

Activity-based Advertising: Techniques and Challenges

Kurt Partridge and Bo Begole

Palo Alto Research Center, Inc., 3333 Coyote Hill Road,
Palo Alto, CA 94304, USA
{kurt,bo}@parc.com

Abstract. Although information technology has transformed advertising dramatically in the past few years, the greatest changes may lie in the future. Researchers in Pervasive and Ubiquitous Computing have spent the last few years prototyping and studying context-aware systems. These systems use sensors on mobile phones and in the infrastructure to build digital models of the world. This technology has not yet achieved commercial success in advertising applications, but has great potential. In this paper we describe three types of activity-based advertising: 1) inferring general interest categories, 2) adapting advertising to the current context, and 3) using contextual histories to predict the future and adapt present advertising to these predictions.

Keywords: Advertising, context-awareness, activity recognition, privacy.

1 Introduction

The purpose of this paper is to clarify the mobile advertising opportunities enabled by a better understanding of human activity. Advertising is, of course, a crucial technology because it finances many information and communications technology products and services. However, it is our opinion that it is understudied within academic institutions because there is little grant money available for its study and because most people experience advertising as consumers, who often find it an annoyance. This is unfortunate, not only because a healthy advertising ecosystem is important not only for finance reasons, but also because research in advertising research can lead to a better experience for end-users and a more efficient system for advertisers.

We view “Pervasive Advertising” as encompassing digital out-of-home advertising and mobile advertising. Both these areas have grown rapidly in the past year. In 2008, out-of-home advertising sales were up 11% according to PQ Media. Mobile advertising revenue rose much faster—an estimated 60%, according to Gartner. These figures compare well to overall global advertising revenue growth of 2.6% (GroupM).

These technologies are distinguished from each other by their delivery channel—out-of-home advertising is displayed on a digital sign in a fixed location whereas mobile advertising is displayed on a mobile phone. But both share the characteristic

with online advertising of being dynamic and adaptable to physical situations. For example, some digital signs can be configured to change with time of day. More recently, digital signs have been outfitted with cameras that can detect demographics such as race, gender, and age group. Although the results are inaccurate, they only need be good enough to make advertisements that are adapted more effectively than those that are not.

Mobile advertisements are also adaptable. Using location-based advertising, advertisements can be delivered when the user enters a region, for example, to deliver coupons for a nearby store.

We believe that through better analysis of contextual data, even better targeting will be possible than location-based advertising offers today. An early example of the kind of analysis that can be done can be found in online behavioral targeting. Such systems consider not only the current content when placing an advertisement, but also the user's browsing *history*. This approach is particularly effective if a page is not commercially oriented, because then advertising can address a consumer's previously demonstrated interests. These approaches must be performed with the consumer's consent, a topic we address more deeply in Section 3.

2 Types of Activity-based Advertising

How, then, can histories of *contextual* data help target advertising more effectively? We propose three types of activity-based advertising.

2.1 Inferring Interest Categories

This type is closest to online behavioral targeting. A user's contextual history, their location in particular, is analyzed to determine the user's interests. For example, if a user frequently goes to Chinese restaurants, then a system may determine that the user prefers Chinese food to other types of cuisine. This fact may not be easily determined through online behavioral targeting, but would be useful for certain advertisers, such as the owners of a newly-opened local restaurant. An example of this approach is Hristova and O'Hare's Ad-e system [1].

Making such a system effective is technically challenging today. A key obstacle is the difficulty of determining the venues that a user visits. GPS resolution is not accurate enough, and does not perform well indoors. Azizyan and Choudhury [2] have investigated ways to improve the accuracy of venue inference by incorporating data from other sensors such as accelerometers, microphones, and light sensors, however this technology has not yet been demonstrated on a wide scale.

Another challenge is that knowing that a person frequently visits a place or a type of place is insufficient. Someone who spends a lot of time in Chinese restaurants might enjoy the food, or they might work there. Hence, any system that uses context to determine interests should have a mechanism to infer and factor out employment. In past work, we have shown that the combination of a person's current location and an inference of whether they are working at that location are generally sufficient for detecting that person's activity outside home environments [3].

2.2 Adapting to Present Context

Another way to use contextual data is to adapt advertising to the current context. Beyond simple location-based advertising, a key component of context is the social situation. Social groups can be detected through a combination of methods, including proximity to others [4]. When a user is in a group, an advertisement that is of interest to many people in the group is more likely to be discussed.

Being in a group can also change how receptive a person is to an interruption. This is important for “push” advertising like SMS messages. When actively engaged in a group, a user might not respond to an interruptive advertisement. However, under other circumstances (say during a conversational lull), the user might be more likely to respond. If these situations can be detected, they could improve the effectiveness of programs in which the consumer has granted permission for push advertising. Our own past work has shown that both timing and content are helpful in improving advertising effectiveness [5].

2.3 Predicting Future Events

Finally, contextual data might be useful for predicting events and situations that would increase the value of particular kinds of advertising in the present moment. For example, a person who regularly eats dinner at different restaurants would have a transportation pattern showing geographical divergence after leaving work. Hence, restaurant advertisements are likely to be valuable before a divergence point, but not afterward [6]. Recent research in predicting future actions from transportation patterns [7] is likely to be useful for these kinds of scenarios.

3 Challenges

Activity-based Advertising faces several challenges. First, the way that advertisers specify and bid on advertisements must be rethought. While keywords have been effective for matching up advertisements with appropriate content, they are not expressive enough to cover the kinds of contextual situations discussed above. However, ad specification and bidding must be simple enough for advertisers to understand how it works, and what adjustments to make to maximize effectiveness.

In addition, the extra complexity of processing contextual information must not overly affect the performance of the ad network placement system. Placement decisions must be made very quickly, and must scale up to millions of users. Portions of such decisions must be precomputed when possible, and the architecture of such a system is likely to be complex and require significant engineering to optimize.

Finally, user privacy must be respected. This means that users must be given control over their data and by default users should not be affected by a system unless they opt-in. Furthermore, parties collecting user data must take steps to make sure that it is securely stored. An additional concern is that someone wishing to access a user’s private data might pose as an advertiser. Working out the complete threat model and requirements for protection is an open research opportunity.

To some, it might seem highly unlikely that a consumer would agree to having so much data about them collected and used for advertising purposes. But consider, however, the popularity of today's online email services that display advertising related to the email contents. Consider also the United Kingdom mobile carrier Blyk, which has over 200,000 subscribers that receive free minutes and text messages in exchange for receiving targeted advertising.

Behaviorally targeted advertising has increased the effectiveness of online advertising in part because there is so much data from which to detect patterns and correlations. In addition, web usage is so prevalent that experimental trials may be conducted quickly. Today, mobile and pervasive technologies are not as widely used as web technologies. Data sets are sparse, and trials have been small. However, smart phones are quickly engulfing the more than 2 billion cell phone device market, bringing mobile data services to more people than the PC market can reach.

4 Activity-based Advertising 25 Years from Now

Over the prior 25 years we've seen the expansion of pervasive network technologies create a global information network through which millions of people interact daily. These interactions generate vast amounts of data from which to mine people's online behavior patterns. Over the coming 25 years, we will see the proliferation of pervasive technologies expand to a point where vast amounts of physical situation and behavior data are available for detecting patterns, allowing for much greater accuracy in the modeling of a person's interests and preferences. We will see the ability for systems to target information at what people *do*, not just what they click on.

References

1. Hristova, N, O'Hare, G.M.P.: Ad-me: Wireless Advertising Adapted to the User Location, Device and Emotions. In Proc. 37th Hawaii Int. Conf. on System Sciences (2004)
2. Azizyan, M., Choudhury, R. R.: SurroundSense: Mobile Phone Localization Using Ambient Sound and Light. In ACM MC2R (2009)
3. Partridge, K., and Golle, P.: On Using Existing Time-Use Study Data for Ubiquitous Computing Applications. In UbiComp (2008)
4. Eagle, N. and Pentland, A.: Reality mining: Sensing complex social systems. Journal of Personal and Ubiquitous Computing (2005)
5. Sala, M., Partridge, K., Jacobson, L., Begole, B.: An Exploration into Activity-Informed Physical Advertising using PEST. In Pervasive (2007)
6. Partridge, K., Greene, D., Begole, B.: Activity-Based Advertising: Auctioning Advertising in Context (2008)
7. Krumm, J., and Horvitz, E.: Predestination: Inferring Destinations from Partial Trajectories. In UbiComp (2006)