IDIL: Exploiting Interdependence to Optimize Multi-Channel Advertising Campaigns

Alessandro Nuara, Nicola Sosio, Francesco Trovò, Maria Chiara Zaccardi, Nicola Gatti, Marcello Restelli
Politecnico di Milano, DEIB, Milano, Italy
{alessandro.nuara, francesco1.trovo, nicola.gatti, marcello.restelli}@polimi.it
{nicola.sosio, mariachiara.zaccardi}@mail.polimi.it

ABSTRACT

In 2017, Internet ad spending reached 209 billion USD worldwide—representing 41% of the global advertising market—, while, e.g., TV ads brought in 178 billion USD. An Internet advertising campaign includes up to thousands of sub-campaigns on multiple channels, e.g., search, social, display, whose bid and daily budget need to be optimized every day, subject to a budget constraint. Such a process is often unaffordable for humans and its automation can be crucial. As also shown by marketing funnel models, the sub-campaigns are usually interdependent, e.g., display ads induce awareness, increasing the number of impressions and conversions of search ads. This interdependence is widely exploited by humans in the optimization process, whereas, to the best of our knowledge, no algorithm takes it into account. In this paper, we provide the first model capturing the sub-campaigns interdependence, designed to guarantee a satisfactory trade-off between accuracy and amount of data used for the optimization phase. We also designed an algorithm, called IDIL, that, employing Granger Causality and Gaussian Processes, learns the model from past data, and returns an optimal stationary bid/daily budget allocation. We provided both theoretical guarantees on the loss of the algorithm w.r.t. the clairvoyant solution, and empirical evidence of the superiority of the proposed algorithm in both realistic and real-world settings w.r.t. previous approaches.

KEYWORDS

Internet Advertising; Granger Causality; Bid/Budget Optimization

1 INTRODUCTION

Since the early stages of the Internet, one of the most remunerative ways to economically exploit this novel media channel is *online advertising*. In 2017 alone, advertising revenues have totalled about 88 billion USD in the US [13] and about 209 billion USD worldwide. The choice of the ads to be displayed and their placement on a webpage are made through auctioning mechanisms [15]. An advertising *campaign* consists of a number (up to thousands) of *sub-campaigns* and a cumulative (per day or month) budget constraint. Each sub-campaign is characterized by an *ad*, a *targeting*, a *bid*, and a *daily budget*, and these last two parameters, crucial for the outcomes of the auctions, can be optimized every day using past performance. Such an optimization process is often unaffordable for humans and its complete, or partial, automation can lead to a significant improvement in revenues.

Nowadays, in addition to the traditional search advertising, other channels can be used, e.g., display and social. The diversification of ad channels is a crucial degree of freedom that one can exploit when setting up an advertising campaign. Indeed, different channels deeply affect each other's performance as Internet users regularly surf from one to another. For instance, Lewis and Nguyen [18] provide empirical evidence that display advertising increases the search activity on a product after the display ads visualization. Kireyev et al. [16] show a similar result between display and search advertising by using the Granger Causality test. The authors also show that this interdependence usually induces delayed dynamics, e.g., an increase in the display advertising impressions can lead to an increase in the conversions of search ads with a delay of some days. The sub-campaigns interdependence is customarily exploited by experts in the field, e.g., setting up sub-campaigns (called *assist*) not providing direct conversions but increasing the number of conversions on the search engine channel. Besides, capturing the interdependence can provide a direct method for comparing and optimizing the performance of sub-campaigns on different channels. Indeed, sub-campaigns on different channels need to be evaluated using different performance metrics, and how to combine them is still an open issue. For instance, display and social ads provide very few conversions compared to search ads but allow search ads to generate a larger number of conversions, and therefore an optimization method based only on the number of conversions might not provide optimal allocations. Although this problem is central in advertising, to the best of our knowledge, no model in the economic literature captures such interdependence.

Related work. Zhang et al. [35] propose an offline joint bid/daily budget optimization algorithm. In addition to neglecting interdependences between sub-campaigns, this work suffers from some weaknesses, like, among the most relevant, the authors assume a specific family of functions describing the relationship between the parameters and do not provide any theoretical bound on the estimation error. Some of these weaknesses have been addressed by Nuara et al. [21], who provide an online joint bid/daily budget optimization algorithm called AdComB-TS. However, this work overlooks the sub-campaigns interdependence.

Another related research field is the study of user behaviors from logging data [23, 27], both on social networks [24], and on search engines [33]. However, these approaches assume to keep

The full version of this work [20], comprehensive of the theorems proofs, will be presented at *The Web Conference* (WWW) in May 2019.

Games, Agents and Incentives Workshops, May 2019, Montreal, Canada

https://doi.org/

track of all the actions of each user perfectly, and this is usually unfeasible for most of the advertising tools used by companies. Less related works concern the optimization of the daily budget [14, 32], bidding strategies in display advertising [17, 29–31, 34], video advertising [10], and targeting [6].

Original contributions. We extend the work by Nuara et al. [21], designing an algorithm based on both learning and optimization techniques that can be adopted for the optimization of real-world Internet advertising campaigns and that, exploiting sub-campaigns interdependence, outperforms the algorithms known so far. To do that, we provide a novel model that, on one side, is expressive enough to capture the interdependences and, on the other side, is simple enough to require few data for the estimation of its unknown parameters. From of the proposed model, we design a data-driven algorithm, called Interdependency Detection, Identification, and Learning (IDIL), which consists of two phases: the *Interdependence Graph Learning Phase* and the *Estimation and Optimization Phase*. In the former phase, the IDIL algorithm learns the sub-campaigns interdependence structure (represented as a graph), identifying the pairs of sub-campaigns with the most significant interdependences by applying the Granger Causality test. In the latter phase, the IDIL algorithm computes the optimal joint bid/daily budget allocation exploiting Gaussian Process modeling [25] and an *ad hoc* dynamic programming procedure.

Finally, we show that neglecting the sub-campaigns interdependence can lead to massive losses even in simple and common scenarios and we theoretically bound the loss of our algorithm. Furthermore, we experimentally evaluate its performance in both realistic and real-world settings, showing the superiority of its performance compared to the previous approaches that neglect the sub-campaigns interdependence.

2 INTERNET ADVERTISING CAMPAIGN

Assume to have an Internet advertising campaign $C = \{C_1, \ldots, C_N\}$, with $N \in \mathbb{N}$, where C_j is the j-th sub-campaign. At day t, we are asked to set for each sub-campaign C_j a bid $x_{j,t} \in [\underline{x}_j, \overline{x}_j]$, and a daily budget $y_{j,t} \in [\underline{y}_j, \overline{y}_j]$, subject to that the daily cumulative budget of all the sub-campaigns cannot exceed $Y \in \mathbb{R}^+$. At day $t+1$, we have the performance of the campaign C at the previous day t, i.e., for every C_j, the tuple $(\widetilde{n}_{j,t}, \widetilde{cl}_{j,t}, \widetilde{co}_{j,t}, \widetilde{c}_{j,t})$, where $\widetilde{n}_{j,t}$ denotes the number of impressions, $\widetilde{cl}_{j,t}$ denotes the number of received clicks, $\widetilde{co}_{j,t}$ denotes the cumulate value of the conversions, and $\widetilde{c}_{j,t}$ denotes the amount of money spent for it. [1]

As aforementioned, both experts in the field of Internet advertising and studies in the Internet economic field, e.g., Kireyev et al. [16] and Hoban and Bucklin [12], demonstrate that impressions, clicks, and conversions of a sub-campaign might be influenced by the same kind of quantities of the other sub-campaigns. We extend the previous studies on the sub-campaigns interdependence, applying the Granger Causality test [11, 28] to two real-world Internet advertising campaigns optimized by an Italian web media agency using the AdComb-TS algorithm [21]. The algorithm, being online, produces policies explorative enough to make the test significant.

[1]We recall that the money spent in one day for a sub-campaign may be different from the daily budget previously allocated.

At first, we test for Granger causality the data collected for 8 months (from 1/1/2018 to 1/8/2018) from an Internet advertising campaign for a financial service of an insurance company: data correspond to $N = 12$ sub-campaigns, on Google AdWords (search), Facebook (social), and Google display with a cumulative budget of $Y = 600$ Euros. The results obtained from the Granger Causality test are shown in Figure 1, where the most significant elements of $(\widetilde{n}_{j,t}, \widetilde{cl}_{j,t}, \widetilde{co}_{j,t}, \widetilde{c}_{j,t})$ are represented as nodes of different colors according to their specific channel (as detailed in the caption of the figure) and the detected interdependences (with a p-value less than 5%) are represented as directed edges. In particular, Figure 1a shows the results when all the sub-campaigns data are aggregated, while Figure 1b focuses on a specific subset of sub-campaigns who share the same targeting (retired people). These results confirm the presence of interdependence between display and search advertising as previously observed in the literature. They also show that social and search advertising are interdependent and that the interdependences may be targeting specific. Moreover, the interdependence between clicks and impressions of the social channels and the impressions of the search one in this specific scenario seems to be more relevant than others, since they appear in both graphs. Furthermore, the Granger Causality test detects that interdependence dynamics between sub-campaigns are delayed up to 2 days.

At second, we test for Granger causality the data collected for 3 months (from 20/7/2018 to 20/10/2018) from an Internet advertising campaign of a different financial product of the same company with about $Y = 1100$ Euros. There are $N = 14$ sub-campaigns belonging to social and search advertising channels. The resulting graph is depicted in Figure 1c (with a p-value less than 5%). As in the previous dataset, many sub-campaigns are subject to interdependence. In particular, in this case, the interdependence phenomenon is only among impressions, suggesting that these can be the most significant in practice. Moreover, differently from the previous case, we distinguish search sub-campaigns into two subclasses which are at different depths in the marketing funnel: branding (orange nodes) or no-branding (red nodes). Finally, the delay of the interdependence dynamics is up to 3.

3 OPTIMIZATION PROBLEM

We provide our optimization problem capturing the sub-campaigns interdependence. For the sake of presentation, we focus on the interdependence between the impressions of different sub-campaigns.[2] Our goal is the maximization of the revenue earned each day from an Internet advertising campaign subject to a cumulative budget constraint. Formally, given a campaign C and a cumulative daily budget of Y, we aim to find, at day t, the value of bid $x_{j,t}$ and the value of daily budget $y_{j,t}$ for every sub-campaign C_j that maximise the revenue by solving the following optimization problem:

$$\max_{x_{j,t}, y_{j,t}} \sum_{j=1}^{N} v_j \; w_j \; n_j(x_{j,t}, y_{j,t}, u_{j,t}) \tag{1a}$$

[2]The use of such quantities is also supported by the experimental results of Section 2, where the interdependence among impressions is the most significant. However, different models, e.g., including the interdependence between the clicks and the conversions, are straightforward extensions of what is proposed in this section.

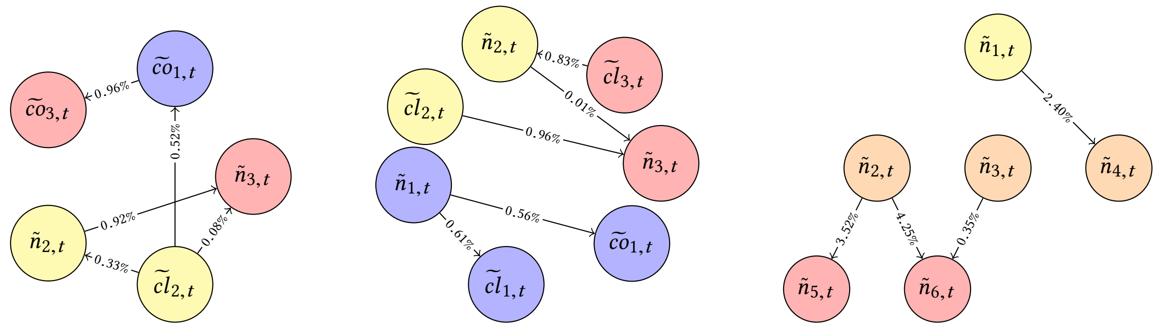

(a) First dataset, overall people **(b) First dataset, only retired people** **(c) Second dataset**

Figure 1: Graphs representing the interdependences of real-world Internet advertising sub-campaigns inferred by Granger Causality test from real data. The numbers on the edges are the p-values (in terms of %) of the Granger Causality test; display ads are depicted in blue, social ads in yellow, and search ads in orange (for branding sub-campaigns) and red (for other search sub-campaigns). The second dataset graph refers to interdependence among impressions $\tilde{n}_{j,t}$ of different sub-campaigns C_j.

$$\text{s.t.} \sum_{j=1}^{N} y_{j,t} \leq Y \tag{1b}$$

$$\underline{x}_j \leq x_{j,t} \leq \overline{x}_j \qquad \forall j \tag{1c}$$

$$\underline{y}_j \leq y_{j,t} \leq \overline{y}_j \qquad \forall j \tag{1d}$$

where $n_j(x_{j,t}, y_{j,t}, u_{j,t})$ is the expected number of impressions given bid $x_{j,t}$, daily budget $y_{j,t}$, and *influence index* $u_{j,t}$, representing the influence of other sub-campaigns towards C_j and computed by using the number of impressions of those sub-campaigns that are interdependent with sub-campaign C_j (see below); w_j and v_j are the click-trough rate and the value per click for the sub-campaign C_j, respectively, and, therefore, $v_j\ w_j\ n_j(x_{j,t}, y_{j,t}, u_{j,t})$ is the revenue provided by sub-campaign C_j.[3] We denote with $(\boldsymbol{x^*}, \boldsymbol{y^*}, \boldsymbol{u^*})$ the optimal solution to the optimization problem.

To model the sub-campaigns interdependence, we define, given an advertising campaign C, an *interdependence graph* $\mathcal{G} := (C, D)$ is a graph in which the adjacency matrix $D = \{d_{ij}\}, D \in \{0,1\}^{N \times N}$ has elements $d_{ij} = 1$ iff the sub-campaign C_i influences the performance of the sub-campaign C_j. We assume that the graph $\mathcal{G}$ is a Directed Acyclic Graph (DAG), i.e., there are no dependency cycles among the sub-campaigns. This assumption is supported by the model of the marketing funnel, in which the majority of the users flows from the top to the bottom, and different advertising channels are positioned at different levels of the funnel. Without loss of generality, we assume that the order over the indices of the sub-campaigns is one of the topological orders induced by $\mathcal{G}$. Given the interdependence graph $\mathcal{G}$, a formal definition of the influence index $u_{j,t}$ is:

$$u_{j,t} := \frac{1}{K} \sum_{i=1}^{j-1} \sum_{h=t-1}^{t-K} d_{ij}\, n_i(x_{i,h}, y_{i,h}, u_{i,h}), \tag{2}$$

where K is a maximum lag order, meaning that users are influenced by ads at most for K consecutive days. Notice that the first sub-campaign C_1, being influenced by no other sub-campaign, has $u_{1,t} = 0$ since the first summation in Equation (2) is over an empty set. The above definition of $u_{j,t}$ is based on the assumption that the increase in the number of impressions provided by a user coming from any sub-campaign influences the number of impressions of C_j in the same way. While this assumption might seem simplistic, it is necessary to keep at a pace the complexity of training the model. Indeed, a more complex model, e.g., where there is a different influence index for every pair of sub-campaigns, might be an option, but this would require an excessively large amount of data for the training of the model, which is not a viable option within the time horizon of the optimization process.

The optimization problem in Equations (1a)-(1d) can be solved using dynamic programming techniques, once all its parameters are known. However, the advertiser does not know the function $n_j(\cdot,\cdot,\cdot)$ that returns the number of impressions for sub-campaign C_j, as well as its click-trough rate w_j and its value per click v_j. Therefore, we resort to learning techniques to produce estimates of these parameters relying on historical data. We assume to have a dataset $Z := \{z_{j,t}\}$ of τ samples that provides, for each day $t \in \{1, \ldots, \tau\}$ and each sub-campaign C_j with $j \in \{1, \ldots, N\}$, the following values: $\boldsymbol{z}_{j,t} := (\tilde{x}_{j,t}, \tilde{y}_{j,t}, \tilde{n}_{j,t}, \widetilde{cl}_{j,t}, \widetilde{co}_{j,t}, \tilde{c}_{j,t})$. This is a tuple with the used bid $\tilde{x}_{j,t}$ and daily budget $\tilde{y}_{j,t}$, the received impressions $\tilde{n}_{j,t}$, clicks $\widetilde{cl}_{j,t}$, values of the conversions $\widetilde{co}_{j,t}$, and costs $\tilde{c}_{j,t}$. We require that the data collected up to day τ to be exploratory enough to properly model the sub-campaigns interdependences.

4 THE IDIL ALGORITHM

The pseudo-code of the IDIL algorithm is provided in Algorithm 1. It requires a dataset Z and two confidence levels $\alpha_{ADF} \in (0,1)$ and $\alpha_{GC} \in (0,1)$ in input. The first phase of the algorithm (Lines 1–8) is called *Interdependence Graph Learning Phase* and is devoted to learning the interdependence graph of the sub-campaigns. The output of this phase is an estimate $\hat{D}$ of the actual adjacency matrix D. The second phase of the algorithm (Lines 9–13) is called *Estimation and Optimization Phase* and is devoted to the estimation of the parameters for each sub-campaign C_j (i.e., $\hat{n}_j(\cdot,\cdot,\cdot), \hat{v}_j, \hat{w}_j$), using

[3] The optimization problem in Equations (1a)–(1d) reduces to the one by Nuara et al. [21] when there is no interdependence, i.e., if $n_j(x_{j,t}, y_{j,t}, u_{j,t}) = n_j(x_{j,t}, y_{j,t})$, for every C_j.

Gaussian Process [25] modeling, and solving the optimization problem in Equations (1a)-(1d), once the parameters have been replaced with their estimates. The outputs of this phase are $(\hat{\boldsymbol{x}}^*, \hat{\boldsymbol{y}}^*, \hat{\boldsymbol{u}}^*)$, i.e., the optimal bid, daily budget, and influence index for each sub-campaign.

4.1 Interdependence Graph Learning Phase

The task of learning $\hat{D}$ is obtained by resorting to the Granger Causality test [11]. This test has been used in many fields to infer the structure among datastreams, e.g., sensor networks by Alippi et al. [1] and by Roveri and Trovò [26], and economics by Calderón and Liu [2]. While in its original formulation the test assumes that the analysed time series are stationary, we rely on a generalization of this test, proposed by Toda and Yamamoto [28], which is suitable for integrated and cointegrated time series.

The basic idea of this approach is to estimate a Vector AutoRegressive model of order $K_{GR} + d_{\max}$ for the vector $(\widetilde{n}_{1,t}, \dots, \widetilde{n}_{N,t})$, where $d_{\max} \in \mathbb{N}$ is the maximum integration order of the time series that we analyse and $K_{GR} \in \mathbb{N}$ is a lag order which is estimated from the data.[4] The use of $K_{GR} + d_{\max}$ lags ensures that the test statistic used in the Granger Causality test for stationary time series has the same asymptotic distribution of the stationary case and, therefore, statistically valid conclusions can be drawn. More specifically, to test if the impressions of the campaign C_i influence the impressions of the campaign C_j, we estimate the parameters a_{jlm}, for each $m \in \{1, \dots, K_{GR} + d_{\max}\}$, of the model:

$$\widetilde{n}_{j,t} = \sum_{l=1}^{N} \sum_{m=1}^{K_{GR}+d_{\max}} a_{jlm}\, \widetilde{n}_{l,t-m} \quad \forall h \in \{1, \dots, N\}$$

and we test for the hypothesis:

$$H_0 : \forall m \in \{1, \dots, K_{Gr}\}\ a_{jim} = 0,$$
$$H_1 : \exists m \in \{1, \dots, K_{Gr}\} \mid a_{jim} \neq 0.$$

The complete description of this procedure is provided by Toda and Yamamoto [28]. The test states that if we reject H_0 there is evidence, with confidence α_{GC}, that the impressions from C_i are influencing those of C_j.

The IDIL algorithm works as follows. For each sub-campaign, we estimate $d_{\max}$ performing the Augmented Dickey Fuller test $\mathsf{ADF}(\tilde{\boldsymbol{n}}_j, \alpha_{ADF})$ on the time series $\tilde{\boldsymbol{n}}_j := (\widetilde{n}_{j,t}, \dots, \widetilde{n}_{j,\tau})$ with confidence α_{ADF} (Lines 1–3), and inferring the time series order adf_j, and, finally, we perform the Granger Causality test on each pair of sub-campaigns (Lines 5–7). The result of this procedure is a matrix $\hat{P}$ containing the p-values of the pairwise tests, which is used to generate a valid estimate of the adjacency matrix $\hat{D} \in \{0, 1\}^{N \times N}$. This operation is performed by $\mathsf{DAG}(\hat{P}, \alpha_{GC})$ (Line 8) by selecting the largest subset S of the p-values $\hat{p}_{ij} < \frac{2\alpha_{GC}}{N(N-1)}$ s.t. the matrix $\hat{D} := \{d_{ij} = 1 \text{ iff } p_{ij} \in S\}$ to correspond to a DAG.[5] This procedure ensures an overall confidence α_{GC} on the Granger Causality test, thanks to the Bonferroni correction for multiple tests, and it avoids that the presence of false positives in the detection of interdependences. Indeed, the edges generated by false positive detections might provide adjacency matrices $\hat{D}$ whose corresponding graph presents cycles, which would compromise the execution of the following optimization procedure.

Algorithm 1 IDIL

Input: dataset Z, confidence α_{ADF}, confidence α_{GC}
Output: optimal bid/budget/new user allocation $(\hat{\boldsymbol{x}}^*, \hat{\boldsymbol{y}}^*, \hat{\boldsymbol{u}}^*)$
▷ *Interdependence Graph Learning Phase*
1: **for** $j \in \{1, \dots, N\}$ **do**
2: $\quad adf_j \leftarrow \mathsf{ADF}(\tilde{\boldsymbol{n}}_j, \alpha_{ADF})$
3: $d_{\max} \leftarrow \max_j \{adf_j\}$
4: $\hat{P} \leftarrow 0$
5: **for** $j \in \{1, \dots, N\}$ **do**
6: $\quad$ **for** $i \in \{j+1, \dots, N\}$ **do**
7: $\quad\quad \hat{p}_{i,j} \leftarrow \mathsf{GCT}(\boldsymbol{n}, i, j)$
8: $\hat{D} \leftarrow \mathsf{DAG}(\hat{P}, \alpha_{GC})$
▷ *Estimation and Optimization Phase*
9: **for** $j \in \{1, \dots, N\}$ **do**
10: $\quad \hat{n}_j(\cdot, \cdot, \cdot) \leftarrow \mathsf{GP}(Z, \hat{D}, j)$
11: $\quad \hat{v}_j \leftarrow \frac{1}{\tau} \sum_{h=1}^{\tau} \frac{\tilde{co}_{j,t}}{\tilde{cl}_{j,t}}$
12: $\quad \hat{w}_j \leftarrow \frac{1}{\tau} \sum_{h=1}^{\tau} \frac{\tilde{cl}_{j,t}}{\tilde{n}_{j,t}}$
13: $(\hat{\boldsymbol{x}}^*, \hat{\boldsymbol{y}}^*, \hat{\boldsymbol{u}}^*) \leftarrow \mathsf{OPT}(\hat{\boldsymbol{n}}, \hat{\boldsymbol{v}}, \hat{\boldsymbol{w}}, \hat{D})$
14: **return** $(\hat{\boldsymbol{x}}^*, \hat{\boldsymbol{y}}^*, \hat{\boldsymbol{u}}^*)$

4.2 Estimation and Optimization Phase

The second phase of the IDIL algorithm exploits predictive models to estimate unknown functions and quantities in the optimization problem defined in Equations (1a)-(1d), and solves it in a dynamic programming fashion with an *ad hoc* procedure.[6]

We use Gaussian Processes (GPs) to compute, for each sub-campaign C_j, the function $\hat{n}_j(x, y, u)$ estimating the expected number of impressions $n_j(x, y, u)$, given the chosen bid x, the allocated budget y, and the influence index u generated by the sub-campaigns influencing the sub-campaign C_j (Line 10). The estimate $\hat{w}_j$ of the click-through rate w_j and the estimate $\hat{v}_j$ of the value per click v_j are the average ratios between the number of clicks and the number of impressions and between the number of conversions and the number of clicks, respectively (Lines 9-12). Finally, the estimated influence index is computed as follows:

$$\hat{u}_{j,t} := \frac{1}{K_{GR}} \sum_{i=1}^{j-1} \sum_{h=t-1}^{t-K_{GR}} \hat{d}_{ij}\, \hat{n}_i(x_{i,h}, y_{i,h}, \hat{u}_{i,h}), \tag{3}$$

where we use K_{GR}, obtained from the Granger Causality test, as an estimate of the actual lag K.

The optimization procedure is an extension of the optimization algorithm by Nuara et al. [21], to handle also campaigns in which the revenue given by a budget allocated to a sub-campaign depends on the budget allocated to other sub-campaigns. The OPT algorithm, presented in Algorithm 2, takes in input the estimates of the

[4] d_{max} can be estimated using the Augmented Dickey Fuller test [3], which requires a confidence level $\alpha_{ADF} \in (0, 1)$, while K_{GR} can be estimated from the dataset Z by standard techniques, see Ozcicek and Douglas Mcmillin [22] for details.

[5] An adjacency matrix $\hat{D}$ identifies a DAG if and only if a depth-first search of the corresponding graph yields no back edges.

[6] For the sake of presentation in what follows we assume that the number of impressions is monotonically increasing in the influence index. A version of the optimization procedure able to handle general cases is discussed in the final part of this section.

Algorithm 2 OPT($\hat{\boldsymbol{n}}, \hat{\boldsymbol{v}}, \hat{\boldsymbol{w}}, \hat{D}$)

Input: estimated adjacency matrix $\hat{D}$, estimated models $\hat{n}_j(\cdot,\cdot,\cdot)$, $\hat{w}_j$, $\hat{v}_j$, discretization of the total budget $\{b_1, \dots, b_B\}$
Output: optimal bid/budget allocation $(\hat{\boldsymbol{x}}^*, \hat{\boldsymbol{y}}^*, \hat{\boldsymbol{u}}^*)$

1: **for** $i \in \{1, \dots, B\}$ **do**
2: $\Pi_{1,i,1} \leftarrow (b_i\ \mathbf{0}_{N-1})$
3: $L_{1,i,1} \leftarrow \hat{v}_1\ \hat{w}_1\ \hat{n}_1(\chi_1, b_i, 0)$
4: $M_{1,i,1} \leftarrow \hat{n}_1(\chi_1, b_i, 0)\ \hat{\boldsymbol{d}}_1$
5: **for** $j \in \{2, \dots, N\}$ **do**
6: **for** $i \in \{1, \dots, B\}$ **do**
7: $c \leftarrow 1$
8: **for** $k \in \{1, \dots, i\}$ **do**
9: $m = |\{\Pi_{j-1,k,h}\}_h|$
10: **for** $h \in \{1, \dots, m\}$ **do**
11: $l \leftarrow \hat{n}_j(\chi_j, b_i - b_k, M_{j-1,k,h}(j))$
12: $\bar{\Pi}_c \leftarrow (\mathbf{0}_{j-1}\ (b_i - b_k)\ \mathbf{0}_{N-j}) + \Pi_{j-1,k,h}$
13: $\bar{L}_c \leftarrow \hat{v}_j\ \hat{w}_j\ l + L_{j-1,k,h}$
14: $\bar{M}_c \leftarrow\ l\ \hat{\boldsymbol{d}}_j + M_{j-1,k,h}$
15: $c \leftarrow c + 1$
16: $c \leftarrow 1$
17: **for** $h \in \{1, \dots, |\{Lt_c\}_c|\}$ **do**
18: **if** $\nexists k \mid \bar{L}_h < \bar{L}_k \wedge \forall p \in \{j+1, \dots, N\} | \bar{M}_h(p) < \bar{M}_k(p)$ **then**
19: $\Pi_{j,i,c} \leftarrow \bar{\Pi}_h$
20: $L_{j,i,c} \leftarrow \bar{L}_h$
21: $M_{j,i,c} \leftarrow \bar{M}_h$
22: $c \leftarrow c + 1$
23: **for** $j \in \{1, \dots, N\}$ **do**
24: $\hat{y}^*_j \leftarrow \max_i \Pi_{N,i,1}(j)$
25: $\hat{u}^*_j$ computed as in Equation (3)
26: $\hat{x}^*_j = x^*_j(\hat{y}^*_j, \hat{u}^*_j)$
27: **return** $(\hat{\boldsymbol{x}}^*, \hat{\boldsymbol{y}}^*, \hat{\boldsymbol{u}}^*)$

adjacency matrix $\hat{D}$, the number of impressions function $\hat{n}_j(\cdot,\cdot,\cdot)$, the click-trough rate $\hat{w}_j$, the value per click $\hat{v}_j$, and a set of available daily budget values $\{b_1, \dots, b_B\}$, which are, for simplicity, evenly spaced in the range $[0, Y]$.

The OPT algorithm uses three structures Π, L, and M defined as follows: $\Pi_{j,i,h}$ is a vector that specifies a partial budget allocation with cumulative budget of b_i among the sub-campaigns $C_1, \dots, C_j$; $L_{j,i,h}$ is the revenue provided by the partial budget allocation $\Pi_{j,i,h}$; $M_{j,i,h}$ is a vector that specifies the value of the influence index of the sub-campaigns $C_{j+1}, \dots, C_N$ provided by the sub-campaigns $C_1, \dots, C_j$ when the partial allocation $\Pi_{j,i,h}$ is used. The third index h in the structures mentioned above is necessary since the algorithm may need to store multiple partial budget allocations for each j and i. More precisely, the set $\{\Pi_{j,i,h}\}_h$ contains Pareto-efficient partial budget allocations, where the optimality criteria are the revenue and the influence indices of campaigns $C_{j+1}, \dots, C_N$. For instance, given two partial budget allocations Π_{j,i,h_1} and Π_{j,i,h_2}, where the former has high revenue and a small number of impressions and the latter *vice versa*, it is not possible to decide which one is the optimal before evaluating their influence on the sub-campaigns $C_{j+1}, \dots, C_N$ and therefore we need to store both.

At first, the algorithm initializes the values of the structures for $j = 1$ (Lines 1–4), corresponding to the allocations of the partial budget to the sub-campaign C_1. For each budget b_i, we allocate it to C_1, formally, $\Pi_{1,i,1} = (b_i, \mathbf{0}_{N-1})$, where $\mathbf{0}_{N-1}$ denotes a null vector of size $N-1$. The sub-campaign C_1, being the first in the topological ordering induced by $\hat{D}$, is not subject to any interdependence from other sub-campaigns. Therefore, the computation of the revenue $\{L_{1,i,1}\}_i$ and the influence index vector $\{M_{1,i,1}\}_i$ is performed using the previously estimated models.[7] The vector $M_{1,i,1}$ is computed as $M_{1,i,1} = n_1(\chi_1, b_i, 0)\ \hat{\boldsymbol{d}}_1$, where $\hat{\boldsymbol{d}}_i$ is the i-th row of the adjacency matrix $\hat{D}$. This means that $M_{1,i,1}(j)$, i.e., the j-th element of $M_{1,i,1}$, is equal to $n_1(\chi_1, b_i, 0)$ if the sub-campaign C_1 influences the campaign C_j and zero otherwise.

For all the $j \in \{2, \dots, N\}$, the algorithm computes the elements of the three structures Π, L, and M using the values previously computed at the $j-1$-th step, in a dynamic programming fashion (Lines 5–22). For each daily budget b_i and for each daily budget $b_k \leq b_i$, we compute the revenue and the influence index provided by the allocation of a daily budget of $b_i - b_k$ to the sub-campaign C_j and the remaining daily budget of b_k to the sub-campaigns $C_1, \dots, C_{j-1}$. We do this by enumerating all the Pareto-efficient partial allocations $\Pi_{j-1,k,1}, \Pi_{j-1,k,2}, \dots$ of the first $j-1$ sub-campaigns, then allocating daily budget $b_i - b_k$ to the sub-campaign C_j and, finally, we evaluate the total revenue $\bar{M}_c$ and the influence indices vector $\bar{L}_c$ provided by the partial allocations obtained, denoted with $\bar{\Pi}_c$ (Lines 9–15).[8] After that, the algorithm discards all the candidate partial allocations which are Pareto dominated (Lines 16-22); see Ehrgott [4] for details on Pareto efficiency and dominance.[9] Finally, the algorithm returns the optimal allocation (Lines 23–27): the optimal budgets $\hat{y}^*_j$ are the elements of $\max_i \Pi_{N,i,1}(j)$; the optimal influence indices $\hat{u}^*_j$ are computed using Equation (3); the optimal bids $\hat{y}^*_j$ are computed using the impressions models $\hat{n}_j(\cdot,\cdot,\cdot)$.

The complexity of the OPT algorithm is $O\left(\sum_j B^{\sum_i \hat{d}_{ij}+2}\right) \leq O\left(N\ B^2\ B^{\max_j \sum_i \hat{d}_{ij}}\right)$ and strictly depends on the maximum indegree of the interdependence graph corresponding to $\hat{D}$. The complexity reduces to that one of the algorithm proposed by Nuara et al. [21] when the sub-campaigns are not interdependent. Notice that capturing only the pairs of sub-campaigns with the most significant interdependence is a crucial issue from a computational point of view since, taking into account all the possible pairs of sub-campaigns, the complexity is bounded by $NB^2 \frac{B^{N+1}-1}{B-1}$, which is intractable when N is large as it happens in real-world applications.

5 THEORETICAL PROPERTIES OF IDIL

We analyse the properties of our problem and those of the IDIL algorithm. Initially, we analyse the suboptimality of any algorithm ignoring the sub-campaigns interdependencies w.r.t. our algorithm, i.e., when the learner uses an adjacency matrix $\hat{D} = 0$, and the real one D is non-null. The following theorem shows that ignoring the sub-campaigns interdependences might be arbitrarily suboptimal.

[7]We define χ_j as the bid that maximise the number of impressions given a budget y and a influence index u, or, formally, $\chi_j := \chi_j(y, u) = \arg\max_x \hat{n}_j(x, y, u)$.
[8]In the pseudo-code, we denoted the number of Pareto optimal allocations at the $j-1$-th row with a budget of b_i with $|\{\Pi_{j-1,i,h}\}_h|$.
[9]Notice that the inequality in Line 18 is designed for settings in which the number of impressions is monotonically increasing in the influence index. However, removing the condition in Line 18, the proposed method also applies to problems without such a monotonicity assumption. This comes at at the cost of storing a larger number of partial allocations in $\{\Pi_{j,i,h}\}_h$.

Theorem 1 (Nuara et al. [20]). *Given the problem of optimizing an advertising campaign C, employing a model $\hat{n}_j(x, y)$ for the number of impressions that ignores the sub-campaigns interdependence may result in an arbitrary large loss in terms of* revenue*, defined as:*

$$R_t = \sum_{j=1}^{N} v_j \, w_j \, n_j(x_{j,t}, y_{j,t}, u_{j,t}). \tag{4}$$

When the model is flexible enough to model the actual process properly, we can bound its error, formally, defined as follows:

Definition 5.1. Given a dataset Z, the *total (estimation) error* is:

$$E_\tau := \sum_{j=1}^{N} \left[v_j w_j n_j(x_j^*, y_j^*, u_j^*) - \hat{v}_j \hat{w}_j \hat{n}_j(\hat{x}_j^*, \hat{y}_j^*, \hat{u}_j^*) \right],$$

where the tuples $(\hat{x}_j^*, \hat{y}_j^*, \hat{u}_j^*)$ are elements of the (stationary) output $(\hat{\boldsymbol{x}}^*, \hat{\boldsymbol{y}}^*, \hat{\boldsymbol{u}}^*)$ of the IDIL algorithm using the estimates of the parameters, and (x_j^*, y_j^*, u_j^*) are elements of the (stationary) output of the IDIL algorithm using the real parameters.

We can show the following:

Theorem 2 (Nuara et al. [20]). *When the expected number of impressions $n_j(\cdot, \cdot, \cdot)$ of every sub-campaign C_j is distributed as a Gaussian Process, the total error between the real revenue and the estimated one using the output of the IDIL algorithm is upper bounded, with a probability of at least $1 - \delta$, as follows:*

$$E_\tau \leq 2Nv^{(\max)} \sqrt{\frac{1}{2\tau} \log \frac{6N}{\delta}} \left(\hat{n}^{(\max)} + \hat{\sigma}_\tau^{(\max)} \sqrt{2 \log \frac{3N}{\delta}} \right)$$
$$+ Nv^{(\max)} \, \hat{\sigma}_\tau^{(\max)} \sqrt{2 \log \frac{3N}{2\delta}},$$

where $\hat{n}^{(\max)} := \max_j \max_{(x,y,u)} \hat{n}_j(x, y, u)$ is the maximum number of estimated expected impressions over all the sub-campaigns, $\hat{\sigma}_\tau^{(\max)} := \max_j \max_{(x,y,u)} \hat{\sigma}_{j,\tau}(x, y, u)$ is the maximum estimated standard deviation, and $v^{(\max)}$ is the maximum value per click.

We remark that Rasmussen and Williams [25] show that, in a generic GP, $\hat{\sigma}_\tau^{(\max)} \to 0$ as $\tau \to \infty$. Therefore, the total error E_τ decreases as the number of samples τ in the training set Z increases.

Our analysis has, so far, focused on the static properties of our problem. However, the scenario we are studying is a dynamical system due to the potentially delayed effects induced by the sub-campaigns interdependence. Therefore, it is crucial to show that, whenever a stationary allocation is used, the dynamics always reach a steady state in finite time and how their length is upper bounded. In this context, a steady state allocation provides a constant number of impressions for each sub-campaign for at least K consecutive days. We can show the following:

Theorem 3 (Nuara et al. [20]). *Using the stationary allocation $(\hat{\boldsymbol{x}}^*, \hat{\boldsymbol{y}}^*, \hat{\boldsymbol{u}}^*)$ we reach a steady state after at most $K\,\Gamma + 1$ days, where K is the maximum lag of the influence index $u_{j,t}$ and Γ is the length of the longest path of the graph $\mathcal{G}$.*

The above theorem states that the more complex the process (e.g., presenting a cascade of interdependences), the more we have to wait to completely remove the effects of a suboptimal allocation.

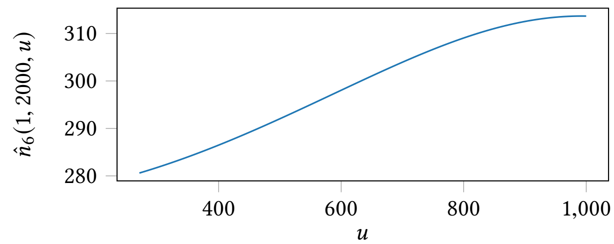

Figure 2: GPs estimation of the number of impressions $\hat{n}_6(1, 2000, u)$ depending on the influence index u.

6 EXPERIMENTAL EVALUATION

We experimentally evaluate the IDIL algorithm in a real-world setting and in a synthetic setting, generated by using a realistic simulator. We compare the revenue R_t produced by IDIL and AdComB-Mean (an off-line version of the algorithm proposed by Nuara et al. [21] neglecting any sub-campaign interdependence).

6.1 Real-world Setting

In this experiment, we rely on the data of the second campaign described in Section 2 to train our model. We recall that the length of the dataset is $\tau = 93$ days (from 20/7/2018 to 20/10/2018), the advertising campaign is composed of $N = 14$ sub-campaigns belonging to both social and search advertising channels. The corresponding estimated interdependence graph is provided in Figure 1c. From 21/10/2018 to 4/11/2018 (15 days), the campaign optimization has been performed by the IDIL algorithm.

When comparing the policies produced by IDIL with those produced by AdComB-Mean (the off-line version of AdComB-TS), the former policies appear more suitable then the latter ones, as a more significant portion of the budget is allocated to social sub-campaigns and branding search sub-campaigns. The interdependence suggested by the Granger Causality Test are confirmed by estimations provided by the GPs. Indeed, in Figure 2, we show the expected value of the prediction provided by GPs of the number of impressions for the sub-campaign C_6 with a bid value of $x = 1$ (i.e., one of the most frequent choice during the training set) and $y = 2000$ (i.e., a budget large enough to capture all the available user for this sub-campaign). The number of impressions increases as the value of the influence index increases, suggesting that a positive correlation between C_2 and C_3 impressions, and C_6 ones exist. However, since in a real setting we cannot exclude the presence of negative interdependence, to compute the optimal allocation with the IDIL algorithm, we remove the condition in Line 18 of Algorithm 2, to be able to provide the optimal allocation even if generic interdependence among sub-campaigns are present. In Figure 3, we show the expected revenue given by optimal policies computed by AdComB-Mean and IDIL for different values of the total budget Y. In this scenario, the exploitation of the sub-campaigns interdependence can lead to a potential revenue increase up to 13%.

In the 15 days of campaign optimization performed by the IDIL algorithm, the number of daily conversions increased by 11% w.r.t. the average of the previous 30 days (the result is compatible with our prediction, given that AdComB-TS/Mean provide very close performance). Although this is a promising result, there is no statistical significance that the IDIL algorithm outperforms in practice AdComB-TS/Mean. Due to the impossibility to directly compare

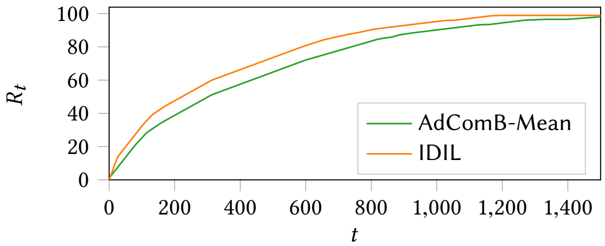

Figure 3: Comparison of the expected revenue R_t given by the AdComB-Mean and IDIL algorithms.

the performance of the two algorithms online (e.g., by using an A/B testing system), we resort to a realistic synthetic environment.

6.2 Synthetic Settings

We evaluate the performance of the IDIL algorithm in two synthetic settings, generated by a realistic simulator, comparing the revenue R_t produced by the following algorithms: IDIL, DA-IDIL (Dependency Aware-IDIL), a variation of the IDIL algorithm *a priori* knowing the dependency matrix D, and AdComB-Mean.

Synthetic Data Generation. The synthetic settings are generated as follows. At day t, each sub-campaign C_j is characterized by the set of the users $S_{j,t} = s_{j,t} \cup \left(\bigcup_{i \neq j} s_{ij,t} \right)$ that could potentially visualize the ad of the sub-campaign C_j. More precisely, we distinguish the set of the users $s_{j,t}$, that would visualize the ad of C_j without having previously visualized the ads of the other interdependent sub-campaigns, from the set of the users $s_{ij,t}$, that would visualize the ad of C_j only after having visualized the ad of C_i. Notice that $s_{ij,t}$ is non-empty only if the sub-campaigns C_i and C_j are interdependent and, more precisely, if $d_{ij} \neq 0$.

The number of users $|s_{j,t}|$ is sampled from $\mathcal{N}(\mu_j, \sigma_j^2)$, i.e., a Gaussian distribution with mean μ_j and variance σ_j^2. Each user in $s_{j,t}$ is characterized by a click probability $p_j^{(cl)}$ and a conversion probability $p_j^{(co)}$ specific for the sub-campaign C_j. Conversely, the number of users $|s_{ij,t}|$ is modeled trough a linear combination of the number of daily impressions $n_{i,t-1}, \ldots, n_{i,t-K}$ (whose generation is described in what follows), where K represents the maximum delay in the interdependence dynamics. Formally, we have that $s_{ij,t} := p_{ij}^{(res)} \sum_{k=1}^{K} \beta_k \, n_{i,t-k}$, where $\beta_k \in [0,1]$ are randomly sampled coefficients and $p_{ij}^{(res)}$ is the probability that a user having visualized ad of C_i is a potential user that may visualize C_j. Each user in $s_{ij,t}$ is characterized by a click probability $p_{ij}^{(cl)}$ and a conversion probability $p_{ij}^{(co)}$.

At each day t, setting the bid/budget pairs on each sub-campaign allows the advertiser to take part to $A_j \leq |S_{j,t}|$ auctions based on the Vickrey-Clarke-Groves mechanism [5, 7–9, 19], in which γ_j available ad slots are allocated to a subset of δ_j advertisers ($\gamma_j \leq \delta_j$). More specifically, each advertiser submits her bid b_h and those with the first γ_j highest values $b_h \, \rho_h$ are allocated in the γ_j slots, where ρ_h is the probability that h-th ad is clicked given it has been observed. The bids b_h of the other ads participating in the auctions are drawn from a truncated Normal distribution $\mathcal{N}(\mu^{(b)}, \sigma^{(b)})$, and the click probabilities ρ_h are uniformly sampled in $[0,1]$. In the case the advertiser wins the m-th auction, the ad gets an impression ($n_{m,j,t} = 1$, otherwise $n_{m,j,t} = 0$). The ad is allocated in a the l-th slot, the ad can be visualized by an user $S_{j,t}$ according to the probability of being observed $p^{(obs)}(l)$. After the impression, the user can click on the ad and generate a conversion according to the click $p_j^{(cl)}$ and conversion $p_j^{(co)}$ probabilities if the user belongs to $s_{j,t}$, and according to the click $p_{ij}^{(cl)}$ and conversion $p_{ij}^{(co)}$ probabilities if the user belongs to $s_{ij,t}$. A click on the ad of C_j provided by the user corresponding to the m-th auction is denoted by $cl_{m,j,t} = 1$ ($cl_{m,j,t} = 0$ otherwise), and imposes a payment of $CPC_{m,j,t}$, as specified by the VCG auction (see [19] for details). The auctions are generated until the daily budget $y_{j,t}$ allocated on the sub-campaign C_j is totally spent, i.e., the total number of auctions A_j is s.t. $\sum_{m=1}^{A_j} CPC_{m,j,t} = y_{j,t}$ or until $A_j = |S_{j,t}|$. Finally, in the case a click happen, the m-th user may convert ($co_{m,j,t} = 1$) or not ($co_{m,j,t} = 0$). The daily impressions, the daily clicks, the daily conversions (assuming unitary value per conversion), and the daily costs are computed as $n_{j,t} = \sum_{m=1}^{A_j} n_{m,j,t}$, $cl_{j,t} = \sum_{m=1}^{A_j} cl_{m,j,t}$, $co_{j,t} = \sum_{m=1}^{A_j} co_{m,j,t}$, respectively. We refer to [20] for the values of the main parameters used in the two synthetic settings in which we test our algorithm.

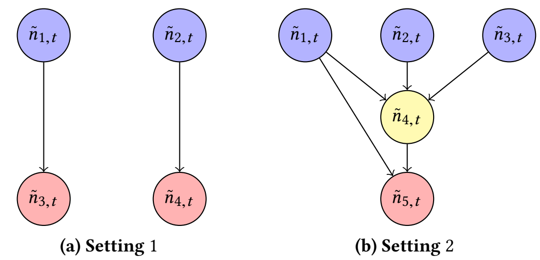

Figure 4: Interdependence graph $\mathcal{G}$ for the two synthetic experimental settings.

Synthetic Setting 1. There are $N = 4$ sub-campaigns, with delayed dynamics of $K = 5$ days, whose interdependence graph is shown in Figure 4a. The longest path of the interdependence graph $\mathcal{G}$ is $\Gamma = 1$. C_1 and C_2 are on the display advertising channel and are targeted to a wide range of daily users, thus generating a large number of daily auctions, but their conversion probability is low. C_3 and C_4 are on the search advertising channel, generating a small number of daily auctions, but their conversion probability is high.

We use $Y = 500$ and $B = 10$ daily budget values evenly spaced in the range $[0, 500]$. The GPs used to estimate the impressions model of the sub-campaigns adopt a squared exponential kernel in which the kernel parameters are chosen as recommended by Rasmussen and Williams [25]. We evaluate the performance of the algorithms with different numbers of samples $\tau \in \{60, 80, 100\}$ in the training set Z. In the first τ days, a uniformly random allocation is used to collect data and, after that, the algorithms compute the optimal solution based on their estimates and then set it.

Results. In Figure 5a, we report the average (over 100 repetitions) revenue R_t produced by the algorithms with a training of $\tau = 100$ samples. From $t = 100$ on, the optimal stationary solution is used.

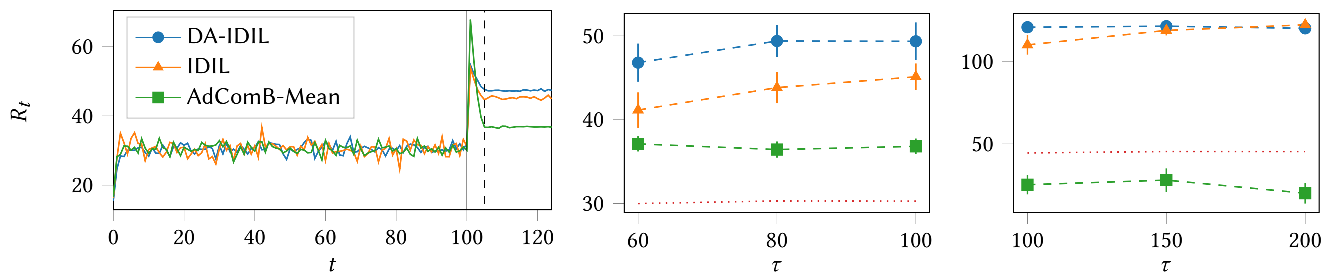

(a) (b) (c)

Figure 5: Results for the Settings 1 and 2. (a) Revenue R_t over time for the Setting 1. (b) Revenue R_t in steady state conditions for different training sizes τ in Setting 1. (c) Revenue R_t in steady state conditions for different training sizes τ in Setting 2. In (b) and (c), the revenue of the random allocation is reported with a dotted magenta line and the vertical lines represent the 95% confidence intervals for the algorithms revenue.

The average revenue of each algorithm peaks at $t = 101$ and, for $t > 101$, decreases by converging to a steady state within $K\Gamma+1 = 6$ days. The peak is generated by the presence of a large number of residual users who have observed display ads during training and who, after $t = 100$, observe search ads. These residual users decrease for $t > 101$ until they reach a steady state. Thus, (temporary) peaks may be achieved with non-stationary policies.

The DA-IDIL algorithm exhibits the best performance, exploiting the *a priori* knowledge of the adjacency graph D. The gap between the revenue produced by the IDIL and DA-IDIL algorithms, due to the estimation error introduced on $\hat{D}$, is sufficiently small, showing that the Granger Causality test used by the IDIL algorithm works well in practice. Instead, the revenue produced by the AdComB-Mean algorithm, neglecting the interdependence among sub-campaigns, is much smaller than that produced by the other two algorithms. This is due to the very different budget allocations chosen by the three algorithms: the IDIL and DA-IDIL algorithms optimally balance the budget on all the sub-campaigns, while the AdComB-Mean algorithm greedily invests the budget only in the search sub-campaigns C_3 and C_4. Interestingly, the performance of the AdComB-Mean algorithm is quite similar to that of the uniformly random allocation used during training.

In Figure 5b, we report the average revenue R_t at the steady-state (averaged over the 100 independent repetitions and over $t \in \{106, \dots, 120\}$) and the 95% confidence intervals as the number of samples τ used for training increases. All algorithms always perform better than the uniformly random allocation. The performance of both the IDIL and DA-IDIL algorithms is significantly better than the one provided by AdComB-Mean (confidence intervals do not overlap). The use of more training samples provides an improvement in terms of steady-state revenue for the IDIL and DA-IDIL algorithms. On the other hand, the performance of the AdComB-Mean algorithm does not benefit from having more samples, which is probably due to the presence of a model bias induced by the fact that it neglects the sub-campaign interdependence.

Synthetic Setting 2. There are $N = 5$ sub-campaigns, whose interdependence graph is shown in Figure 4b. The longest path of the interdependence graph $\mathcal{G}$ is $\Gamma = 2$. C_1, C_2, and C_3 are display sub-campaigns directed to a wide audience and have a low cost per impression, but a low conversion rate. C_4 is a social sub-campaign, whose number of impressions is influenced by the influence index of the display sub-campaigns. Finally, C_5 is a search sub-campaign, whose impressions depend on the influence index of C_1 and C_4. The interdependence among the sub-campaigns occurs within $K = 3$ days and is modeled as in Setting 1. We set a cumulative budget of $Y = 500$ and the budget discretization from the interval $[0, 500]$ with $B = 100$. The number of samples for training is $\tau \in \{100, 150, 200\}$.

Results. In Figure 5c, we report the average (over 100 repetitions and over $t \in \{107, \dots, 120\}$) revenue of the algorithms. The performance of AdComB-Mean is worse than the one of the uniformly random allocation and gets worse as τ increases. This is an empirical confirmation of the statement of Theorem 1, showing that a solution that is optimal without interdependence might perform arbitrarily bad. Conversely, the performance of IDIL and DA-IDIL are significantly larger than that of the uniformly random allocation and increase as the number of samples increases.

Final Remarks. Results obtained in synthetic settings show that this model, relying on a training time which is reasonable for the application, provides a significant improvement in terms of revenue of an Internet advertising campaign. Experts in the marketing field confirmed the feasibility of what proposed in terms of learning time. Conversely, adopting more complex models would most likely result unaffordable in most of the cases, since accurate estimations would require a larger training set and, therefore, excessively long learning periods in real-world scenarios.

7 CONCLUSIONS AND FUTURE WORKS

In this paper, we formalize, for the first time, the problem of optimizing an Internet advertising campaign with sub-campaigns interdependence. We design the IDIL algorithm that, given a set of past observations, models these interdependences and returns an optimal allocation of the bid/daily budget on the sub-campaigns maximizing the revenue. We analyse the properties of the IDIL algorithm both theoretically, providing a bound on the total error, and empirically, showing that it provides revenues on synthetic datasets significantly better than the state of the art in the field.

In the future, we will extend our algorithm to an online framework and test it in a real-world application.

REFERENCES

[1] Cesare Alippi, Manuel Roveri, and Francesco Trovò. 2014. Learning causal dependencies to detect and diagnose faults in sensor networks. In *SSCI*. IEEE, 34–41.

[2] César Calderón and Lin Liu. 2003. The direction of causality between financial development and economic growth. *J DEV ECON* 72, 1 (2003), 321–334.

[3] David A Dickey and Wayne A Fuller. 1979. Distribution of the estimators for autoregressive time series with a unit root. *J AM STAT ASSOC* 74, 366a (1979), 427–431.

[4] Matthias Ehrgott. 2005. *Multicriteria Optimization*. Springer-Verlag, Berlin, Heidelberg.

[5] Gabriele Farina and Nicola Gatti. 2017. Adopting the cascade model in ad auctions: Efficiency bounds and truthful algorithmic mechanisms. *J ARTIF INTELL RES* 59 (2017), 265–310.

[6] Margherita Gasparini, Alessandro Nuara, Francesco Trovò, Nicola Gatti, and Marcello Restelli. 2018. Targeting Optimization for Internet Advertising by Learning Logged Bandit Feedback. In *IJCNN*. 1–8.

[7] N. Gatti, A. Lazaric, M. Rocco, and F. Trovò. 2015. Truthful learning mechanisms for multi-slot sponsored search auctions with externalities. *ARTIF INTELL* 227 (2015), 93–139.

[8] N. Gatti, A. Lazaric, and F. Trovò. 2012. A truthful learning mechanism for contextual multi-slot sponsored search auctions with externalities. In *EC*. ACM, 605–622.

[9] Nicola Gatti, Marco Rocco, Paolo Serafino, and Carmine Ventre. 2018. Towards better models of externalities in sponsored search auctions. *THEOR COMPUT SCI* 745 (2018), 150–162.

[10] S. C. Geyik, S. Faleev, J. Shen, S. O'Donnell, and S. Kolay. 2016. Joint Optimization of Multiple Performance Metrics in Online Video Advertising. In *SIGKDD*. 471–480.

[11] Clive WJ Granger. 1969. Investigating causal relations by econometric models and cross-spectral methods. *ECONOMETRICA* (1969), 424–438.

[12] Paul R Hoban and Randolph E Bucklin. 2015. Effects of internet display advertising in the purchase funnel: Model-based insights from a randomized field experiment. *J MARKETING RES* 52, 3 (2015), 375–393.

[13] IAB. 2017. IAB internet advertising revenue report. https://www.iab.it/wp-content/ uploads/2018/07/IAB-2017-Full-Year-Internet-Advertising-Revenue-Report.REV2_.pdf. (2017). Online; accessed 29 August 2017.

[14] E. M. Italia, A. Nuara, F. Trovò, M. Restelli, N. Gatti, and E. Dellavalle. 2017. Internet Advertising for Non-Stationary Environments. In *AMEC*. 1–15.

[15] M. King, J. Atkins, and M. Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *AM ECON REV* 97, 1 (2007), 242–259.

[16] Pavel Kireyev, Koen Pauwels, and Sunil Gupta. 2016. Do display ads influence search? Attribution and dynamics in online advertising. *INT J RES MARK* 33, 3 (2016), 475–490.

[17] K.-C. Lee, A. Jalali, and A. Dasdan. 2013. Real Time Bid Optimization with Smooth Budget Delivery in Online Advertising. In *ADKDD*. 1–9.

[18] Randall Lewis and Dan Nguyen. 2015. Display advertising's competitive spillovers to consumer search. *QME-QUANT MARK ECON* 13, 2 (2015), 93–115.

[19] Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, et al. 1995. *Microeconomic theory*. Vol. 1. Oxford university press New York.

[20] Alessandro Nuara, Nicola Sosio, Francesco Trovò, Maria Chiara Zaccardi, Nicola Gatti, and Marcello Restelli. 2019. Dealing with Interdependencies and Uncertainty in Multi-Channel Advertising Campaigns Optimization. In *WWW*. to appear.

[21] Alessandro Nuara, Francesco Trovò, Nicola Gatti, and Marcello Restelli. 2017. A Combinatorial-Bandit Algorithm for the Online Joint Bid/Budget Optimization of Pay-per-Click Advertising Campaigns. In *AAAI*. 2457–2463.

[22] Omer Ozcicek and W Douglas Mcmillin. 1999. Lag length selection in vector autoregressive models: symmetric and asymmetric lags. *APPL ECON* 31, 4 (1999), 517–524.

[23] Claudia Perlich, Brian Dalessandro, Troy Raeder, Ori Stitelman, and Foster Provost. 2014. Machine learning for targeted display advertising: Transfer learning in action. *MACH LEARN* 95, 1 (2014), 103–127.

[24] Foster Provost, Brian Dalessandro, Rod Hook, Xiaohan Zhang, and Alan Murray. 2009. Audience selection for on-line brand advertising: privacy-friendly social network targeting. In *SIGKDD*. 707–716.

[25] Carl Edward Rasmussen and Christopher K. I. Williams. 2006. *Gaussian Processes for Machine Learning*. The MIT Press.

[26] Manuel Roveri and Francesco Trovò. 2017. An Ensemble Approach for Cognitive Fault Detection and Isolation in Sensor Networks. *INT J NEURAL SYST* 27, 03 (2017), 1–16.

[27] Ori Stitelman, Brian Dalessandro, Claudia Perlich, and Foster Provost. 2011. Estimating the effect of online display advertising on browser conversion. *ADKDD* 8 (2011).

[28] Hiro Y Toda and Taku Yamamoto. 1995. Statistical inference in vector autoregressions with possibly integrated processes. *J ECONOMETRICS* 66, 1-2 (1995), 225–250.

[29] Francesco Trovò, Stefano Paladino, Marcello Restelli, and Nicola Gatti. 2016. Budgeted multi-armed bandit in continuous action space. In *ECAI*. 560–568.

[30] J. Wang, W. Zhang, and S. Yuan. 2016. Display Advertising with Real-Time Bidding (RTB) and Behavioural Targeting. *CoRR* abs/1610.03013 (2016).

[31] W. Weinan, Y. Rong, J. Wang, T. Zhu, and X. Wang. 2016. Feedback Control of Real-Time Display Advertising. In *WSDM*. 407–416.

[32] J. Xu, K.-C. Lee, W. Li, H. Qi, and Q. Lu. 2015. Smart pacing for effective online ad campaign optimization. In *SIGKDD*. 2217–2226.

[33] Jun Yan, Ning Liu, Gang Wang, Wen Zhang, Yun Jiang, and Zheng Chen. 2009. How much can behavioral targeting help online advertising?. In *WWW*. 261–270.

[34] W. Zhang, S. Yuan, and J. Wang. 2014. Optimal Real-time Bidding for Display Advertising. In *SIGKDD*. 1077–1086.

[35] W. Zhang, Y. Zhang, B. Gao, Y. Yu, X. Yuan, and T.-Y. Liu. 2012. Joint optimization of bid and budget allocation in sponsored search. In *SIGKDD*. 1177–1185.