Forecasting online advertising campaigns in the wild

Radoslaw Nielek[1], Dariusz Parzych[1], Damian Sepczuk[1],
Adam Wierzbicki[1], Jakub Wysoczanski[2]

[1] Polish-Japanese Institute of Information Technology, Koszykowa 86,
02-008 Warsaw, Poland
{nielek, dariusz,parzych, damian.sepczuk, adamw}@pjwstk.edu.pl
[2] Gemius SA, Postepu 18,
02-676 Warsaw, Poland
jakub.wysoczanski@gemius.com

Abstract. Online advertising is a high and a constantly growing business with an elaborate complexity. An important position is occupied by intermediaries, such as ad networks, marketing agencies and companies specialized in delivering software for managing and displaying advertising campaigns. Efficient use of advertising budgets and maximizing profits from ad impressions are the common goal of all players. This goal is usually achieved by manual management of advertising campaigns. Such management requires a very good prediction of Web users' behavior and web site traffic. We have tested the performance of an ad emission simulator for various campaigns and different types of constraints (e.g. capping or targeting). Real data from completed campaigns has been used for the experiment. It has been shown that predicting performance of advertising campaigns that have constraints is considerably more difficult than predicting simple campaigns. Our result applies even if we use a Web traffic sample from the current period (idealized Web traffic prediction) and a simulator that practically emulates ad server operation (the most realistic prediction approach).

Keywords: prediction, ads, web traffic, simulation

1 Introduction

Online advertising is a huge and constantly growing business. In the first half of 2011 revenues reached only in the USA ca. 14.9 bln USD [1]. Europe follows the USA with 23.5 bln USD in 2010 and over 8 bln USD in 2011 only in Germany [2].
Reduced to basics, online advertising is an online scheduling and allocation problem under uncertainty. Several performance measures can be applied to the online advertising system, such as: the number of ad impressions, and the number of clicks (which determine the *Click Through Rate* or CTR); the number of distinct users who have been exposed to a campaign; numbers of distinct users in particular target groups, and so on.

Improving these performance measures is sometimes simple. For example, the simplest ad scheduling algorithms used in practice are simple, greedy algorithms that select a campaign based on its size and priority. A more comprehensive (and complex) approach would be the use of specialized optimization methods to create near-optimal schedules that attempt to maximize some performance measure (or a combination of measures) [3],[4]. However, in practice, campaigns are managed manually by personnel of advertising or marketing agencies. These ad managers make use of short-, and mid-term forecasts of advertising campaigns.

The goal of this paper is to study the prediction of advertising campaigns. The project of creation and implementation of the prediction algorithms were conducted together with AdOcean – innovative company specializing in adserving and ad-measurement technology. The algorithms were tested in real on-line environment which gave unique opportunity to verify the prediction theory.

This task is tricky not just because of the non-stationary nature of Web traffic. A more important problem is the possibility of specifying various constraints on a campaign. These constraints are frequently used in practice, and are therefore an important part of the campaign definition. The constraints and campaign settings (such as the campaign priority) can also be used by the advertisers to change definitions of campaigns during their realization, in order to reach some objective (or to prevent the possibility that a campaign will not be completed on schedule). *Since there exists a fixed limit on the number of impressions (or distinct users) that are available to campaigns, the realization of a campaign is also dependent on competing campaigns that use the same placements*. Therefore, precise predictions for campaigns with complex settings and constraints are an important, but complex problem.

The research presented in this paper bases on a unique dataset obtained from an advertisement campaign management company. This dataset contains the historical data on ad emissions, including detailed data about the Web users. The dataset is sufficient to recreate the history of actual campaign realizations during a specified timeperiod. We have used the data to attempt to forecast campaigns, and to study the following research questions: *is it possible to do online prediction of advertising campaigns with high accuracy? Which campaign settings and constraints have the highest impact on prediction accuracy? How can we improve the currently used campaign prediction procedure?*

The contribution of the paper is the evaluation of the following hypotheses:

The precision of advertising campaign performance prediction for campaigns without constraints is higher than for campaigns with constraints, even when Web traffic is predicted using an oracle (using Web traffic from the forecasted period)

The precision of advertising campaign performance prediction for campaigns with constraints decreases if Web traffic is predicted based on historical data

The precision of advertising campaign performance prediction is strongly influenced by campaign stability, which can be affected by changes in competing campaigns.

Our research is also complemented by a qualitative survey (detailed interviews) with clients of the collaborating company, who are advertisers that use the company's ad campaign management system. The datasets used in this research are available from the authors on request, subject to the agreement of the collaborating company (only non-commercial use is allowed).

A final contribution of this paper is a benchmark for the forecasting of advertising campaigns. New, improved methods of forecasting campaigns may be evaluated *against the results published in this paper in the future*. As our study considers the forecasting of advertising campaigns “in the wild”, based on real data, it creates a benchmark that could be very useful for future lab experiments.

The rest of the paper is structured as follows: in the next section, we describe related work. Section three gives a conceptual background by describing the system of Web advertising. Section four describes the dataset of advertising campaigns used in this research. Section five introduces the simulator used for predicting advertising campaigns. Section six contains an analysis of the correctness of prediction advertising campaign performance and of the practical usage of campaign forecasting by advertisers, based on interviews with clients of the collaborating company. Section seven concludes the paper and discusses future work.

2 Related Work

In the domain of online advertising, authors focus on a system for allocating and pricing display advertisements in contract and auction mechanisms to optimize the profit of an ad publishing company, and design a simple, but near optimal (as is shown by a matching lower bound), online algorithm with provable approximation guarantees [5]. Considering online ad allocation (deciding which ad should be displayed to the user) from the perspective of optimizing the number of ads served to the users, a model has been created as a known distribution over ad slot types. Instead of solving an optimization problem that assumes forecasts will be accurate, this approach solves an optimization problem (or problems) whose solution(s) can serve as a “guide” to a particular robust online selection heuristic [6].

Real-time ad auctions [10] are an alternative method of managing a diverse set of advertising campaigns. However, from our experience with firms operating in the advertising industry the management of campaigns, advertisers usually prefer to use a planning-based approach, which allows them to coordinate their online campaigns with a whole portfolio of marketing campaigns (for example, outdoor campaigns, sales promotions, etc.).

Another problem, encountered among others in predicting clicks on search advertisements, is the sparsity of data at the beginning of a campaign or on new placements, when very little historical data (or none at all) is available. Methods have been proposed to improve click prediction models by mining click behaviour for partial user queries. The click history was aggregated for individual query words, as well as for phrases extracted with a CRF model. The new models show significant improvement in clicks and revenue compared to state-of-the-art baselines trained on several months of query logs [7]. Authors of [8] suggest that the analysis and prediction of the audience’s on-line banner exposure (tracking how the audience is exposed to an on-line campaign), useful in planning emissions of ad more objectively and efficiently, can be done best by looking at its Web site visits. The negative binomial distribution (NBD) model, having long been applied in analysing repeat behaviours, is proposed to serve as a banner ad exposure model.

3. The Web Advertising System

The current system (market and technology) of online advertising is increasingly complex. Apart from the *advertiser*, who wishes to pay for displaying advertisements (also called *creatives*) to *Web users*, and the *publisher* who owns *placements* where creatives can be displayed, an important position is filled by intermediaries, such as ad networks, marketing agencies, and companies which specialize in the management and running of advertising campaigns. Such a campaign *manager* allows its customers (who can be advertisers, ad networks, or marketing agencies) to define *campaigns*, which are detailed schedules for displaying creatives and which can contain complex constraints, such as *targetings*, *cappings* or *keywords*. Based on location, age, sex, previously visited web sites and many more dimensions customers can define a particular subset of Internet users who will see an ad (*targeting*). Defined rules and their interconnections can be very complex and using a lot of information about users past behaviors and browsing habits and is called *behavioral targeting* [9]. The final entity in the system that controls the realization of an advertising campaign by selecting and scheduling ads to be shown to specific users is the *ad server* (or emitter). Virtually all ad servers have an option to limit the number of impressions per user (it is often used for, both, avoiding an annoying number of impressions and assuring a minimum reach) *(capping)*.

A single campaign is usually composed of many creatives that vary not only in size and format but also in transmitted message. Therefore using constraints mentioned in previous paragraph, it is possible to set duration periods of subsequent parts of an ad campaign, assign a precise number of ad impressions or specify the speed for completing the ad impression plan (it is possible to decide either on steady ad impression during the whole campaign or to freely accelerate the completion of an ad impression plan). It is possible to specify the share of individual creatives in an ad impression plan (e.g. one creative is executed three times more often than the remaining ones) and the order of creatives in the schedule. Last but not least, it is possible to specify the priority of a campaign: campaigns with a higher priority are given precedence in the schedule.

The defined campaigns are usually stored by a central management server (owned by the ad manager company), and the actual Web users' browsers communicate (after executing a script on a placement) with the *ad server*. The ad server implements the campaign by choosing the creative to be sent to the browser. The data on all ad impressions and on all ad clicks are recorded in a database that is later used to create detailed statistics for the customer who has defined the campaign. It is also worth to mention that many campaigns go on at the same time on the same web sites. Therefore any changes on even a single creative will influence performances of remaining campaigns.

In this work, we use a unique dataset obtained from an advertisement campaign management company. The company is a provider of Internet advertising solutions and campaign management tools (e.g. ad servers) that are mostly used by media companies, owners of very popular web sites and affiliated networks. The company serves more than 100 clients. All datasets used in this paper come from the main product of the company – the ad server.

The company's system is composed of four layers which are presented on the figure 1. The superior server is responsible for interacting with clients and preparing plans for advertising campaigns. All plans of advertising campaigns are transformed into an impression plan for the next 24 hours, and these plans are asynchronously transmitted (every hour) to ad servers (emitters) responsible for delivering creatives. Every ad server is assigned to a subset of websites and responds to the requests from scripts embedded in web pages (and executed in the browser) by sending an appropriate creative. Ad servers follow delivered impression plans and select creatives to display taking into account constrains and priorities (advertising campaigns can influence each other only on the same emitter). Emitters also collect statistics that are sent back to the superior server every hour.

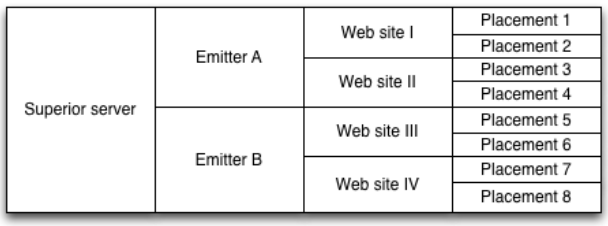

Fig. 1. Ad server design.

The ad servers schedule campaign realizations based on real-time data about ad impressions. This data allows calculating a ratio of the actual impressions to attempted impressions that failed (for example, due to ad blocking). This ratio is taken into account in the scheduling. Advertisers create plans of campaign realizations expressed in the number of impressions, distinct users, etc., for a longer time period (several days). The superior ad server creates schedules for emitters that are updated every hour. An emitter realizes the schedule by first choosing a campaign, then a creative, and finally a placement for this creative. Emitters can also dynamically modify the schedule, for example, if the campaign reaches its target it is stopped immediately.

1.1 Measures of Performance of Advertising Campaign Prediction

Several measures could be applied to measure the correctness of advertising campaign prediction. The simplest one could be the squared error of realized and predicted number of ad impressions over time. Another simple measure is the correlation of the realized and predicted number of ad impressions.

However, for the advertiser, predicting the actual amount of impressions for a given point in time may be less important than then prediction of the sum of impressions in a period of time. Therefore, the difference between the predicted and realized sum of ad impressions over time is another important performance measure.

Finally, due to the large differences in numbers of impressions between various advertising campaigns, comparison of the squared errors between campaigns is difficult. Therefore, a relative absolute error of the predicted and realized number of impressions (absolute error divided by the realized number of impressions) is another valuable performance measure. The same concerns the total sums of impressions.

The following performance measures of advertising campaign prediction have been used in this work:

Relative squared error

$$RSE = \sum_{i=t}^{N}[(E_i - O_i)/E_i]^2/(N-1)\ . \quad (1)$$

Relative absolute deviation

$$RAD = \sum_{i=1}^{N}\frac{|E_i - O_i|}{E_i}/(N-1)\ . \quad (2)$$

N - Number of observations or sum of weights; E_i - Predicted value of case i; O_i - Observed value of case i

Correlation coefficient

$$\boldsymbol{r} = \frac{\sum_{i=1}^{N}(E_i - \bar{E}) * (O_i - \bar{O})}{\sqrt{\sum_{i=1}^{N}(E_i - \bar{E})^2} * \sqrt{\sum_{i=1}^{N}(O_i - \bar{O})^2}}\ . \quad (3)$$

N - Number of observations or sum of weights; E_i - Predicted value of case I; $\bar{E}$- Mean of predicted values; O_i - Observed value of case I; $\bar{O}$ - Mean of observed values

Non-parametric tests.

Because the chosen performance measures may not have normal distributions, we have used non-parametric statistical tests for a comparison of their values. We use the Wald-Wolfowitz runs test, the Mann-Whitney U test, and the Kolmogorov-Smirnov two-sample test.

The Mann-Whitney U test is a nonparametric alternative to the t-test for independent samples. This test assumes that the variable under consideration was measured on at least an ordinal (rank order) scale. The interpretation of the test is essentially identical to the interpretation of the result of a t-test for independent samples, except that the U test is computed based on rank sums rather than means. The U test is the most powerful (or sensitive) nonparametric alternative to the t-test for independent samples; in fact, in some instances it may offer even greater power to reject the null hypothesis than the t-test.

With samples larger than 20, the sampling distribution of the U statistic rapidly approaches the normal distribution. Hence, the U statistic (adjusted for ties) will be accompanied by a z value, and the respective p-value.

4. The Dataset of Advertising Campaigns

The dataset used in this research comes from the AdOcean company and consists of campaign data from over 180 days from 59 emitters. The data includes the definitions and realization history of over 40k individual campaigns, as well as the traffic history of each placement (about 900k impressions on average), as well as information about the cookies of Web users. The data is therefore sufficient to recreate the realization of an advertising campaign. Because of the size of the data (order of 109 impressions for each emitter per month), in practical predictions of campaign realization, only part of the data is used (36 emitters with the highest percentage of stable campaigns).

Studying real data from advertising campaign realization is complicated by the fact that each of these campaigns could be (and frequently was) manually controlled by the customer (e.g. an advertiser). This manual control could modify the entire definition of an advertising campaign. The simulator used in this research can cope with the change of a campaign definition, but such a change completely modifies campaign realization and makes it more difficult to evaluate long-term predictions.

We therefore sought to identify campaigns that have not been modified for at least 7 days. However, because the simulator uses samples of historical Web traffic, the long-term prediction does not work well for placements with a small amount of traffic. We therefore selected emitters with a low degree of variability of placement popularity in order to study the correctness of predictions of individual campaigns. After computing the coefficient of variability for all placement potentials on each emitter, we hand-picked those having the lowest coefficient and total impressions over 4*109. These limitations severely reduced the number of campaigns available for the study; currently we are working on about 4100 campaigns from two different emitters.

The chosen campaigns had a variety of settings, however, for the purpose of statistical comparison, we have identified two largest groups of cappings. This resulted in the creation of three general groups of campaigns: campaigns with no additional settings, campaigns with keywords capping and campaigns with two cappings: keywords capping and geotargetting. Keywords are delivered together with an HTTP request to the emitter, and are independent from the user.

5. The Advertising Campaign Simulator

The prediction of ad impressions based on a history of campaign realization and Web traffic is a complex and challenging task. The reason for this is the non-stationary behavior of both the Web traffic, and of the advertiser who can change campaign definitions. Moreover, the performance of an ad campaigns with constraints (like global cappings) can depend on other campaigns that are run by the same ad server. For these reasons, simple prediction approaches based on time series analysis and fitting are bound to fail, since they depend on the stability of both Web traffic and campaign definitions. A simulation approach can be used instead. The simulator used in this study simply reproduces all the steps of an ad server (it is an ad server emulator), and maintains campaign statistics in the same manner. It is therefore capable of immediately reacting to modifications of campaign definitions, and to changing statistical characteristics of Web traffic. Also, the simulator used in this study is the most realistic approach, as it emulates a real ad server. The drawback of this approach is it’s computational complexity.

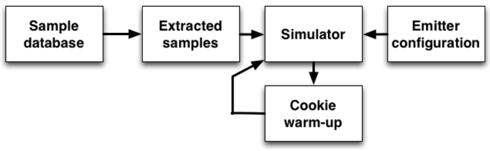

Fig. 2. The campaign simulator.

A simulator receives as an input the advertising campaign definition and the Web traffic history from a specified period for the placements that are controlled by an ad server. The output of the simulator is a prediction of the realization of the advertising campaign for the next period.

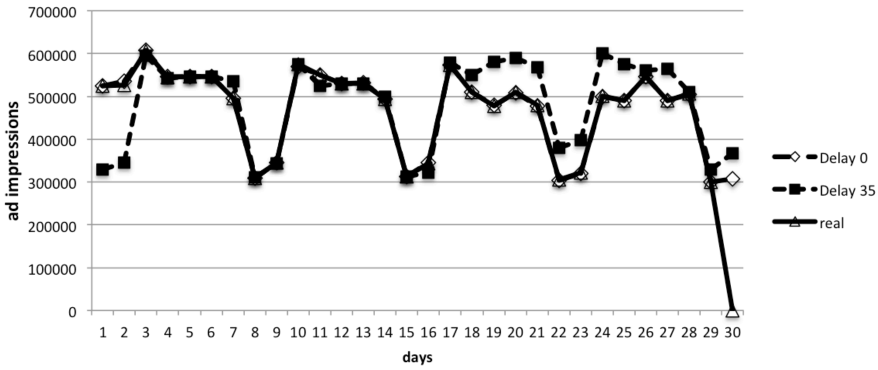

Fig. 3a. A successfully predicted advertising campaign.

Two main steps of the simulator's operation can be easily identified. Every ad is displayed to a single user identified by a unique cookie and placed somewhere on web page. The first step of the simulation consists of a sampling of the available Web traffic history. Because of the size of this data, simulating based on the entire history would be extremely computationally complex. Therefore, up to 1% of the data is sampled using a random selection of cookies. The sampled data can then be used by a Web traffic predictor to create a sample of Web traffic for the next period. The simplest way to create such a sample is to assume that the next period will be statistically similar to the previous one, and to use the sample from the previous period without any modification. Usually, in the simulations considered in this paper, this approach is used, and the traffic sample is based on the previous 35 days of traffic.

Another approach used in these studies that estimates the effect of a better prediction of Web traffic is the use of a traffic sample from the current period, instead of the previous period. Such an "oracle" is not available in practice, but is a benchmark of how well Web traffic can be predicted that is independent of a particular method of Web traffic prediction. In the results in this paper, we compare the results of simulation based on a sample of the previous 35 days (d35) with a sample based on real traffic from the current period (d0).

Using the historical Web traffic, the simulator is also "warmed up" with data about previous campaign realization and on which Web users have been exposed to campaigns. This warm up is important for simulating and predicting the realization of campaign constraints, such as capping or targeting.

The second step of the simulator is essentially an emulation of the ad servers operation on the sampled Web traffic. The simulator repeats the steps of the ad server's scheduling algorithm, given the definitions of advertising campaigns. Like the ad server, the simulator keeps track of the campaign performance. This performance data is the output of the simulator: a prediction of campaign realization.

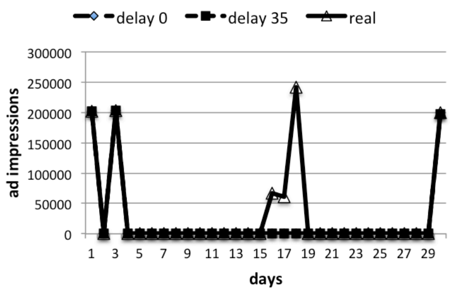

Fig. 3b. Short-term instability of an advertising campaign

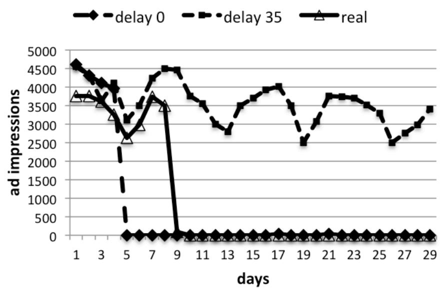

Fig. 3c. Long-term instability of an advertising campaign

6. Analysis of Correctness of Ad Campaign Forecasting

6.1 Typical Examples of Campaign Prediction

Figures 3a–c show the number of ad impressions for three campaigns, along with their predictions based on the past Web traffic (d35) and current Web traffic (d0). Figure 3a shows a typical outcome of the simulator. Prediction based on the current traffic (d0) delivers a very good precision and almost ideally resembles the real ad impressions. Although prediction based on the current traffic can be conducted only after the campaign has finished (ex post), it is very helpful in identifying influence of particular factors. Results of simulation based on past network traffic (35 days old offset) are, as can be expected, slightly worse but the difference is not striking and it is mostly visible for last 10 days where prediction based on historical data overestimates the total number of ad impressions.

Figure 3b shows the effect of short-term Web traffic fluctuations. Such fluctuations cannot be predicted even using a sample of current Web traffic (d0). Such a sample can be too small to capture the fluctuation that occurs only on a particular placement and therefore affects one campaign. However, this does not imply that the overall prediction of campaign correctness is bad; this depends on the frequency of short-term traffic fluctuations.

Figure 3c shows the effect of a long-term change in Web traffic characteristics. This change cannot be predicted using historical data (d35); however, our approach that uses current Web traffic (d0) copes well with this long-term change, which is of sufficient duration to be captured by the sampling.

6.2 Campaign stability

Campaign definitions are subject to dynamic modifications by the advertisers. The creation of new advertising campaigns that compete for the same placements (for example, an increase in priority of competing campaigns). For these reasons, in order to study campaign forecasting, we needed to define and limit our analysis to a set of campaigns that could be considered stable. Our definition of campaign stability is best illustrated on a diagram.

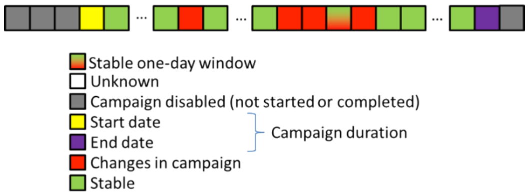

Fig. 4a. States of an advertising campaign

Figures 4a-b show the potential states of an advertising campaign. We consider time windows of one day, and investigate whether the campaign was running at that time and whether or not the campaign settings were modified.

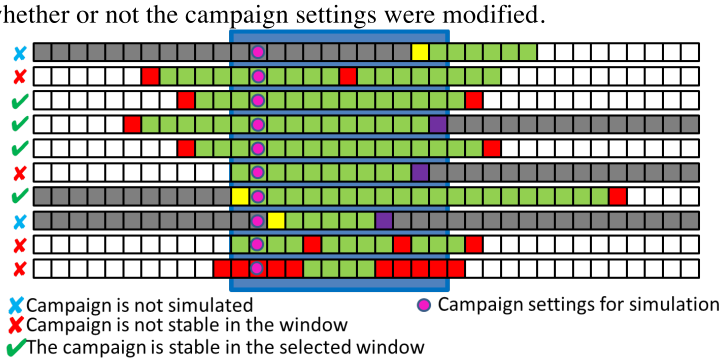

Fig. 4b. Selection of stable campaigns

The selection of stable campaigns was done separately for each emitter and for time windows of 14 days. A campaign is considered stable if it is not modified for a period of minimum 12 days within a time window. Campaigns settings for the simulation are taken from the second day of this period, because if the campaign started on the first day, its settings were usually frequently modified.

6.3 Difficulty of prediction non-stable campaigns

According to the common sense campaigns with constantly changing setups have to be difficult to predict (at least in comparison to stable campaigns). To validate this claim relative absolute deviation and relative square error for stable and non-stable campaigns have been compared and presented in Table 1. Mean error of non-stable campaigns is almost two times bigger than for stable campaigns. Differences are statistically significant.

Table 1. Correctness of predicting campaigns with and without constraints using historical traffic.

Mean of measure	Stable campaigns	Non-stable campaigns
Relative absolute deviation	4.85	8.15
Relative square error	7973.89	112492.32

6.4 Comparison of Prediction of Campaigns With and Without Constraints

The most important question considered in this paper concerned the effect of campaign constraints on the correctness of campaign prediction. This effect is apparent when we compare the correctness of predicting campaigns divided into two groups: campaigns with and without constraints. In the first group (no constraints), we have selected 246 stable campaigns to study long-term prediction. The second group (campaigns with constraints) contained 752 stable campaigns that could have constraints of three types: keywords cappings, or geotargetting or cappings on the number of distinct users (impact).

As can be seen in Table 2 two performance measures, the relative absolute error and the correlation, the correctness of campaign prediction varies strongly (and with statistical significance at $p<0.05$) among four groups. In general constrains make prediction more difficult (constrains on keywords are an exception). All differences presented in Table 2 are statistically significant.

Table 2. Correctness of predicting campaigns with and without constraints using historical traffic.

Median of measure	No constraints	Constraint on no. of distinct users	Constraint on keywords	Constraint on geo-targetting
No. of campaigns	246	230	269	253
Relative absolute deviation	0.96	1.30	0.86	1.32
Relative square error	5.65	17.51	0.81	11.54

6.5 Effect of Idealized Web Traffic Prediction

The difference in the median prediction errors for campaigns without constraints using predictions based on current and historical traffic is large and statistically significant at level p=0.001. Result showed in Table 3 proves the second hypothesis of this paper, that the precision of prediction of advertising campaign performance for campaigns with constraints decreases if Web traffic is predicted based on historical data. This observation also shows a potential for improvement of the prediction of advertising campaigns with constraints: improving Web traffic prediction based on historical data could lead to an improvement of prediction accuracy.

Table 3. Correctness of campaign forecasts depending on used traffic samples (actual vs. past traffic).

Median of measure	Prediction based on actual traffic (d0)	Prediction based on past traffic (d35)
Relative absolute deviation	0.96	1.11
Relative square error	5.65	10.05

6.6 Practical use of campaign forecasting by advertisers

To gain a deeper understanding how clients use predictions to modify campaigns three IDIs[1] were conducted. All interviewed clients are heavy users of ad servers and every year manage few hundreds ad campaigns each. Clients have pointed out that prediction is useful for planning new campaigns but even more important for identifying campaigns which could not be completed because of too strict conditions (actually that what is important is a combination of number of visitors and their demographical features, campaign's settings and other campaigns published on the same placements). Although contracts for displaying ads usually contain a fixed number of impressions, experienced media planners do not react if the deviation from the emission plan is less than 10 to 20% (they do not react at all if the number of impressions is higher than scheduled).

The very first reaction of media planners to the campaign under risk (according to information given during interviews) is the change of priorities – as a consequence this particular campaign is preferred to others. If the priority cannot be changed because is already set to the maximum (or changes to not solve the problem) the next steps considered by media planners are, according to interviews, loosening constrains and adding new placements. Some changes can be done without big consequences for the advertising strategy (e.g. adding new placements), while others require acceptance from clients (e.g. capping). Some restrictions cannot be changed at all (e.g. geo targeting).

Past experiences of interviewed clients show that although many campaigns are at risk, by changing in their parameters almost all can be saved. Typically, the accepted

[1] In-depth interview

margin of error in prediction number of impressions was settled on 10 to 15%, but they also stress that it is much easier to deal with underestimation than overestimation. Interviewed clients use prediction constantly to monitor and modify campaigns, thus, they expect that consequences of every change will be immediately visible in the system. Because all campaigns on the same emitter are interconnected (and strongly depended from each other) such expectation requires a lot of computation which have to be done almost in real time.

7. Conclusions and Future Work

We have found that the addition of constraints has a strong negative impact on the correctness of campaign prediction. This effect is statistically significant in our study. The increase in the number of constraints also decreases the correlation between the predicted and realized numbers of ad impressions.

It is worthwhile to note that the campaign predictions were made in our study using a very realistic simulation of the real ad server. Moreover, the simulations based in one case on an idealized Web traffic prediction that was based on the Web traffic from the same period as the evaluation. This means that the problem with the prediction accuracy of advertising campaigns with constraints most likely cannot be solved easily.

Our study has also demonstrated the importance of campaign stability, understood both as the stability of campaign setting that can be modified at any time by the advertiser, and as a dependence of each campaign on competing campaigns that can be unstable. This dependence will be the object of our future study. The qualitative survey of advertisers has revealed that they mostly use short-term predictions and react to a perceived threat to the realization of their campaigns by changing campaign setting (usually increasing campaign priority). This behavior could result in a feedback among the competing campaigns that continually increase their priorities.

The comparison of prediction correctness using current Web traffic (Web traffic from the same period as the campaign realization) and using historical Web traffic revealed that the correctness of predicting campaigns with constraints decreases further when historical Web traffic is used. This result shows a potential of improvement, since a simple sample of historical Web traffic (using an assumption of stationary distribution) has been used for the Web traffic prediction in our study. Also, for campaigns with constraints, samples should be more frequently refreshed, resulting in a prediction that is based on more current data.

The simulation approach used in this paper seems to be the most promising approach to advertising campaign prediction. Approaches based on analytical models cannot succeed due to a multi-level instability of the problem. The studied simulator creates good predictions for campaigns without constraints. However, for considering campaign optimization, the prediction of campaigns with constraints must be improved. This is currently an open research problem that should be solved in order to apply optimizations of advertising campaigns on the Web.

8. Acknowledgements

This research has been supported by the project "Prediction and reservation of advertising placements on the Internet" (159401/NCBR/12) financed by the Polish National Centre for Research and Development.

9. REFERENCES

1. Internet Advertising Board. IAB Internet Advertising Revenue Report, 2011 First Six Months Results. http://www.iab.net/media/file/IAB-HY-2011-Report-Final.pdf (2011)
2. OVK. OVK Online Report 2011/02 Zahlen und Trend im Ueberblick. http://www.bvdw.org/fileadmin/bvdw-shop/ovk-report-2011-2.pdf (2011)
3. Feldman, J., Mehta, A., Mirrokni, V., Muthukrishnan, S., Offline Optimization for Online Ad Allocation. Fifth Workshop on Ad Auctions, Stanford (2009)
4. Alaei, S., Arcaute, E., Khuller, S., Ma W., Malekian, A., Tomlin, J., Online Allocation of Display Advertisements Subject to Advanced Sales Contracts. s.l. : Proc. ADKDD'09 (2009)
5. Jiang, J., Papavassiliou, S., Enhancing network traffic prediction and anomaly detection via statistical network traffic separation and combination strategies. Computer Communications. 29, 1627--1638 (2006)
6. Sang, A., Li, S.Q., A predictability analysis of network traffic. Computer Networks. 39, 329--345 (2002)
7. Hillard, D., Manavoglu, E., Raghavan, H., Leggetter, C., Cantú-Paz, E., Iyer, R., The sum of its parts: reducing sparsity in click estimation with query segments. Information Retrieval. 14, 315–336 (2011)
8. Huang, C.Y., Lin, C.S., Modeling the audience's banner ad exposure for internet advertising planning. Journal of Advertising. 35, 2, 123–136 (2006)
9. Jaworska, J., Sydow, M., Behavioural Targeting in On-Line Advertising, An Empirical Study. WISE 2008. LNCS, vol. 5175, pp. 62—76. Springer, Heidelberg (2008)
10. Martin, S., Anders, C., Individualized, Real-Time, Interactive E-commerce Auction, United States Patent No: 6,606,607 (2003)