

Automatic Resolution of Normative Conflicts in Supportive Technology Based on User Values

ALEX KAYAL and WILLEM-PAUL BRINKMAN, Delft University of Technology
 MARK A. NEERINCX, Delft University of Technology & TNO Perceptual and Cognitive Systems
 M. BIRNA VAN RIEMSDIJK, Delft University of Technology

Social commitments (SCs) provide a flexible, norm-based, governance structure for sharing and receiving data. However, users of data sharing applications can subscribe to multiple SCs, possibly producing opposing sharing and receiving requirements. We propose resolving such conflicts automatically through a conflict resolution model based on relevant user values such as privacy and safety. The model predicts a user's preferred resolution by choosing the commitment that best supports the user's values. We show through an empirical user study ($n = 396$) that values, as well as recency and norm type, significantly improve a system's ability to predict user preference in location sharing conflicts.

CCS Concepts: • **Human-centered computing** → **Empirical studies in collaborative and social computing**; *User models*; • **Computing methodologies** → **Multi-agent systems**;

Additional Key Words and Phrases: Social media, location sharing, social commitments, normative frameworks, conflict resolution, user values

ACM Reference format:

Alex Kayal, Willem-Paul Brinkman, Mark A. Neerincx, and M. Birna van Riemsdijk. 2018. Automatic Resolution of Normative Conflicts in Supportive Technology Based on User Values. *ACM Trans. Internet Technol.* 18, 4, Article 41 (May 2018), 21 pages.
<https://doi.org/10.1145/3158371>

1 INTRODUCTION

Supportive technology such as personal assistant agents, virtual coaches, location sharing systems, and smart homes have the potential to make our lives more connected, healthy, efficient, and safe. However, research in value-sensitive design and philosophy of technology shows this may come with the risk of demoting other important user values, such as privacy and responsibility (Czeskis et al. 2010; Nihlen-Fahlquist 2013; Nissenbaum 2010). A value is defined in the Cambridge Dictionary as “the importance or worth of something to someone.” Many different values can be distinguished. In particular, Rokeach (1973) published a surveyed list of human values that has become widely used, including, for example, friendship, happiness, and freedom.

Research in philosophy and normative systems (Bench-Capon 2003; van der Weide 2011; Van de Poel 2013), as well as our previous empirical research (Kayal et al. 2014a), observes that values can

This publication was supported by the national Dutch program COMMIT.

Authors' addresses: A. Kayal, W.-P. Brinkman, and M. B. van Riemsdijk, Department of Intelligent Systems, Delft University of Technology, P.O. Box 5031, 2600 GA Delft, The Netherlands; emails: {a.kayal, w.p.brinkman, m.b.vanriemsdijk}@tudelft.nl; M. A. Neerincx, Department of Intelligent Systems, Delft University of Technology, P.O. Box 5031, 2600 GA Delft, The Netherlands and TNO Perceptual & Cognitive Systems, P.O. Box 5031, 3769 ZG Soesterberg, The Netherlands; email: mark.neerincx@tno.nl

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 ACM 1533-5399/2018/05-ART41 \$15.00

<https://doi.org/10.1145/3158371>

be promoted and demoted by (regulatory) norms, such as action guiding statements obligating or prohibiting actions (Hansson 1991). Inspired by this observation, we have put forward the vision that in order to provide improved support for user values, supportive technology should be able to understand and adapt its behavior to diverse and evolving norms of people at runtime—that is, it should be socially adaptive (van Riemsdijk et al. 2015b). This is in contrast to existing supportive technology in which norms are hardwired.

An important challenge that needs to be addressed when making software socially adaptive is dealing with conflicts between norms. New norms can be introduced at runtime, and a situation may arise in which these norms cannot be fulfilled simultaneously. Various methods for detecting, reasoning about, and resolving normative conflicts have already been proposed in the literature (Vasconcelos et al. 2009; Criado et al. 2015; Oren et al. 2008; Ajmeri et al. 2016; Meneguzzi et al. 2015)—for example, scope curtailment (limiting the scope of influence of norms in conflict) and norm ranking, and policies for defining preferences between norms, such as *lex superior* (the norm imposed by the higher power takes precedence) or *lex posterior* (the most recent norm takes precedence).

Since in the context of socially adaptive supportive technology norms originate from users of the system with the aim of guiding the system to provide better support to these users, we argue that the technology should be able to resolve normative conflicts in a way that is aligned with these users' preferences. As a step toward creating technology that can resolve normative conflicts on users' behalf based on their preferences, we study factors that may influence these preferences. Since the underlying motivation for creating this technology is its envisaged improved support for people's values, in this article we specifically focus on how we may use information about people's values to predict their conflict resolution preferences.

The idea we propose in this work is that based on information about (1) how a user ranks the importance of a number of relevant human values within the application domain, and (2) the extent to which specific norms promote these values, the system can resolve the conflict by choosing the norm that best supports fulfillment of the user's values. We call a user ranking of the importance of a set of values a *value profile*.

Taking this idea as the starting point, we provide two main contributions in this article. First, we develop a normative conflict resolution model based on value profiles (Section 3). Second, we show in Sections 4 and 5 through an empirical user study in the domain of mobile location sharing in family life (described in Section 2) that this model can significantly improve a system's ability to predict user preference for resolution of normative conflicts. In addition, we found that other variables, namely recency and norm type (obligation or prohibition), can improve prediction more than user value profiles, and that a combination of all three variables provides the best prediction of user preference. We discuss these findings in Section 6.

2 CASE STUDY

We have selected social data sharing applications, particularly mobile location sharing for families, as our application domain for developing and studying prediction models for user preferences of normative conflict resolution. Allowing parents and children to share their location through mobile technology can support children in exploring their environment, through, for example, helping them go to school on their own, making new friends, and participating in neighborhood events and play dates, as well as increasing parents' awareness of the location of their children. We have chosen this domain since it is well known from the literature that its use can give rise to value tensions (Czeskis et al. 2010; Nihlen-Fahlquist 2013; Vasalou et al. 2012; Hasinoff 2017) (i.e., where promoting certain values comes at the expense of demoting others). Moreover, more and more applications of this type of data sharing and surveillance technology are developed and

used.¹ This makes the investigation of location sharing technology for families not only a means for studying our broader research questions but also relevant for its own sake.

The starting point for the research presented in this article is our previous work (Kayal et al. 2014b), in which we have developed a smartphone app² for family life location sharing based on an exploratory user study (Kayal et al. 2014a). We introduced the idea that social commitment (SC) models—as a specific type of normative model—provide a flexible yet easy-to-use structure to govern sharing and receiving of (location) data. SC models were proposed by Singh (1999, 2008) to describe a commitment between two parties in a sociotechnical system, namely a *debtor* who is committed toward a *creditor* for bringing about a certain proposition, or a *consequent*, when a certain *antecedent* comes to hold.

Our app comes with an interface that allows users to create commitments expressing in which situation which data should and should not be shared and received. For example, a commitment that can be created between a father Bob and his daughter Alice through the app is that Alice should share her location with Bob when Alice is at the park. Once a commitment is created, its behavior follows—broadly speaking—the commitment lifecycle as detailed in Telang and Singh (2011). This means that the app shares and receives location data (if possible) in accordance with the commitments to which the user has subscribed.

Conflicts between commitments may occur (see Ajmeri et al. (2016)) when a user subscribes simultaneously to a number of commitments that may obligate and prohibit the same action (this is called a *prohibition conflict* in van Riemsdijk et al. (2015a)). For example, when one commitment between user A and B specifies that location data from A should be shared with B when A is at the park (e.g., to promote the value safety), and another commitment specifies that this data should not be shared between 3pm and 5pm (e.g., to promote the values privacy and independence), then a conflict occurs when A is at the park between 3pm and 5pm. If this occurs, the app needs to be able to make a decision on which of the two conflicting commitments to satisfy, at the expense of violating the other. The mobile application currently resolves conflicts by selecting the most recent commitment. In this article, we investigate the use of information about users' values for selecting which of two commitments to satisfy.

3 CONFLICT RESOLUTION MODEL

In this section, we present our automatic conflict resolution model for SCs that govern sharing and receiving of data in social platforms. We define a language for creating requests regarding sharing or receiving of location data, and we define the notion of conflict used in this study in Section 3.1. The conflict resolution model is based on the concept of value profiles that we define in Section 3.2, and we present the model for predicting user preference in resolving conflicts in Section 3.3.

3.1 SC Request Language and Conflict Definition

Commitments can be created through the location sharing app described in Section 2 in the following way. The prospective creditor specifies a location sharing request—for example, a parent wants a child to share or not share location in a certain situation—through the graphical interface of the app. This request is sent to the prospective debtor (the child in this example), who can

¹Examples of existing location sharing applications are Life360, Glympse, and wearables such as KizON. Use of these technologies seems to differ across countries. Results from a survey among 920 parents in the UK indicate that the use of location tracking was not prevalent (Vasalou et al. 2012). However, in the United States, the app Life360 is being used by more than 34 million families according to the company (Hasinoff 2017).

²A 3-minute tutorial video (with subtitles) can be seen at <http://bit.do/ePartner>.

decide to accept or decline the request. If the debtor accepts, a commitment is created with the corresponding debtor and creditor, as well as the condition under which data should be shared or received as specified in the request. In the following, we provide the grammar of this language for expressing location sharing requests:

$$\begin{aligned} \langle request \rangle &::= \text{'I want'} \langle debtor \rangle \text{'to'} \langle normType \rangle \langle action \rangle \text{'with/from'} \langle thirdParty \rangle \text{'if'} \langle condition \rangle \\ \langle debtor \rangle &::= \langle individual \rangle \\ \langle normType \rangle &::= \text{'not'} \mid \epsilon \\ \langle action \rangle &::= \text{'share location'} \mid \text{'receive location'} \\ \langle thirdParty \rangle &::= \langle individual \rangle \mid \text{friends} \mid \text{family} \mid \text{others} \mid \text{everyone} \\ \langle condition \rangle &::= \text{'he/she is at'} \langle location \rangle \mid \text{'within'} \langle timePeriod \rangle, \end{aligned}$$

where *debtor* is an individual who forms the target of the commitment that is to be created; *norm type* resolves to either an obligation (empty) or a prohibition (not) of an *action*, which is either to share or receive location information; *third party* is either an individual (e.g., the creator of the request, also known as *creditor* and referred to as “I” in the grammar) or a group of users, such as friends, family members, others, or everyone (i.e., all listed users); and *condition* is either a location or a time period. If the debtor accepts the request, a commitment is created where the debtor and creditor are as indicated previously, the antecedent is the condition, and the consequent is the combination of norm type, action, and third party—the latter can be viewed as the parameter of the action. With some abuse of language, in the following we will sometimes use the term *commitment* to refer to the commitment that is intended to be created through a request.

In the literature on SCs, the consequent typically represents a proposition that the debtor is committed to bringing about. In our case, the consequent represents a sharing or receiving action that should or should not be executed. In line with literature on norms (Balke et al. 2013), we refer to the former as obligations and the latter as prohibitions, which can also be referred to as obligations not to do the action. Furthermore, actions and conditions in our case are specific to the domain of location sharing. We introduce a third party to specify with/from whom data is shared or received, which can be viewed as a parameter of the specified action.

The definition of conflict as introduced in the following underlies the implementation of conflict detection in the application that we employed in the user study presented in this article. Informally, a conflict can occur when two commitments refer to the same debtor, have opposing norm types (i.e., one is an obligation and one is a prohibition), concern the same action (with overlapping third party), and have an overlapping condition. Two conditions overlap when (1) one is a location condition and the other is a time condition (because a person may be at that location at that time), or (2) both are location conditions and they are the same,³ or (3) both are time conditions with overlapping timespan (e.g., “between 8am and 10am” overlaps with “between 9am and 5pm”), denoted as *TimespanOverlap(timespan1, timespan2)*. We use the notation $C.debtor$ to refer to the grammar element “debtor” of commitment C , in correspondence with the grammar defined previously.

Before we define the notion of conflicting commitments formally, we define what we mean by third-party overlap and overlapping conditions. We use $C.condition.type$ to refer to the type of the condition of commitment C (i.e., either place or time period).

Definition 3.1 (Third-party Overlap). Let $C1$ and $C2$ be commitments, and let M be the intersection of the set of third parties of $C1$ and $C2$ (i.e., $M = C1.thirdParty \cap C2.thirdParty$). We define that $C1$ and $C2$ have a third-party overlap, denoted as $Overlap(C1.thirdParty, C2.thirdParty)$, if and only if $M \neq \emptyset$.

³In this study, we assume that locations with different names are geographically different locations.

Definition 3.2 (Overlapping Conditions). Let $C1$ and $C2$ be commitments. We define that $C1$ and $C2$ have a condition overlap, denoted as $Overlap(C1.condition, C2.condition)$, if one of the following two conditions hold:

- (1) $C1.condition.type \neq C2.condition.type$, or
- (2) $C1.condition.type = C2.condition.type$, and
 - (2a) $C1.condition.type = location$ and $C1.condition = C2.condition$, or
 - (2b) $C1.condition.type = timePeriod$ and $TimespanOverlap(C1.condition, C2.condition)$.

Definition 3.3 (Conflict). Let $C1$ and $C2$ be commitments. We define that $C1$ and $C2$ are in conflict if and only if the following conditions hold:

- (1) $C1.debtor = C2.debtor$
- (2) $C1.normType \neq C2.normType$
- (3) $C1.action = C2.action$
- (4) $Overlap(C1.thirdParty, C2.thirdParty)$
- (5) $Overlap(C1.condition, C2.condition)$.

3.2 Value Profiles

Employing users' contextual information has already been established as a viable method to provide more relevant recommendations and a better user experience (Fernández-Tobías et al. 2016; Panniello et al. 2012; Knijnenburg et al. 2012). This, in addition to the link between user values and norms, brought forth the idea of using users' ranking of importance of a number of domain-relevant values as contextual information—to predict their preferred solution if a normative conflict is to occur.

We define a value profile as a user ranking of the importance of a set of values that are relevant in the domain under consideration, which in our case is location sharing in family life. In Kayal et al. (2014a), we have identified several values from Rokeach's survey (Rokeach 1973) as particularly relevant in this domain, namely friendship, family security (here renamed as *safety*), independence, social recognition, and inner harmony. In this domain, social recognition takes shape mainly in the form of friendship, and inner harmony particularly concerns family security. Therefore, and to limit the number of values that users have to rank, in this study we omit social recognition and inner harmony. Moreover, we add the values of responsibility and privacy, as these have been identified in Nihlen-Fahlquist (2013) and Czeskis et al. (2010) as important in this domain and in data sharing in general (Nissenbaum 2010).

We define these values as follows, adapted from Merriam-Webster's dictionary:

- *Friendship (Frnd)*: For you or your family members to build friendships, a social life, and be recognized among others in the social circle
- *Privacy (Priv)*: For you or your family members to be free from unwanted outside intrusion and undesirably shared information
- *Safety (Saf)*: For you or your family members to be free from danger or harm
- *Independence (Ind)*: For you or your family members to be capable of doing what they need to do without another's control or support
- *Responsibility (Res)*: For you or your family members to know and be able to do the tasks that they are expected to do.

To predict users' preference in the resolution of two conflicting commitments, we introduce two types of value profiles: one that provides information about the user's ranking of the importance of a set of values in general and one that provides information about the extent to which a specific

commitment promotes these values. We call the first a *general value profile* (VP_g) and the second a *commitment value profile* (VP_c). By comparing each of the VP_c of two conflicting commitments to the VP_g of a user, one can determine which of the two commitments' profiles is closer to the user's values in the general sense. The idea is then that the commitment closest to the user's VP_g is the commitment that the user would prefer to fulfill in case of conflict.

These two types of value profiles are thus defined as follows, where the domain-relevant values in our case are the five values listed earlier:

- A user's *general value profile* (VP_g): A ranking of the importance of a set of domain-relevant values in the general sense, without any additional context
- A *commitment value profile* (VP_c): A ranking of how an SC promotes a set of domain-relevant values.

In practical terms (i.e., when an application that embodies this conflict resolution model is in use), a user's general value profile will be created as part of the initialization process of the mobile app, whereas commitment value profiles are created by the user every time a request is accepted by the debtor.

3.3 Preference Prediction Model

Our model for predicting which of two conflicting commitments a user will prefer is based on calculating the distance between the value profiles of each of these two commitments and the user's general value profile. The commitment that is closest to the user's general value profile is predicted to be the user's preferred solution for resolving the conflict.

We represent value profiles numerically as vectors. Each element of the vector corresponds to the importance of a particular value—that is, the higher the number, the more important a value is within that profile. For normalization purposes, the sum of the elements of the vector should add up to 1. Thus, a value profile in our case is a five-dimensional vector, representing the relative importance of each of the identified five values relevant in this domain:

$$VP = \langle Frnd, Priv, Saf, Ind, Res \rangle,$$

where $Frnd + Priv + Saf + Ind + Res = 1$.

To illustrate how the model defines the distance between value profiles, consider the values *safety* and *independence*. Let C_1 and C_2 be two commitments, and let VP_g , VP_{c1} , VP_{c2} represent the general value profile, and the value profiles for C_1 and C_2 , respectively. Figure 1 illustrates a partial projection of VP_g , VP_{c1} , VP_{c2} on a two-dimensional plane, showing how close each commitment value profile is to the general value profile. According to this illustration, the safety component in C_1 is closer to its counterpart in the user's VP_g than the safety component of C_2 (i.e., $|Saf_{c1} - Saf_g| < |Saf_{c2} - Saf_g|$). This means that with respect to *safety*, the model predicts that C_1 would be favored over C_2 in case of a conflict. Using the same argument, we can see that according to the value *independence*, C_2 would instead be favored over C_1 .

In this way, we calculate for each of the five values the distance between the general value profile and the value profile of the commitment, defined formally as $|(VP_c - VP_g)|$. We do this for each of two commitments and take the difference between the resulting two vectors to obtain a prediction vector $Pred$:

$$Pred_{c1,c2} = (|(VP_{c2} - VP_g)| - |(VP_{c1} - VP_g)|),$$

for instance, a vector containing five predictive components:

$$Pred_{c1,c2} = \langle Frnd_p, Priv_p, Saf_p, Ind_p, Res_p \rangle.$$




Fig. 1. A two-dimensional representation of safety and independence in VP_g , VP_{c1} , and VP_{c2} .

Each component of this prediction vector represents how close the importance of a certain value in each of C_1 and C_2 is to its importance within the user's value profile. This number thus reflects how much the user is predicted to prefer C_1 over C_2 in a potential conflict with respect to that value. This number can be positive or negative (a negative number means a preference for C_2 over C_1).⁴ A numerical example can be found in Appendix B.

4 USER STUDY

We designed and performed a user study to determine the usefulness of our value-based conflict resolution model for predicting user preferences in resolving conflicts between commitments. The design of the user study was made relatively simple to allow nonexperts on the subject of SCs to perform the required tasks. In this user study, participants were provided with a number of location sharing scenarios in the family life domain. Each of these scenarios ended with a location sharing challenge that required a solution to be created using the SC request language of Section 3.1. The study was designed so that participants were confronted with conflicts between commitments if they were to provide the expected solutions in the scenarios. As part of the study, we elicited users' value profiles and their preferred solution when a conflict occurred. Our aim was to use our conflict resolution model to predict users' preference using information available in their value profiles and compare that prediction to the preference they reported. In this section, we present our hypotheses and research questions and describe the user study in more detail.

Permission from the ethics committee of the university was obtained prior to conducting this user study.

4.1 Hypotheses and Research Questions

Based on the background material and our value profile-based predictive model, we propose the following hypotheses and research question:

⁴In the case of equal VP_{c1} and VP_{c2} , the model will predict an equal user preference for the two conflicting commitments. Within the dataset obtained in the user study related to this work, no such case of an equal prediction was found—although technically possible, it was very unlikely because of the fine-grain and multidimensional method we used for input (see Section 4.3). A reduced grain/dimensionality of the input method would increase the likelihood of equal commitment value profiles happening.




Fig. 2. A Web-style menu representation of the SC request grammar.

- *H1*: People have a preferred resolution when confronted with a conflict between commitments.
- *H2*: Knowledge of people’s general value profiles and commitment value profiles can be used to predict people’s preferred resolutions to conflicts between commitments.
- *RQ3*: Which information within a commitment’s grammatical structure can be used to predict people’s preferred resolutions to conflicts between commitments?

4.2 Material

Research has shown that the Web offers an environment just as powerful as the lab for conducting user studies, with data collected online being of at least similar quality to lab data (Gosling et al. 2004). Results from the two environments have been shown to have high congruence (Krantz and Dalal 2000). An advantage of using the Web is that large numbers of participants can be recruited relatively easily through crowdsourcing platforms. Following this approach, we implemented a Web site containing the tasks that participants had to perform. Participants were recruited through [Microworkers.com](https://www.microworkers.com) and were redirected via a link to our user study Web site.

We have developed a Web-style menu representation of the SC request language to allow participants to create SCs (Figure 2). The menu reflects the user interface of the corresponding smartphone app (see Section 2).

4.2.1 Scenarios and Conflicts. Sixteen scenarios were used in the study, describing fairly common situations within the family life domain, such as children going to school, children playing at a playground, and parents taking their children to meet friends. The origin of these scenarios is rooted in focus group data with members of the target group conducted in Kayal et al. (2014a). A location sharing challenge was presented at the end of each scenario, which participants were asked to solve by creating a data sharing request using the SC request language through the menu in Figure 2. Every scenario was assigned a designated solution (i.e., a specific commitment that we deem to be correct). Scenarios were created such that the commitments forming the designated solutions for each of the 16 scenarios were distributed over 16 possible combinations of norm type (i.e., *obligation* or *prohibition*), action (i.e., *share* or *receive*), third party (i.e., *creditor only* or *any other user or any group*), and condition (i.e., *place* or *time*).

These 16 scenarios were created such that they gave rise to eight conflicting pairs of scenarios. A conflicting pair consisted of 2 scenarios where the commitments forming its two designated solutions would cause a potential conflict according to the definition of conflict of Section 3.1. An example of a pair of scenarios with two designated solutions bearing a potential conflict can be found in Appendix A.

4.2.2 Roles. In each of the 16 scenarios, participants had to assume the role of one of the characters in the scenario—the character is meant to solve the problem in the scenario through creating a commitment with other characters. In 8 of these scenarios, the participant assumed roles of parents, and in the remaining 8 scenarios, participants assumed the roles of 8-year-old primary school




Fig. 3. Values pie chart and legend.

children. A participant assumed a fixed role that did not change throughout their participation. This means that in the scenarios we used, each conflicting pair of commitments was created by the same creditor as impersonated by the participant of the study.

4.2.3 Instructional Videos. Two instructional videos (narrated in English) were created for this user study. The first video provided information regarding the domain, SC request menu, an example task, and the required participant input in case of a normative conflict. Video instructions were customized depending on a participant's assumed role (i.e., parent or child). The second video explained how to operate the values pie chart—the measurement tool we used for elicitation of user values (see Section 4.3).

4.3 Measurement

Although obtaining the relative importance of a mental construct such as human values may be difficult, Carenini and Loyd (2004), Pommeranz et al. (2011), and Hultgren et al. (2014) provide several methods for the visual elicitation of the ranking of a fixed number of user values.

Participants' VP_g s and VP_c s were obtained using a colored pie chart with resizable slices, and a legend relating every slice to a specific value of the five values discussed in Section 3.2 (Figure 3 shows an example of what a pie chart may look like next to the legend). The larger the slice that referred to a certain value, the more important a participant thought this value was in comparison to others, considering the role that they were instructed to play.

If a participant created two successive solutions using the SC request menu whose corresponding commitments were in conflict, a pop-up window showed up at the end of the second scenario. This window displayed information related to the two conflicting solutions. It asked the participant to indicate, from the perspective of his or her character in the scenario (i.e., the creditor of both

Please indicate your preference

It appears that a conflict may occur between the last two agreements that have been created:

Agreement 1: I want Jason to not share check-ins with everyone if he is at the park
 Agreement 2: I want Jason to share check-ins with Mary if it is between 12:00 PM and 07:00 PM

If a conflict occurs, the system must make a choice and comply with only one of the two. Please use the slider below to indicate, from the perspective of mike, how much you favor one agreement over the other.

Fully in favor of Agreement 1 | No preference | Fully in favor of Agreement 2

Proceed

Fig. 4. A pop-up asking the participant to indicate his or her preference, using a slider.

commitments) how much he or she favors one commitment⁵ over the other using a continuous slider (Figure 4).

Note that this setup allowed us to study the conflict resolution preferences of the (same) *creditor* of two conflicting commitments. This setup was chosen since it concerns the most “direct” relation between values and conflict resolution preferences. This is because the commitments originate from the creditor to promote the creditor’s values, and we resolve the conflict by comparing these commitments’ value profiles with the general value profile of the same creditor. Studying how values can be used to predict a debtor’s conflict resolution preferences may require taking into account the debtor’s perspective on the creditors’ value profiles, as well as authority relations between debtor and creditor (in accordance with the conflict resolution criterion *lex superior*) if commitments arise from different creditors. Since this is our first study in this direction, we chose a simple setup.

4.4 Participants

For the study, 400 participants were recruited through Microworkers.com. Participation was open to members living in English-speaking countries (i.e., the United States, Canada, UK, Australia, and New Zealand). Every participant was compensated with \$1, in accordance with the regulations of the crowdsourcing platform.

To ensure the quality of the participant’s contributions, every task contained a quality control question. Only participants who read the task text in full would be able to answer this question correctly. Participants were informed through our terms and conditions that wrong answers to the quality control questions would result in their compensation being canceled and their contribution omitted from the study. Four of the participants did not comply with these regulations, which resulted in a final number of 396 contributing participants. Of these participants, 202 were male and 194 were female, with an age mean of 31.2 years and a standard deviation of 10.8, whereas 156 participants indicated being legal guardians of one or more children.

⁵We used the term *agreement* instead of *commitment* during the experiment for clarity.

4.5 Procedure

Upon reaching the Web site's landing page, participants were instructed to view the first instructional video, customized based on their randomly assigned role. After watching the video, participants were asked to enter demographic information, such as age, gender, and whether they were the legal guardian of one or more children.

Following the landing page, participants were directed to a practice page, where a dummy scenario, a practice SC request menu, a practice values pie chart and colored legend, along with the second instructional video, were presented. Following the practice page, participants were directed to a page that contained the values pie chart and the colored legend. Participants were instructed to "indicate, in the general sense, [their] preference for these five human values" using the pie chart, and to do this "from the perspective of [their] role" (i.e., either a parent or a child) and within the context of family life. This yielded the participant's VP_g .

Next, participants were directed to the scenario pages. In a scenario page, participants could read the scenario text and attempt to solve the location sharing problem in the scenario using the SC request menu. After that, participants indicated how much the solution, namely the request they created, supports each of the five human values if it were accepted using a similar values pie chart, within the context of the scenario and from the perspective of their character in the scenario. This yielded the corresponding commitment's VP_c . Every participant had to solve eight such tasks, dispatched as conflicting pairs but in random order. This meant that every two consecutive scenarios had designated solutions generating a potential conflict, which participants had to manually resolve using the continuous slider in Figure 4.

Finally, after the end of the eighth scenario, participants were directed to a page containing a second value pie chart with a legend, and they were instructed to indicate, once more, their preference for these five human values in the general sense assuming their role and within the context of family life (i.e., their VP_g).

4.6 Data Preparation and Preanalysis

R version 3.2.1 was used for all statistics. Participant demographic data, VP_g (pre and post), assumed role, order of dispatched scenarios, solution to every task, commitments' VP_c , and users' preferences for every conflict resolution were stored.

First, a reliability analysis was conducted for values within VP_g (pre and post) among participants to determine if there was significant change to merit using the average of profiles in further analysis. Results showed a satisfactory⁶ Cronbach's α (Table 1): this means that we can assume consistency among pre- and postexperiment measurements. Therefore, for further analysis, only the premeasurement value of VP_g was used.

We also analyzed to what extent there is agreement among participants regarding how they viewed a commitment's impact on the five values. For this purpose, we performed a reliability analysis among participants for values within VP_c s per scenario, and split across roles (Table 2). Results suggest a high level of consistency between participants in how they viewed a commitment's impact on the five values. This means that it may not be necessary to use the commitment value profiles of individual users, but instead it may suffice to use the average of a number of users' commitment values profiles to predict the preferred solution to a conflict. Relying on consensus data would require less elicitation of users' input. To investigate this, for each commitment that

⁶As an internal reliability measure, Loewenthal (2001) suggests a Chronbach's α of .7 as a threshold for acceptable reliability. However, in a scale of a low (i.e., below 10) number of items (in the case of this article, two: before and after), one may not be able to obtain an acceptable value of α , and thus the threshold may be reduced to .6. In our case, still, one value (privacy) is still considerably below that; however, an average of reliability for all values would still exceed .6.

Table 1. Reliability Analysis for VP_g s

	α
Frnd	.75
Priv	.52
Saf	.63
Ind	.70
Res	.65

Table 2. Reliability Analysis for VP_c s per Scenario, Across Roles

	α	
	$Role_{Parent}$	$Role_{Child}$
Frnd	.97	.97
Priv	.99	.99
Saf	.99	.98
Ind	.94	.98
Res	.97	.91

formed a designated solution to a particular scenario, the average of the VP_c s was calculated for all participants. This created a value profile consisting of the average of all value profiles created for the commitment in that scenario. We call these *consensus value profiles* or VP_{cons} .

Since participants performed multiple tasks, data was transformed to longitudinal form, with two tasks (i.e., one potential conflict between two designated solutions) per row. This way, every row in the data represents a potential conflict and resolution. This generated $396 \cdot (8/2) = 1,584$ rows. In 517 rows, no conflict between commitments was created (i.e., a participant failed to create two conflicting designated solutions). These rows were then dropped, leaving a total of 1,067 rows out of the original 1,584.

To calculate predictions of user preference for conflict resolution (as per the model introduced in Section 3), two different prediction vectors were used. The first one was a *fully personalized* prediction, using a user's VP_g s and the VP_c s they created for every commitment:

$$(1) \text{Pred}_{FP} = (|(VP_{c2} - VP_g)| - |(VP_{c1} - VP_g)|).$$

Although the second was a *semipersonalized* prediction, using a user's VP_g s but replacing their VP_c s with the VP_{cons} . As explained, the latter represents a consensus of opinions rather than users' own opinion of a commitment's impact on values (and hence semipersonalized).

$$(2) \text{Pred}_{SP} = (|(VP_{con2} - VP_g)| - |(VP_{con1} - VP_g)|)$$

Figure 5 represents a visualization of all participants slider values (Figure 4) when presented with a conflict, showing the distribution of users preference for one commitment over another: the closer to the left (right), the more they preferred C1 (C2) respectively, and the closer to the center, the more indifferent they were regarding that conflict. The data shows that participants deviated from the neutral, "no preference" point. This means that people have a preferred resolution when confronted with a conflict between commitments (H1).

For the purpose of analysis and creating prediction models, we translated the slider data into binomial form with the neutral "no preference" as the cutoff point. In other words, measurements




Fig. 5. Histogram of slider data for all participants.

<0 are taken as a preference for the first of two conflicting commitments, and measurements ≥ 0 are taken as preference for the second. The translation to binomial form was performed because resolving an instance of conflict between two commitments means complying with one commitment and violating the other.⁷ The data shows that participants were more likely to favor the second commitment they created (65.5% of total) in case of a conflict, regardless of the content. This means that the order in which commitments were created appears to influence users' preference in case of a conflict. Interestingly, this empirically confirms the relevance of the *lex posterior* policy (see Section 1) for resolving conflicts between norms in the context of supportive technology.

Moreover, the data showed that participants were more likely to favor the commitment created at a scenario where the designated solution contained an *obligation* norm type (62.9% of total). This means that the norm type of a commitment appears to influence users' preference in case of a conflict (RQ3).

In summary, we consider three main factors as possible predictors for user preference in conflict resolution: (1) order, (2) norm type, and (3) user value profiles. For the latter, we consider two options: the fully and the semipersonalized commitment value profile. The order is always taken into account in the prediction models, as it cannot be considered in isolation. This means that we have a total of $2^3 = 8$ possible combinations of predictive factors.

Correspondingly, eight multilevel, linear mixed effect (LME) models were constructed using the *nlme* package of R. Linear models describe relationships in our data between predictive factors and the outcome in terms of a linear formula. LME models contain two types of factors: fixed effects and random effects. Fixed effects are the predictive factors that are within experimental control, in our case norm type and value profiles. Random effects are factors that are outside experimental control, particularly unknown participant-specific factors. Accounting for these in the model is important in case of a repeated measures study, with multiple measurements per participant, as in our case. These measurements are not independent: they are influenced by participant-specific factors that are unknown to the experimenter at the time of the measurement. This could introduce a bias to measurements from an individual participant. A random effects component is added to the model to account for this idiosyncratic variation due to individual differences. This type of model

⁷The choice for obtaining user input via a continuous slider would allow to test for H1, then H2 through translating that input into a binomial form. We opted out of using a three-choice input (i.e., C1, no preference, C2), as it may lead to a more salient choice for "no preference" for participants with a weak preference.

Table 3. Fixed Effects Used in Each of the Eight LME Models

	M_0	M_1	$M_{2.1}$	$M_{2.2}$	$M_{2.3}$	$M_{3.1}$	$M_{3.2}$	$M_{3.3}$
Intercept	x	x	x	x	x	x	x	x
Norm type		x				x	x	x
Group								
Pred _{FP}			x		x	x		x
Pred _{SP}				x	x		x	x

Table 4. Fixed Effect Coefficients for the Eight LME Models

	<i>estM0</i>	<i>estM1</i>	<i>estM2.1</i>	<i>estM2.2</i>	<i>estM2.3</i>	<i>estM3.1</i>	<i>estM3.2</i>	<i>estM3.3</i>
Intercept	.655**	.526**	.658**	.653**	.657**	.541**	.550**	.556**
Norm type		.247**				.226**	.201**	.198**
Pred _{FP}								
Frnd _{FP}			.001**		.001**	.001*		.001*
Priv _{FP}			.000		.000	-.001		-.000
Saf _{FP}			.002**		.001	.001*		.000
Ind _{FP}			.000		.000	-.000		.000
Res _{FP}			.001		.000	.000		.000
Pred _{SP}								
Frnd _{SP}				.001	.000		.001	.000
Priv _{SP}				-.000	-.000		-.001	-.001
Saf _{SP}				.004**	.003**		.002**	.002*
Ind _{SP}				.001	.001		-.000	-.000
Res _{SP}				.003*	.003*		.001	.001
%Prediction	68.4	77.6	72.3	74.4	75.1	77.8	77.0	77.9

* $p < .05$; ** $p < .01$.

containing both fixed and random effects is referred to as a mixed model. Testing the significance of a fixed factor was done by examining the improvement in the models fit on user preference in the conflict resolution data if the model was extended with this fixed effect. For more elaborate introductory explanation, we refer the reader to Winter (2013), as well as Field et al. (2012) for more of a general overview of statistical modeling techniques.

Composition of the models is as follows: in all eight models, the binomial user preference was used as a response (i.e., the output of the prediction), *participant* as a random effect, with an unstructured covariance matrix (i.e., making no assumptions of any relationship between the variances in intraparticipant measurements). Fixed effects (i.e., the predictive factors) used in each of the eight models are shown in Table 3. The intercept concerns the order in which commitments were created, norm type refers to the type of the norm in a commitment (i.e., obligation or prohibition), and the term *Group* used in the table refers to the fact that *Pred_{FP}* and *Pred_{SP}* are each a set of predictive values rather than a single one. Table 4 shows the fixed effect coefficients for the eight LME models.

To determine whether the improvement that a model provides over another model is significant, log-likelihood comparison tests were conducted. Log-likelihood is a measure of fitness of a statistical model. In themselves, log-likelihoods are uninterpretable; however, the difference between the log-likelihood for two models is interpretable as it follows χ^2 distribution, which is a standard measure of difference between expected and observed outcomes. And this can be compared

Table 5. Results of the Log-Likelihood Tests and the R^2 Values

	$\chi^2(R^2)$						
	M1	M2.1	M2.2	M2.3	M3.1	M3.2	M3.3
M0	75.0(.07)**	40.5(.04)**	57.6(.05)**	67.5(.06)**	88.8(.08)**	89.2(.08)**	97.0(.09)**
M1					13.8(.01)*	14.2(.01)*	22.0(.02)*
M2.1				27.0(.03)**			
M2.2				9.9(.01)			
M3.1							8.3(.01)
M3.2							7.9(.01)

* $p < .05$; ** $p < .01$.

to random differences, which means that we can see whether the observed difference is beyond “random chance” and hence significant.

Using log-likelihood comparison tests, a base model can be compared to another model in which fixed effects are added in comparison to the base model (i.e., not all models can be compared in this way). To determine the improvement that each model provided over the base model $M0$, seven log-likelihood comparison tests were conducted with $M0$ and each of the other seven models. To determine the effect of adding value profile predictors over a model containing the norm type predictor, three log-likelihood comparison tests were conducted with $M1$ and each of the $M3.x$ models. To determine whether fully personalized value predictors provided an improvement over the semipersonalized, with and without the presence of the norm type predictor, four log-likelihood comparison tests were conducted, with ($M2.1, M2.3$), ($M2.2, M2.3$), ($M3.1, M3.3$), and ($M3.2, M2.3$). R^2 values were obtained through comparing all of the preceding pairs.

To understand the magnitude (i.e., effect size) of the improvement of one model over another in its ability to explain the data, we also report the pseudo- R^2 (hereafter abbreviated as R^2) values as suggested by Finch et al. (2014). The R^2 value is the percentage of variance in the data that can be explained more by one model than by the other. For example, $R^2 = .1$ (or 10%) means that a model can explain 10% more of the observed outcome than the other model. Cohen (1988) classifies effect size of R for social sciences as small when $R = .1$, medium when $R = .3$, and large when $R = .5$. For R^2 , this means that a value of .01 can be seen as small effect size, .09 as medium, and .25 as large. Table 5 shows the results of the log-likelihood comparison tests and the R^2 values. An explanation (and choice) of the statistical methods used in this analysis can be found in Finch et al. (2014).

5 RESULTS

In Table 4, numbers inside the cells (aside from intercept) are the fixed effect coefficients of the linear model. In other words, in column $estM1$, the slope of .247 represents the norm type’s effect on model $M1$ ’s ability to explain user preference. The intercept is the constant in the linear formula. Since the formula’s 0 to 1 outcome represents our annotated recency of a commitment (i.e., 0 for the first commitment in the conflicting pair and 1 for the second), the intercept on its own here represents a prediction based on recency without any additional predictive factors.

The double asterisk next to that number represents a p value below .01 (see table notes), and thus norm type is considered very significant in this model. Based on this analysis, we can see that column $estM0$ shows that the baseline model (i.e., knowledge of commitment order and participant ID alone) can significantly predict user preference, with 68.4% predicted correctly. Column $estM1$ shows that a significant improvement in prediction can be obtained when adding norm type to the model, with 77.6% of user preferences predicted correctly. Columns $estM2.1$ to $estM2.3$ show

that knowledge of users' value profiles can significantly improve prediction over knowledge of commitment order and participant ID alone, with best prediction out of these three obtained using both $Pred_{FP}$ and $Pred_{SP}$ with 75.1% of the predictions correct. Columns $estM3.1$ to $estM3.3$ show that knowledge of both norm type and users' value profiles can significantly improve prediction over knowledge of commitment order and participant ID alone, with best prediction out of these three obtained using all of norm type, $Pred_{FP}$, and $Pred_{SP}$ with 77.9% of the predictions correct.

In Table 5, the numbers in every cell show the result of a comparison between two models. For example, the values in the uppermost left cell show that additional predictors in $M1$ (namely norm type) affected the ability to explain user preference relatively well ($\chi^2 = 75.0$), with a small to medium effect size ($R^2 = .07$). The double asterisk next to that number represents a p value below .01 (see table notes), and thus the change in prediction ability between $M0$ and $M1$ is very significant. Following this analysis, row $M0$ confirms that each of the seven models with fixed effects provides an improved explanation of user preferences over the base model, particularly with R^2 values suggesting small to a medium effect size (depending on model). Row $M1$ shows that adding value profile predictors to a model containing norm type would have little yet significant improvement in explaining user preferences. Rows $M2.1$ and $M2.2$ show that adding a semipersonalized prediction to a model containing a fully personalized prediction would offer little but significant improvement, if norm type was not included as a predictor. Rows $M2.1$ and $M2.2$ also show that the reverse (i.e., adding a fully personalized prediction to a semipersonalized prediction) would not offer any improvement in prediction. Rows $M3.1$ and $M3.2$ show that no improvement was found in both cases when norm type was included.

6 DISCUSSION

6.1 Hypotheses and Research Question

Regarding hypothesis H1, preanalysis in Section 4.6 has shown that participants were strongly in favor of a resolution for conflicts, as opposed to having no preference for one commitment over another, confirming this hypothesis. The results in Table 4 showed that the most accurate predictors of user preference are certain commitment-relevant information. Within its grammatical structure, norm type was found to be a significant predictor (thus answering RQ3). The table also shows that a commitment's recency was found to be significant as well. Value profiles provided a slight (yet significant) improvement over recency and norm type, with the highest prediction accuracy achieved when using commitment order, norm type, and value profiles all together (thus confirming hypothesis H2). Last, the results of model comparison in Table 5 show that fully personalized value profile predictors do not offer more predictive power than the semipersonalized ones.

6.2 Contributions

The main contributions of this article are (1) development of a conflict resolution model for SCs based on knowledge of user values and (2) a user study showing that this value-based model can be used to automatically solve data sharing conflicts in location sharing platforms. Aside from value profiles, our analysis revealed powerful yet simple and easy-to-obtain information, such as order and norm type, that can be used to significantly increase automatic conflict resolution prediction accuracy. To the best of our knowledge, we are the first to develop a normative conflict resolution model based on user information, particularly user value profiles.

6.3 Limitations

For our user study, we selected five human values relevant to the domain of location sharing in family life to make up the components of value profiles. A more comprehensive list of human

values could be used to provide a wider perspective on the values that users find important and relevant to location sharing commitments.

In the user study, we used 16 scenarios and fixed conflicting commitment pairs. These scenarios and conflicts were based on common family life situations and were rooted in previously collected focus group data (Kayal et al. 2014a). Yet despite our best efforts in selecting and pairing scenarios, and the consistency in results across conflicts, more research is needed to investigate generalizability of our findings to other location sharing scenarios and social data sharing domains.

Moreover, the study was conducted online using a crowdsourcing platform. This means that the conflicts and resolutions were simulated, and participants were in essence actors who simulated both parental and children roles within given scenarios. We therefore cannot immediately assume that real-life location sharing scenarios would generate the same results. Nevertheless, research (Borlund and Schneider 2010) suggests that simulated work tasks produce results that are comparable with real-world behavior. Moreover, working with real-world data has to be balanced against an efficient research approach and ethics justification. Obtaining such data would require the development of entire application, as well as asking participants to use the application over a long period of time, to evaluate only one element of the system. Using a gamification mechanism known as abstractions (Kapp 2012) was therefore more justified. With abstractions, participants were only exposed to a simplification of the situation by removing less relevant factors (e.g., a parent actually going to his or her office) while also making cause and effect clearer with time being sped up (e.g., participants did not wait for the period of a school day for the second scenario). The advantage of conducting a study in this controlled type of setting is a strong internal validity. Because we have control over the variables (Robson 2002), we are in a good position to attribute the observed effects to our manipulations instead of potential biases that may come from confounding variables in a field study. Although field studies have higher external validity, confounding variables (i.e., variables outside of experimental control) would make findings less generalizable—thus, we opted for a setup with a strong internal validity, with a view of conducting further research in a field setup.

In addition, using a crowdsourcing platform limited participation to those who chose to perform that task out of personal interest. This limits the generalizability of our findings beyond interested parties for the time being. Furthermore, the label referring to the more recent commitment made was always displayed on the right side of the slider, leaving our findings in regard to recency vulnerable to visual bias—although the effect of recency is more documented in the literature (Howard and Kahana 2002).

Moreover, and since we needed to ensure that participants were able to fully understand our scenarios, participation was limited to English-speaking countries only (i.e., primarily western cultures). Different cultures may, on average, rank their values differently. At first glance, this would not affect how the prediction model works—the model uses a user value profile(s), and sometimes community value profiles, to generate a prediction for that specific user. Different users within one community rank their values differently as well, and there is no reason to expect that the model will be less capable of predicting individual user preferences if it uses user value profiles and community value profiles from another culture that is equally as homogeneous as western culture. However, if we were to collect community profiles from various cultures and use their average in the semipersonalized prediction, then this may negatively affect the prediction accuracy for semipersonalized predictions.

6.4 Proposed Future Work

Our main finding is that our results provide evidence that values are a relevant factor influencing users' preferences regarding normative conflict resolution. This is important in light of our

overall aim of creating supportive technology that better supports people's values by adapting to their norms, for two main reasons: (1) it provides empirical evidence for the link between norms and values, which underlies our vision of socially adaptive supportive technology, and (2) if we can improve our understanding of the relation between values and normative conflict resolution preferences through further research, this may allow us to improve the predictive power of our conflict resolution models, leading to supportive technology that better supports people's values.

Improving our understanding of the relation between values and normative conflict resolution preferences involves also studying other factors that are (potentially) relevant for conflict resolution and their interaction with values. In this study, we have already identified two other factors (recency and norm type). We expect that a third important factor is the nature of the relation between debtor and creditor (e.g., an authority relation).⁸ This is particularly relevant when considering a debtor's conflict resolution preferences in case of conflicting commitments toward different creditors. It will be interesting to investigate if for these other factors we can also identify accompanying values as the underlying factor (e.g., respect for authority). Moreover, more research is required to investigate more involved interplay between values of different users. For example, debtors may take into account their own values, as well as their perception of the values of creditors, in establishing conflict resolution preferences.

In support of efforts to acquire a better understanding between values and normative conflict resolution preferences, we believe that an interesting next step would be to investigate other ways of obtaining value profile information. In this study, we asked participants to provide this information directly via a pie chart, and it would be interesting to obtain that profile indirectly, for example, through behavioral information or sensor data to investigate if such information can lead to better predictions. Furthermore, in our study, the starting point was a predetermined set of relevant values. It would be interesting to integrate and further develop value elicitation techniques, such as techniques for eliciting which values are important in the context of particular applications (Pommeranz et al. 2011).

Finally, this user study was conducted in a simulated setting with all-adult actors who simulated both parental and child roles within given scenarios in the location sharing domain. Conducting this research in a field setting with both parents and children with a location sharing mobile app would be necessary to confirm that our findings carry over to use of the technology in real life. An important challenge to consider when performing a field study with technology that automatically makes decisions on users' behalf, as in our case, is how to balance automatic decision making and user control over the application's behavior. Although our predictive models have good accuracy, this does not necessarily mean that users will easily accept an application that automatically resolves their conflicts. Moreover, it will be interesting to investigate the generalizability of our results to other social data sharing settings. If we obtain evidence that this is the case, it supports our broader vision of developing socially adaptive supportive technology.

APPENDIXES

A EXAMPLE SCENARIO PAIR AND CONFLICTING DESIGNATED SOLUTIONS

A.1 Scenario A

Mary is an 8-year-old child, and Paul is her father. Paul wants to find out when Mary arrives at the park. She is going there on her own for the first time, and Paul is worried. You are Paul; use the menu below to construct an agreement to find out when Mary arrives at the park.

Designated solution: I want *Mary* to *share* her location with *me* if *she's at the park*.

⁸We thank one of the anonymous reviewers for highlighting this.

A.2 Scenario B

Mary is an 8-year-old child, and Paul is her father. Paul has a meeting between 3pm and 5pm, and will be very busy during that time. But Mary checks in frequently all day long and shares with everybody. You are Paul; use the menu below to construct an agreement to ensure that Mary does not notify you with her location during your meeting.

Designated solution: I want *Mary* to *not share* her location with *me* if *it's between 3pm and 5pm*.

A.3 Conflict Between Designated Solutions

Using the definition of conflict of Section 3.1, we can see that the two agreements have the same debtor, opposite norm types, same action, overlapping third party, and possibly overlapping conditions. We therefore conclude that a conflict may occur (e.g., to share or not to share Mary's location if she enters the park between 3pm and 5pm).

B A NUMERICAL EXAMPLE DEMONSTRATING THE CONFLICT RESOLUTION MODEL

Assume that a user provided the following VP_g , VP_{c1} , VP_{c2} consisting of

$$VP_g = \langle Frnd_g = .1, Priv_g = .2, Saf_g = .3, Ind_g = .2, Res_g = .2 \rangle$$

$$VP_{c1} = \langle Frnd_{c1} = .1, Priv_{c1} = .2, Saf_{c1} = .2, Ind_{c1} = .4, Res_{c1} = .1 \rangle$$

$$VP_{c2} = \langle Frnd_{c2} = .2, Priv_{c2} = .1, Saf_{c2} = .5, Ind_{c2} = .1, Res_{c2} = .1 \rangle.$$

We then calculate the prediction that each value gives—for example, the distance of friendship component in the first commitment value profile to its counterpart in the user's general value profile: $|Frnd_{c1} - Frnd_g| = 0$.

The distance of friendship component in the second commitment value profile to its counterpart in the user's general value profile: $|Frnd_{c2} - Frnd_g| = .1$.

The friendship component of the first commitment's value profile is closer to its counterpart in the user's general value profile. This means that the value of friendship predicts that the user will prefer C_1 , or

$$Pred_p = |Frnd_{c2} - Frnd_g| - |Frnd_{c1} - Frnd_g| = .1 - 0 = .1.$$

Following the same example, we get the following prediction vector:

$$Pred_{c1,c2} = \langle Frnd_p = .1, Priv_p = .1, Saf_p = .2, Ind_p = -.1, Res_p = 0 \rangle.$$

This means that the values of friendship, privacy, and safety predict that the user will prefer C_1 , the value of independence predicts that the user will prefer C_2 , and the value of responsibility predicts no preference.

ACKNOWLEDGMENTS

The authors would like to thank Ursula Beer for her assistance during the user study and the anonymous reviewers for their comments, which helped us in making substantial improvements to the article.

REFERENCES

- Nirav Ajmeri, Jiaming Jiang, Rada Chirkova, Jon Doyle, and Munindar P. Singh. 2016. Coco: Runtime reasoning about conflicting commitments. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI'16)*. 17–23.
- Tina Balke, Célia da Costa Pereira, Frank Dignum, Emiliano Lorini, Antonino Rotolo, Wamberto Vasconcelos, and Serena Villata. 2013. Norms in MAS: Definitions and related concepts. In *Normative Multi-Agent Systems, Dagstuhl Follow-Ups*, Vol. 4. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Wadern, Germany, 1–31.

- Trevor Bench-Capon. 2003. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation* 13, 3, 429–448.
- Pia Borlund and Jesper W. Schneider. 2010. Reconsideration of the simulated work task situation: A context instrument for evaluation of information retrieval interaction. In *Proceedings of the 3rd Symposium on Information Interaction in Context (IIX'10)*. ACM, New York, NY, 155–164.
- Giuseppe Carenini and John Loyd. 2004. ValueCharts: Analyzing linear models expressing preferences and evaluations. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI'04)*. ACM, New York, NY, 150–157.
- J. Cohen. 1988. *Statistical Power Analysis for the Behavioral Sciences*. Lawrence Erlbaum Associates.
- Natalia Criado, Elizabeth Black, and Michael Luck. 2015. A coherence maximisation process for solving normative inconsistencies. *Autonomous Agents and Multi-Agent Systems* 30, 4, 640–680. DOI: <http://dx.doi.org/10.1007/s10458-015-9300-x>
- Alexei Czeskis, Ivayla Dermendjieva, Hussein Yapit, Alan Borning, Batya Friedman, Brian Gill, and Tadayoshi Kohno. 2010. Parenting from the pocket: Value tensions and technical directions for secure and private parent-teen mobile safety. In *Proceedings of the 6th Symposium on Usable Privacy and Security (SOUPS'10)*. ACM, New York, NY, 15:1–15:15.
- Ignacio Fernández-Tobías, Matthias Braunhofer, Mehdi Elahi, Francesco Ricci, and Iván Cantador. 2016. Alleviating the new user problem in collaborative filtering by exploiting personality information. *User Modeling and User-Adapted Interaction* 26, 2–3, 221–255.
- A. Field, J. Miles, and Z. Field. 2012. *Discovering Statistics Using R*. SAGE Publications.
- W. H. Finch, J. E. Bolin, and K. Kelley. 2014. *Multilevel Modeling Using R*. Taylor & Francis.
- Samuel D. Gosling, Simine Vazire, Sanjay Srivastava, and Oliver P. John. 2004. Should we trust Web-based studies? A comparative analysis of six preconceptions about Internet questionnaires. *American Psychologist* 59, 2, 93–104.
- S. Hansson. 1991. Norms and values. *Critica* 23, 67, 3–13.
- Amy Adele Hasinoff. 2017. Where are you? Location tracking and the promise of child safety. *Television and New Media* 18, 6, 496–512.
- Marc W. Howard and Michael J. Kahana. 2002. A distributed representation of temporal context. *Journal of Mathematical Psychology* 46, 3, 269–299.
- Alina Huldgtren, Pascal Wiggers, and Catholijn M. Jonker. 2014. Designing for self-reflection on values for improved life decision. *Interacting with Computers* 26, 1, 27–45.
- Karl M. Kapp. 2012. *The Gamification of Learning and Instruction: Game-Based Methods and Strategies for Training and Education*. Pfeiffer & Company.
- Alex Kayal, Willem-Paul Brinkman, Rianne Gouman, Mark A. Neerincx, and M. Birna van Riemsdijk. 2014a. A value-centric model to ground norms and requirements for ePartners of children. In *Coordination, Organization, Institutions, and Norms in Agent Systems IX*. Springer-Verlag.
- Alex Kayal, Willem-Paul Brinkman, Hanna Zoon, Mark A. Neerincx, and M. Birna van Riemsdijk. 2014b. A value-sensitive mobile social application for families and children. In *Posters, Demos, Late-Breaking Results and Workshop Proceedings of the 22nd Conference on User Modeling, Adaptation, and Personalization Co-Located With the 22nd Conference on User Modeling, Adaptation, and Personalization (UMAP'14)*.
- Bart P. Knijnenburg, Martijn C. Willemsen, Zeno Gantner, Hakan Soncu, and Chris Newell. 2012. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction* 22, 4, 441–504.
- John H. Krantz and Reeshad Dalal. 2000. Validity of Web-based psychological research. In *Psychological Experiments on the Internet*, M. H. Birnbaum (Ed.). Academic Press, San Diego, CA, 35–60.
- Kate Miriam Loewenthal. 2001. *An Introduction to Psychological Tests and Scales*. UCL Press, London, England.
- Felipe Meneguzzi, Odinaldo Rodrigues, Nir Oren, Wamberto W. Vasconcelos, and Michael Luck. 2015. BDI reasoning with normative considerations. *Engineering Applications of Artificial Intelligence* 43, C, 127–146.
- J. Nihlen-Fahlquist. 2013. Responsibility and privacy—ethical aspects of using GPS to track children. *Children and Society* 29, 1, 38–47.
- Helen Nissenbaum. 2010. *Privacy in Context: Technology, Policy and the Integrity of Social Life*. Stanford University Press, Stanford, CA.
- Nir Oren, Michael Luck, Simon Miles, and Timothy J. Norman. 2008. An argumentation inspired heuristic for resolving normative conflict. In *Proceedings of the International Workshop on Coordination, Organisations, Institutions, and Norms in Agent Systems*.
- Umberto Panniello, Alexander Tuzhilin, and Michele Gorgoglione. 2012. Comparing context-aware recommender systems in terms of accuracy and diversity. *User Modeling and User-Adapted Interaction* 24, 1, 35–65.
- Alina Pommeranz, Christian Detweiler, Pascal Wiggers, and Catholijn Jonker. 2011. Elicitation of situated values: Need for tools to help stakeholders and designers to reflect and communicate. *Ethics and Information Technology* 14, 4, 285–303.
- C. Robson. 2002. *Real World Research: A Resource for Social Scientists and Practitioner-Researchers* (2nd ed.). Blackwell Publishing, Malden, MA.
- M. Rokeach. 1973. *The Nature of Human Values*. Free Press, New York, NY.

- Munindar Singh. 1999. An ontology for commitments in multiagent systems. *Artificial Intelligence and Law* 7, 1, 97–113.
- Munindar P. Singh. 2008. Semantical considerations on dialectical and practical commitments. In *Proceedings of the 23rd National Conference on Artificial Intelligence—Volume 1 (AAAI'08)*. 176–181.
- Pankaj R. Telang and Munindar P. Singh. 2011. Specifying and verifying cross-organizational business models: An agent-oriented approach. *IEEE Transactions on Services Computing* 5, 3, 305–318.
- Ibo Van de Poel. 2013. Translating values into design requirements. In *Philosophy and Engineering: Reflections on Practice, Principles and Process*. Springer Netherlands, 253–266.
- T. van der Weide. 2011. *Arguing to Motivate Decisions*. Ph.D. Dissertation. Utrecht University.
- M. Birna van Riemsdijk, Louise Dennis, Michael Fisher, and Koen V. Hindriks. 2015a. A semantic framework for socially adaptive agents: Towards strong norm compliance. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS'15)*. 423–432.
- M. Birna van Riemsdijk, Catholijn M. Jonker, and Victor Lesser. 2015b. Creating socially adaptive electronic partners: Interaction, reasoning and ethical challenges. In *Proceedings of the 14th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'15)*. 1201–1206.
- Asimina Vasalou, Anne-Marie Oostveen, and Adam N. Joinson. 2012. A case study of non-adoption: The values of location tracking in the family. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW'12)*. ACM, New York, NY, 779–788.
- Wamberto W. Vasconcelos, Martin J. Kollingbaum, and Timothy J. Norman. 2009. Normative conflict resolution in multi-agent systems. *Autonomous Agents and Multi-Agent Systems* 19, 2, 124–152.
- Bodo Winter. 2013. A very basic tutorial for performing linear mixed effects analyses. arXiv:1308.5499.

Received November 2016; revised August 2017; accepted November 2017