Nova

Value-based Negotiation of Norms

Aydogan, Reyhan; Kafali, Özgür; Arslan, Furkan; Jonker, Catholijn M.; Singh, Munindar P.

DOI

[10.1145/3465054](https://doi.org/10.1145/3465054)

Publication date

2021

Document Version

Final published version

Published in

ACM Transactions on Intelligent Systems and Technology

Citation (APA)

Aydogan, R., Kafali, Ö., Arslan, F., Jonker, C. M., & Singh, M. P. (2021). Nova: Value-based Negotiation of Norms. *ACM Transactions on Intelligent Systems and Technology*, 12(4), [45].
<https://doi.org/10.1145/3465054>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Nova: Value-based Negotiation of Norms

REYHAN AYDOĞAN, Özyeğin University, Turkey, and Delft University of Technology, The Netherlands

ÖZGÜR KAFALI, University of Kent, UK

FURKAN ARSLAN, Özyeğin University, Turkey

CATHOLIJN M. JONKER, Delft University of Technology, The Netherlands and Leiden University, The Netherlands

MUNINDAR P. SINGH, North Carolina State University, USA

Specifying a normative multiagent system (nMAS) is challenging, because different agents often have conflicting requirements. Whereas existing approaches can resolve *clear-cut* conflicts, tradeoffs might occur in practice among alternative nMAS specifications with no apparent resolution. To produce an nMAS specification that is acceptable to each agent, we model the specification process as a negotiation over a set of norms. We propose an agent-based negotiation framework, where agents' requirements are represented as *values* (e.g., patient safety, privacy, and national security), and an agent revises the nMAS specification to promote its values by executing a set of norm revision rules that incorporate ontology-based reasoning. To demonstrate that our framework supports creating a *transparent* and *accountable* nMAS specification, we conduct an experiment with human participants who negotiate against our agent. Our findings show that our negotiating agent reaches better agreements (with small *p*-value and large effect size) faster than a baseline strategy. Moreover, participants perceive that our agent enables more *collaborative* and transparent negotiations than the baseline (with small *p*-value and large effect size in particular settings) toward reaching an agreement.

CCS Concepts: • **Computing methodologies** → **Multi-agent systems**; **Knowledge representation and reasoning**; • **Security and privacy** → *Security requirements*;

Additional Key Words and Phrases: Sociotechnical systems, conflicting requirements, human-agent negotiation

ACM Reference format:

Reyhan Aydoğan, Özgür Kafalı, Furkan Arslan, Catholijn M. Jonker, and Munindar P. Singh. 2021. Nova: Value-based Negotiation of Norms. *ACM Trans. Intell. Syst. Technol.* 12, 4, Article 45 (July 2021), 29 pages. <https://doi.org/10.1145/3465054>

Singh was partially supported by the US National Science Foundation under grant IIS-1908374.

Author's addresses: R. Aydoğan, Özyeğin University, Nişantepe Mah. Orman Sok. No:34–36, Alemdağ, Çekmeköy, İstanbul, 34794, Turkey, Delft University of Technology, Van Mourikbroekmanweg 6, Delft, 2628XE, The Netherlands; email: reyhan.aydogan@ozyegin.edu.tr; Ö. Kafalı, University of Kent, School of Computing, Canterbury, Kent, CT2 7NF, UK, R.O.; email: Kafali@kent.ac.uk; F. Arslan, Özyeğin University, Nişantepe Mah. Orman Sok. No: 34–36, Alemdağ, Çekmeköy, İstanbul, 34794, Turkey; email: furkan.arslan@ozu.edu.tr; C. M. Jonker, Delft University of Technology, Van Mourikbroekmanweg 6, Delft, 2628XE, The Netherlands, Leiden University, Niels Bohrweg 1, Leiden, 2333CA, The Netherlands; email: c.m.jonker@tudelft.nl; M. P. Singh, North Carolina State University, Department of Computer Science, Raleigh, NC, 27695-8206, USA; email: singh@ncsu.edu.

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.

© 2021 Association for Computing Machinery.

2157-6904/2021/07-ART45 \$15.00

<https://doi.org/10.1145/3465054>

1 INTRODUCTION

A sociotechnical system is one where multiple agents, as autonomous social entities, interact with each other via technical components Chopra et al. [2014] and Kafalı et al. [2016, 2020]. We understand a sociotechnical system as a **normative multiagent system (nMAS)** that is governed via norms [Singh 2013] that regulate interactions among autonomous agents. Specifying an nMAS is nontrivial when the participating agents' requirements conflict with each other.

Example: Consider a tradeoff scenario regarding safety, patient privacy, and national security in a hospital setting. Among three stakeholders, the hospital administration values safety over the other two, the legal department values patient privacy over the other two, and the government agency values national security over the other two. If we represent this setting as an nMAS, then there are three requirements: The hospital administration has the requirement “ R_1 : Physicians should treat patients”; the legal department has the requirement “ R_2 : Physicians should not disclose patients' sensitive information without their consent”; and the government agency has the requirement “ R_3 : Government officers may access patients' information.” In a medical emergency [HHS 2014], R_1 and R_2 conflict with each other, because the norms that regulate the sharing of patients' sensitive information among physicians might hinder treatment efficiency. If the patient is involved in a national security case, then R_2 and R_3 conflict with each other as physicians might need to disclose patients' information to government officers.

Problem in brief: A transparent and accountable [AIHLEG 2019] specification of an nMAS should take into account the tradeoffs between agent requirements as exemplified in the above setting. Relevant approaches that aim to tackle this problem can be grouped into three lines of research. One, approaches to specifying nMASs [Chopra et al. 2014; Kafalı et al. 2016, 2020] adopt a *central* (requirements engineering) perspective to meet agent preferences but lack automated support to avoid or resolve potential conflicts among preferences. Two, centralized conflict resolution approaches [Ajmeri et al. 2016; Günay and Yolum 2013; Kollingbaum et al. 2006] can eliminate clear-cut contradictions but lack reasoning about tradeoffs from a decentralized perspective where all agents are equally represented. Three, norm negotiation [Boella et al. 2009; Dijkstra et al. 2007] seeks to achieve social goals and resolve conflicts but omits preferences or bidding strategies. In general, to accommodate all agents' preferences and the potential tradeoffs among them, we need an automated approach to support specifying an nMAS by reasoning about norms and enabling agents to make concessions by revising one or more of the norms.

Approach: We present **Norms and Values (Nova)**, an agent-based automated negotiation framework for collaboratively specifying an nMAS. An *offer* is an nMAS specification represented as a set of norms. The satisfaction of agent requirements contributes toward the achievement of social values [Carnevale and Probst 1998]. For example, the hospital administration's requirement R_1 contributes to *patient safety*, the legal department's requirement R_2 contributes to *privacy*, and the government agency's requirement R_3 contributes to *national security*. We propose value-based negotiation, a variation of interest-based negotiation [Fisher et al. 1983], to generate an offer that corresponds to a preferred nMAS specification for an agent. Agents revise the norms in an offer by following a set of rules based on ontology reasoning and evaluating how the norms contribute to their preferred values via a dynamically generated *norm-value graph*.

Nova supports desired principles of responsible AI [AIHLEG 2019]: transparency through traceability of negotiation steps and explainability of agents' reasoning via norm revision rules and accountability via normative semantics. Nova enables mutually acceptable outcomes via equal participation of all agents. To demonstrate how Nova supports these principles in a practical setting, we implement a negotiator agent that follows a time-dependent concession bidding strategy, and conduct a human participant experiment pertaining to a healthcare privacy use case.

Table 1. Norm Syntax

Norm	$\rightarrow$	$t(\text{AG}, \text{AG}, \text{Expr}, \text{Expr})$
t	$\rightarrow$	$c \mid p \mid a$
Expr	$\rightarrow$	$\text{true} \mid \phi \mid \neg\text{Expr} \mid \text{Expr} \wedge \text{Expr}$

Contributions: (i) An agent-based negotiation framework for the specification of nMASs that supports transparency and accountability, (ii) a set of heuristic rules for revising norms based on ontology reasoning, (iii) a value-based concession bidding strategy that adapts its offers at runtime based on opponent’s behavior without predefined utility functions, and (iv) a prototype negotiator tool and ontology that enables human–agent interactions in healthcare privacy setting.

Structure: Section 2 introduces the relevant technical background. Section 3 describes the Nova framework. Section 4 presents the evaluation of the Nova framework. Section 5 reviews related work. Section 6 concludes the article.

2 BACKGROUND

We introduce the foundational concepts that are used in our article: norms, ontologies, and agent-based automated negotiation. These concepts constitute the prior work on top of which we frame our contribution.

2.1 Norms

An nMAS supports interaction across two levels: (i) among the agents regarding the applicable norms and (ii) between the agents and the technical components of the underlying nMAS.

Definition 2.1 characterizes a directed norm, adopted from Singh [2013]. Below, $\mathcal{R}$ is the set of agent roles, $\mathcal{P}$ is the set of propositional expressions ($\text{Expr} \in \mathcal{P}$), and $\mathcal{N}$ is the set of norms that can be constructed using elements from $\mathcal{R}$ and $\mathcal{P}$. Table 1 presents the syntax of a norm in Nova. Note that we adopt distinct type styles for AGENTS and propositions.

Definition 2.1. A norm, $n = t(\text{SBJ}, \text{OBJ}, \text{ant}, \text{con})$, represents a directed, conditional relationship between two agents, where t, its type, is one of $\{c, p, a\}$ (denoting commitment, prohibition, authorization; respectively); $\text{SBJ} \in \mathcal{R}$ is its subject; $\text{OBJ} \in \mathcal{R}$ is its object; $\text{ant} \in \mathcal{P}$ is its antecedent; and $\text{con} \in \mathcal{P}$ is its consequent.

There are two main underlying intuitions for this conception of norms. First, a norm is directed and the directionality provides information on who is accountable to whom given the norm. Second, the antecedent of a norm brings it into force and the consequent of a norm completes it. We describe these intuitions for the specific norm types.

We consider three type of norms. Figure 1 demonstrates the lifecycle of each norm type. Typically, each norm starts its lifecycle as conditional (i.e., initial state), unless its antecedent is already true. If the antecedent is true, then the norm starts its lifecycle as detached. That is, when a norm is first created (or activated) among two agents, its satisfaction requirement is conditional upon the antecedent condition. Only after the antecedent condition holds, the norm is “in force,” meaning that its subject is now committed to, or prohibited by, or authorized by its object. Below, we demonstrate the norm states and the transitions among them with examples of each norm type.

A commitment (indicated by c) means that its subject commits to its object to bringing about the consequent if the antecedent holds. For example, the hospital administration prefers a commitment $n_1 = c(\text{PHYSICIAN}, \text{HOSPITAL_ADMINISTRATION}, \text{true}, \text{treat})$ for R_1 . According to this norm, the PHYSICIAN is committed to the HOSPITAL_ADMINISTRATION to treating patients. Since n_1 ’s antecedent is true, n_1 is *detached*. In other word, this commitment is unconditional. When treat

Fig. 1. Lifecycle of norms. Double rectangles represent terminal states (i.e., norm's lifecycle ends).

holds (i.e., the transition labeled “con” in Figure 1 occurs), n_1 is satisfied. Note that a commitment can be satisfied before being detached, i.e., its consequent holds before its antecedent holds. If treatment fails to hold, then n_1 is violated. PHYSICIAN is accountable to HOSPITAL_ADMINISTRATION for n_1 's violation.

A prohibition (indicated by p) means that its subject is prohibited by its object from bringing about the consequent if the antecedent holds. For example, the legal department prefers a prohibition $n_2 = p(\text{PHYSICIAN}, \text{HOSPITAL_ADMINISTRATION}, \text{true}, \text{disclose_information})$ for R_2 . Accordingly, the PHYSICIAN is prohibited by the HOSPITAL_ADMINISTRATION from disclosing patient information. Similarly to the example commitment above, n_2 is already *detached* (unconditional). When *disclose_information* holds, n_2 is violated. PHYSICIAN is accountable to HOSPITAL_ADMINISTRATION for n_2 's violation. If the consequent never occurs, then n_2 is satisfied.

An authorization (indicated by a) means that its subject is authorized by its object for bringing about the consequent if the antecedent holds. For example, the government agency prefers an authorization $n_3 = a(\text{GOVERNMENT_OFFICER}, \text{HOSPITAL_ADMINISTRATION}, \text{case}, \text{access_patient_information})$ for R_3 . Accordingly, the GOVERNMENT_OFFICER is authorized by the HOSPITAL_ADMINISTRATION to access patient information if a case is presented regarding the need for such information. When *case* holds (i.e., the transition labeled “ant” in Figure 1 occurs), n_3 is *detached*. For some propositions, it is reasonable to assume inherent deadline conditions, whose passing renders the proposition false. For example, if there is a specified deadline for the government agency to present a case to access to patient information and the deadline passes (i.e., the transition labeled “never ant” in Figure 1 occurs), then n_3 is *expired*. Note that norms whose antecedents are true never expire. When *access_patient_information* holds, n_3 is satisfied. Similarly to a commitment, an authorization can be satisfied before being detached, i.e., its consequent holds before its antecedent holds. When n_3 is *detached* but the consequent never occurs (i.e., the transition labeled “never con” in Figure 1 occurs), n_3 is violated. Here HOSPITAL_ADMINISTRATION is accountable to GOVERNMENT_OFFICER for n_3 's violation. Notice that an authorization makes its object accountable, meaning that the party who granted an authorization to another has to support the latter to exercise that authorization. This notion of authorization is a major departure from the common notion of permissions, which can never be violated [Singh 2013; Von Wright 1999].

We do not explicitly model norm deadlines in this article. Previous works have addressed deadlines for the antecedent and consequent using temporal languages such as the Event Calculus [Chesani et al. 2013; Kafalı and Torroni 2012] or CTL [Alechina et al. 2013], where deadlines can be expressed via discrete time constraints. We follow the formalisation of Kafalı et al. [2020] for norms, wherein an event (ev) that brings about an antecedent or consequent condition may have a complement ($\overline{\text{ev}}$), which describes the *nonoccurrence* of the event. Such complement events may be used to express deadline conditions. For example, *present* means that the government agency has presented its case within the specified time interval, then $\overline{\text{present}}$ means that the government agency has failed to present its case during that time interval.

Simple permissions can be represented via technical mechanisms such as role-based access control [Marinovic et al. 2014], whereas authorizations also support accountability for permissions by holding the authorizer accountable, i.e., when the authorizee cannot perform the authorized action, then the authorizer violates the authorization [Von Wright 1999].

Note that the term “norm” in our work follows the legal literature, i.e., the Hohfeldian approach regarding rights and duties [Hohfeld 1919] and Von Wright’s deontic logic [Von Wright 1999], as opposed to the social literature. That is, a norm here is not merely a widely accepted practice or a convention but a representation of some deontic expectations that one agent holds regarding another. Definition 2.1 conceptualizes a norm in this sense, drawing upon works that adopt norms to represent agent requirements [Kafali et al. 2016, 2017, 2020; Singh 2013].

Definition 2.2 formalizes how norms can be inferred from other norms by comparing their antecedents and consequents [Kafali et al. 2017].

Definition 2.2. Let $n, n' \in \mathcal{N}$ be two norms of the same type and with the same subject and object. Then, n is stronger than n' iff n is detached whenever n' is detached; and n' is satisfied whenever n is satisfied.

Note that if n is stronger than n' , then n' is weaker than n . If n is stronger than n' and n' is stronger than n , then the only way this can happen is that n and n' are identical norms [Kafali et al. 2017].

We extend Definition 2.2 with heuristics and ontology reasoning for norm revision, e.g., to generate counteroffers.

2.2 Ontology Reasoning

An *ontology* [Guarino 1998] formalizes knowledge about a domain. We focus on the *is-a* (or taxonomy) and *has-a* (or property) relations to express knowledge in a specific domain. The *is-a* relation denotes that a concept is a specialisation of another concept. For example, in the hospital domain, a *Physician* is a *Hospital Staff*. Such relations can be used to describe role hierarchies. Properties connect concepts and can be described via the *has-a* relation. For example, in the hospital domain, a *Physician* has a *treatment* for a *Patient*. That is, we specify the property *hasTreatment* for the concept *Physician* and the range of the property is another concept *Patient*.

Domain-specific values are explored in the economic literature [Van Raaij and Verhallen 1994] and together with applicable norms they affect how autonomous agents behave in an nMAS. For example, in the healthcare setting, values such as patient safety and privacy are relevant. We denote the set of values as $\mathcal{V} = \{\text{patient_safety}, \dots, \text{privacy}\}$.

The set of roles is denoted as $\mathcal{R} = \{\text{PHYSICIAN}, \dots, \text{PATIENT}\}$ and the set of atomic propositions is denoted as $\mathcal{P} = \{\text{case}, \dots, \text{disclose_information}\}$. As a result, the set of norms is denoted by $\mathcal{N} = \{\text{c}(\text{PHYSICIAN}, \text{HOSPITAL_ADMINISTRATION}, \text{true}, \text{treat}), \dots, \text{a}(\text{GOVERNMENT_AGENCY}, \text{HOSPITAL_ADMINISTRATION}, \text{case}, \text{access_patient_information})\}$. The sets $\mathcal{R}$, $\mathcal{P}$, $\mathcal{N}$, and $\mathcal{V}$ together constitute the concepts in the ontology, whereas the set $\mathcal{O}$ constitutes the properties.

Three such properties are particularly useful for our development. Given a proposition $\phi \in \mathcal{P}$ and a value $v \in \mathcal{V}$, ϕ promotes v , if $\text{pro}(\phi, v) \in \mathcal{O}$. For example, the ontology property $\text{pro}(\text{consent}, \text{privacy})$ means that obtaining consent from a patient to access their data promotes patient privacy. Given a proposition $\phi \in \mathcal{P}$ and a value $v \in \mathcal{V}$, ϕ demotes v if $\text{dem}(\phi, v) \in \mathcal{O}$.

Given two propositions $\phi_1, \phi_2 \in \mathcal{P}$ and a value $v \in \mathcal{V}$, ϕ_1 is preferred over ϕ_2 with respect to v if $\text{pref}(\phi_1, \phi_2, v) \in \mathcal{O}$. Consider two types of consent that promote privacy. The ontology property $\text{pref}(\text{written_consent}, \text{verbal_consent}, \text{privacy})$ means that written consent is preferred over verbal consent with respect to privacy, because written consent stays on record and can also be implemented in technical access control mechanisms. However, preference is reversed with respect to

patient safety as obtaining verbal consent from the patient in an emergency is easier than obtaining written consent, thus the ontology property $\text{pref}(\text{verbal_consent}, \text{written_consent}, \text{patient_safety})$.

The public ontology that stores unified role, proposition, and value names is shared among all negotiating agents to ensure agent interoperability. Moreover, each agent has a private ontology for representing individual preferences, e.g., how propositions promote values. Note that contradicting elements in $\mathcal{O}$, such as the inclusion of both $\text{pro}(\phi, v)$ and $\text{dem}(\phi, v)$ or the inclusion of both $\text{pref}(\phi_1, \phi_2, v)$ and $\text{pref}(\phi_2, \phi_1, v)$, can be eliminated using ontology validation checks.

2.3 Automated Negotiation

In general negotiations are about a finite set of n issues $\mathcal{I} = \{1, 2, \dots, n\}$. Each issue $i \in \mathcal{I}$ has a range $\mathcal{D}_i$ of possible instantiations. An outcome is a complete assignment to the set of issues, i.e., an element of the Cartesian Product of the ranges of instantiations per issue. Formally, the set of all possible outcomes is defined as $\mathcal{O} = \mathcal{D}_1 \times \mathcal{D}_2 \times \dots \times \mathcal{D}_n$. The interaction among agents during an automated negotiation is governed by the **stacked alternating offers protocol (SAOP)** [Aydoğan et al. 2017], which extends alternating offers protocol [Rubinstein 1992] to more than two negotiating parties. In SAOP, agents have a shared deadline to reach an agreement and one of the agents initiates the negotiation with an offer (first round). In each turn, the agent receiving an offer can *accept* the current offer, make a *counteroffer*, or *end* the negotiation without an agreement. Turn-taking continues until success (agreement is reached: all agents accept) or failure (no agreement: any agent ends the negotiation or the deadline is reached). An agent negotiates by reasoning about its user's preference profile, often represented via a utility function, as well as taking its opponent's offers into account. In general, an agent does not know its opponent's preferences or its opponent's negotiation strategies. In some works, agents learn their opponent's preferences or strategies [Baarslag et al. 2016] by using machine learning techniques such as Version Space [Aydoğan and Yolum 2012], Bayesian learning [Hindriks and Tykhonov 2008], regression [Williams et al. 2013], or heuristics such as the frequency model [Tunali et al. 2017].

In addition to the uncertainty about the opponent, agents may have uncertainty about their own preferences especially when their preferences are represented by a qualitative model such as a CP-net [Boutilier et al. 2004]. Since CP-nets often lack total ordering, heuristic approaches for negotiating via CP-nets have been introduced [Aydoğan et al. 2015].

The article introducing CP-nets [Boutilier et al. 2004] explains how CP-nets induce preference graphs, and for the details about this procedure we refer to Boutilier et al. [2004]. For this article, it is enough to understand how to read a preference graph induced by a CP-net. A preference graph is a graph $\langle \mathcal{O}, \mathcal{E} \rangle$, where $\mathcal{E}$ is a set of edges between elements of $\mathcal{O}$. Having an edge (o, o') , from o to o' means that o' is preferred over o . A property of preference graphs induced from CP-nets is that such a direct edge can only exist between two nodes that differ only by one "flip," i.e., a change of instantiation in only one issue. Formally, suppose that $o = \langle o_1, o_2, \dots, o_n \rangle$, and $o' = \langle o'_1, o'_2, \dots, o'_n \rangle$, then o and o' differ by one flip iff $\exists! 1 \leq i \leq n: o'_i \neq o_i \wedge \forall j \neq i: o'_j = o_j$.

The root node is the least desired outcome, whereas the sink node is the most preferred outcome. Agents can compare possible outcomes only if there is a path between them. In this article, we use the heuristics that the longer the path from the root to an outcome is, the more preferred is the outcome by the agent [Aydoğan et al. 2015]. Therefore, the depth of a node represents to what extent the corresponding offer is preferred. Recall that the depth of an outcome in the graph corresponds to *the length of the longest path from the root node to that node*.

3 NOVA FRAMEWORK

We now introduce heuristic rules for revising norms that apply the strength relation and ontology reasoning and an automated negotiation strategy based on norm revision.

3.1 Norm Revision

In NOVA, an agent generates a counteroffer by revising the norms contained in the opponent's most recent offer. We define three revision rules using the underlying domain ontology and the strength relation (Definition 2.2). Definition 3.1 describes how the subject or object (i.e., actors) of a norm can be revised.

Definition 3.1. Let $\text{SBJ}, \text{SBJ}', \text{OBJ}, \text{OBJ}' \in \mathcal{R}$. Then, $t(\text{SBJ}', \text{OBJ}, \text{ant}, \text{con})$ and $t(\text{SBJ}, \text{OBJ}', \text{ant}, \text{con})$ are actor revisions of $t(\text{SBJ}, \text{OBJ}, \text{ant}, \text{con})$.

An agent can perform an actor revision based on how a norm promotes or demotes a value and using the taxonomic relations among the roles (see Definition 3.7).

Next, we describe how a norm promotes or demotes a value. To do so, we first describe how a propositional expression (see Table 1) promotes and demotes a value, and how an expression is preferred over another with respect to a value.

The first condition of Definition 3.2 tests for membership of a promotes property in the ontology for atomic propositions. The second and third conditions describe negation and conjunction of propositions; respectively. Definition 3.3 is described similarly for the demotes relation.

Definition 3.2. An Expr promotes a value $v \in \mathcal{V}$, denoted $\text{promotes}(\text{Expr}, v)$, iff

- (1) $\text{Expr} = \phi$, $\text{pro}(\phi, v) \in \mathcal{O}$; or
- (2) $\text{Expr} = \neg \text{Expr}'$, $\text{demotes}(\text{Expr}', v)$; or
- (3) $\text{Expr} = \text{Expr}_1 \wedge \text{Expr}_2$, $\text{promotes}(\text{Expr}_1, v)$, $\text{promotes}(\text{Expr}_2, v)$.

Definition 3.3. An Expr demotes a value $v \in \mathcal{V}$, denoted $\text{demotes}(\text{Expr}, v)$, iff

- (1) $\text{Expr} = \phi$, $\text{dem}(\phi, v) \in \mathcal{O}$; or
- (2) $\text{Expr} = \neg \text{Expr}'$, $\text{promotes}(\text{Expr}', v)$; or
- (3) $\text{Expr} = \text{Expr}_1 \wedge \text{Expr}_2$, $\text{demotes}(\text{Expr}_1, v)$, $\text{demotes}(\text{Expr}_2, v)$.

Definition 3.4 describes the preference relation for expressions. The first condition of Definition 3.4 tests for membership of a preference property in the ontology for atomic propositions. That is, the base case of the preference relation among propositional expressions is the ontology property for describing preference among for atomic propositions, $\text{pref}(\phi_1, \phi_2, v)$. The second and third conditions describe negation of propositions. For example, if an agent prefers written consent over verbal consent with respect to privacy, then the agent also prefers written consent over “no verbal consent” (i.e., the negation of verbal consent) with respect to privacy (second condition). However, verbal consent is preferred over no written consent (i.e., the negation of written consent) with respect to privacy (third condition). The fourth and fifth conditions handle conjunction of propositions.

Definition 3.4. An Expr is preferred over an Expr' with respect to a value $v \in \mathcal{V}$, denoted $\text{preferred}(\text{Expr}, \text{Expr}', v)$, iff

- (1) $\text{Expr} = \phi$, $\text{Expr}' = \phi'$, $\text{pref}(\phi, \phi', v) \in \mathcal{O}$; or
- (2) $\text{Expr}' = \neg \text{Expr}''$, $\text{preferred}(\text{Expr}, \text{Expr}'', v)$; or
- (3) $\text{Expr}' = \neg \text{Expr}''$, $\text{preferred}(\text{Expr}'', \text{Expr}, v)$; or
- (4) $\text{Expr} = \text{Expr}_1 \wedge \text{Expr}_2$, $\text{preferred}(\text{Expr}_1, \text{Expr}', v)$, $\text{preferred}(\text{Expr}_2, \text{Expr}', v)$; or
- (5) $\text{Expr}' = \text{Expr}_1 \wedge \text{Expr}_2$, $\text{preferred}(\text{Expr}, \text{Expr}_1, v)$, $\text{preferred}(\text{Expr}, \text{Expr}_2, v)$.

We are now ready to describe how a norm promotes (Definition 3.5) or demotes (Definition 3.6) a value.

Definition 3.5. Given a value $v \in \mathcal{V}$ and a norm $n = t(\text{SBJ}, \text{OBJ}, \text{ant}, \text{con}) \in \mathcal{N}$, n promotes v , denoted $\text{promotes}(n, v)$ iff

- $t \in \{c, a\}$, $\text{promotes}(\text{con}, v)$, and $\text{preferred}(\text{con}, \text{ant}, v)$; or
- $t = p$, $\text{demotes}(\text{con}, v)$, and $\text{preferred}(\text{ant}, \text{con}, v)$.

Definition 3.6. Given a value $v \in \mathcal{V}$ and a norm $n = t(\text{SBJ}, \text{OBJ}, \text{ant}, \text{con}) \in \mathcal{N}$, n demotes v , denoted $\text{demotes}(n, v)$ iff

- $t \in \{c, a\}$, $\text{demotes}(\text{con}, v)$, and $\text{preferred}(\text{ant}, \text{con}, v)$; or
- $t = p$, $\text{promotes}(\text{con}, v)$, and $\text{preferred}(\text{con}, \text{ant}, v)$.

Consider the following. Commitment $c(\text{STAFF}, \text{HOSPITAL_ADMINISTRATION}, \text{shared_computer}, \text{logout})$ means that hospital staff should log out from their accounts after using shared computers. This commitment promotes privacy, because logging out from an account promotes privacy and is preferred over using shared computers with respect to privacy. Prohibition $p(\text{PHYSICIAN}, \text{HOSPITAL_ADMINISTRATION}, \neg\text{treat}, \text{access_patient_information})$ means that physicians should not access the information of a patient whom they are not treating. This prohibition demotes patient safety, because accessing a patient's information promotes patient safety and is preferred over not treating the patient with respect to patient safety.

Along with the promotes and demotes relations, a negotiator agent can use the taxonomic relations in the ontology to choose new roles for its norms, e.g., replace the subject of an authorization with a more general role to authorize a wider set of roles for a specific functionality. Definition 3.7 describes how an agent can assign new roles using a norm preference relation, $n' >_v n$, meaning that norm n' is preferred over norm n with respect to value v .

Definition 3.7. Given two roles $R, R' \in \mathcal{R}$: R' is a specialization of R (R is a generalization of R'), a value $v \in \mathcal{V}$, and two norms $n = t(R, \text{OBJ}, \text{ant}, \text{con})$, $n' = t(R', \text{OBJ}, \text{ant}, \text{con}) \in \mathcal{N}$, then

- $\text{promotes}(n, v)$ iff $n >_v n'$;
- $\text{demotes}(n, v)$ iff $n' >_v n$.

Consider an offer consisting of a norm $n = a(\text{EVERYONE}, \text{HOSPITAL_ADMINISTRATION}, \text{true}, \text{access_patient_information})$. Authorizing everyone to access patient information demotes privacy. Thus, a norm n' with a specialized subject, e.g., $a(\text{GOVERNMENT_AGENCY}, \text{HOSPITAL_ADMINISTRATION}, \text{true}, \text{access_patient_information})$, is preferred by the hospital administration: $n' >_{\text{privacy}} n$. Such taxonomic relations among agent roles can be extracted from a given ontology (using *is-a* relations).

Definition 3.8 describes how the antecedent or consequent (i.e., propositional expressions) of a norm can be revised. A negotiator agent can use the taxonomic relations in the ontology as well as the strength relation between norms (Definition 2.2) to choose new propositional expressions.

Definition 3.8. Let $\text{ant}, \text{ant}', \text{con}, \text{con}' \in \mathcal{P}$, $\text{ant} \not\equiv \text{ant}'$, and $\text{con} \not\equiv \text{con}'$. Then, $t(\text{SBJ}, \text{OBJ}, \text{ant}', \text{con})$ and $t(\text{SBJ}, \text{OBJ}, \text{ant}, \text{con}')$ are proposition revisions of $t(\text{SBJ}, \text{OBJ}, \text{ant}, \text{con})$.

For example, AG_1 can revise authorization $a(\text{AG}_1, \text{AG}_2, \text{written_consent}, \text{access_records})$ to a stronger authorization $a(\text{AG}_1, \text{AG}_2, \text{written_consent}, \text{access_records} \vee \text{share_records})$. The revised norm authorizes AG_1 to perform an additional functionality. Or, AG_1 can revise the norm with authorization $a(\text{AG}_1, \text{AG}_2, \text{verbal_consent}, \text{access_records})$, where AG_1 prefers verbal consent over written consent. Agents can reason about such preference relations among antecedents and consequents using Definition 3.4.

Definition 3.9 describes how the type of a norm can be changed. Type revision is useful both by itself and in combination with another revision.

Definition 3.9. Let $t, t' \in \{c, p, a\}$ and $t \neq t'$. Then, $t'(\text{SBJ}, \text{OBJ}, \text{ant}, \text{con})$ is a type revision of $t(\text{SBJ}, \text{OBJ}, \text{ant}, \text{con})$.

The notion of revision is defined inductively as follows.

Definition 3.10. Let $n_i, n_j \in \mathcal{N}$, then n_j is a revision of n_i , denoted by $\text{revision}(n_j, n_i)$ iff

- n_j is an actor, proposition, or type revision of n_i ; or
- $\exists n \in \mathcal{N} : \text{revision}(n_j, n)$ and $\text{revision}(n, n_i)$.

Note that revision is a transitive relation. The agent's negotiation strategy determines how (i.e., in what order and number of times) the agent applies these three rules to a norm that is received as part of an offer from another agent.

3.2 Value-based Negotiation

NOVA uses value-based negotiation to enable agents with conflicting requirements to agree on norms. Agents reason about values to generate offers, each a set of norms (Definition 3.11), without the need for predefined utility functions. Negotiation between agents is governed by SAOP.

Definition 3.11. $\mathcal{B} = \mathcal{P}(\mathcal{N})$ is the set of all potential offers the negotiators can make. We use $b \in \mathcal{B}$ with or without subscripts to denote variables over $\mathcal{B}$.

According to the proposed value-based negotiation approach, an offer is evaluated without having an explicit utility function or a preference model to compute the ranking of offers. That is, an agent in NOVA evaluates an offer by considering to what extent the norms in the offer *promote* or *demote* the agent's values. Furthermore, an offer is generated dynamically with respect to the opponent's most recent offer. We present an incremental offer-generation approach and introduce a set of *improving steps* inspired by "improving flips" in CP-nets [Aydoğan et al. 2015; Boutilier et al. 2004]. Our improving steps follow the norm revision rules defined in Section 3.1. Next, we describe the improving steps using two agents: AG_1 receives an offer from agent AG_2 , and AG_1 's aim is to improve the offer with respect to AG_1 's values. Note that an improving step is carried out from an agent's perspective. One agent's improvement may not (in many cases will not) be an improvement for another agent. The level of improvement for one agent determines the level of concession for the other, which is essential for coming to an agreement. In this work, three types of improving steps are introduced as follows:

- (a) Actor-based improving step: Let $n \in b_{AG_2}$ be a norm that AG_1 receives. Then, AG_1 can make an improving step with respect to a value $v \in \mathcal{V}$ by revising n to n' , where $n' \succ_v n$.
- (b) Proposition-based improving step: Let $n = t(\text{SBJ}, \text{OBJ}, \text{ant}, \text{con}) \in b_{AG_2}$ be a norm that AG_1 receives. Then, AG_1 can make an improving step with respect to a value $v \in \mathcal{V}$ by revising n to $n' = t(\text{SBJ}, \text{OBJ}, \text{ant}', \text{con}')$, where
 - (1) $\text{promotes}(n, v)$ and n' is stronger than n ; or
 - (2) $\text{demotes}(n, v)$ and n' is weaker than n ; or
 - (3) $\text{preferred}(\text{ant}', \text{ant}, v)$; or
 - (4) $\text{preferred}(\text{con}', \text{con}, v)$.
- (c) Type-based improving step: Let $n \in b_{AG_2}$ be a norm that AG_1 receives. If $\text{demotes}(n, v)$, where $v \in \mathcal{V}$, then AG_1 can make an improving step with respect to v by revising n to n' , where n' is a type revision of n , such that $\text{promotes}(n', v)$.

Note that there might be cases where we cannot compute whether a norm promotes (Definition 3.5) or demotes (Definition 3.6) a value. This is due to the preference relation for propositional

expressions (Definition 3.4), e.g., the consequent of the norm contains a conjunction of two propositions, where one proposition promotes the value and the other proposition demotes the value. However, in practical settings, agents following our improving steps would not generate such norms. In case an agent receives such a norm, the agent does not need to compute whether the norm promotes or demotes the value—the agent can use the third or fourth conditions of the “Proposition-based improving step.”

The aforementioned steps improve a given offer by replacing one of the components (i.e., subject, antecedent, consequent, or type) of one of the norms in the offer. By applying a sequence of improving steps, we dynamically construct a norm-value graph that can be used to generate offers by following a bidding strategy. In Nova, a norm-value graph is an acyclic weighted directed graph where nodes represent candidate offers, edges represent improving steps, and weights on edges represent scores of corresponding improving steps. If there is an edge from node b_i to node b_j , then it means that there is an improving step l from b_i to b_j . That is, b_j is preferred over b_i with respect to a value v . The weight of that edge denotes to what extent making this improving step improves the agent’s offer score (Definition 3.12).

ALGORITHM 1: Agent decision-making in Nova.

Data: b_{op} is the most recent offer; t_d is the deadline; t_c is the current time; G is the norm-value graph
Result: Negotiation action $a = \{\text{Offer}(b_{ag}), \text{Accept}\}$

```

1  if  $G = \emptyset$  or  $b_{op} \notin G$  then
2     $G.\text{nodes.add}(b_{op});$ 
3     $\text{expand-nova-graph}(G, b_{op});$ 
4  end
5   $b_{ag} = \text{generate-offer}(G, b_{op}, t_c, t_d);$ 
6  if  $(b_{op}.\text{score} \geq b_{ag}.\text{score})$  then
7    return  $\text{Accept};$ 
8  else
9    return  $\text{Offer}(b_{ag});$ 
10 end

11 Function  $\text{expand-nova-graph}(G, b):$ 
12    $\text{step-types} = [\text{actor-revision}, \text{predicate-revision}, \text{norm-revision}];$ 
13   foreach  $\text{type } s \in \text{step-types}$  do
14      $b' = b.\text{improve-step}(s);$ 
15     if  $b'.\text{score} > 0$  and  $b' \notin G$  and not  $\text{cause-cycle}(G, b, b')$  then
16        $G.\text{nodes.add}(b');$ 
17        $G.\text{edges.add}(b, b');$ 
18        $\text{expand-nova-graph}(G, b');$ 
19     end
20   end

```

Algorithm 1 describes how a Nova agent constructs its norm-value graph by revising the norms in its opponent’s offer and decides its next action accordingly. When the agent receives an offer for the first time, the norm-value graph does not yet exist (Line 1). The agent creates the norm-value graph by applying a sequence of improving steps to promote its values (Lines 2–4). The overview of applying a sequence of improving steps are given in Lines 12–19. To ensure there are no cycles in the norm-value graph, the agent does not add any edge that results in a cycle even though doing so via an improving step may promote a value of the agent.

If the norm-value graph already exists, then there are two possibilities. One, the opponent's most recent offer is a node in the norm-value graph, enabling the agent to evaluate the received offer. Two, the opponent's most recent offer is not a node in the norm-value graph, preventing the agent to evaluate the received offer. In this case, the agent expands the norm-value graph by applying a sequence of improving steps to link the opponent's offer to an existing branch of the norm-value graph (Lines 2–4). Note that the agent's norm-value graph can always be expanded to include the opponent's current offer and that offer can be linked to an existing branch in the norm-value graph using our improving steps. There are two cases as follows:

- if opponent's current offer is better for the agent than the root node in the norm-value graph, then the opponent's current offer must already be contained in the agent's norm-value graph, since the opponent's current offer is reachable from the root node using the improving steps; or
- if opponent's current offer is worse for the agent than the root node in the norm-value graph, then the norm-value graph can be expanded such that the opponent's current offer can be linked to one of the (better) nodes in the norm-value graph using the improving steps.

Once the norm-value graph is expanded, the agent generates its offer using a bidding strategy (Line 5). According to the proposed approach, the bid generation is explained in Algorithm 2. By comparing its next offer and opponent's most recent offer, the agent decides whether or not to accept the opponent's offer (Lines 6 and 7). The bid score estimation is explained in Definition 3.13. We adopt the most-commonly used acceptance strategy AC-Next [Jonker et al. 2017], where the agent accepts its opponent's offer if the utility of the opponent is greater than or equal to the utility of that offer. If it does not accept the opponent's offer, then it makes a counteroffer (Lines 8 and 9). Now, we describe how an agent can reason about its norm-value graph.

When we apply an improving step ι to a given offer b and obtain $\iota(b)$, we draw an edge (labeled ι) from b to $\iota(b)$. That is, a child node in the preference graph is preferred over its parent node with respect to a value. Note that the importance of values may vary according to the role of the agent. A rational agent would make improving steps that promote its most important values. For example, authorizing a government officer to access a patient's record promotes "national security"; however, it demotes "privacy." Accordingly, we introduce a scoring function for improving steps. The score of an edge ι from offer b to offer $\iota(b)$ is calculated based on the weights of the agents' values. If the child offer $\iota(b)$ is preferred over parent offer b with respect to a value v , then it increases the improving score of the node by the weight of value v . To the contrary, if b is preferred over $\iota(b)$ with respect to a value v' , then we reduce the improving score of the offer by the weight of v' . We calculate the overall score S_i^b as the sum of individual scores.

Definition 3.12. For an agent AG, let w_v^{AG} denote the weight of value v for AG, such that $w_v^{AG} \geq 0$ for each v and $\sum_{v \in V} w_v^{AG} = 1$. The score of an improving step ι with respect to value $v \in V$ is calculated as follows:

$$s_i^b(v) = \begin{cases} +w_v^{AG} & \text{if } \iota(b) >_v b \\ -w_v^{AG} & \text{if } b >_v \iota(b) \\ 0 & \text{otherwise} \end{cases}.$$

The overall improving score of offer b produced by applying improving step ι is $S_i^b = \sum_{v \in V} s_i^b(v)$, where $S_i^b > 0$.

Consider an improving step, ι_{access} , from an offer consisting of one norm $\{a(\text{GOVERNMENT_AGENCY, HOSPITAL_ADMINISTRATION, true, share_patient_information})\}$ to $\{a(\text{GOVERNMENT_OFFICER,$

HOSPITAL_ADMINISTRATION, true, share_patient_information)) made by the hospital administration. Regarding privacy, the latter offer authorizes only a single government officer to share patient information and is preferred over the former offer that authorizes the whole government agency. Regarding national security, the former offer is preferred over the latter. Therefore, the score of this improving step, s_{access} , is calculated as $(w_{\text{privacy}} - w_{\text{national_security}})$.

We define a scoring function to evaluate offers according to the norm-value graph of the agent. The score of a given offer (a node in the graph) is calculated by summing the score of each improving step (following Definition 3.12) on the path from the source node(s) (i.e., root node) to the node containing that offer. Note that instead of using the term “root node,” we refer it as “source node,” since there might be multiple nodes, which do not have any parents in the norm-value graph. If there are multiple paths, then the score of the path whose overall score is maximum, is considered the final score, as mentioned below.

Definition 3.13. Let λ denote a path of improving steps from the source node to a node b . Then, the score for the given path is $U_\lambda = \sum_{i \in \lambda} S_i$. For any $b \in \mathcal{B}$, let Ω_b denote the set of all paths from the root node to b . Then, the score of node b is $U_b = \max_{\lambda \in \Omega} U_\lambda$.

Figure 2 depicts part of a norm-value graph. The root node denotes that the hospital administration unconditionally authorizes everyone to share patient records. From the point of the hospital administration, specializing the subject (in this example the authorizing government agency instead of everyone) will revise this norm regarding the privacy value. Alternatively, the root node can be revised by adopting more restrictive propositions (e.g., introducing a condition that the patient is involved in a criminal case). The score of those improving steps are denoted by S_{i_1} and S_{i_2} respectively. Similar improving steps can be applied for these two outcomes and consequently the outcome, $a(\text{GOVERNMENT_AGENCY}, \text{HOSPITAL_ADMINISTRATION}, \text{involve_case}, \text{share_patient_records})$ can be reached. Let's estimate the overall score of this newly constructed outcome according to Definition 3.13. That is, the sum of the path score calculation is done for each possible path from the root node to the current node, and the maximum score is considered as the final score. For this simplified example, there are two paths from the root node; thus, the score would be estimated as $\max(S_{i_1} + S_{i_4}, S_{i_2} + S_{i_5})$.

Following Aydoğan et al. [2015], we deploy a bidding strategy that searches through the norm-value graph as each improving step lets the agent reach a preferred outcome with respect to a value. In the norm-value graph, the opponent's first offer is one of the source nodes. When we apply all possible improving steps, the offers in the sink nodes become the most preferred outcomes. Since a longer path from the source node(s) indicates more improving steps, we expect to have more preferred outcomes in nodes that have greater depth in the graph. Accordingly, we propose a path-length-based bidding strategy that takes the remaining negotiation time and the opponent's most recent offer into account. Algorithm 2 presented this duration-based concession strategy. The agent initially has a tendency to make an offer at the maximum depth (i.e., the most preferred offer is at a sink node) and makes a concession by means of decreasing the depth of its offers over time. To decide the amount of concession, the agent estimates a target depth based on the remaining time (Line 1), which is similar to calculating the target utility in time-based concession strategies [Faratin et al. 1998]. That is, the target depth decreases gradually, computed as $\frac{\max_d \times t_c}{t_d}$, where $\max_d$ corresponds to the maximum depth of the sink nodes in the norm-value graph, t_c denotes current negotiation time, and t_d is the negotiation deadline. Among the nodes with the target depth, the nodes with the highest scores (Definition 3.13) are considered candidate offers (Lines 2–4). Among those offers, the agent chooses the offer that is the most similar to the opponent's most recent offer (Line 5). We adopt relative proximity, shortly RP similarity [Aydoğan and Yolum 2007] to compute similarity, which is initially set to one and reduced by a predefined

Fig. 2. Example norm-value graph (simplified).

amount every time a node is visited over the path. As shown in Lines 7–11, we estimated the weighted sum of the similarities for each component. The details of the similarity estimation can be found in Aydoğan and Yolum [2007].

ALGORITHM 2: Duration-based concession strategy.

Data: b_{op} , t_d , t_c and G is the norm-value graph

Result: b_{ag} is the agent's offer

```

1  $d_{target} \leftarrow \frac{Max_{depth} * t_c}{t_d}$ ;
2  $List_{bid} \leftarrow \text{retrieve-bids-at-depth}(G, d_{target})$ ;
3  $Max_{score} \leftarrow \text{getMaxScore}(G, List_{bid})$ ;
4  $List_{max-bids} \leftarrow \text{filter-bids-with-score}(G, List_{bid}, Max_{score})$ ;
5  $b_{ag} \leftarrow \text{argmax}\{\text{similar}(G, b_{op}, b_{ag}) \mid b_{ag} \in List_{max-bids}\}$ ;
6 Function  $\text{similar}(G, b_{op}, b_{ag})$ :
7    $sim_{sub} \leftarrow \text{RP-Similarity}(b_{ag}.subject, b_{op}.subject, G)$ ;
8    $sim_{ant} \leftarrow \text{RP-Similarity}(b_{ag}.antecedent, b_{op}.antecedent, G)$ ;
9    $sim_{con} \leftarrow \text{RP-Similarity}(b_{ag}.consequent, b_{op}.consequent, G)$ ;
10   $sim_{type} \leftarrow \text{RP-Similarity}(b_{ag}.normtype, b_{op}.normtype, G)$ ;
11  return  $(W_{sub} * sim_{sub}) + (W_{ant} * sim_{ant}) + (W_{cons} * sim_{cons}) + (W_{type} * sim_{type})$ ;
    
```

4 EVALUATION

To demonstrate that Nova supports creating transparent and accountable nMAS specifications via collaborative negotiations among the stakeholders, we conduct an experiment with human participants. First, we describe our use case. Second, we introduce our experiment setting and associated metrics. Third, we present our results and discuss their significance.

4.1 Use Case

We adopted a tradeoff scenario between two values: patient privacy and national security. Our scenario involves negotiation between the hospital administration and the government agency, inspired by the “Apple vs FBI” case [Hack 2016] and other similar application domains [Dijkstra et al. 2007]. The hospital administration prioritizes patient privacy over national security, and the government agency the reverse.

Figure 3 demonstrates the steps of a simple negotiation between the government agency and the hospital administration. First, the government agency makes an initial offer: The whole government agency is authorized to access patient records without further restrictions (i.e., the antecedent is true). Second, the hospital administration makes a counteroffer in response to the government agency’s offer: The hospital administration makes one type-based improving step—replace the initial authorization with a prohibition. Third, the government agency makes a counteroffer in response to the hospital administration’s offer: The government agency makes one type-based improving step—replace the prohibition back with the initial authorization—along with a concession on the antecedent of the authorization that restricts access to patient information only when there is a case of national security. Whereas this change is considered a concession with respect to national security (preferred by the government agency), it is an overall improving step (all values considered). Fourth, the hospital administration makes a counteroffer in response to the government agency’s offer: The hospital administration makes one actor-based improving step to the authorization—replace the subject with a more specific agent role (government officer).

4.2 Human-Subject Study

To realize the above scenario, we implemented a Web-based negotiator that enables human-agent negotiations. Details of the Web-based negotiator are in Appendix A. We recruited 16 participants (computer science students) and obtained their informed consent. The participants negotiate on behalf of the government agency. The government agency values national security more than patient privacy, and thus prefers to access patients’ information whenever possible with minimal restrictions. To the contrary, the hospital administration (enacted by our agent) prefers to restrict the government agency’s access to patients’ information under certain conditions, e.g., consent from the patient. We randomly assign the participants to two settings and instruct them to negotiate the following outcomes. *Challenging* setting (seven participants) requires participants to reach a norm where the government agency as a whole agency or a single government officer from the government agency should be able to access patients’ information without any restrictions (e.g., patient consent). *Concession* setting (nine participants) requires participants to reach a norm where it would be sufficient if a single government officer from the government agency can access patients’ information and the government agency is willing to present to the hospital administration that the patient is involved in a national security case. We compare the performance of our agent against a baseline strategy that starts with the offer at the maximum depth (max_d) and concedes one step at a time (i.e., $max_d - 1$, $max_d - 2$, ...) irrespective of the remaining time (similar strategy to Sequential Search Strategy in Aydoğan et al. [2010]). If there are multiple offers at the same depth, then it chooses one of them randomly. After each negotiation is over, the participants fill in a short questionnaire to rate the agreement and report their perceptions about Nova. The participants reported that the instructions to carry out the experiment were clear (means are 6.07/7 and 6/7 for challenging and concession settings, respectively). Details of the participant instructions and the post-experiment questionnaire are in Appendix B.

Table 2 summarizes our results. Our main hypothesis is that Nova performs better than the baseline. We test this hypothesis separately for each of the two settings and five metrics. We

Fig. 3. Sample negotiation turns.

perform two-tailed t -tests for social welfare and number of rounds. We perform Wilcoxon's rank-sum tests [Hollander and Wolfe 1999] for the user rating, collaboration score and transparency score. Wilcoxon's test (w) has a power advantage over the t -test for qualitative data, e.g., Likert scale. Moreover, we provide descriptive statistics beyond p -values. To measure the effect size [Grissom and Kim 2012] for the difference between means, we compute Hedges' g (recommended for measuring the magnitude of the difference for small sample sizes).

We compute social welfare by taking the average of the agent (the hospital administration and the government agency) scores of an agreement. To calculate the hospital administration's score, we normalize the score computed from a norm-value graph to 0–1. To calculate the government agency's score, we compare the agreement with the preference ordering from the participants' instructions (normalized to 0–1). A desired outcome for a negotiating party yields an individual score of 0.5–1. When both parties are considered, a social welfare above 0.5 is desirable. Nova achieves this threshold whereas baseline is below the threshold as three participants from each setting did not reach an agreement with the baseline. Although the differences in means in both settings are not significant at the 95% confidence level, p -values are small. Moreover, the effect sizes are large for both settings. We also report the number of rounds where Nova reaches better agreements faster than the baseline. However, the differences in means in both settings are not significant at the 95% confidence level. The effect size is medium for the challenging setting and small for the concession setting.

Table 2. Summary of Results

	Challenging				Concession			
	Nova	Baseline	Significance	Effect size	Nova	Baseline	Significance	Effect size
Average social welfare	0.50	0.28	$p = 0.08$	$g = 1.07$	0.50	0.33	$p = 0.09$	$g = 0.91$
Average number of rounds	8.71	11	$p = 0.27$	$g = 0.61$	8.56	9.78	$p = 0.37$	$g = 0.23$
Average user rating	6.43	3.71	$w = 7$	$g = 0.71$	7.78	4.56	$w = 3$	$g = 1.09$
Average collaboration score	4.50	3.21	$w = 6$	$g = 0.83$	4.78	3.33	$w = 5$	$g = 0.82$
Average transparency score	4.86	4.71	$w = 7$	$g = 0.07$	4.78	4.00	$w = 5$	$g = 0.48$

Social welfare ranges over [0–1] (higher is better). Number of rounds ranges over [1–20]. User rating ranges over [0–10] (higher is better). Collaboration and transparency scores are measured from a seven-point Likert scale (higher is better).

For user rating in the challenging setting, the w value is greater than the critical value for seven subjects at the 95% confidence level for Wilcoxon’s rank-sum tests. Thus, we conclude that the difference in means is not significant at the 95% confidence level. However, when we inspected individual ratings, we noticed that this is primarily due to one participant giving a rating of 1 to Nova and 10 to the baseline although the participant reached similar agreements in both cases. For the concession setting, the w value is less than the critical value for nine subjects at the 95% confidence level. Thus, we conclude that users gave significantly higher ratings to their agreements with Nova than the baseline. Moreover, the effect sizes are large for both settings. For collaboration score, we asked the participants to rate two statements, “The opponent was sensitive to my values” and “The opponent collaborated with me to reach an agreement” and computed the average of their responses. The difference in means is not significant for the challenging setting, whereas Nova is significantly more collaborative than the baseline in the concession setting. Moreover, the effect sizes are large for both settings. For transparency score, we asked the participants to rate the statement “The opponent negotiated in a way that made sense to me,” for which positive responses mean that the participants were able to understand the agent’s strategy. We conclude that Nova is significantly more transparent than the baseline in the concession setting, but no significant difference is observed in the challenging setting. The effect size is small for the challenging setting and medium for the concession setting.

In addition, we inspected individual negotiation traces to gather further insight into how participants interacted with Nova and whether the resulting agreements support accountability of various roles. We selected traces that are at least five rounds long and resulted in an agreement with an authorization (participant’s desired outcome). We noticed a common trend among participants who negotiated with the baseline. Such negotiations often lasted for 10–15 rounds with parties exchanging counteroffers that alternate between authorizations and prohibitions, and either not reaching an agreement or ending in prohibition agreements. To the contrary, a majority of negotiations against Nova ended in authorization agreements. One participant took 10 rounds to negotiate against the baseline following the above trend (no agreement), whereas the same participant reached an authorization agreement with Nova in 12 rounds. Another participant who negotiated against Nova took 14 rounds in total, and unlike most participants who accept the agent’s first authorization offer, this participant negotiated for five additional rounds to revise the authorization. We noticed that such negotiations resulted in finer control, i.e., only the required role is authorized via actor revision or additional antecedent conditions are specified via

proposition revision. Such interactions suggest that Nova enables concessions from both parties, thus corroborating our collaboration scores and providing accountability for specific roles in certain situations.

5 RELATED WORK

We now review the relevant literature grouped into the following major categories.

Automated Negotiation. The field of automated negotiation focuses on the design and evaluation of negotiation strategies for bidding [Baarslag et al. 2015], opponent modeling [Baarslag et al. 2016], and acceptance [Fatima et al. 2014], which makes the field challenging and open for future improvement [Baarslag et al. 2017]. The domains of negotiation are used strategically, to test negotiation strategies in domains of different complexity and size [Marsá-Maestre et al. 2014; Vahidov et al. 2017]. The **Automated Negotiating Agents Competition (ANAC)** [Jonker et al. 2017] employs predefined utility functions. We considered using ANAC agent strategies as benchmarks in our experiments. But, adapting those agents to our domain is not trivial—ANAC agents work with real-valued utilities normalized to $[0, 1]$, which would require an individual utility score for each norm element in an offer. In contrast, we generate offers dynamically via computing our norm relationships without the need for individual scores. We are not aware of any existing agent implementation that performs norm negotiation.

In our work, node scores in the norm-value graph (unlike utility functions) are dynamically computed by reasoning over an ontology of norms, values, and value preferences. This approach is relevant for situations in which agents do not know their preferences in advance [Fatima et al. 2014; Rahwan et al. 2003]. For instance, negotiating agents playing the Diplomacy game do not have predefined utility functions; the utilities of offers depend on the current state of the game [de Jonge et al. 2018]. Baarslag and Kaisers [2017] improve agents' preference models by obtaining feedback from users during the negotiation. Whereas Baarslag and Kaisers focus on how and when to improve the underlying preference model efficiently, we focus on dynamic evaluation of offers on the basis of the agents' values. One advantage of working with a norm-value graph is that agents' values and requirements are stable. Thus, an agent can represent the same user in repeated negotiations for various nMAS specifications. The agent needs only to know how propositions relate to values in a specific nMAS, which is captured in the ontology.

Deontic Logic and Normative Systems. Previous work often consider norms as *expected* social properties [Criado et al. 2013] that are typically enforced in an agent society via sanctions [Nardin et al. 2016]. Sergot [2013] formalize normative relations among agents as Hohfeldian legal concepts such as duties and rights. King et al. [2015] propose a hierarchical governance model for normative systems with a multitier architecture, where institutions on a higher level of the hierarchy govern institutions on lower levels. We build on these conceptual normative frameworks, and focus on legal norms [Hohfeld 1919], to provide a negotiation process for the design of nMAS. Including sanctions in Nova would be an interesting addition, that would provide a way for agents to associate norm violations with concrete follow-up actions.

Deontic logic theories and frameworks have been the basis for most normative multiagent systems. With respect to agency, deontic logic [Horty 2001] describes what the agents are permitted to do, obliged to do, and prohibited from doing. Although Von Wright introduced deontic logic in 1951, he later realized that his construction of permissions as duals of obligations was not effective [Von Wright 1963]. Our approach toward norms is aligned with Von Wright's later thinking, especially as explained in Von Wright [1999]. A key feature of our approach is that it does not treat authorization as a dual of commitment but rather, in Von Wright's term, as "prohibition of self."

Conflict Resolution. Santos et al. [2017] provide a review of previous techniques for the detection and resolution of normative conflicts, where a conflict is described as a situation where satisfying one norm violates another norm. One such technique is unification, where the components of a prohibition are checked whether they overlap with the components of an obligation or permission. Another type of conflict might occur among norms with the same deontic concept, such as two norms obliging the same agent to perform actions that cannot be performed at the same time. Beirlaen and Straßer [2014] propose an abstract multiagent deontic logic for dealing with obligation-obligation and obligation-permission conflicts. Benzmüller et al. [2018] propose a practical computational tool for legal and ethical reasoning regarding norms. They focus on three types of reasoning: compliance checking deals with the current state being compliant with a set of norms, consistency checking deals with the set of norms being consistent with each other, and entailment checking deals with norms following a stated regulation. Such reasoning can be provided via classic theorem provers, and a case study regarding legal norms from the General Data Protection Regulation is provided. In contrast with compliance checking and conflict resolution approaches, where clear-cut conflicts can be resolved, we provide a means for agents to deal with tradeoffs scenarios that arise among requirements. That is, rather than verifying whether a set of norms satisfies a set of requirements, Nova enables agents to come up with a set of norms (starting from an initial set of norms) that is acceptable to all the involved agents given the requirements of the nMAS.

Interest-Based Negotiation. Interest-based negotiation [Fisher et al. 1983] is the closest to Nova. Interest-based negotiation is based on the idea that agents' underlying concerns may differ from the issues they negotiate. In our work, the issues being negotiated (i.e., norms) are directly related to such concerns (i.e., values) via the ontology. Rahwan et al. [2003] propose an interest-based negotiation approach where an agent evaluates an offer based on its impact on satisfying the agent's goal. Boella et al. [2009] propose a protocol for negotiating norms and sanctions with respect to agents' goals in social settings, e.g., online multiplayer games. To deal with the underlying concerns of the negotiating parties, we adopt a value-based approach rather than reasoning about agents' goals. Our novelty comes from the systematic derivation of norm preference relations based on individual norm components and the underlying domain ontology, accompanied by a practical bidding strategy. Kollingbaum et al. [2006] focus on eliminating conflicts among negotiating parties. We deal with conflicts that are not as clear-cut as the those in Kollingbaum et al. [2006], since agents in our setting need concede to come to an agreement when tradeoffs are involved. Whereas they propose a qualitative conflict resolution approach, our work introduces a scoring function to evaluate the offers dynamically according to the values of the agent.

Norm Synthesis. Other authors discuss normative conflicts in the context of norm synthesis. Günay and Yolum [2013] discuss the feasibility of commitments, i.e., whether it is possible to satisfy all (existing and prospective) commitments of an agent. They formulate feasibility as a constraint satisfaction problem. Vasconcelos et al. [2009] propose a conflict resolution method, called *norm curtailment*, that manipulates the constraints associated with norms, e.g., reduce the scope of a prohibition to avoid conflict with an obligation. Zhang et al. [2016] discuss probabilistic commitments in open environments with uncertainty. Ajmeri et al. [2016] propose Coco, a formalism to express and reason about conflicting commitment instances at runtime, and dominance among them. Coco employs Answer Set Programming to compute nondominated commitment instances and uses [Alechina et al. 2013] framework to determine compliance of actions with nondominated commitment instances. In contrast to resolving normative conflicts by adopting one norm or another, a Nova agent can evaluate an offer using our proposed scoring function by analyzing the components of a norm and then decide which components to revise based on the agent's values. Therefore, the agents can agree on new norms that were not considered originally.

Kafalı and Yolum [2016] propose an approach for monitoring an agent's interactions to determine whether the agent is progressing as expected. An agent verifies whether its expectations are satisfiable by its current state. In their setting, agents are cooperative. In contrast, in our setting agents' requirements and expectations might be competing. Moreover, norms and normative conflicts have been investigated in the context of verification. El-Menshawy et al. [2015] survey research on verification, including model checking to monitor the satisfaction or violation of commitments. Giannikis and Daskalopulu [2011] take a centralized perspective where reasoning on contracts are performed by a single entity that has a global view of the environment.

In the context of agents cooperatively enacting business protocols, problems might arise and the agents need to diagnose the cause of the problems in a timely manner, with the aim of solving the problems either individually or cooperatively [Moubaidin and Obeid 2013]. To verify whether a business process (modelled with the standard BPMN notation, in which tasks can be related to each other via different relations) is compliant with its requirements, that are set out as obligations, Governatori [2013] proposes a conceptual framework to model normative requirements and different types of obligations (punctual and persistent). When checking for compliance, Governatori assumes that the obligations are already in force as well as the tasks that make up the process, to reason about violations and their possible compensations.

Argumentation. Dijkstra et al. [2007] present a multiagent framework that regulates the information exchange between agents through negotiation dialogues. They define a set of negotiation policies to determine how the agents act during the negotiation. For instance, if the opponent's offer includes an action that is prohibited by the agent, then the agent automatically rejects the offer. If the offer is not compatible with the agent's own interests, then the agent looks for additional conditions and makes a counteroffer with such conditions. That is, by querying its knowledge base, the agent finds the minimal set of conditions to add to the current offer so that the offer becomes compatible with its interests. Whereas Dijkstra et al.'s negotiation policy relies on logic and argumentation systems to generate justified or defensible arguments, our approach is based on reasoning on preferences over agents' values and heuristics for improving offers by also taking the negotiation deadline into account. Furthermore, they do not compute a preference relation over offers and accept any offer that is compatible with the agent's own interests, similar to a constraint satisfaction problem.

Argumentation frameworks have been applied in the context of legal norms [van der Torre and Villata 2014] as well as social norms [Modgil and Luck 2008]. Modgil and Luck [2008] differentiate between individual (desires) and normative goals of agents, and use BDI agents and dialogues based on Dung's abstract argumentation framework [Dung 1995] to resolve conflicts among norms. van der Torre and Villata [2014] use smoking regulations as a motivating example to reason about deontic detachment; they provide a formal grammar for legal norms. Whereas such frameworks can resolve clear-cut conflicts among norms via argumentation based attacks on individual norms and finally adopting one of the initially stated norms, NOVA is capable of generating new (components of) norms via negotiation among agents. Riveret et al. [2019] combine probabilistic argumentation and reinforcement learning to characterize various agent types, e.g., self-interested or norm compliant. They measure values of arguments via a utility function, which the agents can automatically learn by employing reinforcement learning. Extending NOVA with reinforcement learning would be an interesting future direction to investigate a series of negotiations where an agent can learn from previous interactions.

Self-organization. As a result of our negotiation process among the agents, new norms might emerge for the multiagent system in design. In that regard, NOVA is relevant to self-organizing multiagent systems. Morris-Martin et al. [2019] provide a literature survey on norm emergence

and convergence. They describe norms as expected behaviour patterns that are acceptable to the society as a whole. However, since societies change, norms need to be adapted accordingly. Norm emergence is the result of a bottom-up approach, where agents learn *appropriate* actions based on the observation of their previous interactions with other agents. Nova supports this type of norm emergence, both at the design time where a multiagent system is first built as a result of the negotiation among the involved agents, and at maintenance where agents might continue to (re)negotiate due to additional changes in their environment.

Picard et al. [2015] propose Agamemnon, a decentralized governance process for adaptive multiagent systems. They adopt roles, social goals and norms for the multiagent system, and reactive and proactive processes for agent reasoning. Agents can reorganize to adopt new contracts in their environment. Their implementation is based on previously proposed frameworks (such as Moise, Jason, Cartago, and Jacamo) and simulation experiments are provided on a smart parking scenario. Whereas their concept of a norm is tightly coupled with a system requirement (e.g., norms state quality of service related non-functional requirements in their scenario), their governance model—especially the governance policy and adaptation components that are responsible for monitoring, evaluating, and refining the system—can be adopted as a natural next step after the design of the nMAS using Nova.

Aldewereld et al. [2016] discuss group norms and propose a compliance based reasoning mechanism for agents that are involved in a group norm. They differentiate between actors who should achieve the goal stated by the norm, “responsibles” who are accountable for the norm (i.e., they are sanctioned if the norm is violated), and addressees who might be actors or responsibles depending on the context of the norm. Whereas our norm formalization covers the above roles with the subject and the object, and describe which individual agent is accountable to which other agent and to what, extending Nova with group norms and adopting the negotiation process to involve individual accountability in a given group norm would be an interesting direction for future work.

6 CONCLUSIONS AND FUTURE WORK

We presented Nova, an automated negotiation framework that tackles nMAS specification via norm revision. Our main contribution is the handling of tradeoffs among stakeholder requirements via automatic revision of norms using the domain ontology. In previous work, as explained in Section 5, normative conflicts are resolved by adopting one norm or another. However, in Nova, an agent can evaluate an offer using our proposed scoring function by analyzing the components of a norm and then decide which components to revise based on the agent’s values. A negotiation between two agents does not end with one of them adopting the other agent’s norm, but instead with them coming up with a new norm that may include completely new components that are not included in the initial norms from either agent. The agents use the domain ontology to reason about such norm components and create revisions accordingly. An example of this process is given in Figure 3—note that the agreed upon norm is not necessarily the “best” norm that either agent had in mind initially, but a result of their respective concessions.

Various circumstances might affect the behavior of the agent and how it is perceived by its opponent. The starting point for the generation of the agent’s norm-value graph is the opponent’s initial offer. If the opponent offers the least desired outcome for the agent, then it might be time consuming to generate the entire norm-value graph depending on the outcome space. In addition, it leads the agent to concede slowly against the opponent and explore more options. Therefore, the agent would be perceived less collaborative. If the opponent’s initial offer is relatively good for the agent, then the search space will be narrowed down. The agent will still make the best offers for itself but the opponent will perceive the agent more collaborative. As future work, we are planning to analyze the behavior of the agent against different human strategies.

We acknowledge several directions for future work. Offers can be extended to include additional nMAS components beyond norms such as technical mechanisms and domain assumptions. Current experimental setup can be extended to involve additional agents (more than two) and norms (more than one) in the negotiation process. A number of additional baseline strategies can be included to compare against altruistic, selfish, and random agents. Such comparisons may involve agent-agent simulations with large number of runs to achieve statistical significance of results. Moreover, a large-scale human participant experiment can be conducted drawing upon our findings, which may include human–human negotiations to compare against. Dsouza et al. [2013] present a negotiation approach where negotiation participants can ask about their opponent’s private objectives in a repeated negotiation setting. Their results show that disclosing such objectives leads to more beneficial agreement outcomes for the participants. Inspired from this work, we can extend our agent strategy to provide supporting information, e.g., disclosing the values they prioritize, why the given offer is not acceptable for them.

APPENDICES

A WEB-BASED NEGOTIATOR

We developed a web-based negotiator tool, called *Nova Negotiation Simulator* (Figure 4) to test our use-case scenario with human participants.

Fig. 4. Web-based negotiator.

A.1 Negotiator Agent Selection Page

In the beginning of the simulation the participants should fill his/her name and surname and select an negotiation agent. There are two types agent in our simulation. The “*opponent-1*” represents “*baseline agent*,” and the second type of agent is “*opponent-2*” represents the “*nova agent*.” In order for the participants to evaluate the two different strategies in an unbiased manner, the information about which agent applied which strategy was hidden from them.

A.2 Negotiation Process

After the negotiator agent selection, the negotiation process between the participant and the agent begins. The participant makes the first bid according to the preferences order given to him. To construct a bid, there are four drop-downs at the bottom of the screen (Figure 5). From these drop-downs, the participants should select a subject, a norm, a consequent, and an antecedent in, respectively. These selections will be converted to a bid representation internally that the agent can compute. These selections are internally converted into a norm representation (Definition 2.1).

Fig. 5. Bid construction.

Fig. 6. Accept the opponent's last offer or produce a counteroffer.

Then the negotiator agent can process and evaluate this proposed norm. According to proposed norm, the agent either propose a new bid or accept the participant's offer.

If the agent proposes a new bid, then the participant will observe the new bid via chat screen. After that, the two selection buttons will appear at the bottom of the screen for the participant to choose either accepting the bid of the agent or continuing the negotiation (Figure 6).

A.3 Evaluation Page

The negotiation continues until one of the negotiators accept a bid or 5 minutes deadline ended. After negotiation process ended the evaluation page will be shown. The participant will evaluate satisfaction level of the agreement by using stars in the middle of the screen. Then the experiment will be over for this setup.

B MATERIALS NEEDED FOR THE EXPERIMENTS

Next, we describe the instructions given to the participants in the challenging setting (Group₁) and the concession setting (Group₂).

B.1 Nova Experiment Instructions for Group₁—Challenging Setting

In this experiment, we will be simulating the negotiation between a **hospital administration (HA)** and a **government agency (GA)** to reach a norm for regulating the access to patients' data.

The HA has *privacy concerns*. Therefore, it may try to restrict the GA's access to patients' data under certain conditions, e.g., written consent from the patient or warrant, and so on, and not willing to easily permit the GA to access patients' data. To the contrary, the GA values *national security more than patient privacy*, thus will try to access patients' data whenever possible with *minimal restrictions*.

Group 1 Role Profile: Group 1 will play *the role of the GA* and try to get a norm from the HA that provides flexible access to patients' data. During the negotiation, the bids of the Group 1 will consist of the following items:

- For **“Subject”**: You should prefer *Government Agency* to *Government Officer* for authorization (i.e., may). In addition, you will be pretty okay with *“Everyone”* when it is authorization. When the norm is prohibition (i.e, should not), your preferences are opposite.
- For **“Norm”**: You strictly prefer *“may”* over *“should not.”*
 - *“may”*: authorize the selected subject to perform the selected actions [What] under the selected conditions [Condition];
 - or *“should not”*: prohibit the selected subject from performing the selected actions [What] under the selected conditions [Condition].
- For **“What”**:
 - You prefer *“modify”* to *“share”* and *“share”* is preferred over *“access”*. (modify > share > access)
 - Patient records cover both personal information and medical treatment data. Personal information is more useful for the case rather than medical treatment data. Therefore, you prefer *patient records* to *personal information* and prefer *personal information* to *medical history*. (patient records > personal information > medical history)
- For **“When”**: The preferences on when conditions are listed below:
Always > Patient is involved in a criminal case > Duty on Patient Case > Case concerns national security > Police has warrant > Patient gives consent.

The goal of the participants of Group 1 is to reach an agreement in 5 minutes and the minimum acceptance condition specified below. Although this outcome would be acceptable for you, you would prefer to get a stronger authorization from the HA.

Minimum acceptable norm for Group₁: The GA as a *whole department* or a single **government officer (GO)** from the GA should be able to access any patients' data *without further restrictions* (e.g., no need to obtain consent or warrant). Note that modifying or sharing data naturally involves accessing the data.

Steps for the Experiment:

- (1) Please study the scenario carefully and when you are ready follow the link to the application:
- (2) Enter your name and choose "Opponent-2." Please recall that you have only 5 minutes to complete your negotiation. When you click "Start Negotiation" button, negotiation starts. It is recommended to check your time before clicking the Start Negotiation button.
- (3) After your negotiation, please fill in the questionnaire form.
- (4) Now you will do another negotiation. Enter your name and choose "Opponent-1." Please recall that you have only 5 minutes to complete your negotiation. When you click "Start Negotiation" button, negotiation starts. It is recommended to check your time before clicking the Start Negotiation button.
- (5) After your negotiation, please fill in the questionnaire form.

B.2 Nova Experiment Instructions for Group₂—Concession Setting

In this experiment, we will be simulating the negotiation between a HA and a GA to reach a norm for regulating the access to patients' data.

The HA has *privacy concerns*. Therefore, it may try to restrict the GA's access to patients' data under certain conditions, e.g., written consent from the patient or warrant, and so on, and not willing to easily permit the GA to access patients' data. To the contrary, the GA values *national security more than patient privacy*, thus will try to access patients' data whenever possible with *minimal restrictions*.

Group 2 Role Profile: Group 2 will play *the role of the GA* and try to get a norm from the HA that provides flexible access to patients' data. During the negotiation, the bids of the Group 2 will consist of following items:

- For "**Subject**": You should prefer *Government Agency* to *Government Officer* for authorization (i.e., may). In addition, you will be pretty okay with "*Everyone*" when it is authorization. When the norm is prohibition (i.e., should not), your preferences are opposite.
- For "**Norm**": You strictly prefer "*may*" over "*should not*."
 - "*may*": authorize the selected subject to perform the selected actions [What] under the selected conditions [Condition];
 - or "*should not*": prohibit the selected subject from performing the selected actions [What] under the selected conditions [Condition].
- For "**What**":
 - You prefer "*modify*" to "*share*" and "*share*" is preferred over "*access*". (modify > share > access)
 - Patient records cover both personal information and medical treatment data. Personal information is more useful for the case rather than medical treatment data. Therefore, you prefer *patient records* to *personal information* and prefer *personal information* to *medical history*. (patient records > personal information > medical history)

- For **“When”**: The preferences on when conditions are listed below:
Always > Patient is involved in a criminal case > Duty on Patient Case > Case concerns national security > Police has warrant > Patient gives consent.

The goal of the participants of Group 2 is to reach an agreement in 5 minutes and the minimum acceptance condition specified below. Although this outcome would be acceptable for you, you would prefer to get a stronger authorization from the HA.

Minimum acceptable norm for Group₂: It would be sufficient even if a single GO from the GA may access patients’ data and the GA will be willing to present a case on national security to the hospital (e.g., Patients is involved in a criminal case, case concerns national security) but *cannot* provide consent or warrant.

Steps for the Experiment:

- (1) Please study the scenario carefully and when you are ready follow the link to the application: <https://pasta-framework.firebaseio.com/>
- (2) Enter your name and choose *“Opponent-2.”* Please recall that you have only *5 minutes* to complete your negotiation. When you click *“Start Negotiation”* button, negotiation starts. It is recommended to check your time before clicking the Start Negotiation button.
- (3) After your negotiation, please fill in the questionnaire form.
- (4) Now you will do another negotiation. Enter your name and choose *“Opponent-1.”* Please recall that you have only *5 minutes* to complete your negotiation. When you click *“Start Negotiation”* button, negotiation starts. It is recommended to check your time before clicking the Start Negotiation button.
- (5) After your negotiation, please fill in the questionnaire form.

B.3 Questionnaire

- Q₁: The Name of The Participant
 Q₂: Which group did participant play for? It can be either Group1 or Group2
 Q₃: Which opponent did the participant negotiate with? It can be either Opponent 1 or Opponent 2
 Q₄: Is the participant confident about his/her negotiation skills? (Rating between 1 and 7)
 Q₅: Were the instructions provided to the participant for the experimental negotiation clear? (Rating between 1 and 7)
 Q₆: Was it easy to decide what to offer during the negotiation? (Rating between 1 and 7)
 Q₇: Select bidding strategy were you used during the negotiation.
 O₁: Time-based Concession (i.e., starting with the most desired offer and conceding over time)
 O₂: Boulware (i.e., insisting on my best offer but starting to concede when deadline is approaching)
 O₃: Tit for Tat (i.e., mimicking my opponent’s moves and acting accordingly)
 O₄: Hardheaded (i.e., insisting on my best offer and never conceding)
 O₅: Random Dancer (i.e., making my bids randomly)
 O₆: Conceder (i.e., starting with a reasonable offer and conceding over time)
 Q₈: Select acceptance strategy were you used in the negotiation:
 O₁: I did not accept my opponent’s offer
 O₂: I accepted my opponent’s offer since it was as good as my current offer.
 O₃: I accepted my opponent’s offer since it met my minimum acceptance condition.

- O₄: I accepted my opponent's offer even though it did not meet my minimum acceptance condition
- Q₉: Did I take the opponent's values into account (such as patient privacy, or hospital reputation) when making your offers? (Yes or No question)
- Q₁₀: The opponent negotiated in a way that made sense to me. (Rating between 1 and 7)
- Q₁₁: The opponent was sensitive to my values such as national security. (Rating between 1 and 7)
- Q₁₂: The opponent collaborated with me to reach an agreement. (Rating between 1 and 7)
- Q₁₃: Select acceptance strategy were you used in the negotiation:
- O₁: Time-based Concession (i.e., starting with the most desired offer and conceding over time)
 - O₂: Boulware (i.e., insisting on my best offer but starting to concede when deadline is approaching)
 - O₃: Tit for Tat (i.e., mimicking my opponent's moves and acting accordingly)
 - O₄: Hardheaded (i.e., insisting on my best offer and never conceding)
 - O₅: Random Dancer (i.e., making my bids randomly)
 - O₆: Conceder (i.e., starting with a reasonable offer and conceding over time)
- Q₁₄: Select acceptance strategy were you used in the negotiation:
- O₁: She did not accept my offer
 - O₂: It was as good as her current offer.
 - O₃: It met her minimum acceptance condition.
 - O₄: Even though it did not meet her minimum acceptance condition.
- Q₁₅: My opponent was negotiating like a human. (Rating between 1 and 7)

REFERENCES

- AIHLEG. 2019. Ethics Guidelines for Trustworthy AI. (2019). High-Level Expert Group on Artificial Intelligence. Retrieved from https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419.
- Nirav Ajmeri, Jiaming Jiang, Rada Y. 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. <https://www.ijcai.org/Abstract/16/010>.
- Huib Aldewereld, Virginia Dignum, and Wamberto W. Vasconcelos. 2016. Group norms for multi-agent organisations. *ACM Trans. Auton. Adapt. Syst.* 11, 2 (2016), 1–31.
- Natasha Alechina, Mehdi Dastani, and Brian Logan. 2013. Reasoning about normative update. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*. 20–26.
- Reyhan Aydoğan, David Festen, Koen V. Hindriks, and Catholijn M. Jonker. 2017. Alternating offers protocols for multi-lateral negotiation. In *Modern Approaches to Agent-based Complex Automated Negotiation*, Katsuhide Fujita, Quan Bai, Takayuki Ito, Minjie Zhang, Fenghui Ren, Reyhan Aydoğan, and Rafik Hadfi (Eds.). Number 674 in Studies in Computational Intelligence. Springer, Cham, 153–167. https://doi.org/10.1007/978-3-319-51563-2_10
- Reyhan Aydoğan, Tim Baarslag, Koen V. Hindriks, Catholijn M. Jonker, and Pinar Yolum. 2015. Heuristics for using CP-nets in utility-based negotiation without knowing utilities. *Knowl. Inf. Syst.* 45, 2 (2015), 357–388.
- Reyhan Aydoğan, Nuri Taşdemir, and Pinar Yolum. 2010. Reasoning and negotiating with complex preferences using CP-nets. In *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis*. Springer, Berlin, 15–28.
- Reyhan Aydoğan and Pinar Yolum. 2007. Learning consumer preferences using semantic similarity. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*. IFAAMAS, 1293–1300.
- Reyhan Aydoğan and Pinar Yolum. 2012. Learning opponent's preferences for effective negotiation: An approach based on concept learning. *Autonom. Agents Multi-Agent Syst.* 24, 1 (2012), 104–140.
- Tim Baarslag, Enrico H. Gerding, Reyhan Aydoğan, and M. C. Schraefel. 2015. Optimal negotiation decision functions in time-sensitive domains. In *Proceedings of the 2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT'15)*, Vol. 2. IEEE, 190–197.
- Tim Baarslag, Mark J. C. Hendriks, Koen V. Hindriks, and Catholijn M. Jonker. 2016. Learning about the opponent in automated bilateral negotiation: A comprehensive survey of opponent modeling techniques. *Auton. Agent Multi Agent Syst.* 30, 5 (2016), 849–898.

- Tim Baarslag and Michael Kaisers. 2017. The value of information in automated negotiation: A decision model for eliciting user preferences. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems*. IFAAMAS, 391–400.
- Tim Baarslag, Michael Kaisers, Enrico H. Gerding, Catholijn M. Jonker, and Jonathan Gratch. 2017. When will negotiation agents be able to represent us? The challenges and opportunities for autonomous negotiators. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17)*. International Joint Conferences on Artificial Intelligence Organization, 4684–4690. <https://doi.org/10.24963/ijcai.2017/653>
- Mathieu Beirlaen and Christian Straßer. 2014. Non-monotonic reasoning with normative conflicts in multi-agent deontic logic. *J. Logic Comput.* 24, 6 (2014), 1179–1207.
- Christoph Benz Müller, Xavier Parent, and Leendert van der Torre. 2018. A deontic logic reasoning infrastructure. In *Proceedings of the Conference on Computability in Europe*. Springer, 60–69.
- Guido Boella, Patrice Caire, and Leendert van der Torre. 2009. Norm negotiation in online multi-player games. *Knowl. Inf. Syst.* 18, 2 (2009), 137–156.
- Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. 2004. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res.* 21 (2004), 135–191.
- Peter J. Carnevale and Tahira M. Probst. 1998. Social values and social conflict in creative problem solving and categorization. *J. Pers. Soc. Psychol.* 74, 5 (1998), 1300.
- Federico Chesani, Paola Mello, Marco Montali, and Paolo Torroni. 2013. Representing and monitoring social commitments using the event calculus. *Autonom. Agents Multi-Agent Syst.* 27, 1 (Jul. 2013), 85–130.
- Amit K. Chopra, Fabiano Dalpiaz, F. Başak Aydemir, Paolo Giorgini, John Mylopoulos, and Munindar P. Singh. 2014. Protos: Foundations for engineering innovative sociotechnical systems. In *Proceedings of the 22nd IEEE International Requirements Engineering Conference (RE)*. IEEE Computer Society, 53–62. <https://doi.org/10.1109/RE.2014.6912247>
- Natalia Criado, Estefania Argente, Pablo Noriega, and Vicente Botti. 2013. MaNEA: A distributed architecture for enforcing norms in open MAS. *Eng. Appl. Artif. Intell.* 26, 1 (Jan. 2013), 76–95.
- Dave de Jonge, Tim Baarslag, Reyhan Aydoğan, Catholijn M. Jonker, Katsuhide Fujita, and Takayuki Ito. 2018. The challenge of negotiation in the game of diplomacy. In *Proceedings of the 6th International Conference on Agreement Technologies*. Springer International Publishing, 100–114.
- Pieter Dijkstra, Henry Prakken, and Kees de Vey Mestdag. 2007. An implementation of norm-based agent negotiation. In *Proceedings of the 11th International Conference on Artificial Intelligence and Law*. ACM, 167–175.
- Sohan Dsouza, Ya'akov K. Gal, Philippe Pasquier, Sherief Abdallah, and Iyad Rahwan. 2013. Reasoning about goal revelation in human negotiation. *IEEE Intell. Syst.* 28, 2 (2013), 74–80. <https://doi.org/10.1109/MIS.2011.93>
- Phan Minh Dung. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.* 77, 2 (1995), 321–357.
- Mohamed El-Menshaw, Jamal Bentahar, Warda El Kholy, Pinar Yolum, and Rachida Dssouli. 2015. Computational logics and verification techniques of multi-agent commitments: Survey. *Knowl. Eng. Rev.* 30, 5 (Nov. 2015), 564–606.
- Peyman Faratin, Carlos Sierra, and Nick R. Jennings. 1998. Negotiation decision functions for autonomous agents. *Robot. Auton. Syst.* 24, 3 (1998), 159–182.
- Shaheen Fatima, Sarit Kraus, and Michael Wooldridge. 2014. *Principles of Automated Negotiation* (1st ed.). Cambridge University Press, New York, NY.
- Roger Fisher, William L. Ury, and Bruce Patton. 1983. *Getting to Yes: Negotiating Agreement without Giving in* (3rd ed.). Penguin Books, New York.
- Georgios K. Giannikis and Aspasia Daskalopulu. 2011. Normative conflicts in electronic contracts. *Electr. Commerce Res. Appl.* 10, 2 (2011), 247–267.
- Guido Governatori. 2013. Business process compliance: An abstract normative framework. *Inf. Technol.* 55, 6 (2013), 231–238.
- Robert J. Grissom and John J. Kim. 2012. *Effect Sizes for Research: Univariate and Multivariate Applications*. Routledge, Abingdon-on-Thames.
- Nicola Guarino. 1998. *Formal Ontology in Information Systems* (1st ed.). IOS Press, Amsterdam.
- Akın Günay and Pinar Yolum. 2013. Constraint satisfaction as a tool for modeling and checking feasibility of multiagent commitments. *Appl. Intell.* 39, 3 (Oct. 2013), 489–509.
- Michael Hack. 2016. The implications of Apple's battle with the FBI. *Netw. Secur.* 2016, 7 (2016), 8–10.
- HHS. 2014. Bulletin: HIPAA Privacy in Emergency Situations. (2014). United States Department of Health and Human Services (HHS). Retrieved from <http://www.hhs.gov/ocr/privacy/hipaa/understanding/special/emergency/>.
- Koen V. Hindriks and Dmytro Tykhonov. 2008. Opponent modelling in automated multi-issue negotiation using bayesian learning. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*. IFAAMAS, 331–338.

- Wesley Newcomb Hohfeld. 1919. *Fundamental Legal Conceptions as Applied in Judicial Reasoning and Other Legal Essays*. Yale University Press, New Haven, CT.
- Myles Hollander and Douglas A. Wolfe. 1999. *Nonparametric Statistical Methods*. Wiley, New York.
- John F. Horty. 2001. *Agency and Deontic Logic*. Oxford University Press, Oxford, UK.
- Catholijn M. Jonker, Reyhan Aydoğan, Tim Baarslag, Katsuhide Fujita, Takayuki Ito, and Koen V. Hindriks. 2017. Automated negotiating agents competition (ANAC). In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI'17)*. AAAI Press, 5070–5072.
- Özgür Kafalı, Nirav Ajmeri, and Munindar P. Singh. 2016. Revani: Revising and verifying normative specifications for privacy. *IEEE Intell. Syst.* 31, 5 (Sep. 2016), 8–15. <https://doi.org/10.1109/MIS.2016.89>
- Özgür Kafalı, Nirav Ajmeri, and Munindar P. Singh. 2017. Kont: Computing tradeoffs in normative multiagent systems. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI'17)*. AAAI, 3006–3012.
- Özgür Kafalı, Nirav Ajmeri, and Munindar P. Singh. 2020. DESEN: Specification of sociotechnical systems via patterns of regulation and control. *ACM Trans. Softw. Eng. Methodol.* 29, 1 (Feb. 2020), 7:1–7:50. <https://doi.org/10.1145/3365664>
- Özgür Kafalı and Paolo Torroni. 2012. Exception diagnosis in multiagent contract executions. *Ann. Math. Artif. Intell.* 64, 1 (Jan. 2012), 73–107.
- Özgür Kafalı and Pinar Yolum. 2016. Pisagor: A proactive software agent for monitoring interactions. *Knowl. Inf. Syst.* 47, 1 (Apr. 2016), 215–239.
- Thomas C. King, Tingting Li, Marina De Vos, Virginia Dignum, Catholijn M. Jonker, Julian Padget, and B. Birna van Riemsdijk. 2015. A framework for institutions governing institutions. In *Proceedings of the 14th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS'15)*. IFAAMAS, 473–481.
- Martin J. Kollingbaum, Timothy J. Norman, Alun Preece, and Derek Sleeman. 2006. Norm refinement: Informing the re-negotiation of contracts. In *Proceedings of the International Workshop on Coordination, Organization, Institutions and Norms in Agent Systems*. IFAAMAS, 46–51.
- Srdjan Marinovic, Naranker Dulay, and Morris Sloman. 2014. Rumpole: An introspective break-glass access control language. *ACM Trans. Inf. Syst. Secur.* 17, 1 (Aug. 2014), 2:1–2:31.
- Ivan Marsá-Maestre, Miguel A. López-Carmona, Mark Klein, Takayuki Ito, and Katsuhide Fujita. 2014. Addressing utility space complexity in negotiations involving highly uncorrelated, constraint-based utility spaces. *Comput. Intell.* 30, 1 (2014), 1–29.
- Sanjay Modgil and Michael Luck. 2008. Argumentation based resolution of conflicts between desires and normative goals. In *International Workshop on Argumentation in Multi-agent Systems*. Springer, 19–36.
- Andreas Morris-Martin, Marina De Vos, and Julian Padget. 2019. Norm emergence in multiagent systems: A viewpoint paper. *Auton. Agents Multi-Agent Syst.* 33, 6 (2019), 706–749.
- Asma Moubaidin and Nadim Obeid. 2013. On formalizing social commitments in dialogue and argumentation models using temporal defeasible logic. *Knowl. Inf. Syst.* 37, 2 (2013), 417–452.
- Luis G. Nardin, Tina Balke-Visser, Nirav Ajmeri, Anup K. Kalia, Jaime S. Sichman, and Munindar P. Singh. 2016. Classifying sanctions and designing a conceptual sanctioning process model for socio-technical systems. *Knowl. Eng. Rev.* 31, 2 (Mar. 2016), 142–166.
- Gauthier Picard, Camille Persson, Olivier Boissier, and Fano Ramparany. 2015. Multi-agent self-organization and reorganization to adapt M2M infrastructures. In *Proceedings of the 9th International Conference on Self-Adaptive and Self-Organizing Systems*. IEEE, 91–100.
- Iyad Rahwan, Liz Sonenberg, and Frank Dignum. 2003. Towards interest-based negotiation. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems*. IFAAMAS, 773–780.
- Régis Riveret, Yang Gao, Guido Governatori, Antonino Rotolo, Jeremy Pitt, and Giovanni Sartor. 2019. A probabilistic argumentation framework for reinforcement learning agents. *Auton. Agents Multi-Agent Syst.* 33, 1 (2019), 216–274.
- Ariel Rubinstein. 1992. The electronic mail game: Strategic behavior under “almost common knowledge.” In *Knowledge, Belief, and Strategic Interaction*, Cristina Bicchieri and Maria Luisa Dalla Chiara (Eds.). Cambridge University Press, Cambridge, UK, 317–326.
- Jéssica S. Santos, Jean O. Zahn, Eduardo A. Silvestre, Viviane T. Silva, and Wamberto W. Vasconcelos. 2017. Detection and resolution of normative conflicts in multi-agent systems: A literature survey. *Auton. Agents Multi-agent Syst.* 31, 6 (2017), 1236–1282.
- Marek Sergot. 2013. Normative positions. In *Handbook of Deontic Logic and Normative Systems*, Dov Gabbay, John Horty, Ron van der Meyden, Xavier Parent, and Leendert van der Torre (Eds.). College Publications, London, 353–406.
- Munindar P. Singh. 2013. Norms as a basis for governing sociotechnical systems. *ACM Trans. Intell. Syst. Technol.* 5, 1 (Dec. 2013), 21:1–21:23. <https://doi.org/10.1145/2542182.2542203>
- Okan Tunali, Reyhan Aydoğan, and Victor Sanchez-Anguix. 2017. Rethinking frequency opponent modeling in automated negotiation. In *Proceedings of the International Conference on Principles and Practice of Multi-Agent Systems*. Springer International Publishing, 263–279.

- Rustam Vahidov, Gregory E. Kersten, and Bo Yu. 2017. Human-agent negotiations: the impact agents' concession schedule and task complexity on agreements. In *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS'17)*. HICSS Press, 412–420.
- Leendert van der Torre and Serena Villata. 2014. An ASPIC-based legal argumentation framework for deontic reasoning. In *Proceedings of the 5th International Conference on Computational Models of Argument*. IOS Press, 266–421.
- W. Fred Van Raaij and Theo M. M. Verhallen. 1994. Domain-specific market segmentation. *Eur. J. Market.* 28, 10 (1994), 49–66.
- Wamberto W. Vasconcelos, Martin J. Kollingbaum, and Timothy J. Norman. 2009. Normative conflict resolution in multi-agent systems. *Auton. Agents Multi-Agent Syst.* 19, 2 (Oct. 2009), 124–152.
- Georg Henrik Von Wright. 1963. *Norm and Action: A Logical Enquiry*. Humanities Press, New York.
- Georg Henrik Von Wright. 1999. Deontic Logic: A personal view. *Ratio Juris* 12, 1 (Mar. 1999), 26–38. <https://doi.org/10.1111/1467-9337.00106>
- Colin R. Williams, Valentin Robu, Enrico H. Gerding, and Nick R. Jennings. 2013. Iamhaggler2011: A gaussian process regression based negotiation agent. In *Complex Automated Negotiations: Theories, Models, and Software Competitions*. Springer, Berlin, 209–212.
- Qi Zhang, Edmund H. Durfee, Satinder P. Singh, Anna Chen, and Stefan J. Witwicki. 2016. Commitment semantics for sequential decision making under reward uncertainty. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI'16)*. International Joint Conferences on Artificial Intelligence Organization, 3315–3323.

Received June 2020; revised February 2021; accepted May 2021