

AdMotional: Towards Personalized Online Ads

Manfred Meyer, Markus Balsam, Arlo O'Keeffe, Christian Schlüter, Jens Fendler

University of Applied Sciences Gelsenkirchen

Münsterstr. 265

46397 Bocholt, Germany

+49 2871 2155 910

{meyer,balsam,okeeffe,schluter,fendler}@admotional.org

ABSTRACT

AdMotional is a research project aiming at achieving a win-win situation for online advertisers and web users alike by optimizing the campaign selection process and creating personalized ads. This results in increased campaign performance for advertisers, and in more relevant and thus less annoying ads for consumers. We give a general overview and present the system architecture, before describing the main components in greater detail. We also introduce the learning and optimization component and strategies, before concluding with a summary and brief outlook into future developments.

Categories and Subject Descriptors

H.3.5 [Information Storage and Retrieval]: Online Information Services – *Commercial services, Web-based services*

K.4.4 [Computers and Society]: Electronic commerce

General Terms

Algorithms, Performance, Design, Economics, Human Factors

Keywords

AdMotional, Ad Serving, Emotional Targeting, Inductive Learning, Online Advertising, Personalization, Rule Based Ad Optimization

1. INTRODUCTION

In typical online ad scenarios neither website publishers nor advertisers contact each other, but both rely on a third party: an advertising network. Advertisers, and in turn advertising networks, have a natural interest in high performance of their campaigns. This demand drives a constant search for new or enhanced marketing forms (e.g. viral marketing), ad formats (e.g. rich-media ads), and channels (e.g. in-video ads). However the identification of such only leads to short-term advantages over competitors, and is thus not sustainable in itself. The AdMotional project addresses this issue by aiming for *long-term* enhancements of online ad performance through a) targeting optimization, and b) personalization of ads taking contextual, emotional, and other

aspects into account [1]. “Targeting” refers to the selection of ads for a particular audience while matching underlying campaign requirements [14], while “personalization” describes the process of customizing ads in order to make them more appealing to a particular consumer (user) in a given situation. Within our project these two aspects are equally covered. We present the overall system architecture in chapter 2, before covering the details of our targeting and personalization mechanisms. In chapter 3 we discuss our targeting process, which exploits existing targeting strategies in combination with a new emotional dimension, allowing for campaign selections based on consumers' moods. This step is further extended by moving from the traditional (campaign-based) targeting to an additional more fine-grained (ad-based) targeting – aiming for a single, best-matching, advertisement. After selection, specific customization points are identified as the basis for individual customization of the ad media, resulting in a personalized ad: potentially unique for every individual consumer and online scenario, as further explained in chapter 4. The design of the system's feedback component is presented in chapter 5. This machine learning module complements the overall system by constantly monitoring and analyzing ad performance in an attempt to derive rules for not only optimizing the targeting and personalization processes, but also to inform ad designers of the most influential factors to be considered. We conclude with a brief summary and outlook into future developments in chapter 6.

2. SYSTEM OVERVIEW

Our project is not an ad server in itself, but was designed as a loosely coupled system to only facilitate the campaign/ad selection and generation process upon requests from existing ad servers. From an external point of view, the system is fully embedded within a surrounding ad server. Yet as the system heavily depends on advertiser parameters and campaign details (statistics, campaign history, etc.), overall ad delivery speed will greatly benefit from hosting in close vicinity to the surrounding ad infrastructure. The remainder of this chapter takes a closer look into the system's internal modules and operations, while focusing on steps relevant to AdMotional only.

Upon initial page requests from consumers, web servers respond with HTML content containing URLs to scripts on the ad server. During the HTML parsing process the consumer's browser requests these additional URLs. Ad servers then typically pre-process these requests (e.g. to deliver high-priority campaigns directly if necessary), extract session information, and send separate requests for ads to the AdMotional system. In addition to information as present in the initial consumers' requests (original URL, browser type, language, etc.) the ad server may provide further information in support of our system, such as priority

parameters or a browsing history as extracted from consumers' session data.

Based on this enriched request our system first selects the most relevant campaign and ad, before identifying ad customization points (e.g. banner size, background color, font size) and dynamically creating a personalized ad for the consumer. The media format to be used is specified in the selected campaign. While currently focusing on image generation in JPG and PNG formats as well as HTML creation, the underlying ad description language is sufficiently powerful to describe other (dynamic) formats such as Adobe Flash or PDF as well. The generated media (or respective URL) is finally returned to the ad server for immediate delivery to requesting consumers' browsers. The overall process is roughly illustrated in figure 1.

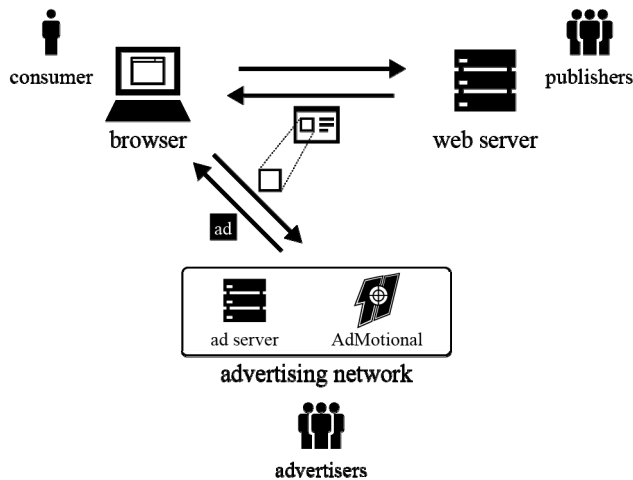


Figure 1: The general ad serving process

For campaign selection and personalization to be sufficiently accurate, we combine a wide variety of internal as well as external data sources. These range from the consumers' request/session data, additional (possibly aggregated) information from the ad server and underlying campaign details to external data sources. The latter includes e.g. location-based services for geo-targeting or weather information to name just a few. We combine all these sources to not only guide the conventional targeting process, but to further derive a consumer model which includes an emotional dimension, taking as much of the consumers' situational information into account as possible. While non-emotional information is used to initially reduce the number of applicable campaigns, consumer models significantly influence the final selection of the most relevant campaign and ad.

This model further supports the population of customization points with specific data instances such as colors or font styles. Figure 2 summarizes the different steps within the selection and generation process based on a single consumer request.

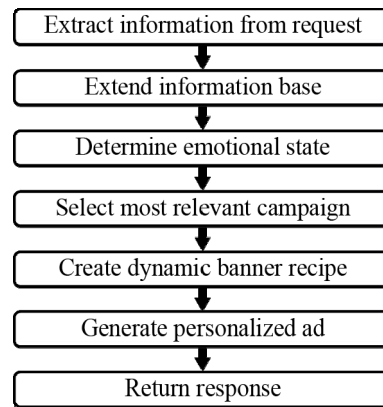


Figure 2: The internal process

On the architectural level, our system is composed of five individual components as shown in figure 3. Despite a relatively tight coupling of components, clear separation of concern allows for simplified improvements and extensions.

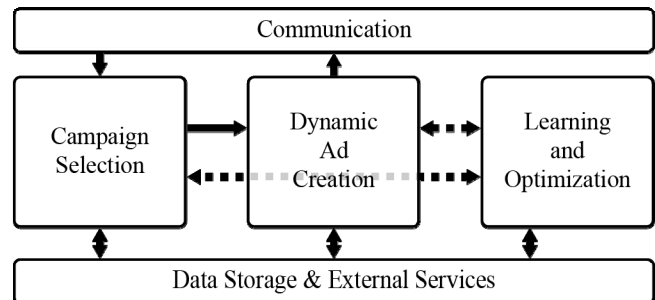


Figure 3: The system architecture

The "Communication" component provides the external interface to the ad server. It is responsible for providing high performance data exchange using a custom TCP protocol, but could be adjusted to provide e.g. HTTP-based adapters for legacy ad server systems.

"Data Storage & External Services" provides access to all persistent or external data sources, as well as to the underlying database module to keep track of (internal) statistical and historic data.

Specific campaigns and ads are evaluated within the "Campaign Selection". This in turn passes its results to the "Dynamic Ad Creation", where the ad media is generated and personalized based on consumer model and results from the "Learning & Optimization" component. While final ads are delivered to the ad server, results from the creation process are passed to the learning and optimization component for analysis of ad performance in comparison to similar ads (see chapter 5). Once sufficient information about an ad's performance is available, the learning and optimization component feeds back optimization information to the selection and/or creation process.

3. CAMPAIGN SELECTION

The Campaign Selection component has a twofold responsibility: first to collect and aggregate targeting information, and second to evaluate this information to select the most appropriate campaign for the requesting user.

As accurate targeting is the result of a combination of heterogeneous targeting methods – each utilizing diverse information – flexibility has been a major design goal, particularly within the campaign selection. We currently integrate contextual, behavioral, and situational targeting strategies [4]. Contextual targeting focuses on information about the requested web page, its keywords and text, while behavioral targeting utilizes historical information and thus focuses on users' browsing behavior [13,2]. Situational targeting relies on location and temporal information in evaluating campaign appropriateness. Having combined the above techniques, and to further enhance the overall targeting quality, we propose the concept of *emotional* targeting. We define emotional targeting as an ad selection process based on a rudimentary model of users' emotions and anticipated feelings, as being derived relevant (partially aggregated) information from consumers' requests, external services, as well as internal deduction techniques and heuristics.

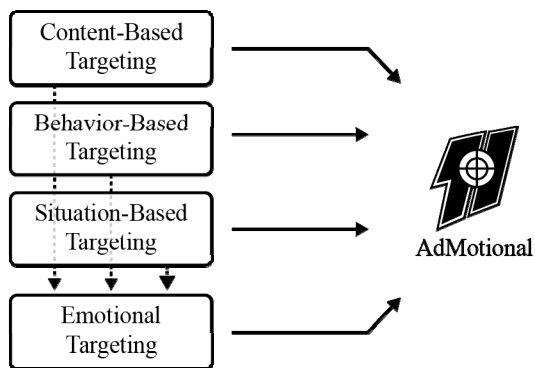


Figure 4: The targeting strategies

The knowledge foundation of our emotional classification system is populated through a modular and extensible information collection process. Results from traditional targeting methods are combined using a weighted, score-based approach. For this purpose, EmotionML [9] is being evaluated as a standardized language to model and further process emotional states in a flexible way. We also investigate different options of embedding rule-based concepts in the selection process. First experiments with the integration of a rule engine (e.g. Take/Mandarax [3,10]) in the final decision step to increase performance are promising. This approach further allows for simplified, asynchronous interfacing between Optimization and Selection/Generation components, particularly if the execution of standardized knowledge base representations (such as RuleML [12]) in the selection process prove to provide the required performance.

4. DYNAMIC AD CREATION

Rather than using pre-rendered ads, AdMotional requires advertisers to provide *templates* with customization points. Customization points are basically ad skeletons including rules and sets of alternative data instances or value generation strategies. During ad creation, customization points are identified and evaluated to produce individual ads (see figure 5 for a simple example).




Figure 5: An example of two ads generated by one template

While this concept demands an increased initial design effort from advertisers, it effectively integrates a potentially large number of situational, personalized ads into a single design, thus reducing the number of ads to be created manually. Yet the approach bears potential for conflict: on the advertisers' side, a release of control over some ad parameters is necessary to achieve a high degree of personalization. A more critical constraint arises from the fact that advertising networks typically limit the time allowed for ad servers to deliver ads, while further aiming for low-bandwidth solutions to meet distributed environments' requirements [7]. These issues are overcome by implementing different strategies for ad creation and evaluation at run-time.

Rather than creating banners directly, the campaign selection module produces *recipes* which are interpreted by separate ad creation modules. We currently provide three such modules, based on two different creation techniques. The first two modules use graphic libraries [6,5] to produce pixel images. While producing high quality ads for given resolutions, this causes relatively high traffic on advertising networks. To reduce the bandwidth impact, images are cached on external servers (e.g. within content delivery networks). The third module creates HTML5 code utilizing canvas elements for embedding in the target website. This strategy meets low bandwidth requirements in advertising networks as image elements can reside with (external) advertisers. Yet cross-browser issues as known from website development still need to be addressed. We are further evaluating a new module for the first strategy, using a separate image creation server. One of the main challenges currently addressed for all ad creation modules is the automatic generation of attribute values. With typed customization points being defined, we might be able to select specific values such as background color automatically, based on consumers' current emotions [8]. As these auto-generated values are expected to have a high impact on ad performance, they are also subject to analysis and constant optimization by the learning and optimization component.

5. LEARNING COMPONENT

To continuously improve the performance of individually targeted ads, we have added a learning and optimization component. Two performance types need to be differentiated in this regard, depending on advertisers' intentions. On one hand so called "performance campaigns" aim for consumers' immediate responses (i.e. clicks on the ad media). "Branding campaigns" on the other hand intend to strengthen advertisers' brands, and are expected to result in increased business turnover in the long run.

While the success of performance campaigns (e.g. simple ad clicks, completed contact forms, or effective online sales) is easily measured and fed back into the optimization component, data about branding campaigns' success is not as readily available and would involve active participation from individual advertisers. We therefore only focus on the optimization of performance campaigns, based on ad server feedback in different categories as related to consumer actions immediately following ad impressions.

The learning and optimization component runs asynchronously to the ad selection and creation process. Successful ad impressions are analyzed as to determine presence and impact of particular targeting criteria as well as customization point values used. Given sufficient statistical support on ad performance, delivered ads are first clustered according to similarity (i.e. product type, media format, etc.). Within these clusters we identify a sub-cluster of high performance ads proven to yield exceptional results. We then iterate over the remaining (low performance) ads within the similarity cluster, and identify particular targeting and customization point dimensions with the greatest differences to the center of the high performance ad cluster. We can now define individual ad modification rules, suggesting to automatically adjust particular ad criteria during the selection and generation process based on experience with this and similar ads. As these candidate criteria are hardly independent (e.g. foreground vs. background color), we are considering the application of a set of design constraints resulting in criteria groups – rather than individual criteria – to be jointly adjusted. The generated rules are represented using RuleML [12] before being serialized and injected into the Take/Mandarax rule compiler [10] within the selection and creation modules. This approach allows for a direct object notation, while not only ensuring a smooth integration with other AdMotional components, but also allowing for efficient rule evaluation as well as simplified knowledge inspection by advertisers.

Ad modification rules as described above each address one particular ad. However, we anticipate a significant number of rules to state similar adjustments – possibly across similarity clusters. As such redundancy in rules has a negative impact on rule evaluation performance, an inductive learning component is triggered upon every addition to the rule base. This component tries to induce a higher-level “abstract ad modification rule”, substituting sets of individual rules, and thus leading to performance improvements during the ad creation process. Moreover, these abstract modification rules constitute a qualitative new level of knowledge as they embody empirically proven rules towards the design of more effective online ads.

6. CONCLUSION

Considering today's situation in the area of online marketing as well as related research projects, personalization is one of the most important topics aiming to maximize campaign performance. Selecting the most appropriate campaign, and in turn dynamically creating customized ads further depends on an extensive data basis which needs to be efficiently evaluated. For this purpose we have presented the AdMotional system which combines existing targeting strategies with the new dimension of emotional targeting. We not only utilize the underlying emotional model for the selection, but also exploit its potential in the personalization of template-based ads. We also presented the learning and optimization component, which integrates a feedback cycle in order to continuously optimize ad performance,

and to derive knowledge about critical factors in ad design for particular audiences.

Future work will mostly focus on the development of a more fine-grained emotional model, as well as on the identification of new emotional indicators for web users. This will further lead to the necessity of re-evaluating options for dynamic ad creation. We will also continue to work in the area of feedback-based ad optimization techniques, and fine-tune the existing modules. Finally we will investigate the potential of AdMotional in other areas of online advertising, such as Retargeting [11].

7. ACKNOWLEDGMENTS

The authors would like to thank the entire AdMotional team at our industry partner The AdInsiders GmbH in Oberhausen, at FH Gelsenkirchen and at the Life & Brain institute at the University of Bonn.

The AdMotional project is being supported by a grant from the



EUROPEAN UNION
Investment in our future
European regional
development fund

European Union within the ZIEL2
program 2007-2013 (EFRE),
IKT.NRW (FKZ 29 00 341 02).

Images used under Creative Commons license and in courtesy of <http://www.flickr.com/photos/anirudhkoul/3784809456/> and <http://www.flickr.com/photos/mostrauss/4691973470/>

8. REFERENCES

- [1] AdMotional Project Web Site. University of Applied Sciences Gelsenkirchen, Germany. <http://www.admotional.org>.
- [2] De Bock, K., and Van den Poel, D. 2010. Predicting Website Audience Demographics for Web Advertising Targeting Using Multi-Web Site Clickstream Data. *Journal Fundamenta Informaticae – Intelligent Data Analysis in Granular Computing*, Vol. 98 Nr. 1, IOS Press. Amsterdam.
- [3] Dietrich, J., Hiller, J., and Schenke, B. 2007. Take - A Rule Compiler for Derivation Rules. In *Proc. RuleML 2007*, pp. 134-148.
- [4] Goldfarb, A. and Tucker, C. 2011. Online Advertising. In: *Advances in Computers*, Vol. 81, Zerkowitz, M. (Ed.), Elsevier.
- [5] GraphicsMagick Image Processing System <http://www.graphicsmagick.org/>.
- [6] Java Advanced Imaging (JAI) API. <http://java.sun.com/javase/technologies/desktop/media/jai/>.
- [7] Johnson, D.W., and Jankun-Kelly, T.J. 2008. A scalability study of web-native information visualization. In *Proc. Graphics Interface (GI'08)*, Canadian Information Processing Society, Toronto (Canada).
- [8] Pitsch, N. 2011. *Preliminary Results from an FMRT study investigating the influence of colors on a person's mood*. Internal Report, AdMotional and Life & Brain, Bonn <http://www.admotional.org/pub/docs/results-fmrt-2011.pdf>.
- [9] Schröder, M. (Ed.) 2010. *Emotion Markup Language (Emotion ML) 1.0*. W3C Working Draft 29 July 2010. <http://www.w3.org/TR/2010/WD-emotionml-20100729/>.
- [10] TAKE – a Java rule compiler. <http://code.google.com/p/take/>.
- [11] Tucker, C. 2010. The Economics Value of Online Customer Data. In *Proc. Joint WPISP-WPIE Roundtable 'The*

Economics of personal Data and Privacy: 30 Years after the OECD Privacy Guidelines, OECD.

- [12] Wagner, G., Antoniou, G., Tabet, S., and Boley, H. 2004. The Abstract Syntax of RuleML – Towards a General Web Rule Language Framework. In Proc. *IEEE/WIC/ACM Int'l. Conf. on Web Intelligence (WI'04)*, IEEE, Washington D.C.
- [13] Yan, J., Liu, N., Wang, G., Zhang, W., Jiang, Y., and Chen, Z. 2009. How much can behavioral targeting help online advertising? In Proc. *WWW 2009*, 261-267, Madrid, Spain.
- [14] Zeff, R.L., and Aronson, B. 1999. *Advertising on the Internet*. John Wiley, New York.