias-variance* trade-off.

median of every component. As we are interested in detecting accidental clicks we only focus on the *first* component of each ad. Using the second and third component to study short and long clicks, respectively, is something we leave for future work.

For example, if we fit the data to a mixture of three log-normal distributions we can represent the first component by a random variable distributed as a log-normal with parameters $\theta_1 = (\mu_1, \sigma_1^2)$. In general, for any random variable $Z \sim \text{Ln}\mathcal{N}(\mu, \sigma^2)$, or equivalently $\ln(Z) \sim \mathcal{N}(\mu, \sigma^2)$, we can derive the following:

- $E[Z] = e^{\mu + \sigma^2/2}$ (where E denotes the expected value);
- $\text{Median}(Z) = e^\mu$.

We therefore estimate a per-ad threshold of dwell time for detecting accidental clicks using either the expected value or the median of the first component—the latter being more robust to the presence of outliers—by letting $\mu = \mu_1$ and $\sigma^2 = \sigma_1^2$ in the equations above. Finally, to obtain an overall estimate of dwell time threshold of accidental clicks across all the ads, we compute the mean or the median of the individual per-ad estimates.

In this section, we described a two-stage approach for computing thresholds of dwell time from observed data, to detect accidental clicks on ads. Next, we present experimental results when these thresholds are deployed within a large ad network.

4 Experiments

We conduct two experiments on multiple datasets of ads served by a large ad network, codenamed *Gemini*, on several Yahoo mobile apps. We focus on mobile apps as these are where accidental clicks are more likely to happen [12,27]. We only consider ads with at least 100 clicks to increase the confidence of the estimates of our thresholds.

In the first experiment, we choose one *pivoting* mobile app to estimate the dwell time threshold of accidental clicks, which is then used to identify accidental clicks on other (two in our case) mobile apps. To protect sensitive information, we refer to the former as **App 1** and to the other two as **App 2** and **App 3**, respectively. The experiment is performed on two 1-month datasets, which we refer to \mathcal{D}_1 and \mathcal{D}_2 , respectively. Each consists of a random sample of around 10,000 ads and 6.5M clicks, unevenly distributed on the three mobile apps. In the second experiment, instead of having one pivoting mobile app used to estimate a *single* dwell time threshold of accidental clicks, we generate a threshold for each mobile app. The dataset used is a random sample from three weeks worth

of data containing 120,000 ads and 70M clicks, hereinafter called \mathcal{D}_3 .

With the first approach—using a pivoting app—the aim is provide a single estimate of dwell time threshold of accidental clicks, using the app with the highest volume of traffic. The second approach provides one threshold per app, which is more flexible and accounts for the effect of different user experience, population, and operating systems (i.e., Android vs. iOS).

4.1 Data preprocessing

In both experiments, we remove outliers by discarding clicks with dwell time greater than 600s.⁶ We also apply a logarithmic transformation to all the observations. This allows us to fit the log-transformed data to a mixture of Gaussian distributions, as this is the same as fitting the original data to a mixture of log-normals.

The logarithmic transformation emphasises differences between small values of dwell time, whereas it smooths the same differences when they happen between larger values of dwell time. In this way, we are able to capture relative differences instead of absolute ones. To identify accidental clicks a difference of 1 s between two small values of dwell time, such as 2 and 3 s, is more important than the same difference between, say, 101 and 100 s.

Our approach is not tailored to any specific family of parametric distributions, and we select a mixture of log-normal distributions purely because this gives us the best (smallest) AIC on average over all the ads.

4.2 Single threshold from a pivoting app

In this first setting, we consider (log-transformed) observations of dwell time from the ads clicked on our pivoting mobile app—**App 1**—collected both from \mathcal{D}_1 and \mathcal{D}_2 . Then, we fit each set of ad dwell time observations to a mixture of Gaussian distributions. Figure 2 presents three examples of ads, each one fitted to mixture of one, two, and three Gaussian distributions. Except for the ad shown in Fig. 2a, the other two exhibit a first component centred around very small value of dwell time (around $e^{7.7} \approx 2.2$ s), which likely represents accidental clicks.

Table 3 shows how ads in the two datasets fit to models having one, two, and three components, respectively. For both, the vast majority of ads (82.5 and 65.4%, respectively) fit to exactly three components. These ads are the ones we focus on to detect the dwell time threshold of accidental clicks, since their corresponding landing pages contain all the three categories of clicks described in Sect. 1. Since our

⁶ A threshold already used in previous work [16].




Fig. 2 Examples of ads clicked on the pivoting App 1 which fit to one, two, and three components

Table 3 Percentage of ads clicked on the pivoting App 1 which fit to one, two, and three components

	1 Comp (%)	2 Comps (%)	3 Comps (%)
\mathcal{D}_1	1.0	16.5	82.5
\mathcal{D}_2	2.9	31.7	65.4

aim is to “isolate” accidental clicks happening on our pivoting app, we concentrate on the first component. According to our conjecture, this would capture users clicking on an ad by mistake, or simply returning to the publisher site without actually landing on the advertiser’s page.

For each ad, we compute an estimate of the dwell time threshold using the median of its first fitted component, as this is more robust to the variance. In fact, we observe that the variance increases going from the first to the second and finally to the third component. Intuitively, this reflects the variability of dwell time on each click category: Dwell times of accidental clicks are expected to differ less between each other than what would be the case with short and long clicks.

To obtain an overall estimate of the threshold t_{acc} (an estimate derived from all the per-ad estimates), we propose two strategies: (i) the mean of all the per-ad medians; (ii) the median of all the per-ad medians. The first one results in a generally higher threshold, which implies considering “accidental” a larger number of clicks. The second estimate is more “conservative” and usually generates a smaller value of the threshold.

In Fig. 3, we plot the distribution of per-ad medians of the first component computed from the pivoting App 1 on \mathcal{D}_1 and \mathcal{D}_2 . The median of all those medians (the blue dashed line) seems more suitable than the mean (the red dashed line), as it perfectly aligns with the “peak” we are interested in, which sits around $2.1 \div 2.2$ s. For business confidentiality, we do not disclose the percentage of accidental clicks for each of

the considered apps, but we can report that this percentage is stable over the two datasets, once the thresholding strategy is fixed. Anecdotally, this percentage was shown to change using a dataset from a different time period, as the result of a change of the user interface on a specific app.

4.3 Multiple per-app thresholds

In the previous setting, we reported results obtained when using a pivoting mobile app to compute a single threshold of dwell time, which in turn can be used to identify accidental clicks for other apps. In this section, we discuss another approach, which was rolled out in production. Instead of computing a single threshold from one pivoting app, we generate a dwell time threshold of accidental clicks for each app, individually. By doing so, we are also accounting for the effect of different user experiences or user populations on the different mobile apps and operating systems (i.e., Android vs. iOS). We use a default threshold value for apps with not enough observations of dwell time.

In Fig. 4, we plot the empirical cumulative distribution function (eCDF) of the thresholds computed on each week of the dataset (\mathcal{D}_3). We observe that the distribution of thresholds remains stable over time. The median of all the per-app thresholds identified with this approach results in a value of 2.1 s, which aligns with the threshold generated using the pivoting strategy.

In Fig. 5, we separate the eCDFs of the thresholds obtained from apps on Android and iOS. We see that there are differences, as the median values are now ≈ 2.44 and ≈ 1.80 s, respectively. In general, thresholds of accidental ad clicks on Android apps are more right-skewed than those computed for iOS apps, thereby suggesting that thresholds on Android are somewhat less dependent from the app which they are computed on.

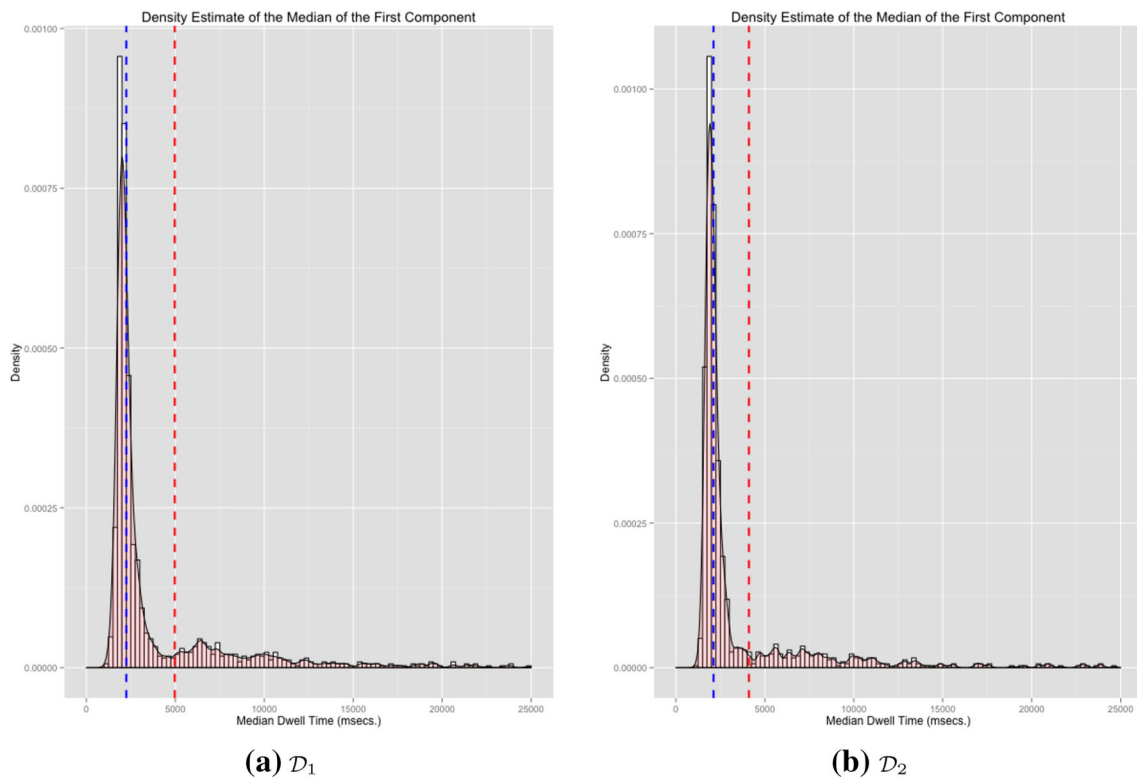


Fig. 3 Distributions of per-ad medians computed from the pivoting App 1 (median vs. mean)

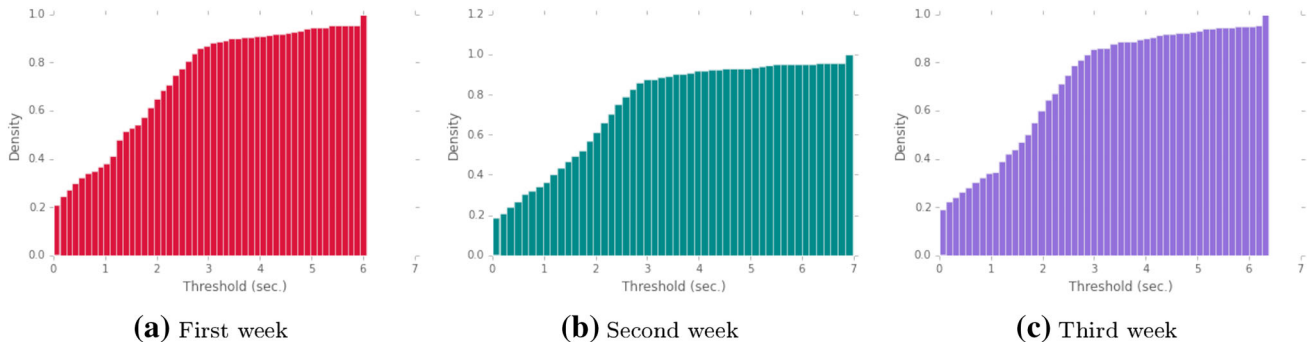


Fig. 4 Empirical CDF of thresholds of accidental ad clicks on mobile apps computed on \mathcal{D}_3

We presented two strategies to calculate dwell time thresholds of accidental clicks. Both strategies rely on the same data-driven approach, i.e., estimating the parameters of a mixture of three dwell time distributions and computing aggregated statistics (e.g., the median) on the first component. The two strategies differ in the (historical) dwell time data used to fit the mixtures. One uses observations of dwell time *only* from a single app, and hence provides the threshold for apps with few historical dwell time observations, who may be less popular than the pivoting app or have just entered the market. The other provides a per-app threshold, which can be used when there are multiple apps with a sufficiently large number of dwell time observations. In addition, this strategy

considers the impact of different user experience on those apps.

In the next two sections, we present two use cases where our proposed data-driven approach for identifying accidental clicks was deployed.

5 Use case I: discounting accidental clicks when billing advertisers

With a mechanism for detecting accidental clicks, an ad network may simply discard all accidental clicks when billing the advertiser. This can, however, severely impact revenue

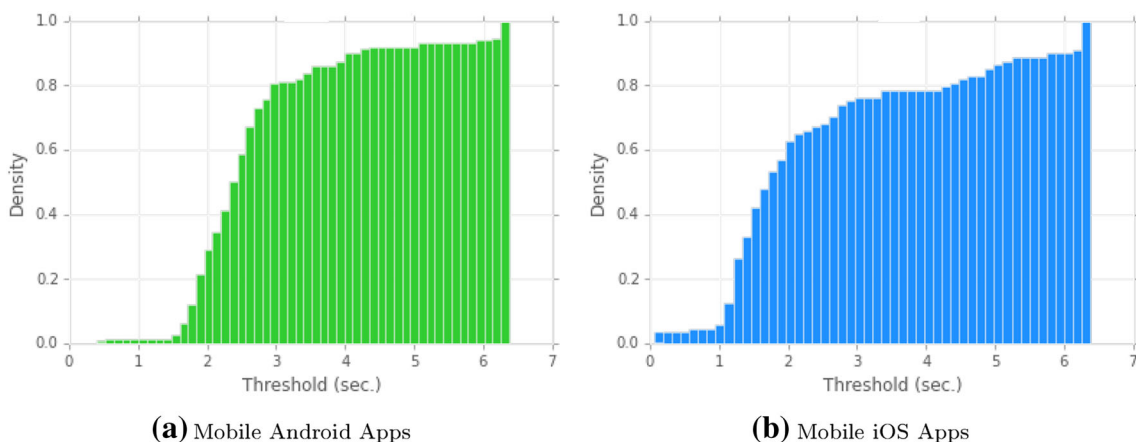


Fig. 5 Empirical CDF of thresholds of accidental ad clicks computed on \mathcal{D}_3 across different platforms

for both the ad network and the publisher, at least in the short term. For example, in our dataset \mathcal{D}_3 , we saw that the top-3 most revenue-losing apps account for $\approx 72.4\%$ of the overall revenue loss for all the apps under consideration.⁷ It is therefore important to control how much revenue loss is acceptable, hence looking for a trade-off between accounting for accidental clicks (satisfying the advertisers) and containing revenue loss (satisfying the ad network and publishers). We present a smooth method to *discount* the price of accidental clicks, instead of discarding them, so that advertisers are not fully charged for those clicks.

5.1 Smooth discounting strategy

One of the main attractions of ad networks is scale; advertisers have access to a large number of impressions and reach a wide audience with a single buy. However, not all the apps in the network perform equally. The advertiser is then faced with the problem of either selecting which apps to bid on or adjusting the bids by app. Both cases create extra friction for the advertiser. The algorithmic discounting we present below addresses this problem by adjusting the cost per click on each app of the network, such that the return on investment (ROI) for the advertiser is the same across all the apps.

We assume the existence of a *pivoting* app from which we estimate the dwell time threshold of accidental clicks t_{median} , such as discussed in Sect. 4.2. Let $C_{i,\text{pivot}}$ be the set of clicks observed on ad i , which have been impressed on the pivoting app on a fixed time window. Moreover, let $\text{clicks}_{i,\text{pivot}} = |C_{i,\text{pivot}}|$ be the total number of observed clicks on i . Therefore, $C_{i,\text{pivot}}^{\text{non-acc}} = \{c \in C_{i,\text{pivot}} \mid \tau(c) > t_{\text{median}}\}$ —where $\tau(c)$ is the dwell time of c —is the set of *non-accidental* clicks on i identified using t_{median} on the same time window, and $\text{clicks}_{i,\text{pivot}}^{\text{non-acc}} = |C_{i,\text{pivot}}^{\text{non-acc}}|$ is the total number of

non-accidental clicks on i . Similarly, for any other app, we define $\text{clicks}_{i,\text{app}}$ and $\text{clicks}_{i,\text{app}}^{\text{non-acc}}$.

An advertiser may associate a *value* to each click on ad i , referred to as vfc_i . This corresponds to the amount of money the advertiser would like to earn from a click on i , *independently* of the source (app, in this case) where such click occurs. Under a CPC cost model, there is a maximum amount of money the advertiser is willing to pay for having ad i impressed and clicked, denoted by cpc_i . We define the advertiser ROI for ad i on the generic app as:

$$\frac{\text{vfc}_i \times \text{clicks}_{i,\text{app}}^{\text{non-acc}}}{\text{cpc}_i \times \text{clicks}_{i,\text{app}}} \tag{8}$$

where the numerator is the total value earned by the advertiser considering only valid—non-accidental—clicks on app, and the denominator is the total cost the advertiser would pay for *all* the clicks on ad i occurred on app. If we knew what is the *true* non-accidental click rate of the app (NACR_{app}), we could rewrite Eq. 8 as:

$$\frac{\text{vfc}_i \times \text{clicks}_{i,\text{app}}^{\text{non-acc}}}{\text{cpc}_i \times \text{clicks}_{i,\text{app}}} \approx \frac{\text{vfc}_i \times \text{NACR}_{\text{app}}}{\text{cpc}_i} \tag{9}$$

Indeed, $\frac{\text{clicks}_{i,\text{app}}^{\text{non-acc}}}{\text{clicks}_{i,\text{app}}}$ is the MLE estimate of the true NACR_{app} . We require the app we chose as pivot not only to be the one from which we can accurately estimate t_{median} but also the best performing app with the highest NACR, i.e., the highest proportion of valid ad clicks overall. This is because we want the pivoting app to be the “benchmark” against which we compare all the other apps of the network.

The advertiser ROI calculated on any app of the network should be ideally equal to that of the pivoting app:

$$\frac{\text{vfc}_i \times \text{NACR}_{\text{app}}}{\text{cpc}_i} = \frac{\text{vfc}_i \times \text{NACR}_{\text{pivot}}}{\text{cpc}_i} \tag{10}$$

⁷ The actual revenue loss is not shown due to business confidentiality.

Moreover, since the value that the advertiser would get from a click on ad i (vfc_i) is independent of the source, we can rewrite Eq. 10 as:

$$\frac{\text{NACR}_{\text{app}}}{\text{cpc}_i} = \frac{\text{NACR}_{\text{pivot}}}{\text{cpc}_i} \quad (11)$$

We may observe that $\frac{\text{NACR}_{\text{app}}}{\text{NACR}_{\text{pivot}}} \leq 1$, since $\text{NACR}_{\text{pivot}}$ is the highest among all the apps by design. For Eq. 11 to be satisfied, we define $\text{cpc}'_{i,\text{app}} \leq \text{cpc}_i$ as the *adjusted* cost of each accidental click on ad i , specific to app:

$$\text{cpc}'_{i,\text{app}} = \frac{\text{NACR}_{\text{app}} \times \text{cpc}_i}{\text{NACR}_{\text{pivot}}} \quad (12)$$

The intuition is that the cost of an accidental click on i on a generic app ($\text{cpc}'_{i,\text{app}}$) should be obtained by discounting the price for a valid click (cpc_i) proportionally to its relative value with respect to the best performing pivoting app of the network ($\frac{\text{NACR}_{\text{app}}}{\text{NACR}_{\text{pivot}}}$), which is exactly our *discount factor*.

This strategy does not discount accidental ad clicks on the pivoting app itself. This is because the pivoting app is chosen as the one with the smallest accidental click rate, thus likely with little need to apply a discount factor. Nonetheless, we may also decide to account for accidental clicks on the pivoting app, especially if its click value performance deteriorates. Various strategies may be deployed. For example, we can monitor the number of accidental clicks on the pivoting app and if it eventually exceeds some established threshold, we can apply to those accidental clicks a default discounting strategy by choosing among one of the discount factors computed for the other apps.

5.2 Estimating non-accidental click rate

To implement our proposed discounting strategy, we must accurately estimate the (binomial) proportion of non-accidental clicks NACR of an app. The most straightforward way is to use maximum likelihood estimate, namely a single-point estimate $\widehat{\text{NACR}}_{\text{MLE}} \approx \text{NACR}$, which is the overall number of non-accidental ad clicks divided by the total number of ad clicks observed during a specific time window: $\widehat{\text{NACR}}_{\text{MLE}} = \frac{\sum_{i \in \text{ads}} \text{clicks}_i^{\text{non-acc}}}{\sum_{i \in \text{ads}} \text{clicks}_i}$. This estimate, however, is not robust when we have an app with a low number of observations. To overcome this, we compute the confidence interval for $\widehat{\text{NACR}}_{\text{MLE}}$.

There exist several ways to compute a confidence interval for an estimate \hat{p} of a binomial proportion p .⁸ The normal approximation interval is the simplest and most common approach, and assumes the distribution of error of a binomially distributed observation to be Gaussian. This is computed

as $\hat{p} \pm z \sqrt{\frac{1}{n} \hat{p}(1 - \hat{p})}$, where \hat{p} is the proportion of successes in a Bernoulli trial process estimated from the statistical sample, z is the $1 - \frac{1}{2}\alpha$ percentile of a standard normal distribution, α is the error percentile and n is the sample size.

The normal approximation however does not always work. Several competing formulas are available that perform better, especially for situations with a small sample size and a proportion very close either to 0 or 1. The choice will depend on how important it is to use a simple and easy-to-explain interval versus the desire for better accuracy. As such, the Agresti–Coull interval [7] is another approximate binomial confidence interval, which is more robust than the normal approximation interval. Given X successes in n Bernoulli trials, it defines the following quantities: $\tilde{n} = n + z^2$ and $\tilde{p} = \frac{1}{\tilde{n}} (X + \frac{1}{2}z^2)$. Then, a confidence interval for p is given by: $\tilde{p} \pm z \sqrt{\frac{1}{\tilde{n}} \tilde{p}(1 - \tilde{p})}$, where z is the $1 - \frac{1}{2}\alpha$ percentile of a standard normal distribution, as before.

5.3 Comparing non-accidental click rates

The proposed discounting strategy requires computing the ratio of non-accidental click rates between the app of interest and the pivoting app. If the estimate of NACR is not a single-point estimate such as $\widehat{\text{NACR}}_{\text{MLE}}$ but is a confidence interval, one way to compare two NACR estimates is to take the ratio of their *upper confidence bounds*:

$$\frac{\text{NACR}_{\text{app}}}{\text{NACR}_{\text{pivot}}} = \frac{\text{ucb}(\text{NACR}_{\text{app}})}{\text{ucb}(\text{NACR}_{\text{pivot}})} \quad (13)$$

where ucb is the upper confidence bound computed by the Agresti–Coull interval defined at the end of previous section.

We already stated that the pivoting app is assumed to be the one with the highest NACR . However, for our discounting strategy to be robust it should also account for the case when an app is overperforming in terms of “click value” the pivoting app⁹. In such a case, we would like the discount factor $\frac{\text{NACR}_{\text{app}}}{\text{NACR}_{\text{pivot}}}$ to be greater than 1 only when we have a degree of confidence in it. One way to implement this is to require the overperforming app’s *lower confidence bound* (lcb) being greater than the pivot’s. We can therefore modify Eq. 13 to:

$$\frac{\text{NACR}_{\text{app}}}{\text{NACR}_{\text{pivot}}} = \max \left[\frac{\text{lcb}(\text{NACR}_{\text{app}})}{\text{ucb}(\text{NACR}_{\text{pivot}})}, \min \left(\frac{\text{ucb}(\text{NACR}_{\text{app}})}{\text{ucb}(\text{NACR}_{\text{pivot}})}, 1 \right) \right] \quad (14)$$

⁸ In this setting, $p = \text{NACR}$ and $\hat{p} = \widehat{\text{NACR}}_{\text{MLE}}$.

⁹ This may happen if the same pivoting app has been running for long and a new, better performing app slightly overtakes it.

We make the following observations for Eq. 14. First, the ratio of non-accidental click rate will be greater than 1 only if the lower confidence bound of the app is greater than the upper confidence bound of the pivot. Second, in case of a large sample with nonzero valid clicks the ratio of non-accidental click rate will converge to the ratio of single-point estimate (MLE). Third, in case of a small sample size, the ratio of non-accidental click rate will be close to 1 indicating that we do not have enough data to suggest that the app of interest is any different from the pivoting app. Fourth, there is still a minimum number of clicks needed for Eq. 14 to produce reliable results.

When we have enough confidence that the ratio of non-accidental click rates between an app and the current pivot is greater than 1, we can update the pivot with that app and compute the discount factors using the new app as the new benchmark.

Next, we discuss through an example the impact on revenue of this smooth discounting strategy once implemented.

5.4 The impact of smooth discount factors

We compute the discount factors for accidental ad clicks on the two datasets \mathcal{D}_1 and \mathcal{D}_2 , described in Sect. 4.1. We consider all the ads impressed and clicked on all three apps, App 1, App 2, and App 3. To increase the confidence in our estimates, we discard ads with less than 40 clicks on *each* app. We select the pivoting app as the one with the highest NACR, estimated either via MLE or with the Agresti–Coull estimator. In both cases, App 1 is chosen.

Table 4 shows the discount factors computed using t_{median} as the dwell time thresholds of accidental clicks ($t_{\text{median}} = 2.1\text{ s}$ on \mathcal{D}_1 ; $t_{\text{median}} = 2.2\text{ s}$ on \mathcal{D}_2), and two different ratio of estimates of NACR, one obtained with MLE and the other using Agresti–Coull in combination with Eq. 14. Each row shows how much an advertiser should be charged for one accidental click on an ad shown on the app indicated by that row depending on the estimator used, providing that a valid (non-accidental) click on the same ad would cost 1. For example, looking at \mathcal{D}_2 if an ad click on App 2 originally costs 1\$ to the advertiser, any accidental ad click on the same app

Table 4 Discount factors computed using t_{median} , and MLE and Agresti–Coull estimates of NACR

		Discount factors	
		MLE	Agresti–Coull
\mathcal{D}_1	App 2	0.72	0.79
	App 3	0.64	0.73
\mathcal{D}_2	App 2	0.66	0.75
	App 3	NA (not enough obs.)	NA (not enough obs.)

will instead cost 0.75\$ after discounting using the Agresti–Coull estimate of NACR from Eq. 14.

We observe that the discount factors are comparable across the two datasets when computed using the same strategy. Moreover, larger discounting happens when generated from a single-point MLE of non-accidental click rate; the discounting is smaller when computed using the Agresti–Coull confidence interval. Depending on how aggressive the discounting has to be, one or the other approach may be chosen.

At the beginning of this section, we discussed how with dataset \mathcal{D}_3 the potential revenue loss that would result from fully discarding accidental clicks when billing advertisers was too high. Now using Eq. 12, the discount factor defined in Eq. 14, and using App 1 as pivoting app, the revenue drop is reduced by about 73.1%; allowing for advertisers to save money on likely less valuable clicks, while controlling for revenue impact for the ad network and publishers.

6 Use case II: filtering out accidental clicks when training an ad click model

Many ad networks improve the logic behind their ad serving algorithm through machine-learned models. At each ad request, these models provide a ranked list of ads to serve to maximise the overall *expected revenue*, eCPM.¹⁰ The eCPM is an estimate of the truly observed CPM, computed for each ad i as $\text{eCPM}_i = \text{cpc}_i \times \text{eCTR}_i$. Each cpc_i (or bid) is the price an advertiser is willing to pay for buying an impression, whereas eCTR_i is the estimate of the click-through rate of ad i (CTR_i). Estimating eCPM_i means estimating CTR_i for all ads i in the ad network inventory. Machine-learned models achieve exactly this task, i.e., they are trained on historical datasets of ad clicks to predict CTR from a feature vector representations of serviceable ads.

Training models on datasets containing a “large enough” ratio of accidental clicks may overestimate an ad CTR. This is because the estimated CTR becomes “inflated” with accidental clicks, eventually leading to the selection of irrelevant ads to serve. Filtering out accidental clicks when training machine-learned models may provide a more accurate selection of ads, resulting in higher revenue for the ad network and the publisher.

Using our data-driven approach to identify accidental clicks, we compute a threshold for a large number of Yahoo mobile apps (Yahoo News, Yahoo Mail, etc). We compute per-app thresholds of accidental click, as we have a sufficiently large number of dwell time observations for each app.

For each app, we then filter out clicks that are below the corresponding threshold for that app. The filtered clicks

¹⁰ CPM stands for *cost per mille (impressions)* and indicates the earnings gained every thousand ad impressions sold.

are not used to train the ad click model. We refer to this model as `accidental_click`. The model where no filtering of accidental clicks is our baseline mode, denoted as `baseline`. We setup an online A/B testing experiment, where a fraction of Yahoo Gemini incoming ad traffic is split between a *control* bucket and a *variation* bucket. More specifically, the A/B test affects about 1.8% of the overall ad serving traffic on the Yahoo apps considered. The traffic served by the control bucket is handled by the `baseline` ad click model, while variation bucket dispatches traffic to the `accidental_click` model. We thus compare the performance of the two models by measuring both CTR and CPM (click-through rate and revenue).

A significant lift is obtained using the `accidental_click` model of +3.9% compared to the `baseline` model (significance $\alpha = 0.01$ and p value $\ll 0.01$ using one-tailed two-proportion z-test). This means that an ad click model trained on a cleaner dataset (i.e., without accidental clicks) leads to a better estimation of ad CTR. This trained click model is better at predicting ads that are more likely to be clicked, i.e., more relevant to users, as it relies on ads that were clicked by users with an intent. Similarly, we observe a statistically significant lift in CPM of +0.2% (level of significance $\alpha = 0.01$ and p value $\ll 0.01$), which is partly due to the lift in CTR.

In this section, we showed that our data-driven approach for identifying accidental clicks is effective as a preprocessing step to training machine-learned ad click models, which ad networks leverage to estimate CTR to rank ads at serving time.

7 Related work

Various works investigated the role of dwell time on web pages. Liu et al. [18], who modelled dwell time on web pages using a Weibull distribution [20], found that web browsing exhibits a significant “negative ageing” phenomenon, suggesting that some initial screening has to be passed before a page is examined in detail by a user. They also demonstrated that dwell time distributions can be predicted purely using low-level web page features. We extend this work—focussing on ad (mobile) web pages—by proposing a model of dwell time based on a mixture of distributions instead of Weibull. This allows us to capture three categories of ad clicks, *accidental*, *short*, and *long*, where the focus of this paper is on accidental clicks.

Kim et al. [15] presented a method to explain dwell time on search engine result pages. They estimate dwell time distributions for SAT (satisfied) or DSAT (dissatisfied) clicks for different click segments and use them to derive features to train a click-level satisfaction model. Yi et al. [29] use item-level dwell time as a proxy to quantify how likely a content

item is relevant to a particular user in a recommender system. Furthermore, Yin et al. [30] show how to enrich the user-vote matrix by converting dwell time on items into users’ “pseudo votes” and then improve recommendation performance. All these works show that considering dwell time leads to improved decision-making tasks. In our work, our task is the identification of accidental clicks.

In the context of sponsored search e.g., [5,11,21,25,26] and display advertising e.g., [2,3,14,22], studies have mostly focused on predicting ad performance in terms of CTR. An ad with a high ad CTR is considered to perform well, since it indicates that the ad attracts users, who click on it. CTR, however, does not account for the post-click experience, that is how users experience the ad landing page. Dwell time on ad landing page has proven to be a good proxy of the post-click experience [4,16], reflecting the assumption that the longer the time user spends on the ad landing page the higher the chance he or she “converts” (e.g., by purchasing an item, registering to a mailing list, etc.), or simply the more likely is for the user to build an affinity with the brand [5,22]. These are cases of clicks that bring values to advertisers. Our work also looks at how dwell time relates to conversion rate, confirming what was shown in [10] that the former is a significant predictor of the latter, and hence a good proxy for measuring the “value” of a click. We use dwell time as our proxy of the ad click value, and proposed a data-driven methodology to identify accidental clicks, i.e., ad clicks with very short dwell time that are not only valueless to advertisers, but also to machine-learning models trained to predict ad CTR.

Other studies have investigated click “value” in the context of online advertising, mostly to contrast with fraudulent activities perpetrated by dishonest advertisers and/or web publishers like *click spam* [8,28]. To the best of our knowledge, this is the first work using dwell time to identify accidental clicks, in combination with a data-driven approach that can be applied to other domains, where it is important to quantify the value of a click.

8 Conclusions

In this paper, we propose a data-driven method to identify accidental clicks. An accidental click happens when a user that click on an ad, likely by mistake, is redirected to the ad landing page and bounce back without having seen the page. This type of clicks happens often on ads impressed on mobile devices.

We collect empirical dwell time observations from several Yahoo mobile apps for a large number of ads. We decompose the distribution of dwell time into a *mixture of components*, with each component corresponding to a click category: *accidental*, *short*, and *long*. Representative statistics for the first component of each ad are then further aggregated to provide

an overall estimate of the dwell time threshold of accidental clicks.

We assess the validity of our method when this is applied on two use cases. First, we describe a technique that estimates a *smooth* discounting factor, so that accidental clicks are not fully charged nor totally discarded at billing time. This allows for a trade-off between advertiser's satisfaction and potential revenue loss. Experiments conducted on different Yahoo mobile apps confirm that thresholds found are stable over time, and revenue loss can be mitigated by around 73.1% using our discounting strategy compared to not charging at all accidental clicks. Second, we demonstrate that an existing machine-learned ad click model used at serving time leads to better online performance if trained on datasets where accidental clicks are removed using our data-driven approach. We observe a positive and statistically significant lift on CTR and CPM (+3.9 and +0.2%, respectively) with the model trained without accidental clicks.

As future work, we plan to look at the two other components of the mixture models, so as to estimate thresholds—again in a data-driven manner—for short and long clicks. The former are clicks that suggest that the user had a negative post-click experience, whereas the latter is an indication of a positive post-click experience. Being able to do so per-app would remove the often ad hoc setting of dwell time thresholds. In addition to that, we would also like to investigate if the same (or similar) methodology proposed in this work could be used to assess the engagement of users with *any* section of a web page or a mobile app (i.e., not only with ads shown). Having a rigorous, data-driven methodology to classify content on the basis of the time users spent on it (i.e., dwell time) might be useful to providers, who could in turn make better decision on which of their assets they should invest more.

Acknowledgements The authors would like to thank Michal Aharon and Marc Bron for their support in setting up the online A/B test, which allowed them to deploy and assess their approach on a second use case, i.e., the ad click model.

Compliance with ethical standards

Conflict of interest The authors declared that they have no conflict of interest.

References

1. Akaike, H.: Information theory and an extension of the maximum likelihood principle. In: Petrov, B.N., Csaki, F. (eds.) *Selected Papers of Hirotugu Akaike*, pp. 199–213. Springer, Budapest, Hungary (1973)
2. Azimi, J., Zhang, R., Zhou, Y., Navalpakkam, V., Mao, J., Fern, X.: Visual appearance of display ads and its effect on click through rate. In: Chen, X-w, Lebanon, G., Wang, H., Zaki, M.J. (eds.) 21st ACM International Conference on Information and Knowledge Management, CIKM'12, pp. 495–504. Maui, HI, USA (2012)
3. Barajas, J., Akella, R., Holtan, M., Kwon, J., Flores, A., Andrei, V.: Dynamic effects of ad impressions on commercial actions in display advertising. In: Chen, X-w, Lebanon, G., Wang, H., Zaki, M.J. (eds.) 21st ACM International Conference on Information and Knowledge Management, pp. 1747–1751. Maui, HI, USA (2012)
4. Barbieri, N., Silvestri, F., Lalmas, M.: Improving post-click user engagement on native ads via survival analysis. In: Chen, X-w., Lebanon, G., Wang, H., Zaki, M.J. (eds.) 21st ACM International Conference on Information and Knowledge Management, CIKM'12, pp. 1747–1751. Maui, HI, USA (2012)
5. Becker, H., Broder, A., Gabrilovich, E., Josifovski, V., Pang, B.: What happens after an ad click?: Quantifying the impact of landing pages in web advertising. In: Cheung, D.W-L., Song, Il-Y., Chu, W.W., Hu, X., Lin, J.J. (eds.) *Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM 2009*, pp. 57–66. Hong Kong, China, 2–6 November 2009
6. Bishop, C.M.: *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, New York (2006)
7. Brown, L.D., Cai, T.T., DasGupta, A.: Interval estimation for a binomial proportion. *Stat. Sci.* **16**(2), 101–117 (2001)
8. Daswani, N., Stoppelman, M.: The anatomy of clickbot. A. In: Provos, N. (ed.) *First Workshop on Hot Topics in Understanding Botnets, HotBots'07*, p. 11. Cambridge, MA, USA, (2007)
9. Elkan, C.: Mixture Models. <http://cseweb.ucsd.edu/~elkan/250Bwinter2011/mixturemodels.pdf>. Accessed 4 Mar 2010
10. Goldman, M., Rao, J.M.: Experiments as instruments: heterogeneous position effects in sponsored search auctions. <https://ssrn.com/abstract=2524688>. Accessed 20 Nov 2014
11. Grbovic, M., Djuric, N., Radosavljevic, V., Silvestri, F., Baeza-Yates, R., Feng, A., Ordentlich, E., Yang, L., Owens, G.: Scalable semantic matching of queries to ads in sponsored search advertising. In: Perego, R., Sebastiani, F., Aslam, J.A., Ruthven, I., Zobel, J. (eds.) *Proceedings of the 39th International ACM, SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2016*, pp. 375–384. Pisa, Italy, 17–21 July 2016
12. Jacobson, A.: Preventing accidental clicks for a better mobile ads experience. <https://adwords.googleblog.com/2016/05/preventing-accidental-clicks-for-better-mobile-ads.html>. Accessed 6 May 2016
13. James, G., Witten, D., Hastie, T., Tibshirani, R.: *An Introduction to Statistical Learning: With Applications in R*. Springer, Berlin (2014)
14. Kae, A., Kan, K., Narayanan, V.K., Yankov, D.: Categorization of display ads using image and landing page features. *Proceedings of the Third Workshop on Large Scale Data Mining: Theory and Applications*, pp. 1:1–1:8. San Diego, California
15. Kim, Y., Hassan, A., White, R.W., Zitouni, I.: Modeling dwell time to predict click-level satisfaction. In: Carterette, B., Diaz, F., Castillo, C., Metzler, D. (eds.) *Seventh ACM International Conference on Web Search and Data Mining, WSDM 2014*, pp. 193–202. New York, NY, USA, 24–28 February 2014
16. Lalmas, M., Lehmann, J., Shaked, G., Silvestri, F., Tolomei, G.: Promoting positive post-click experience for in-stream yahoo gemini users. In: Cao, L., Zhang, C., Joachims, T., Webb, G.I., Margineantu, D.D., Williams, G. (eds.) *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1929–1938. Sydney, NSW, Australia, 10–13 August 2015
17. Lindsay, B.G.: *Mixture models: theory, geometry and applications*. In: *NSF-CBMS Conference Booktitle in Probability and Statistics*, Pennsylvania State University (1995)
18. Liu, C., White, R.W., Dumais, S.: Understanding web browsing behaviors through Weibull analysis of dwell time. In: Crestani,

- F., Marchand-Maillet, S., Chen, H.-H., Efthimiadis, E.N., Savoy, J. (eds.) *Proceeding of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2010*, pp. 379–386. Geneva, Switzerland, 19–23 July 2010
19. Mclachlan, G., Peel, D.: *Finite Mixture Models*, 1st edn. Wiley, Hoboken (2000)
 20. Papoulis, A., Pillai, S.U.: *Probability, Random Variables, and Stochastic Processes*, 4th edn. McGraw-Hill Higher Education, New York (2002)
 21. Raghavan, H., Hillard, D.: A relevance model based filter for improving ad quality. In: Allan, J., Aslam, J.A., Sanderson, M., Zhai, C.X., Zobel, J. (eds.) *Proceedings of the 32nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2009*, pp. 762–763. Boston, MA, USA, 19–23 July 2009
 22. Rosales, R., Cheng, H., Manavoglu, E.: Post-click conversion modeling and analysis for non-guaranteed delivery display advertising. In: Adar, E., Teevan, J., Agichtein, E., Maarek, Y. (eds.) *Proceedings of the Fifth International Conference on Web Search and Web Data Mining, WSDM 2012*, pp. 293–302. Seattle, WA, USA, 8–12 February 2012
 23. Schlattmann, P.: *Medical Applications of Finite Mixture Models. Statistics for Biology and Health*. Springer, Berlin (2009)
 24. Schwarz, G.: Estimating the dimension of a model. *Ann Stat* **6**, 461–464 (1978)
 25. Sculley, D., Malkin, R.G., Basu, S., Bayardo, R.J.: Predicting bounce rates in sponsored search advertisements. In: Elder IV, J.F., Fogelman-Soulié, F., Flach, P.A., Zaki, M.J. (eds.) *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1325–1334. Paris, France (2009)
 26. Sodomka, E., Lahaie, S., Hillard, D.: A predictive model for advertiser value-per-click in sponsored search. In: Schwabe, D., Almeida, V.A.F., Glaser, H., Baeza-Yates, R.A., Moon, S.B. (eds.) *22nd International World Wide Web Conference, WWW '13*, pp. 1179–1190. Rio de Janeiro, Brazil, 13–17 May 2013
 27. Stewart, C., Hoggan, E., Haverinen, L., Salamin, H., Jacucci, G.: An exploration of inadvertent variations in mobile pressure input. In: Churchill, E.F., Subramanian, S., Baudisch, P., O'Hara, K. (eds.) *Mobile HCI '12, Proceedings of the 14th international conference on Human-computer interaction with mobile devices and services*, pp. 35–38. San Francisco, CA, USA, 21–24 September 2012
 28. Stone-Gross, B., Stevens, R., Zarras, A., Kemmerer, R., Kruegel, C., Vigna, G.: Understanding fraudulent activities in online ad exchanges. In: Thiran, P., Willinger, W. (eds.) *Proceedings of the 11th ACM SIGCOMM Internet Measurement Conference, IMC '11*, pp. 279–294. Berlin, Germany (2011)
 29. Yi, X., Hong, L., Zhong, E., Liu, N.N., Rajan, S.: Beyond clicks: dwell time for personalization. In: Kobsa, A., Zhou, M.X., Ester, M., Koren, Y. (eds.) *Eighth ACM Conference on Recommender Systems, RecSys '14*, pp. 113–120. Foster City, Silicon Valley, CA, USA, 06–10 October 2014
 30. Yin, P., Luo, P., Lee, W.-C., Wang, M.: Silence is also evidence: interpreting dwell time for recommendation from psychological perspective. In: Dhillon, I.S., Koren, Y., Ghani, R., Senator, T.E., Bradley, P., Parekh, R., He, J., Grossman, R.L., Uthurusamy, R. (eds.) *The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013*, pp. 989–997. Chicago, IL, USA, 11–14 August 2013