International Journal of Computer Science and Applications,
©Technomathematics Research Foundation
Vol. 8, No. 2, pp. 59 – 80, 2011

ADMOTIONAL: TOWARDS PERSONALIZED ONLINE ADS

MANFRED MEYER

University of Applied Sciences Gelsenkirchen, Campus Bocholt
Münsterstr. 265, 46397 Bocholt, Germany
meyer@admotional.org

MARKUS BALSAM

University of Applied Sciences Gelsenkirchen, Campus Bocholt
Münsterstr. 265, 46397 Bocholt, Germany
balsam@admotional.org

ARLO O'KEEFFE

University of Applied Sciences Gelsenkirchen, Campus Bocholt
Münsterstr. 265, 46397 Bocholt, Germany
okeeffe@admotional.org

CHRISTIAN SCHLÜTER

University of Calgary, Department of Computer Science, Agile Software Engineering Group
2500 University Drive NW, Calgary, Alberta, T2N 1N4, Canada
email@schlueternetz.de

AdMotional is aiming at achieving a win-win situation for online advertisers and web users alike through optimization of the campaign selection process and the creation of individually personalized ads. This approach is expected to result in increased campaign performance for advertisers, and in more relevant – and thus less annoying – ads for consumers. We present a general overview of the system architecture before describing the main components in detail. For campaign selection, AdMotional combines and enhances state-of-the-art targeting technologies by a novel concept of emotional targeting. The key aspect of ad personalization furthermore demands for the dynamic creation of the ad itself for which different approaches are presented and evaluated. We then introduce the learning and optimization component and strategies before concluding with a summary and brief outlook into future developments.

Keywords: AdMotional; Ad Serving; Emotional Targeting; Learning; Online Advertising; Personalization; Rule Based Ad Optimization

1. Introduction

Online advertising uses the Internet for the purpose of delivering marketing messages in order to strengthen brands and/or attract customers. Examples include text ads that appear on search engine results pages, banner (display) ads, in-text ads, or rich-media ads appearing on web pages, portals, or applications. Online advertising dates back to 1994 when the first online ad (banner) was sold. Since then, the online advertising industry has continued to grow, enduring the setback of the dot-com bust in 2001, now standing as a $65 billion industry worldwide [Shanahan and Kurra (2011)]. In 2010, advertisers in

Germany spent more than 5 billion Euro advertising on websites, approximately 22 percent of all advertising spending in Germany [OVK (2011)]. The proportion of advertising that is done online is expected to continue to increase over time as the accessibility of web-based services and content relentlessly expands and web access becomes available to more people, through more devices, such as mobile phones and Internet-enabled TV [Evans (2008)].

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 advertising networks have a mutual interest in consistently achieving high performance in their campaigns. This demand drives a permanent 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). Yet the identification and exploitation of such typically only leads to short-term advantages over competitors and is thus not sustainable in itself. The AdMotional project addresses this issue by identifying methods for long-term enhancements of online ad performance, taking contextual, emotional and other aspects into account. We first focus on the initial targeting process, second on run-time personalization of ad design, and third on automatic improvements thereof through a self-learning optimization component.

In the remainder of this paper we briefly review some related work in chapter 2 and present the overall system architecture in chapter 3, before covering the details of our targeting and personalization mechanisms. In chapter 4 we discuss the AdMotional targeting process, exploiting 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 the 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 5. The design of the system's self-learning feedback component is then presented in chapter 6. It complements the overall system by constantly monitoring and analyzing ad performances 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. Preliminary results of conducted system performance tests as well as user studies regarding the relationship between colors, visuals and emotions are presented in chapter 7. We conclude with a brief summary and outlook into future developments in chapter 8.

2. Related Work

Targeting refers to the selection of ads for a particular audience while matching underlying campaign requirements [Zeff and Aronson (1999)], whereas personalization describes the process of customizing ads in order to make them more appealing to consumers (users) in a given situation. Within our project these two aspects are equally covered.

While many different targeting approaches have been studied by various researchers, e.g. [De Bock and Van den Poel (2010)], [Yan *et al.* (2009)], and some of them have already become the core of modern ad servers [Plummer *et al.* (2007)], the personalization of online ads based on consumers' data has only recently been addressed in the context of Retargeting [Helft and Vega (2010), Lambrecht and Tucker (2011)]. However, personalization in this case is limited to presenting products the user has previously shown some interest in. Also, consumers' moods are regarded an important aspect in advertising business since long [Holbrook and Batra (1987)] but have not yet been explored in the context of targeting online ads.

Performance optimization of online ad campaigns has also been studied, e.g. by [Liu *et.al.* (2007)] who presented interesting empirical results on optimizing online campaigns in response to various user data, e.g. weather, location and age. However, this approach focuses on the campaign selection only, while AdMotional also includes optimization of the rules being applied for dynamically creating the appropriate online ad.

Although there are some approaches in literature aiming at a more personalized selection of online ads based on both context and user profiles, e.g. AD ROSA [Kazienko and Adamski (2004)], or focusing on intelligent ad creation, e.g. [Badali *et. al.* (2010)], we present an integrated and flexible approach bringing together state-of-the-art and innovative targeting methods for campaign selection with dynamic, rule-based creation of personalized ads in a framework that furthermore supports optimization of ad-creation rules based on click performance.

3. System Overview

AdMotional does not provide an ad server by 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 (see figure 1). Yet, as the system heavily depends on advertiser parameters and campaign details (statistics, campaign history, etc.) the 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 specifically relevant to AdMotional.

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, consumers’ browsers request 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 present in the initial consumers' requests (original URL, browser type, language, etc.), the ad server may provide further information to aid our system, such as priority parameters or 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. The media format to be used is specified in the selected campaign.

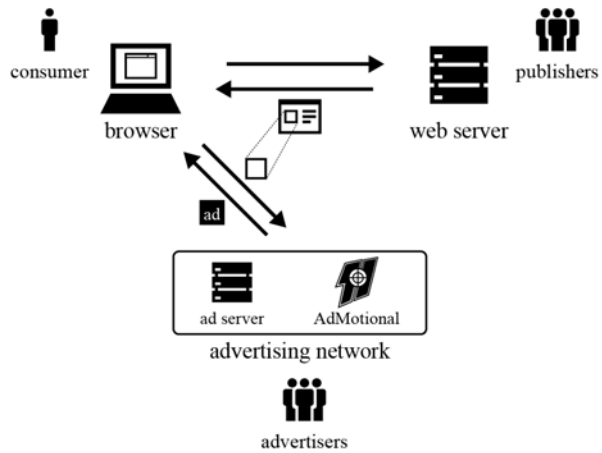

Figure 1: The general ad serving process.

While currently focusing on image generation in JPG and PNG formats, as well as HTML creation, the underlying ad description language (see chapter 5) is sufficiently powerful to describe other (dynamic) formats such as Adobe Flash or PDF. 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, while figure 2 summarizes the different steps within the selection and generation modules based on a single consumer request.

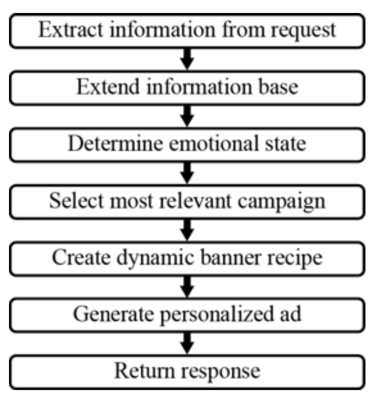

Figure 2: The internal process.

Based on the request parameters provided by the ad server, additional information may be gathered, e.g. geo location or weather data, which extend the information base for the targeting process. Information about content (topic) of the recently visited web sites is then used to determine the emotional state of the user before all available information is taken into account for evaluating available campaigns and selecting the most relevant one, i.e. the one with the highest score. From the selected campaign, a banner recipe is generated, which provides specific pre-defined, static data necessary for the creation of a personalized ad in response to the given ad request. However, the time-consuming task of creating the ad itself may be performed concurrently while a response containing a link to the ad's location on an image server is being returned to the ad server in order to meet time requirements.

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:

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, for example, 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 ad media are generated and personalized based on consolidated consumer models and results from the “Learning & Optimization” component.

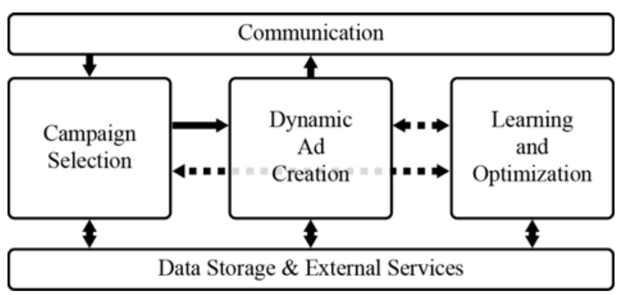

Figure 3: The system architecture

While rendered ads are delivered back to the ad server, all relevant data from the creation process is forwarded to the learning and optimization component for analysis of ad performance in comparison to similar ads (see chapter 6). 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.

4. 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 ad.

4.1. *State-of-the-art targeting strategies*

In order to achieve a high targeting accuracy we combine multiple targeting strategies, which can be categorized as content-based, behavior-based, and situation-based (figure 4). In order to illustrate the strategies and their contribution to the campaign selection, we assume a situation with a 25 year old male web user with recent session history showing dating sites and music stores, who is currently reading a review of the latest Yeah Yeah Yeahs album on a music blog. Additional information gathered includes the location of the user (Berlin), the time (Friday afternoon), the current weather at the user's location (rainy), and the result from an advertiser's web-service about an upcoming Yeah Yeah Yeahs concert in Berlin within the next two weeks.

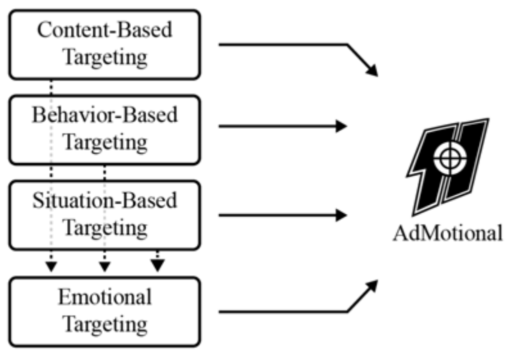

Figure 4: The targeting strategies.

Channel targeting is the first content-based strategy. It displays ads based on the page's overall topic. As opposed to contextual targeting, explained later on, the overall topic is determined manually and assigned statically to the whole site or certain areas. This coarse-grained classification is prone to human error and cannot take page-specific content into account.
Example: Since only the overall topic - in this case, music - is taken into consideration, the user receives an unspecific music-related ad, e.g. for an online music store.

Keyword targeting gives advertisers the opportunity to assign sets of keywords to their ads. To place the best-matching ad on a particular web page, its content is parsed for keywords and the ad that matches the most of them is shown. Due to polysemy and homonymy of keywords, this method is prone to inaccuracy and may even be harmful for an advertiser as it is, for example, not able to determine if the keyword is used in a negative context.
Example: Keyword targeting is able to detect the band name 'Yeah Yeah Yeahs' inside the page's content. Hence, an ad for the band's merchandising store may be displayed. Since the context is not determined (e.g. a negative review) , displaying the ad may be counterproductive or even harm the band's image.

Geo targeting utilizes location information about the user [Plummer *et al.* (2007)]. This information is generally provided as latitude and longitude values. The accuracy of these values depends on the location source, which may be a user's IP address, triangulation of nearby Wi-Fi access points, or a GPS satellite signal. The latter is especially relevant since most modern mobile devices ship with integrated GPS receivers. Location information enables the restriction of advertisements to certain areas, creating opportunities for local companies.
Example: Geo targeting detects that the user is located in Berlin; an ad for a recently opened coffee shop close to the user's location is displayed.

Contextual targeting focuses on information about the requested web page [Plummer *et al.* (2007)]. In addition to the utilization of static information contained in page meta data, the content is analyzed dynamically by applying classification methods: e.g. using support vector machines [Schölkopf and Smola (2002)], or clustering algorithms, in order to determine pages' semantic context. Besides providing means for accurate targeting, this contextual information may further serve as a basis for protecting brands from being advertised in negatively annotated contexts.
Example: Contextual targeting could determine that the article is a music review of the Yeah Yeah Yeahs album and that it is generally positive. Hence, an ad may be displayed which links to the album being offered in the label's online store.

Behavioral targeting utilizes historical information of the users' browsing behavior [Yan *et al.* (2009)][De Bock and Van den Poel (2010)]. Based on the history of visited websites, users' areas of interest as well as socio-demographic parameters such as age or gender are extrapolated by utilizing statistical and heuristic procedures. A study based on a survey of twelve ad networks reveals that, first, advertising rates are significantly higher for behavioral targeting as advertisers are willing to pay more if their ad is supposed to better match the consumer's interests, and, second, behavioral targeting is more successful than standard advertising, since more relevant ads lead to a higher conversion rate [Beales (2010)].
Example: Analysis suggests, with a high confidence, that the user is about twenty-five years old, male and is interested in dating and music. Thus an ad for a music-related social online community is displayed.

Situational targeting evaluates information about the users' current situation and aims to extend geo targeting by evaluating not only the location but also information such as time, weather, and cultural events taking place nearby the user's location.
Example: The user is located in Berlin on a rainy afternoon in the fall. This situation leads to an ad for a nearby movie theater.

Combining the above techniques to further enhance the existing targeting strategies, we propose the concept of *emotional targeting*.

4.2. *Emotional targeting*

Actions and decisions as made by persons are significantly affected by their current emotional (or sentic[a]) state. Due to this well-recognized fact, the advertising industry, for example, exploits these relations to convey positively annotated stimuli together with their advertising messages. Such annotations traditionally consist of specific themes, images, colors, and text elements. We attempt to reliably extrapolate the users' emotional states through an analysis of available parameters. In the ad customization step, these emotional states – together with the underlying emotional model – are used to generate emotion-aware personalized ads.

4.2.1. Emotional model

It should be noted that currently no commonly used formal definition of emotions has been agreed on. While this may be partly due to lack of consensus within the topic, as suggested by Schmidt-Atzert (1996), our further discussions in this paper shall be based on the following definition of Meyer *et al.* (2001), who defines emotions as

- current mental states;
- having a certain quality, intensity and duration;
- object oriented;
- showing in the reaction triad of [Lazarus *et al.* (2001)] (subjective, behavioral and physiological aspects).

For the classification of emotions we employ the semantic differential model as described by [Osgood *et al.* (1957)] with three input dimensions: evaluation, potency and activity. Evaluation describes whether the emotion is positive or negative. Potency describes whether one feels strong or weak while experiencing the emotion. Activity refers to the amount of arousal associated with the emotion. An example for a very active emotion is excitement in contrast to satisfaction.
We chose the model mainly due to available data for text analysis [Heise (2004)]. While this model is not very specific in terms of targeting, it is sufficiently simple as to facilitate an accurate analysis.

4.2.2. Analysis

For determining the emotional values of web sites, it has to be distinguished between multi-topic sites (e.g. a news site) and single-topic sites (e.g. a dating site). Thus we first categorize sites into one of the two categories, which lead to varying determination processes that are described below. Further, the content of the site is processed similar to [Kazienko and Adamski (2004)]: A crawler indexes all pages of the sites and extracts relevant terms. During this extraction, terms without emotional significance are filtered. But rather than using this basis for clustering sites into thematic groups, we analyze the content to derive the emotional value. In contrast to existing approaches for assigning

[a] Used by [Clynes (1978)] to avoid negative connotations with the word emotion ([Picard (1995)])

terms with static emotional values, we include users' profiles in our emotional dictionary. Among other factors, this allows for a rich spectralanalysis of gender and thus increases the personalization.
For a site s_h with a single topic (having pc_h pages) we calculate the emotional value as follows:

$$e_{hk}^{su} = \frac{e}{pc_h} * \sum_{i=1}^{pc_h} e_{ik}^{pu} * w_{ih}^{ps}. \tag{1}$$

When analyzing a site with multiple topics, each page is treated individually and this first step is therefore skipped. The weight of the page on the site w_{ih}^{ps} is calculated according to the so-called PageRank algorithm described by [Page (1997)]. The emotional value of a single page (containing tc terms) e_{ik}^{pu} is calculated as follows:

$$e_{ik}^{pu} = \frac{1}{tc_1} * \sum_{j=1}^{tc_1} e_{jk}^{tu} * w_{jk}^{tp} . \tag{2}$$

Where w_{jk}^{tp} is the weight of the term t_j on the page p_k (see [Kazienko and Adamski (2004)] equation (1)).
The value e_{jk}^{tu} is extracted from an *emotional dictionary*, which maps terms and profile information to three-dimensional emotional vectors. As an initial baseline, we use data supplied by [Heise (2004)], which linearly maps terms to emotional values. However, in contrast to existing linear mappings, the proposed emotional dictionary for AdMotional results in higher accuracy through the inclusion of profile information, but requires a fairly large *a priori* emotional vocabulary to provide the required overall accuracy.
Example: While the words "sports" or "soccer" on a web page may well induce different or even contrary emotions depending on the predicted user's profile, e.g. male 20-25 years vs. female 60+ years, state-of-the-art linear approaches are unable to reflect and make use of this fact for targeting purposes.

While first tests on selected web pages show promising results, the main issue remains the lack of sufficient data for extensive analysis. As AdMotional defines the emotional dictionary as a general framework for representing user-profile-specific emotional values of words (or concepts), available or future results from various studies can be later-on incrementally added to this dictionary, while words without different profile-specific entries are still mapped to their linear emotional value, as supplied by [Heise (2004)].
The final result is a three-dimensional vector describing the emotional value of either a site or a single page, depending on the user's profile. This value is then used to determine the user's current emotional situation. For this, we analyze the time series [Hamilton (1994)] of the user's history as follows:

$$e_k^u = \sum_{x=0}^{sc} e_{kx}^{us} * w_{kx}^{us}. \tag{3}$$

Where w_{kx}^{us} is dependent on the time since the visit. Additional to the time series analysis, we use an authored rule-based evaluation of the user's history.
Example: The emotional value of a visit to a dating site is dependent on the time spent on the site. A long visit suggests many interactions and therefore should be rated higher than

a short visit, which, in contrast, may result from an empty inbox and no chats, leading to a more negative emotion. This final emotional value is then used in the ad selection and creation components.

Further research is still necessary for the evaluation of entire sentences, since the emotional value can differ significantly, especially when negations are used. Furthermore, the emotional dictionary needs to be extended to increase the profile information available and thus improve the accuracy of the emotional values derived.

4.3. *System integration*

To efficiently and flexibly integrate the multitude of strategies, we represent the targeting information as parameters. These parameters can be separated in these two categories:

(1) Parameters regarding the user, his behavior, and current situation. These include age, gender, the history of visited websites, and location-based information.

(2) Parameters regarding the website the ad is supposed to be displayed on, including the URL and topic of the website, as well as the keywords.

The source of these parameters may be internal or external. The inclusion of external data sources allows for later extension and utilizing the intelligence of third-party tools or suppliers (e.g. specialized on behavioral targeting).

These external data sources include services for content classification and behavioral targeting, as well as location-based services for geo targeting or weather information. The integration of external sources is realized via web services.

This approach enables the inclusion of as much information as possible in the selection and personalization increasing the accuracy of these processes and also reduces the implementation necessary within the system.

4.4. *Selection process*

The campaign selection process is based on the described parameters spanning an n-dimensional space where both the request and the campaigns are placed.

The campaigns are placed based on targeting rules composed by the advertiser on a set of selected parameters in order to define their target groups as well as target publisher websites. For instance, advertisers may specify their campaigns to be exclusively targeted at customers aged 20 or over. In this case, the resulting campaign effectively represents an (n-1)-dimensional subspace of our parameter space.

The request is placed based on the acquired parameters inside the search space where the m best-matching campaigns are selected.

The remaining campaigns are evaluated using weighted rules, with individual rule weights being subject to optimization within the learning component (see chapter 6). For this purpose, we currently use a score-based approach resting on the inference of targeting rules incorporating a derivation rule compiler from the Mandarax[b] project and applying the assigned weights. This approach further allows for simplified, asynchronous interfacing between optimization and selection/generation components, particularly if the

[b] Relational Programming for the JVM. http://code.google.com/p/mandarax/

execution of standardized knowledge base representations (e.g. RuleML [Wagner *et al.* (2004)]) in the selection process proves to provide the required performance.
Performance is a crucial factor, particularly in the campaign selection process. In order to save execution time, we cache the first step of the campaign selection based on a hash of the request parameters. For subsequent requests, the hash value is calculated initially and if there has been a similar request the previously calculated *m* campaigns are evaluated further.

4.5. *Issues*

In this section, some current issues in the context of online advertising, as well as of the implementation, are addressed.

4.5.1. Economic constraints

As the main goal is to serve users with personalized ads, the campaign selection process currently does not take economic criteria, such as contractual agreements among advertisers, publishers, and the ad server network into account. As such, constraints cannot be neglected in a production environment; they are currently addressed in a preceding step within the ad server, which supplies the system with a preselected set of appropriate campaigns.
While such a two-phase selection mechanism effectively reduces the targeting accuracy, the necessity of actively fulfilling contractual obligations forced the introduction of this limitation. It would be ideal if such economic indicators could be considered directly within the campaign selection, but since this was not the primary research goal, this may only be addressed in future developments.

4.5.2. Accuracy

Accuracy is the most important issue concerning the campaign selection component and the advertising industry in general. All targeting methods which are used in the system are prone to inaccuracy. Contextual classification or geo targeting may not always return correct results. For example, location detection based on the user's IP address may detect the wrong location if a proxy server is used. For behavioral, situational, and emotional targeting, the threat of inaccuracies is even higher as these methods are based on predictions. Since targeting resting on inaccurate or incorrect parameters may even be less effective than displaying an ad without any user specific targeting the issue of accuracy needs to be addressed. The learning and optimization component provides the means to detect ineffective targeting rules, which may be due to low significance or insufficient accuracy of their parameters.

4.5.3. Privacy

User privacy is another important issue when it comes to personalized online advertisements. In the EU, for example, legal regulations have been made that bar ad networks from storing personal information without users’ consent [European Parliament

(2002)]. Studies have shown that these regulations have had a significant impact on the overall performance of campaigns that relied on mere behavioral targeting as users tend to not agree to the utilization of their personal data [Goldfarb and Tucker (2011)]. In order to avoid an impairment of targeting accuracy without disregarding existing laws and regulations, personal data is stored only anonymizingit, e.g. by deleting information such as IP addresses or applying a hashing algorithm on the confidential parts. [Toubiana *et al.* (2010)] proposes an architecture that sources behavioral profiling and targeting out to the user's browser and thus eliminates the need for storing personal data on the server side. Although this architecture does not seem sufficient to fully carry a targeting as complex as AdMotional's into execution, it might be used as a basis to let single privacy-critical targeting modules take part in the user's browser. By combining targeting methods, AdMotional avoids relying on personal data exclusively and is able to deliver targeted ads even when there is only sparse personal data.

4.5.4. Relevance

It has yet to be determined if all available targeting criteria are relevant for our targeting and hence for the overall campaign selection process. In order to keep the number of parameters and, correspondingly, the size of requests, as small as possible, identifying and rejecting unnecessary criteria where applicable is another responsibility of the learning and optimization component. For this purpose, the component evaluates past requests with their respective parameters and the success of the displayed ad and, if there is no correlation between success and the parameter, the parameter can be discarded.

5. Dynamic Ad Creation

AdMotional's dynamic ad creation component is a novelty in online advertisement. It was introduced to allow for a high degree of ad customisation and – eventually – truly personalised ads. While we currently only use the dynamic ad creation features for ad customisation, it could be employed equally well for related online advertising tasks such as Retargeting. [Helft and Vega (2010), Lambrecht and Tucker (2011)].

Traditional banner ad processes require advertisers to provide advertising networks with ready-to-use banner instances, accompanied by targeting information supporting the selection process. Within these processes, customization was limited to simple text overlays, possibly guided by information about users' geo location. Besides a resulting – mostly poor – visual design due to the inherent split in the ad creation process, dynamic capabilities of such ads are severely limited. Instead of depending on existing, pre-rendered ads, AdMotional promotes live ad composition at request-time, based on user-specific customizations. The customization points used in this approach form ad *skeletons* consisting of rules and sets of alternative data instances or value generation strategies. Such instances can be any rich media content, and include pixel images, vector shapes, text templates and various attributes, such as color or relative position within the ad.

Yet, in order to introduce this dynamic approach in the market, significant changes within advertisers' processes are initially required, as we shall briefly outline in the following section.

5.1. *Changing the role of the advertiser*

The process underlying the facilitation of custom ad creation as outlined above requires comprehensive changes within the processes of advertisers. Using AdMotional, advertisers need to provide the network with a template consisting of set of data instances and a rule-set for the ad-medium generation, rather than with pre-rendered static ads. The rule-set encapsulates the combinatorial complexity of the parameters customizing the ad for specific user groups.

The overall ad layout is defined using XML. It mainly consists of two parts: a list of possible elements to be placed on the ad, and a rule set defining how and under which conditions these elements are applied to the template. The elements defined in the list further contain all attributes necessary for specific appearance (e.g. color, font) and transformations (e.g. positioning, scaling) within ads. The rules in the defined set consist of an element selector identifying one or more ad elements to be customized, a customization trigger (condition), and one or more customization commands detailing how particular ad elements need to be adjusted. As, among other things, variables and nested structures are fully supported, the DSL used for the three rule parts has sufficient expressive power to also allow for complex real-world ad-generation scenarios.

The ad templates are compiled and cached when first added to the system to make them available for high-performance processing when a request is received. We provide an XSD file[c] that clearly defines available ad elements, their properties, and the technical commands to apply them on an ad. Although this file gets more complex with increasing numbers of element and rule attributes, it clearly defines an API for both the advertisers and the generation component. Rather than using the complex file that defines the entire functionality, XSLT can be used for creating XSD files representing common ad designs, minimizing the effort for marketing agencies by defining the required transformation to an AdMotional-compatible description.

While this concept demands an increased initial design effort from advertisers compared to creating static ads, 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. However, using this approach we see a potential for conflict as advertisers would have to render parts of their control over ads' appearances to the AdMotional system to achieve a high degree of personalization. A sophisticated implementation and a successful introduction phase for this new approach of dynamic advertisements is required to introduce this technique to the market and to achieve acceptance among advertisers.

[c] Available at http://www.admotional.org/schema/acdc/banner-2.0-system.xsd

5.2. *The generation process*

AdMotional provides a common interface for pluggable ad generation modules creating customized ads. Currently three different implementations are available: creating PNG pixel images, HTML code, and Flash movies. Being a reference implementation, these modules are currently limited to the production of static, image-like banners, rather than producing video-style multi-media content. However, as has been shown by Bayles (2002), ad success of animated banners so far does not significantly outperform static image banners. We thus don't see the limited capabilities of our ad creation modules as a drawback regarding the overall validity of our (proof of concept) results.

The generation process consists of several steps which are executed within sub components. As shown in figure 5, these components provide performance and feature optimization points, thereby promote extensibility of the overall system.

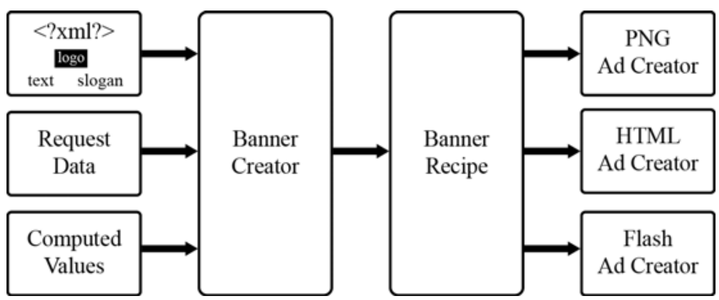

Figure 5: The Ad Generation Process.

Initially, the entire set of request parameters has been computed and is made available to the banner-creation module. A specific campaign and a related banner have been selected and it is now customized based on three categories of input data: An XML document containing possible banner elements, user-defined variable values, and a rule- set to define how the banner should be created. The current request data is provided to evaluate the rules with the user-specific values. Additionally, access to automatically computed variable values is granted. This is an extremely powerful capability, since it can be extended using techniques identified by ongoing research. After evaluating the studies we conducted with our project partners (see chapter 7), we implemented a provider to choose a background color that matches the user's current emotion. The system is kept generic so that new findings such as [Danaher (2003)] can be integrated at any time.

A banner creator instance processes this input data by executing the rules, creating a set of customized and positioned banner elements that make up the final advertisement. While some implementations may focus on a direct high-performance generation of the medium, our implementation creates an interim representation of the customized ad, a "banner recipe". This allows us to create and evaluate several ad creators that process the temporary banner elements to a final ad instance that can be presented in the user's browser.

Example: In the case of the sample user checking out the latest reviews about the Yeah Yeah Yeahs, the same banner template may result in the ad shown on the left in figure 6, while it may lead to the banner on the right when executed with slightly different profile data, e.g. age 60+ and surfing on a classical music web site:

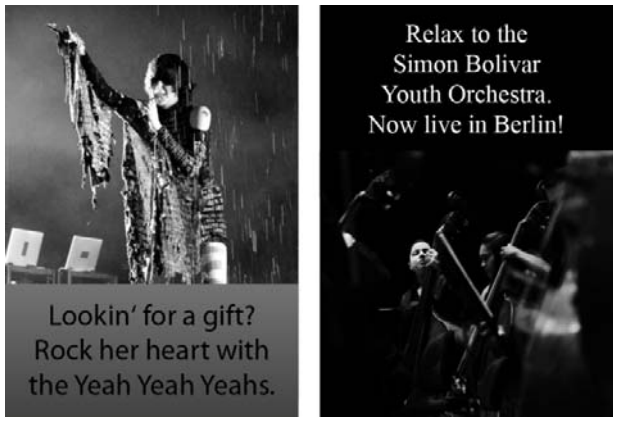

Figure 6: An example of two ads generated by the same template.

5.3. *Available implementations*

We currently provide three ad-generating modules, based on two different creation techniques. The first two modules use graphic libraries, GraphicsMagick[d] and Java Advanced Imaging API[e], to produce raster 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. The generation of HTML and CSS code also requires less computational power than the image file generation.

This approach also faces several challenges. First of all, the appearance of the ad medium is subject to many factors. The medium has to be defined in a sandbox-like environment so that the publisher's website cannot affect the ad's layout and vice versa. Since the rendering of HTML elements is done on the client side, they are subject to all well-known cross-browser issues that are considered the second-most time-consuming issue in web application development, after ensuring security [Rode *et al.* (2002)]. Research to overcome the problems in this area is still being conducted (e.g., [Choudhary *et. al.*

[d] http://www.graphicsmagick.org
[e] http://java.sun.com/javase/technologies/desktop/media/jai/

(2010)], [Mesbah and Prasad (2011)]); at this point only very basic features are supported by all browsers and processed in the same way.

6. Learning Component

To continuously improve the performance of individually targeted ads, a learning and optimization component has also been developed where two performance types need to be differentiated depending on advertisers' intentions. While so called "performance campaigns" aim for consumers' immediate responses (i.e. clicks on the ad media), so called "branding campaigns" intend to strengthen advertisers' brands, and are expected to result in increased business turnover in the long run [Wang *et. al.* (2002)].

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 in order to allow extensive analyses. Results are fed back to the system via web services. The system has two optimization targets: improving the campaign selection and improving the visual representation:

6.1. *Optimizing campaign selection*

For improving the campaign selection, we use success feedback from the ad server. For every successful display situation, we increase the weight of the used targeting rules. The approach is similar to reinforcement learning [Sutton and Barto (1998)]. If parameters are set in the situation that currently have no corresponding targeting rule, these are added to the campaign. This enables the discovery of new relevant targeting groups. For this, a small percentage of requests are randomly selected for modification during campaign selection. While this obviously reduces the overall amount of specifically targeted ads, the knowledge gain is expected to yield higher click-rates in the long run. Additionally, campaign settings allow exclusion from this experimental optimization, in case advertisers want their specific targeting rules to be applied without any change.

6.2. *Optimizing ad creation*

The goal of the visual optimization is to find general statements about a better (more successful) ad design. 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 customization point dimensions with the greatest differences to the center of the high performance ad cluster. It is then possible to define individual ad modification rules, suggesting automatic adjustments to particular ad criteria during the generation process, based on experience with the current ad and those that are similar. As these candidate criteria are hardly independent (e.g. foreground vs. background color), a set of design constraints are defined, resulting in criteria groups – rather than individual criteria – to be jointly adjusted.

Example: Revisiting the sample ad that had been generated for a young, emotionally activated Yeah Yeah Yeahs fan in Berlin, it can be assumed that, over time, similar requests have to be answered by the system. Unless blocked by the advertiser through specific campaign settings, the system may start modifying selected design parameters within the banner recipe in order to improve the ad performance. This may lead to a layout modification (e.g. text above image instead of image above text) and a change in background color for the text part. Figure 7 shows an automatically modified version of the ad shown on the left in figure 6.

Figure 7: A modified ad with layout and background color modification.

The generated rules are represented using RuleML [Wagner *et. al.* (2004)] before being serialized and injected into the creation modules. This approach allows for a direct object notation, while not only ensuring a smooth integration with other components, but also allowing presentation as simplified knowledge for inspection by advertisers.

Ad modification rules as described above each address one particular ad. However, a significant number of rules must be anticipated 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.

Example: For the sample Yeah Yeah Yeahs campaign, such an abstract ad modification rule may state (but has not yet been observed, however) that in the context of dating and for an emotionally activated user of age 25, ads with pink background color generally perform better compared to those using a standard blue background.

7. Preliminary Results

At this stage of the project, preliminary results about internal performance tests of the image creation modules and the entire request processing within the system are available. Live-tests that incorporate actual user-data and feedback for the learning component are currently being prepared. In addition to the test results, research results of user studies determining the relationship between emotions, colors and image types are presented.

7.1. *System performance test*

The AdMotional component was tested for performance of (a) the entire system in terms of response times and (b) for the comparison of the described ad creators.
A user was simulated by a multithreaded test environment, sending 20 generated random requests to the system. Aiming for a possibly realistic scenario, the requests were sent in a varying period of 1 to 20 seconds. The environment allowed the generation of user-specific histories using partly related target websites in order to build up a user story.

	GraphicsMagick Ad Creator		HTML5 Ad Creator	
Users	Ad Creator	Request Total	Ad Creator	Request Total
1	7 (SD = 1)	17 (SD = 3)	3 (SD = 3)	12 (SD = 2)
100	11 (SD = 3)	22 (SD = 6)	8 (SD = 6)	17 (SD = 4)
10000	10 (SD = 3)	24 (SD = 6)	8 (SD = 6)	17 (SD = 4)

Table 1: Preliminary performance test results.

These results are based on a campaign selection using a small number of campaigns (n=10), each containing two or three banner templates. This is a realistic value for the prospective environment where the ad server limits the number of candidate campaigns.
A common performance goal for current ad servers is 320ms per request [Internet World Business (2010)], including all necessary steps from initial Javascript execution to the delivery of the actual ad medium. Aiming for this performance goal, first integration tests with the ad server of our research partners revealed that the timeframe for the AdMotional component must not exceed 20ms. Table 1 shows the comparison of the system performance for 1, 100 and 10000 users. The processing time and the standard derivations are shown in milliseconds. Depending on the ad creator implementation, AdMotional can serve dynamically personalized ads within the commonly accepted timeframe. The system is able to scale to a realistic load.

7.2. *The relationship between colors, images and emotions*

User studies were conducted in order to implement a number of providers for automatically computed values that can be used for designing ads. In a first step, questions about generated ads had to be answered in order to determine color preferences in general. A surprising result is the little preference for the color red compared to the strong positive effect of yellow. Visual effects such as gradient backgrounds revealed that their usage does not significantly increase the user's interest. Figure 8 shows the cumulated preference for selected colors (in percentage, multiple choices allowed).

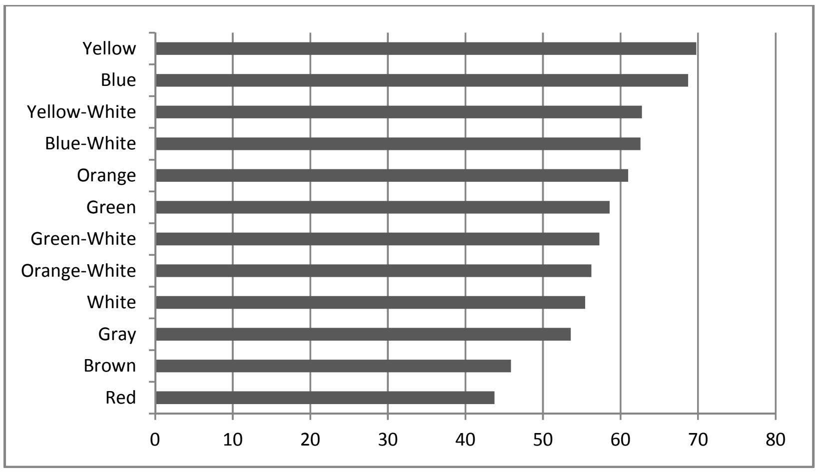

Figure 8: Empirical preferences for selected colors.

The second part of the user study showed that the choice of colors only has very little influence on the user if other visuals such as images are displayed in the ad. In that case, the interplay between image and color has to be investigated as the color only supports the message of other visuals.

The third and final finding is based on a categorization of images that were used in advertisements. It was proven that depending on the ad topic, different image types have significant influence on the campaign performance. While most results seem obvious (for example, visuals of humans are positively connected with family-related ads), the study also revealed interesting preferences such as landscape visuals for entertainment and lifestyle ads – rather than actual product images.

The results of all user studies were combined, forming a matrix showing common color and image type preferences of users, depending on the content topic. AdMotional's learning component is able to track the dependencies of user emotions and website contexts, and based on the presented research results, optimizing the selection and ad designing processes over time. This also offers the possibility of conducting more user studies on the relationship between colors/visuals and emotions in the near future.

8. Conclusion

Considering today's environment of online marketing and related research, personalization appears to be one of the most important and promising approaches to optimizing campaign performance. Selecting the most appropriate campaign and, in addition, dynamically creating customized ads, depends on rich, extensive databases that need to be designed to allow for efficient evaluation. It is for these reasons that the AdMotional system was developed and presented: it combines existing targeting strategies with a new dimension of emotional targeting. It not only utilizes established targeting strategies for its selection process, but also exploits their potential in the personalization of template-based ads, enhancing personalization. Additionally, it employs a powerful 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 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 lead to the necessity of re-evaluating options for dynamic ad creation. Recent research on intelligent ad resizing [Badali *et. al.* (2010)] may open new perspectives in dynamic ad creation, as ad zones on web pages may then be automatically filled with best-fitting re-scaled ads tailored to the specific user's state and situation, as described in the various targeting aspects discussed in this paper. Directions for ongoing efforts include the evaluation and fine-tuning of our feedback-based ad optimization techniques in long-term real-life tests.

Lastly, given that AdMotional has developed along a modular and highly flexible architecture, its potential for use in other areas of online advertising, such as Retargeting [Tucker (2010)], is an avenue worthy of investigation.

Acknowledgements

The authors would like to thank the entire AdMotional team at our industry partner The AdInsiders GmbH in Bocholt, at FH Gelsenkirchen and at the Life & Brain institute at the University of Bonn. Special thanks to Jens Fendler for his contributions to the project, to an earlier version of this paper and for proofreading parts of it. The anonymous reviewers also provided valuable remarks for further improving the paper. Thanks also to Neil Tallim for his help polishing the final version.

The AdMotional project is being supported by a grant from the European Union within the ZIEL2 program 2007-2013 (EFRE), IKT.NRW (FKZ 29 00 341 02).

EUROPEAN UNION
Investment in our future
European regional development fund

References

Badali, A.; Aarabi, P.; Appel, R. (2010). Intelligent Ad Resizing. In Proceedings 19th International Conference on World Wide Web (WWW'10), Raleigh, USA 2010, pp. 1053-1054

Bayles, M. (2002). Designing online banner advertisements: should we animate?. In Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves (CHI '02). ACM, New York, NY, USA, pp. 363-366.

Beales, H. (2010). The Value of Behavioral Targeting. Network Advertising Initiative (NAI) http://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf

Choudhary, S.R.; Versee, H.; Orso, A. (2010). WEBDIFF: Automated identification of cross-browser issues in web applications. In Proceedings IEEE International Conference on Software Maintenance (ICSM), Timisoara, Romania.

Clynes, D. M. (1978). Sentics - The Touch Of Emotions. Anchor Press/Doubleday, New York, USA.

Danaher, P. J. (2003). "Factors Affecting Online Advertising Recall: A Study of Students" Journal of Advertising Research 43.3 (2003): 252-267.

De Bock, K.; 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, Netherlands.

European Parliament (2002). Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector. Official Journal L 201, 31 July 2002, pp. 37-47.

Evans. D. (2008). Economics of the Online Advertising Industry. Review of Network Economics, Vol. 7, No. 3, September 2008

Goldfarb, A.; Tucker, C. (2011). Privacy regulation and online advertising. In: Management Science, Vol. 57, Jan. 2011, pp. 57-71.

Hamilton, J. D. (1994). Time Series Analysis. Princeton University Press, USA.

Heise, D. R. (2004). Affect Control Theory, Indiana University, Bloomington, Indiana: http://www.indiana.edu/~socpsy/ACT/.

Helft, M.; Vega, T. (2010). Retargeting Ads Follows Surfers to Other Sites. New York Times, August 29, 2010.

Holbrook, B.; Batra, R. (1987). Assessing the Role of Emotions as Mediators of Consumer Responses to Advertising. Journal of Consumer Research, Vol. 14, No. 3, pp. 404-420

Internet World Business (2010). Adserver Benchmark Vol. 13 (21.06.2010). Online edition: http://www.internetworld.de/Heftarchiv/2010/Ausgabe-13-2010/Adserver-Benchmark

Kazienko, P.; Adamski, M. (2004). Personalized Web Advertising Method. In: W. Nejdl and P. De Bra (Eds.), AH 2004, LNCS 3137, pp. 146–155. Springer, Berlin, Germany.

Lambrecht, A.; Tucker, C. (2011). When does Retargeting Work? Timing Information Specificity. Social Science Research Network http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1795105

Lazarus, R. S.; Averill, J. R.; Opton, E. M. (1970). Toward a cognitive theory of emotion. In: M. B. Arnold (Ed.), Feelings and emotions. pp. 207-232. Academic Press, New York, USA.

Liu, W.; Zhong, S.; Chaudhary, M.; Kapur, S. (2007). Online Advertisement Campaign Optimization. In Proceedings IEEE International Conference on Service Operations and Logistics, and Informatics (SOLA2007), Philadelphia, PA (USA) 2007

Mesbah, A.; Prasad, M. R. (2011). Automated cross-browser compatibility testing. In Proceedings International Conference on Software Engineering (ICSE '11). Waikiki, Honolulu, HI, USA.

Meyer, W.-U.; Schützwohl, A.; Reisenzein, R. (2001). Einführung in die Emotionspsychologie, pp. 22-40. Huber, Bern, Switzerland.

Osgood, C.E.; Suci, G.; Tannenbaum, P. (1957). The measurement of meaning. University of Illinois Press, Urbana, IL, USA.

OVK (2011). OVK Online-Report 2011/02. Bundesverband Digitale Wirtschaft (BVDW) e.V., Düsseldorf, Germany

Page, L. (1997). Method for node ranking in a linked database. US Patent No. 6,285,999. USA.

Picard, R.W. (1995). Affective Computing, MIT Press, Cambridge, MA, USA.

Plummer, J.; Rappaport, S.; Hall, T.; Barocci, R. (2007). The Online Advertising Playbook. John Wiley & Sons, New York, USA.

Rode, J.; Rosson, M. B.; Perez-Quinones, M. A. (2002). The Challenges of Web Engineering and Requirements for Better Tool Support. Technical Report TR-05-01, Computer Science, Virginia Tech. http://eprints.cs.vt.edu/archive/00000649/

Schmidt-Atzert, L. (1996). Lehrbuch der Emotionspsychologie. Kohlhammer. Stuttgart, Germany.

Schölkopf, B.; Smola, A. (2002). Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond (Adaptive Computation and Machine Learning), MIT Press, Cambridge, USA

Shanahan, J.; Kurra, G. (2011). Digital Advertising: An Information Scientist's Perspective. In: Melucci, M.; Baeza-Yates, R. (eds.): Advanced Topics in Information Retrieval (The Information Retrieval Series), vol. 33, Springer, pp. 209-237

Sutton, R.; Barto A. (1998). Reinforcement Learning: An Introduction. MIT Press, Cambridge, MA, USA.

Toubiana, V.; Narayanan, A.; Boneh, D.; Nissenbaum, H.; Barocas, S. (2010). Adnostic: Privacy preserving targeted advertising. In 17th Annual Network & Distributed System Security Symposium. San Diego, CA, USA.

Tucker, C. (2010). The Economics Value of Online Customer Data. In Proceedings Joint WPISP-WPIE Roundtable 'The Economics of personal Data and Privacy: 30 Years after the OECD Privacy Guidelines', OECD.

Wagner, G.; Antoniou, G.; Tabet, S.; Boley, H. (2004). The Abstract Syntax of RuleML – Towards a General Web Rule Language Framework. In Proceedings IEEE/WIC/ACM Int'l. Conf. on Web Intelligence (WI'04). IEEE, Washington D.C, USA.

Wang, C.; Zhang, P.; Choi, R.; D'Eredita, M. (2002). Understanding Consumers Attitude Toward Advertising. In Proc. Eighth Americas Conference on Information Systems, pp. 1143-1148

Yan, J.; Liu, N.; Wang, G.; Zhang, W.; Jiang, Y.; Chen, Z. (2009). How much can behavioral targeting help online advertising? In Proc. Intern. World Wide Web Conference (WWW) 2009, pp. 261-267, Madrid, Spain.

Zeff, R. L.; Aronson, B. (1999). Advertising on the Internet. John Wiley, New York, USA.