

# INFORMED CONSENT AND THE DESIGN OF RECOMMENDER SYSTEMS

**Batya Friedman**

Department of Computer Science

Colby College

Waterville, ME 04901 USA

+1 707 983 6030

[batya@saber.net](mailto:batya@saber.net)

(This text was adapted from Friedman, B., 1999, *Network browsers and informed consent*. Unpublished manuscript.)

## STATEMENT OF PROBLEM

Informed consent is an important value to cultivate in our online interactions and communities. Informed consent refers to the principle that individuals should be informed about the harms and benefits from participating in a particular activity, and should have a say in whether or not they participate. Pragmatically, informed consent underlies a climate of trust, and is a critical protection for privacy.

Recommender systems -- as described in this workshop call -- “automate the process of passing experience from a community of users to an individual user who is about to make a decision.” To achieve this goal, recommender systems typically determine what data to collect from users, how to store that data, when to access that data, whether or not and how to aggregate that data, and when and how to present that data to other users. In addition to all this, recommender systems also determine (a) what users *know* about these activities and (b) what *say* users have about the use of these activities. This point is crucial. For in most any account of informed consent, a consentee must be aware of relevant knowledge of what is being requested (i.e., informed) and then say that the request can be fulfilled (i.e., consent). I want to say more about what it means to be informed and what it means to consent.

## WHAT COUNTS AS “INFORMED”?

Exactly how much information and of what sort does a recommender system need to provide in order to satisfy the principle of informed

consent? While there are no hard and fast rules about what counts as “informed”, we can gain insight from the guidelines established for informed consent in the context of research on human subjects.

Here several dimensions are relevant:

- Potential harms and risks to participants should be made explicit.
- Purpose/benefits of the intervention should be made explicit.
- If information is collected, who will have access to it, what will be collected, how long will it be archived, what it will be used for, and how the anonymity of the individual will be protected should be made explicit.
- Items 1 –3 should be accomplished in language that is clear and will be easily understood by the participant.

The broad principle requires recommender systems to inform the individual sufficiently about potential risk and benefits so that the individual can make a reasoned and reasonable judgment about whether or not to consent. For example, the recommender system would need to make explicit what information was being collected about the user, how the user’s anonymity would be protected, what the information would be used for, and how long the information would be maintained. Technical language should be avoided in favor of simple, meaningful communication.

### **WHAT COUNTS AS “CONSENT”?**

Consent carries with it three conditions: First, the individual must have a reasonably clear opportunity to accept or decline. For example, opportunities to consent (or decline) that are buried under layers of menus, hidden in obscure locations, or require elaborate sequences of key presses are at best, marginally viable. Thus, before collecting any data, recommender systems should visibly and openly request the user’s consent.

Second, if what an individual is informed about does not reasonably match what the individual experiences then consent has been obtained under false pretenses. In the game of football, for example, participants consent to being bumped, bashed and smashed and, typically, are not disappointed. However, if by entering a game of football, a participant became the target of gunfire, we would call “foul” claiming that gunfire was not a reasonable expectation of how the participant understood the ground rules for football. Consenting to the game did not include consenting to gunfire. In the context of recommender systems this might occur as follows: If a recommender system solicits a user’s consent under one set of conditions (e.g., to be used by Web Site A to help inform other users of Web Site A) and then allows this information to be used under other conditions (e.g., makes this information available to Web Site B to help inform users of Web Site B).

Third, the opportunity to consent or decline must be genuine – that is, not coerced. The canonical case of coercion would occur if someone pointed a gun to your head and said: “I’ll shoot you if you don’t accept this cookie.” However, less obvious forms of coercion abound. For example, if the information or services you need can only be obtained online and all such sites require the user to contribute data to a recommender system then the user has no genuine choice. The user must go along with a recommender system if the user is to obtain the information or services needed.

In sum, valid consent must be obtained in a straightforward manner, non-coerced, and with adequate information.

### **IMPLIED CONSENT – ISN’T THAT ENOUGH?**

As a society, we have a great deal of experience with implicit consent – situations in which by virtue of voluntarily entering into the situation we consent to the activities that are known to

occur in that context. To continue with the football example: While in normal interactions I do not consent to being bumped, bashed, or smashed by another person, when I step onto the football field in football garb and enter the “game”, I implicitly consent and accept being bumped, bashed or smashed by the other players -- who have similarly implicitly consented to accept such treatment from me.

In the context of web interactions, one could try to argue that by virtue of visiting a site the user has implicitly consented to the activities sponsored by that site. After all, no one forced the user to visit that site or, for that matter, to engage in any web interactions.

However, this argument encounters problems with two of the three conditions delineated above: reasonable expectation and coercion. Taking each in turn. For implied consent to hold, what the individual imagines that he or she has consented to by virtue of participation must match reasonably well with what the individual experiences. In the context of web interactions, we can ask what do users expect of web sites? Do they expect web sites to generate profiles of their interactions? Do they expect these profiles to be stored over long periods of time and slowly elaborated? Do they expect their profiles to be presented in some format to other users? While typical users probably understand that “sites may build a user profile about them to aid with billing” and that in some vague way the profile may be stored and accessible when next they visit the site, most users would not assume that the site would volunteer information about their past activities at the site to other users. And, in fact, for most purposes of visiting a site, this sharing of information about users to other users is not necessary. Now it is possible that users might assume that a site could do other things, if those things were to benefit the user – advanced features, so to speak. Recommender systems could be portrayed in this light as an advanced feature that allows a site to collect information about you to provide to other users and, in turn, provides you with information based on the activities of other users. But, users would also expect that those special features would be advertised and well known. Thus, recommender systems as currently implemented fail to meet the criteria of reasonable expectations necessary for implicit consent.

In addition to reasonable expectation, implicit consent requires an absence of coercion. On the

surface, visiting a web site that employs a recommender system appears to meet this criterion. However, closer examination reveals problems here as well. For participation to be voluntary, a user would need to have viable alternatives to achieve the same ends. To the extent that some activities are now only available through the web (e.g., accessing certain job announcements only visiting a web site with a recommender system if they are to access these goods. The question becomes is there a diverse enough group web sites – some without recommender systems -- such that users have the option to choose a site that reasonably matches the activities to which they wish to consent. If that was the case, then there could be leverage for implicit consent. However, historically, we have seen extraordinary similarity among competing web sites with respect to recommender systems, user control, and informed consent.

Taking both problems together, the argument for implicit consent carries little force.

## **CONCLUSION**

Recommender systems have the potential to benefits for users. When used well, they can provide people with tailored information to support decision-making. However, recommender systems that run roughshod over informed consent in the process of generating their user database risk violating user's privacy and undermining user's trust in web interactions. I believe that our design of recommender systems would profit from a systematic and cohesive analysis of online informed consent. In that spirit, I offer my participation in the workshop on interacting with recommender systems.