

# Selecting Effective Means to Any End: Futures and Ethics of Persuasion Profiling

Maurits Kaptein<sup>1</sup> and Dean Eckles<sup>2</sup>

<sup>1</sup> Technical University of Eindhoven, Eindhoven, the Netherlands.

<sup>2</sup> Stanford University, Stanford CA 94305.

**Abstract.** Interactive persuasive technologies can and do adapt to individuals. Existing systems identify and adapt to user preferences within a specific domain: e.g., a music recommender system adapts its recommended songs to user preferences. This paper is concerned with adaptive persuasive systems that adapt to individual differences in the effectiveness of particular means, rather than selecting different ends. We give special attention to systems that implement persuasion profiling — adapting to individual differences in the effects of influence strategies. We argue that these systems are worth separate consideration and raise unique ethical issues for two reasons: (1) their end-independence implies that systems trained in one context can be used in other, unexpected contexts and (2) they do not rely on — and are generally disadvantaged by — disclosing that they are adapting to individual differences. We use examples of these systems to illustrate some ethically and practically challenging futures that these characteristics make possible.

## 1 Introduction

You are just finishing up your Christmas shopping. You surf to an online bookstore and look around for books your family members might like. While you are not so much attracted by the content, as you don't share your sister's chicklit affection, you do intend to buy a great present. Luckily, the bookstore provides you with lots of options to base your choices on. Some books are accompanied by ratings from users, and some are sold at a special discount rate just for Christmas. There is also a section of books which is recommended by famous authors, and there are several bestsellers — that is, books that many people apparently chose.

You unknowingly spend more time looking at the books that are recommended by famous authors than the books presented with the other messages. In fact, the present you buy was identified as a famous authors' selection. The online store takes you through the checkout process. Since you frequent the store there is no need to specify your details. Your account is recognized, and in just two clicks a great gift is purchased.

As it turns out, this is not the first time you have been a *sucker for authority*: while being presented with persuasive attempts — yes, all of the messages on the bookstore are presented to sell more books — you hardly ever buy the special

discount books (you would be a sucker for *scarcity*) or the bestsellers (*consensus*) and frequently end up with a book that in one way or another is endorsed by a relevant authority. In the context of book sales alone, authority as a persuasive strategy can be implemented in a host of ways (e.g., selections by authors, critics' book reviews). You are more likely to buy books supported by this strategy — and its many implementations — than those supported by implementations of other strategies.

Based on your online behavior since you first signed in to the online bookstore, the company is able to estimate the effects of different influence strategies. The bookstore knows you listen to relevant authorities and experts more than friends or just the anonymous majority. And, in comparison to the average effects of these strategies, the positive effect of authoritative sources is larger for you. This latter point holds true irrespective of the context: though for some attitudes and behavior, authorities are more or less persuasive on average, across these contexts they are more persuasive for *you* than they for others. Perhaps you know this too, or perhaps you don't. But others can now know this also: the online bookstore sells this information — your *persuasion profile*<sup>3</sup> — for additional income. In this case, your persuasion profile has been sold to a political party.

In the run-up to the next election you receive mailings to vote for a particular candidate. A number of arguments by influential political commentators and esteemed, retired politicians in the door-to-door mailing changes your attitude about the local party from indifference to approval. That is, the authority figures in the mailing have done a good job of changing your attitudes in favor of their political party. Your neighbor received a similar leaflet, although hers seems to stress the fact that everyone else in the neighborhood votes for this specific party. None of the authority arguments that persuaded you appear on her personalized copy of the mailing.

## 1.1 Overview

This article discusses the future development and the ethical implications of adaptive persuasive technologies, especially those that develop and employ persuasion profiles — profiles that specify estimates of the effects of particular influence strategies on an individual. That is, we highlight a subclass of adaptive persuasive technologies that selects influence strategies for use based on individuals' profiles. These profiles are constructed based on individuals' previous responses to implementations of this same strategy and other available demographic and psychographic information. In particular, we describe how the former type of information — past responses to implementations of influence strategies — can be aggregated into persuasion profiles that may be generalizable to contexts other

---

<sup>3</sup> Fogg has used the term 'persuasion profiling' in lectures since 2004. We found it an apt and evocative name for these adaptive systems. Fogg [personal communication] has indicated that its meaning in this article is generally consistent with his prior use. This introductory example is similar to one in Fogg [9].

than those in which the data was collected, as illustrated in the scenario above. In order to highlight both the similarities and differences to longstanding practices, our examples include persuasion processes in human–human interaction and interactions with both adaptive and non-adaptive persuasive technologies.

The adaptive persuasive technologies that are the focus of this paper have two features that distinguish them from such adaptive persuasive technologies more generally. First, they are end-independent: a persuasion profile constructed in one domain for some end (or goal) can be applied in other domains for other ends or within the same domain for different ends. Second, unlike other adaptive persuasive technologies, such as recommender systems, they do not benefit from or are even disadvantaged by disclosing the adaptation.

We have been motivated to write this paper in anticipation of negative responses to the scenarios we describe here and to our empirical research in this area. Our colleagues have often responded to drafts of this and other papers with concern about the future of these technologies. We too are concerned: we selected the term ‘persuasion profiling’ precisely to evoke careful consideration and concern. But we see substantial positive potential for adaptive persuasive technologies, and we hope that readers entertain the idea that using persuasion profiling may sometimes be the most ethical course of action.

Finally, we revisit each of the examples using Berdichevsky and Neuenchwander’s [3] decision tree and principles, one proposal for how to evaluate designers’ ethical responsibility when creating persuasive technologies.

## 2 Adaptive persuasive technologies

One can describe adaptive technologies in many ways, and profiles or models of users have been employed and studied in some of these systems. These notions however do not unambiguously translate into the domain of persuasive technologies. In this section we distinguish different kinds of adaptive persuasive technologies and define the associated concepts. As with the rest of the paper, we illustrate these with examples.

Book recommendations by, e.g., Amazon are often spot on: the book is something the user does not own, but would like to. These recommendations are provided by a *recommender system* [17]. Recommender systems frequently have an aim to persuade people (to, e.g., buy books), and thus are reasonably counted among the larger category of adaptive persuasive technologies. That is, we define *adaptive persuasive technologies* as technologies that aim to increase the effectiveness of some attitude or behavior change by responding to the behavior of (and other information about) their individual users.<sup>4</sup> We intend this to be an inclusive and non-controversial definition of persuasive technologies that adapt to their *users*. We continue by distinguishing between two kinds of adaptive persuasive technologies.

---

<sup>4</sup> Like Fogg’s [8] definition of persuasive technology, this definition depends on the intentions of the designer or deployer of the technology, rather than success [7].

## 2.1 Ends and means adaptation

We distinguish between those adaptive persuasive technologies that adapt the particular *ends* they try to bring about and those that adapt their *means* to some end.

First, there are systems that use models of individual users to select particular *ends* that are instantiations of more general target behaviors. If the more general target behavior is book buying, then such a system may select which specific books to present.

Second, adaptive persuasive technologies that change their *means* adapt the persuasive strategy that is used — independent of the end goal. One could offer the same book and for some people show the message that the book is recommended by experts, while for others emphasizing that the book is almost out of stock. Both messages may be true, but the effect of each differs between users.

### – *Example 2. Ends adaptation in recommender systems*

Pandora is a popular music service that tries to engage music listeners and persuade them into spending more time on the site and, ultimately, purchase music. For both goals it is beneficial for Pandora if users enjoy the music that is presented to them by achieving a match between the music offering to individual, potentially latent music preferences. In doing so, Pandora adaptively selects the end — the actual song that is listened to and that could be purchased, rather than the means — the reasons presented for the selection of one specific song.

The distinction between *end*-adaptive persuasive technologies and *means*-adaptive persuasive technologies is important to discuss since adaptation in the latter case could be domain independent. In *end* adaptation, we can expect that little of the knowledge of the user that is gained by the system can be used in other domains (e.g. book preferences are likely minimally related to optimally specifying goals in a mobile exercise coach). *Means* adaptation is potentially quite the opposite. If an agent expects that a person is more responsive to authority claims than to other influence strategies in one domain, it may well be that authority claims are also more effective for that user than other strategies in a different domain. While we focus on novel means-adaptive systems, it is actually quite common for human influence agents adaptively select their means.

### – *Example 3. Means adaptation by human influence agents*

Salespeople adapt their pitches to the audience. After spending some time observing the salesman in a used-car garage, it is clear that the same car — a total wreck which definitely was a bad-buy — was sometimes presented as either “*a great ride which your friends would look up to*” or “*one of the safest cars according to the NHTSA*”. The car that was being advertised was the same in both cases, but the salesman had judged the recipient likely to respond better to one description or the other.

### 3 Influence strategies and implementations

Means-adaptive systems select different means by which to bring about some attitude or behavior change. The distinction between adapting means and ends is an abstract and heuristic one, so it will be helpful to describe one particular way to think about means in persuasive technologies. One way to individuate means of attitude and behavior change is to identify distinct *influence strategies*, each of which can have many implementations. Investigators studying persuasion and compliance-gaining have varied in how they individuate influence strategies: Cialdini [5] elaborates on six strategies at length, Fogg [8] describes 40 strategies under a more general definition of persuasion, and others have listed over 100 [16].

Despite this variation in their individuation, influence strategies are a useful level of analysis that helps to group and distinguish specific influence tactics. In the context of means adaptation, human and computer persuaders can select influence strategies they expect to be more effective than other influence strategies. In particular, the effectiveness of a strategy can vary with attitude and behavior change goals. Different influence strategies are most effective in different stages of the attitude to behavior continuum [1]. These range from use of heuristics in the attitude stage to use of conditioning when a behavioral change has been established and needs to be maintained [11]. Fogg [10] further illustrates this complexity and the importance of considering variation in target behaviors by presenting a two-dimensional matrix of 35 classes behavior change that vary by (1) the schedule of change (e.g., one time, on cue) and (2) the type of change (e.g., perform new behavior vs. familiar behavior). So even for persuasive technologies that do not adapt to individuals, selecting an influence strategy — the means — is important. We additionally contend that influence strategies are also a useful way to represent individual differences [9] — differences which may be large enough that strategies that are effective on average have negative effects for some people.

– *Example 4. Backfiring of influence strategies*

John just subscribed to a digital workout coaching service. This system measures his activity using an accelerometer and provides John feedback through a Web site. This feedback is accompanied by recommendations from a general practitioner to modify his workout regime. John has all through his life been known as authority averse and dislikes the top-down recommendation style used. After three weeks using the service, John’s exercise levels have *decreased*.

### 4 Persuasion profiles

When systems represent individual differences as variation in responses to influence strategies — and adapt to these differences, they are engaging in *persuasion profiling*. Persuasion profiles are thus collections of expected effects of different influence strategies for a specific individual. Hence, an individual’s persuasion

profile indicates which influence strategies — one way of individuating means of attitude and behavior change — are expected to be most effective.

Persuasion profiles can be based on demographic, personality, and behavioral data. Relying primarily on behavioral data has recently become a realistic option for interactive technologies, since vast amounts of data about individuals’ behavior in response to attempts at persuasion are currently collected. These data describe how people have responded to presentations of certain products (e.g. e-commerce) or have complied to requests by persuasive technologies (e.g. the DirectLife Activity Monitor [12]).

Existing systems record responses to particular messages — implementations of one or more influence strategies — to aid profiling. For example, Rapleaf uses responses by a users’ friends to particular advertisements to select the message to present to that user [2]. If influence attempts are identified as being implementations of particular strategies, then such systems can “borrow strength” in predicting responses to other implementations of the same strategy or related strategies. Many of these scenarios also involve the collection of personally identifiable information, so persuasion profiles can be associated with individuals across different sessions and services.

## 5 Consequences of means-adaption

In the remainder of this paper we will focus on the implications of the usage of persuasion profiles in *means*-adaptive persuasive systems. There are two properties of these systems which make this discussion important:

1. *End-independence*: Contrary to profiles used by *end*-adaptive persuasive systems the knowledge gained about people in *means*-adaptive systems can be used independent from the end goal. Hence, persuasion profiles can be used independent of context and can be exchanged between systems.
2. *Undisclosed*: While the adaptation in *end*-adaptive persuasive systems is often most effective when disclosed to the user, this is not necessarily the case in *means*-adaptive persuasive systems powered by persuasion profiles. Selecting a different influence strategy is likely less salient than changing a target behavior and thus will often not be noticed by users.

Although through the previous examples and the discussion of adaptive persuasive systems these two notions have already been hinted upon, we feel it is important to examine each in more detail.

### 5.1 End-independence

Means-adaptive persuasive technologies are distinctive in their end-independence: a persuasion profile created in one context can be applied to bringing about other ends in that same context or to behavior or attitude change in a quite different context. This feature of persuasion profiling is best illustrated by contrast with end adaptation.

Any adaptation that selects the particular *end* (or goal) of a persuasive attempt is inherently context-specific. Though there may be associations between individual differences across context (e.g., between book preferences and political attitudes) these associations are themselves specific to pairs of contexts. On the other hand, persuasion profiles are designed and expected to be independent of particular ends and contexts. For example, we propose that a person’s tendency to comply more to appeals by experts than to those by friends is present both when looking at compliance to a medical regime as well as purchase decisions.

It is important to clarify exactly what is required for end-independence to obtain. If we say that a persuasion profile is end-independent then this does not imply that the effectiveness of influence strategies is constant across all contexts. Consistent with the results reviewed in section 3, we acknowledge that influence strategy effectiveness depends on, e.g., the type of behavior change. That is, we expect that the most effective influence strategy for a system to employ, even given the user’s persuasion profile, would depend on both context and target behavior. Instead, end-independence requires that the *difference between the average effect of a strategy for the population and the effect of that strategy for a specific individual is relatively consistent across contexts and ends*.<sup>5</sup>

**Implications of end-independence** From end-independence, it follows that persuasion profiles could potentially be created by, and shared with, a number of systems that use and modify these profiles. For example, the profile constructed from observing a user’s online shopping behavior can be of use in increasing compliance in saving energy. Behavioral measures in latter two contexts can contribute to refining the existing profile.

Not only could persuasion profiles be used across contexts within a single organization, but there is the option of exchanging the persuasion profiles between corporations, governments, other institutions, and individuals. A market for persuasion profiles could develop [9], as currently exists for other data about consumers. Even if a system that implements persuasion profiling does so ethically, once constructed the profiles can be used for ends not anticipated by its designers.

Persuasion profiles are another kind of information about individuals collected by corporations that individuals may or have effective access to. This raises issues of data ownership. Do individuals have access to their complete persuasion profiles or other indicators of the contents of the profiles? Are individuals compensated for this valuable information [14]? If an individual wants to use Amazon’s persuasion profile to jump-start a mobile exercise coach’s adaptation, there may or may not be technical and/or legal mechanisms to obtain and transfer this profile.

---

<sup>5</sup> This point can also be made in the language of interaction effects in analysis of variance: Persuasion profiles are estimates of person–strategy interaction effects. Thus, the end-independence of persuasion profiles requires not that the two-way strategy–context interaction effect is small, but that the three-way person–strategy–context interaction is small.

## 5.2 Non-disclosure

Means-adaptive persuasive systems are able and likely to not disclose that they are adapting to individuals. This can be contrasted with end adaptation, in which it is often advantageous for the agent to disclose the adaptation and potentially easy to detect. For example, when Amazon recommends books for an individual it makes clear that these are personalized recommendations — thus benefiting from effects of apparent personalization and enabling presenting reasons why these books were recommended. In contrast, with means adaptation, not only may the results of the adaptation be less visible to users (e.g. emphasizing either “*Pulitzer Prize winning*” or “*International bestseller*”), but disclosure of the adaptation may reduce the target attitude or behavior change.

It is hypothesized that the effectiveness of social influence strategies is, at least partly, caused by automatic processes. According to dual-process models [4], under low elaboration message variables manipulated in the selection of influence strategies lead to compliance without much thought. These dual-process models distinguish between central (or systematic) processing, which is characterized by elaboration on and consideration of the merits of presented arguments, and peripheral (or heuristic) processing, which is characterized by responses to cues associated with, but peripheral to the central arguments of, the advocacy through the application of simple, cognitively “cheap”, but fallible rules [13]. Disclosure of means adaptation may increase elaboration on the implementations of the selected influence strategies, decreasing their effectiveness if they operate primarily via heuristic processing. More generally, disclosure of means adaptation is a disclosure of persuasive intent, which can increase elaboration and resistance to persuasion.

**Implications of non-disclosure** The fact that persuasion profiles can be obtained and used without disclosing this to users is potentially a cause for concern. Potential reductions in effectiveness upon disclosure incentivize system designs to avoid disclosure of means adaptation.

Non-disclosure of means adaptation may have additional implications when combined with value being placed on the construction of an accurate persuasion profile. This requires some explanation. A simple system engaged in persuasion profiling could select influence strategies and implementations based on which is estimated to have the largest effect in the present case; the model would thus be engaged in passive learning. However, we anticipate that systems will take a more complex approach, employing active learning techniques [e.g., 6]. In active learning the actions selected by the system (e.g., the selection of the influence strategy and its implementation) are chosen not only based on the value of any resulting attitude or behavior change but including the value predicted improvements to the model resulting from observing the individual’s response. Increased precision, generality, or comprehensiveness of a persuasion profile may be valued (a) because the profile will be more effective in the present context or (b) because a more precise profile would be more effective in another context or more valuable in a market for persuasion profiles.



These later cases involve systems taking actions that are estimated to be non-optimal for their apparent goals. For example, a mobile exercise coach could present a message that is not estimated to be the most effective in increasing overall activity level in order to build a more precise, general, or comprehensive persuasion profile. Users of such a system might reasonably expect that it is designed to be effective in coaching them, but it is in fact also selecting actions for other reasons, e.g., selling precise, general, and comprehensive persuasion profiles is part of the company’s business plan. That is, if a system is designed to value constructing a persuasion profile, its behavior may differ substantially from its anticipated core behavior.

## 6 Ethical considerations

We have illustrated variations on adaptation in persuasion processes through four examples. We now turn to ethical evaluation of these examples, with particular attention to means-adaptive persuasive systems.

Persuaders have long stood on uneasy ethical grounds [3]. From the more recent beginning of the study of interactive persuasive technologies, researchers and practitioners have questioned the ethics of developing and deploying of persuasive systems. Several attempts to develop frameworks or principles for the ethical evaluation of persuasive systems have been undertaken [e.g., 3, 8].

Berdichevsky and Neuenschwander [3] present a decision tree for ethical evaluation of persuasive technologies and the moral responsibility of system designers. This decision tree identifies how intent, predictability of outcomes, and ethical judgment interact to determine the proposed judgement and response to the system designers. According to this decision tree, the system designer is (a) praiseworthy if the outcome is intended and good; (b) not responsible if the outcome is unintended and not reasonably predicable, or if the outcome is reasonably predictable and good but not intended; and (c) otherwise at fault and blameworthy. Table 1 presents our application of this decision tree to the four extended examples in this paper, including the introductory example. Berdichevsky and Neuenschwander [3] additionally offer eight principles that they regard as heuristics that could be justified within rule-based consequentialism.

We do not intend to make a final judgment about the ethical aspect of using persuasive profiles. Instead, we are employing the method of *reflective equilibrium* [15], in which one iteratively moves between principles and intuitive responses to individual cases to reach justified ethical judgements. Thus, we see the application of Berdichevsky and Neuenschwander’s [3] principles and decision tree as but one move in reaching a stable and justified understanding of the ethics of adaptation in persuasion. To this end, we invite readers to participate in the generation and evaluation of examples at <http://persuasion-profiling.com>.

In the first example, the relevant, intended outcome is increased votes for a particular candidate, which has an unknown ethical status. Depending on this outcome, the designers of the personalized mailing system for the political campaign may be praiseworthy or blameworthy, according to the decision

Example	Intended?	Predictable?	Outcome?	Judgement
1 (Introduction): Undisclosed persuasion profiling to influence voting	Yes	Yes	(Ethical)	(Praise)
2. Disclosed end adaptation for sales	Yes	Yes	(Ethical)	Praise
3. Undisclosed means adaptation for sales	Yes	Yes	Unethical	Blame
4. Influence strategies backfire	No	(Yes)	Unethical	(Blame)

**Table 1.** Ethical evaluation of the examples used in this article according to Berdichevsky and Neuenschwander [3]. The entries in parentheses are provisional values used for discussion.

tree. However, Berdichevsky and Neuenschwander’s [3, p. 52] sixth principle is that systems should disclose the “the motivations, methods, and intended outcomes, except when such disclosure would significantly undermine an otherwise ethical goal.” Thus, the proposed response to this example would, assuming the goal is ethical, be determined by how much the disclosure of adaptation would reduce the effectiveness of the personalized mailing system.

In Example 2 one relevant, intended outcome is purchasing music; additional intended outcomes include discovering and enjoying new music. Although sales as a goal of the persuasive intent may be debatable, for the present discussion, we consider it so accepted in Western society that we do not enter this outcome as unethical. Assuming the sales outcome is not unethical, then the designers are praiseworthy in proportion to the value of this and the other intended outcomes. Due to the disclosed nature of the end adaptation, the system is consistent with the sixth principle, assuming that the designers’ motivations are made sufficiently clear to users.

Example 3 is on its face unethical and the salesman is blameworthy, as he convinces someone to make a poor decision in purchasing a car not worth buying. This outcome is intended and unethical, so the decision tree yields the same judgement. Furthermore, this means adaptation is undisclosed, violating the sixth principle; however, it is unclear whether this makes the salesman *more* blameworthy than if the same outcome occurred with disclosure. Nonetheless, undisclosed means adaptation when achieving the same outcome with disclosure would not be praiseworthy strikes us as a practice that should be avoided.

In Example 4 the relevant outcome is an unintended decrease in exercise. While the system is designed to increase exercise levels — and the designers would likely be praiseworthy when successful, the system’s failure to adapt to John’s aversion to authority arguments leads to lowered exercise level for John. Is this unintended and otherwise unethical outcome reasonably predictable? While it may not be reasonably predictable that John in particular would response this way, we propose that it is in aggregate: it is reasonably predictable that some — and potentially many — users will have a negative response to this influence strategy. Thus, in contrast to Examples 1 and 3, *failure to use persuasion profiling (or other means adaptation) is unethical* in this case. This example illustrates the faults of a hypothetical principle according to which persuasion

profiling is unethical and to be avoided in general.

However, one can not move directly from this point to the stronger conclusions that the system designers should implement persuasion profiling. In particular, the implications of end-independence and non-disclosure are of importance in Example 4. When implementing persuasion profiling, actions taken by the system selected to improve to the profile — while perhaps not maximizing the primary, ethical outcome — warrant ethical judgement as well. Additionally, users’ possible inability to inspect, modify, or use their persuasion profiles may present ethical concerns that would nonetheless motivate not implementing persuasion profiling, at least in particular ways.

## 7 Limitations

We have aimed to bring about consideration of ethical issues associated with means-adaptive persuasive technologies, and this is partially motivated by these systems’ end-independence. So it is reasonable to ask whether such end-independence is possible and likely to be widespread in the future. We think the answer is “yes”. Though we cannot yet offer exhaustive, empirical evidence, this answer is provisionally justified by two considerations. First, psychologists have developed personality constructs that predict differences in attitude change processes and outcomes across different contexts. For example, need for cognition predicts differences in the processing of persuasive messages, whether these messages are about policy proposals or consumer products [4]. Second, as mentioned in section 4, people are exposed to an increasing number of influence attempts, and these attempts and their responses are recorded and available for training models. We expect that this will support the identification of stable individual differences in processing of and responses to influence attempts. We are currently engaged in empirical research testing this and related hypotheses.

## 8 Conclusion

In this article we defined persuasion profiling as adapting to individual differences in the effects of influence strategies. Adaptive persuasive systems that adapt their *means* rather than *ends* enables constructing models of users (e.g., persuasion profiles) that are end-independent and can and will likely be obtained and used while being undisclosed to the user.

Our focus here has been to define this area of research and describe why the futures and ethics of adaptive persuasive systems are worth careful consideration. In particular, we hope we have illustrated both some concerns presented by persuasion profiling, but also the potential for ethical use of these techniques.

## Bibliography

- [1] Aarts, E.H.L., Markopoulos, P., Ruyter, B.E.R.: The persuasiveness of ambient intelligence. In: Petkovic, M., Jonker, W. (eds.) *Security, Privacy and Trust in Modern Data Management*. Springer (2007)
- [2] Baker, S.: Learning, and profiting, from online friendships. *BusinessWeek* 9(22) (May 2009)
- [3] Berdichevsky, D., Neunschwander, E.: Toward an ethics of persuasive technology. *Commun. ACM* 42(5), 51–58 (1999)
- [4] Cacioppo, J.T., Petty, R.E., Kao, C.F., Rodriguez, R.: Central and peripheral routes to persuasion: An individual difference perspective. *Journal of Personality and Social Psychology* 51(5), 1032–1043 (1986)
- [5] Cialdini, R.: *Influence: Science and Practice*. Allyn & Bacon, Boston (2001)
- [6] Cohn, D.A., Ghahramani, Z., Jordan, M.I.: Active learning with statistical models. *Journal of Artificial Intelligence Research* 4, 129 – 145 (1996)
- [7] Eckles, D.: Redefining persuasion for a mobile world. In: Fogg, B.J., Eckles, D. (eds.) *Mobile Persuasion: 20 Perspectives on the Future of Behavior Change*. Stanford Captology Media, Stanford, California (2007)
- [8] Fogg, B.J.: *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann (2002)
- [9] Fogg, B.J.: Protecting consumers in the next tech-ade, U.S. Federal Trade Commission hearing (Nov 2006), [http://www.ftc.gov/bcp/workshops/techade/pdfs/transcript\\_061107.pdf](http://www.ftc.gov/bcp/workshops/techade/pdfs/transcript_061107.pdf)
- [10] Fogg, B.J.: The behavior grid: 35 ways behavior can change. In: Chatterjee, S., Dev, P. (eds.) *Proc. of Persuasive Technology 2009*. p. 42. ACM (2009)
- [11] Kaptein, M., Aarts, E.H.L., Ruyter, B.E.R., Markopoulos, P.: Persuasion in ambient intelligence. *Journal of Ambient Intelligence and Humanized Computing* 1, 43 –56 (2009)
- [12] Lacroix, J., Saini, P., Goris, A.: Understanding user cognitions to guide the tailoring of persuasive technology-based physical activity interventions. In: *Proc. of Persuasive Technology 2009*. vol. 350, p. 9. ACM (2009)
- [13] Petty, R.E., Wegener, D.T.: The elaboration likelihood model: Current status and controversies. In: Chaiken, S., Trope, Y. (eds.) *Dual-process theories in social psychology*, p. 4172. Guilford Press, New York (1999)
- [14] Prabhaker, P.R.: Who owns the online consumer? *Journal of Consumer Marketing* 17, 158–171(14) (2000)
- [15] Rawls, J.: The independence of moral theory. *Proceedings and Addresses of the American Philosophical Association* 48, 5–22 (1974)
- [16] Rhoads, K.: How many influence, persuasion, compliance tactics & strategies are there? (2007), <http://www.workingpsychology.com/numbertactics.html>
- [17] Schafer, J.B., Konstan, J.A., Riedl, J.: E-commerce recommendation applications. *Data Mining and Knowledge Discovery* 5(1/2), 115–153 (2001)