Building Persuasion Profiles in the Wild: Using Mobile Devices as Identifiers.

Maurits Kaptein
Eindhoven University of Technology / Philips Research
Den Dolech 2, 5600MB, the Netherlands. m.c.kaptein@tue.nl

ABSTRACT
Tailoring — presenting the right message at the right time — has long been identified as one of *the* core opportunities of persuasive systems. In this paper we describe a scenario in which an adaptive persuasive system which identifies users by the Bluetooth key of their mobile phone is used to promote energy savings. By describing this simplistic system and its possible implementation we identify several key-criteria of adaptive persuasive systems.

Author Keywords
Persuasive Technology, Influence strategies

ACM Classification Keywords
H.1.2 User/Machine Systems: Software psychology.

INTRODUCTION
CHI2010 attendees were presented with a choice on entering the conference hotel: A large revolving door provided access to the hotel while next to it was a sliding door—some things simply do not fit through a revolving door. With the air conditioning in full operation revolving doors are efficient at keeping the heat in. Sliding doors, however, are not. To help save energy a paper-sign was put up: *"Please take the revolving door"*. A brief observation proved the paper-sign to be effective just over half the time: 60% of the visitors took the revolving door. This scenario, the "Revolving Door Problem", offers a framework to describe *adaptive persuasive systems*. By further elaborating this scenario and exploring a solution we describe the neccesities and difficulties that arise when designing adaptive persuasive systems.

The Promises of Persuasive Technology
There are three reasons why employing a persuasive system might be more effective than the current paper-sign: (1) Persuasive technologies function as social actors and can use social influence strategies, (2) they can be context aware, and (3) they can adapt to individual users [5, 8]. While the paper-sign is probably located at the right place and at the right time—when visitors make their choice—the current version does not implement social influence strategies and does not adapt to its users.

Social Influence Strategies
Cialdini [2] shows how small changes to messages—such as the message on the door—can increase their effectiveness. For example, a message in a hotel room asking guests to "*reuse their towels*" compared to a message stating "*Join your fellow citizens in helping to save the environment*" led to a difference in towel re-usage of 28.4% [7]. To structure thes types of messages Cialdini [2] identifies six *social influence strategies*: *Authority*, *Consensus*, *Reciprocity*, *Liking*, *Scarcity*, and *Commitment*. The message in the towel re-usage example implements the *Consensus* strategy: people act like other people do. A message (e.g.) stating that "*The general manager of this hotel requests you to re-use...*" would implement the Authority strategy. These social influences strategies can easily be used to improve upon the effectiveness of the paper-sign.

The final promise of persuasive technologies however—adapting influence attempts to individuals—will require some kind of interactive system. While *adaptation* of persuasive strategies to responses by users is mentioned early on in the literature on persuasive technologies Fogg [5, e.g.] we are unaware of any actual implementations.

Individual Differences
There is growing evidence that individuals differ in their responses to influence strategies: Constructs like Need For Cognition [1] predict the response of individuals to the usage of social influence strategies. More concretely, Kaptein et al. [9] show that usage of influence strategies for individuals who are low susceptible to these strategies can lead to backfiring: for a portion of participants in their study compliance to a request was lower when the social influence strategy was presented. Next to this overall tendency to respond to influence strategies, some individuals seem more likely to respond to one specific strategy—e.g. an authority argument—while others are more influenced by implementations of other strategies. Cialdini et al. [3] shows that there are sizable and stable individual differences in people's responses to the commitment strategy. Similar results have been obtained when looking at the consensus strategy: Self-reported

susceptibility to this strategy highly correlates with behavioral responses to this strategy [10].

These individual differences in susceptibility to different persuasive strategies imply that persuasive systems should personalize the way in which they attempt to influence individuals. Such a class of systems, which we call *adaptive persuasive systems*, are an unexplored area in that we still need to understand how to model, design and build these systems. This paper takes a concrete but simple example that encapsulates the quintessence of this problem to discuss how to address these challenges.

SOLVING THE REVOLVING DOOR PROBLEM?

Returning to the revolving door problem, let us consider what is involved in implementing an adaptive persuasive system. We need to (A) identify the visitors entering the lobby—minimally by giving each a unique ID, and (B) measure the effectiveness of a presented message. The Bluetooth key of visitor's mobile phone could be used for identification [11]. This will capture around 12% of the visitors entering the lobby. This same identification method can also be used to measure the effectiveness of each persuasive attempt: One Bluetooth scanner next to the revolving door and one next to the sliding door could determine which entrance was used by the current visitor. Based on this knowledge about the visitor and records of earlier decisions a message implementing the right influence strategy can be selected. In the remainder of this paper, we focus on the mechanism by which these strategies can be selected.

Suppose we have only two messages to show, one implementing the authority strategy—*"The general manager of this hotel urges you to..." (A)*—and one implementing the consensus strategy—*"80% of our visitors always use...etc." (B)*. The system then needs a mechanism to choose the message that is most likely to be effective for the current visitor. It is intuitive that for a *new* visitor the system should present the message which has lead to the highest compliance for other, previously observed, visitors. If this message is successful then there is no need to try different messages on subsequent visits. However, when the selected message is not effective, it might become attractive to present another message on the next visit. This decision logically depends on the initial succes probabilities of the messages under consideration, the variance of effectiveness of messages between visitors, and the number of succes's or failures observed for the current visitor. A collection of estimates of the effectiveness of different influence strategies for an individual is called a *Persuasion Profile* and can be used to select the most-likely-to-be effective message on a next visit.

Formalizing the Adaptation Problem

The probability of a single visitor taking the revolving door on multiple occasions can be regarded a binomial random variable $B(n, p)$ where n denotes the number of approaches the visitor has made to the doors and p denotes the probability of success: the probability of taking the revolving door. Given M messages one can compute for each individual, for each message, probability $p_m = k_m/n_m$ where k_m is the number of observed successes after representation of message m, n_m times to a specific visitor. It makes intuitive sense to present a visitor with the messages with the highest p_m.

For a large number of observations N of one visitor this would make perfect sense. However, this will not inform a decision for a newly observed visitor. For a *new* visitor one would present the message m for which p_m is maximized over previously observed visitors[1]. Actually—given Stein's result [4]—for *every* user a weighted average of the p_m for an individual user and those of other users—one where the estimated $\widehat{p}_m$ for an individual is "shrunk" toward the population mean—will provide a better estimate than an estimate based on observations of a single visitor alone. E.g., if the authority message is effective 70% of the time over all visitors and only 30% percent of the time for the specific visitor under consideration, the best estimate of the (real) effectiveness of the authority message $\widehat{p}_A$ for this visitor is a weighted average of these two.

Adapting to Individuals

To include both the known effectiveness of a message for others, and a specific visitors previous responses to that same message, into a new estimate of message effectiveness, p_m, we use a Bayesian approach. A common way of including prior information in a binomial random process is to use the Beta-Binomial model [12]. The Beta $Beta(\alpha, \beta)$ distribution functions as a conjugate prior to the binomial. If we re-parametrize the beta distribution as follows

$$\pi(\theta|\mu, M) = Beta(\mu, M)$$

where $\mu = \frac{\alpha}{\alpha+\beta}$ and $M = \alpha + \beta$, then the expected value of the distribution is given by: $E(\theta|\mu, M) = \mu_m$. In our scenario this represents the expected probability of a successful influence attempt by a specific message. The certainty of this estimated success probability is represented by:

$$Var(\theta|\mu, M) = \sigma^2 = \frac{\mu(1-\mu)}{M+1}$$

After specifying the probability of success μ_m of message m and the certainty about this estimate σ^2_m we can treat this as our prior expectancy about the effectiveness of a specific message and update this expectancy by multiplying it by the likelihood of the observation(s) to obtain the distribution of our posterior expectation:

$$\begin{aligned} p(\theta|k) &\propto l(k|\theta)\pi(\theta|\mu, M) \\ &= Beta(k + M\mu, n - k + M(1-\mu)) \end{aligned}$$

[1]This is assuming the error costs—the effects of presenting the wrong message—are equal for each message.

The newly obtained Beta distribution, $B(\mu, M)$, functions as our probability distribution with a new point-estimate of the effectiveness of the presented message given by:

$$E(\theta|k) = \frac{k + M\mu}{n + M}$$

Decision Rule

The Beta-Binomial model described above allows us to estimate the effectiveness of message m, include prior knowledge, and update these estimates based on new observations. A individual's persuasion profile would be a record of both the expected success, μ_m, and the certainty, σ_m^2 of different influence strategies.

To determine which message to present next, one could pick the message which has the highest μ_m. However, if σ_m^2 is large this decision might not be feasible given that the difference between effectiveness estimates might not be significant. To address this we can choose to show the message with the highest estimate when this estimate is "certain enough"—in the binomial case only once sufficient observations are obtained. In uncertain situations we can randomly present one of the H messages which have the highest estimates out of the total set of estimates of M messages. This decision rule would avoid presenting each new visitor with only the single most effective message when responses to messages are variant.

Because the Beta distribution is not necessarily symmetrical the variance σ_m^2 provides and inadequate starting point to compute confidence intervals. This problem can be solved using simulations: By generating a number of draws from the specified Beta distribution and computing (e.g.) the 20^{th} and 80^{th} percentiles one can compute a empirical confidence interval. The above described decision rule for $M = 2$ would then result in:

$$M_{selected} = \begin{cases} 1 & \mu_1 > Perc(80)_2 \\ 2 & \mu_2 > Perc(80)_1 \\ Rand(1,2) & \text{otherwise} \end{cases}$$

Thus, if the estimated effectiveness of a message 1, $\widehat{p}_1 = \mu_1$, is higher than the 80^{th} percentile of message 2, $Perc(80)_2$, the system presents message one.[2] If the confidence interval of two messages overlap the system could randomly present one of these two.

Simulations

To explore the presented Beta-Binomial approach in the $M = 2$ scenario we simulated a dataset presenting different visitors observed at multiple points in time. The simulated data describes the message success of two different messages for four different groups of visitors with 20 visitors each on 50 approaches to the doors. The four groups represent (1) *general insusceptible visitors*—those that respond favorable to only 10% of the message which implement strategy A and 50% to strategy B, (2) *susceptible visitors*, $A = 40\%, B = 90\%$, (3) *visitors susceptible to message* B, $A = 10\%, B = 90\%$, and (4) *visitors susceptible to message* A, $A = 90\%, B = 10\%$. Table 1 shows an excerpt of the simulated data. Based on these simulated data we first compute our population estimates of message effectiveness for each message: $\widehat{p}_A = 0.38$, $\widehat{p}_B = 0.58$. Thus, message B—the consensus message—was most effective.

	Type	User	Occasion	Mes. A	Mes. B
1	1	1	1	0	0
2	1	1	2	0	0
3	1	1	3	0	1
..	..	..	..	..	..
..	..	..	..	..	..
1000	4	20	50	1	0

Table 1. Overview of the simulated data for the 4 different user groups. Columns *Mes. A* and *Mes. B* represent the success of the influence message at that point in time.

Next, we simulate for each visitor, each occurrence at the doors. We select the message as specified by our decision rule and record the (simulated) outcome. Next, we update our expectancy for the selected message and iterate through all occurrences. To ensure a flexible starting point for each user we set the prior variance of each estimate at the first encounter to be high: $\sigma_A^2 = \sigma_B^2 = 0.05$.[3]

Figure 1 shows for four users—one out of each group—in separate panels, the estimated probability of success of the two messages (left and right side of each panel). In the upper left panel—representing a *general insusceptible visitor*—convergence to message B, whose estimated effect is presented on the right side of the upper left panel, is slow: it takes about 40 observations before B is consistently estimated to be the "best" message. With higher compliance and/or larger differences in effectiveness of the two strategies convergence is much faster. The bottom right of figure 1 shows a user from the *visitors susceptible to message* A group. For this user after 10 observations strategy A is correctly identified as the most successful strategy.

Limitations of the proposed solution

There are a number of drawbacks of the proposed Beta-Binomial solution to create adaptive persuasive systems. Besides the fact that when the number of strategies grows the number of necessary occasions for convergence will increase, there are three more fundamental issues which are not addressed by this algorithm. First, while including prior information based on other users, the algorithm described here does not use a shrunken estimate on each occasion: After including the initial knowledge of the behavior of other visitors the model is specific for an individual

[2] The 80^{th} percentile is an arbitrary choice.

[3] One could estimate this variance based on the between-visitor variance.

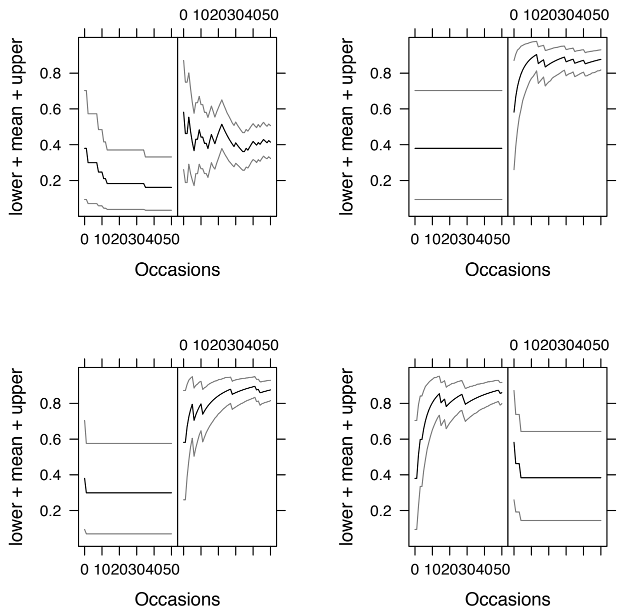

Figure 1. Progression of point estimates of the effects of two messages on four different users (the four panels). Within each panel the left side shows the estimated effect of message A, including in gray its 80% confidence interval, and the right shows the estimates for message B. A horizontal section in the estimates of message A indicates that at that point in time the message B was shown and updated.

visitor. While this provides quick adaptation there is no opportunity to adapt estimates based on changing population wise trends. Second, since the estimates for the effectiveness of the strategies are treated independently there is no way to of "borrowing strength" [6] based on correlations with other strategies. Both of these concerns could be addressed using a multilevel approach. Finally, the proposed model provides no method of including prior believes about the distribution of visitor profiles over a population.

CONCLUSIONS

We identified two core necessities of adaptive persuasive systems: a means to identify users and a means to measure effectiveness of persuasive attempts. Furthermore, we highlighted a number of challenges associated with the design of these systems. The presented Beta-Binomial solution is lightweight and functions well in simulations with only two messages. More elaborate algorithms which are (1) variant to changing population trends,(2) allow for relationships between strategies, and (3) enable us to include prior beliefs about user profiles should be explored. Given the current state of social science literature on influence strategies we believe that persuasive technologies *should* tailor the influence strategies they use to their users. We described one possible—but limited—implementation of such a system. This, and other, implementations should now be tested empirically.

Mobile devices—as used in our scenario—provide a core opportunity to serve as an identifier for adaptive persuasive technologies. Currently we are operating a system, like the one described here, in real-life and we would like to share our experiences building and deploying this system during the CHI 2011 PINC workshop.

References

[1] Cacioppo, J. T. and Petty, R. E. (1982). The need for cognition. *J. of Pers. and Soc. Psy.*

[2] Cialdini, R. (2001). *Influence, Science and Practice.* Allyn & Bacon, Boston.

[3] Cialdini, R. B., Trost, M. R., and Newsom, J. T. (1995). Preference for consistency: The development of a valid measure and the discovery of surprising behavioral implications. *J. of Pers. and Soc. Psy.*, 69:318–328.

[4] Efron, B. and Morris, C. (1975). Data analysis using Stein's estimator and its generalizations. *Journal of the American Statistical Association*, 70(350):311–319.

[5] Fogg, B. J. (2002). *Persuasive Technology: Using Computers to Change What We Think and Do.* Morgan Kaufmann.

[6] Gelman, A. and Hill, J. (2007). *Data Analysis Using Regression and Multilevel/Hierarchical Models.* Cambridge University Press.

[7] Goldstein, N. J., Cialdini, R. B., and Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *J. of Cons. Res.*, 35(3):472–482.

[8] Kaptein, M., Aarts, E. H. L., Ruyter, B. E. R., and Markopoulos, P. (2009a). Persuasion in ambient intelligence. *Journal of Ambient Intelligence and Humanized Computing*, 1:43—56.

[9] Kaptein, M., Lacroix, J., and Saini, P. (2010). Individual differences in persuadability in the health promotion domain. In Ploug, T., Hasle, P., and Oinas-Kukkonen, H., editors, *Persuasive Technology*, pages 94—105. Springer Berlin / Heidelberg.

[10] Kaptein, M. C., de Ruyter, B., Markopoulos, P., and Aarts, E. (2009b). Persuading you: Individual differences in susceptibility to persuasion. In *12th IFIP TC 13 International Conference - INTERACT*, pages 24–28, Uppsala, Sweden. ACM Press.

[11] Kostakos, V. (2008). Using bluetooth to capture passenger trips on public transport buses. *CoRR*, abs/0806.0874.

[12] Wilcox, R. R. (1981). A review of the beta-binomial model and its extensions. *J. of Educ. Stat.*, 6:3–32.