# Computers That Negotiate on Our Behalf: Major Challenges for Self-sufficient, Self-directed, and Interdependent Negotiating Agents

Tim Baarslag<sup>1</sup>, Michael Kaisers<sup>1</sup>, Enrico H. Gerding<sup>2</sup>, Catholijn M. Jonker<sup>3</sup>, Jonathan Gratch<sup>4</sup>

 Intelligent and Autonomous Systems Group Centrum Wiskunde & Informatica Amsterdam, The Netherlands
Agents, Interaction and Complexity Group University of Southampton Southampton, UK
<sup>3</sup> Interactive Intelligence Group Delft University of Technology Delft, The Netherlands
<sup>4</sup> Institute for Creative Technologies University of Southern California Marina del Rey, CA, USA

Abstract. Computers that negotiate on our behalf hold great promise for the future and will even become indispensable in emerging application domains such as the smart grid, autonomous driving, and the Internet of Things. Much research has thus been expended to create agents that are able to negotiate in an abundance of circumstances. However, up until now, truly autonomous negotiators have rarely been deployed in real-world applications. This paper sizes up current negotiating agents and explores a number of technological, societal and ethical challenges that autonomous negotiation systems are bringing about. The questions we address are: in what sense are these systems autonomous, what has been holding back their further proliferation, and is their spread something we should encourage? We relate the automated negotiation research agenda to dimensions of autonomy and distill three major themes that we believe will propel autonomous negotiation forward: accurate representation, long-term perspective, and user trust. We argue these orthogonal research directions need to be aligned and advanced in unison to sustain tangible progress in the field.

# 1 Introduction

Negotiation, the process of joint decision making, is pervasive in our society [35]. Whenever actors meet and influence each other to forge a mutually beneficial agreement, a form of negotiation is at work [76].

Negotiation arises in almost every social and organizational setting, yet many avoid it out of fear or lack of skill and this contributes to income inequality [9], political gridlock [34] and social injustice [26]. This has led to an increasing focus on the design of autonomous negotiators capable of automatically and independently negotiating with others. This interest has been spurred since the beginning of the 1980s with the work of early pioneers such as Smith [66] and Sycara [67].

Automated negotiation research is fueled by a number of benefits that computerized negotiation can offer, including better (win-win) deals, and reduction in time, costs, stress and cognitive effort on the part of the user. Moreover, autonomous negotiation will soon become not just desired but *required* in instances where the human scale is simply too slow and expensive. For instance, with the world-wide deployment of the smart electrical grid and the must for renewable energy sources, flexible devices in our household will soon (re-)negotiate complex energy contracts automatically. Another example is the rise of the Internet of Things (IoT), which will introduce countless smart, interconnected devices that autonomously negotiate the usage of sensitive data and make trade-offs between privacy concerns, price, and convenience.

To properly fulfill its representational role in an ever-dynamic environment, a negotiation agent has to balance and adhere to different aspects of autonomous behavior, including self-reliance and the capability and freedom to perform its actions, while at the same time remaining interdependent in its joint activity with the user. While many successes have been achieved in advancing various degrees of autonomy in negotiating agents, it is readily apparent that fully-deployed and truly autonomous negotiators are still a thing of the future. Continued development will be required before agents will be able to forge even mundane agreements such as the personalized renewal of our energy or mobile phone contracts. This begs the obvious question: what is still lacking currently and what is needed for autonomous negotiators to be able to fulfill their promise?

This paper discusses the challenges and upcoming application domains for (almost) entirely autonomous negotiation on people's behalf, extending the vision set out in [8]. We describe the technological challenges associated with these future domains and provide a roadmap towards full autonomy, together with stops along the way, highlighting what we deem important solution concepts for enabling future autonomous negotiation systems. As a basis for our discussion, we provide a unifying view of autonomous negotiation based on three orthogonal dimensions of autonomy that research has focused on so far: being self-sufficient, self-directed, and interdependent. We argue that automated negotiation opportunities of tomorrow are calling for a combined effort to address these three pillars of a negotiator's autonomy.

This paper does not aim to survey all research or challenges in the field comprehensively, but rather presents pointers to what we consider important focal points for autonomous negotiation, now and in the future. We pinpoint and elaborate on the following major challenges for autonomous negotiation:

- 1. Domain knowledge and preference elicitation;
- 2. Long-term perspective; and
- 3. User trust and adoption.

Lastly, this paper also pays homage to the 2001 landmark publication by Jennings et al. [42] and asks what has happened, 16 years later, with the prospects and challenges of automated negotiation. We examine which main challenges have been addressed, and

which stay relevant in a world that offers more opportunities for automated negotiation than ever before.

# 2 The Autonomy Diagonal of Negotiation

Autonomous negotiation is more than just *automated* negotiation; it is the freedom to negotiate independently. Rather than being uni-dimensional, autonomy incorporates at least two components [14]: *self-sufficiency* (the capability of the actor to take care of itself) and *self-directedness* (the freedom to act within the environment and the means to reach goals). Following [44] we distinguish a third dimension called *support for interdependence* – being able to work with others and influence and be influenced by team members.<sup>5</sup>



**Fig. 1.** By and large, negotiation research can be clustered around one of the three main orthogonal dimensions of autonomy: *self-sufficiency*, *self-directedness*, and *interdependence*. The efforts of the three need to be integrated to arrive at truly autonomous negotiators that can progress along the *autonomy diagonal*.

<sup>&</sup>lt;sup>5</sup> Note that the notion of autonomy is notoriously difficult to capture (see [44] for an overview). We are concerned here with those aspects especially relevant for negotiation and for their autonomy in relation to their environment; an alternative, more self-contained definition, for example, is an agent's ability to generate its own goals [51].

We can distinguish three strands of research in automated negotiation that each cluster around one of the three dimensions of autonomy (Figure 1):

# Self-sufficient: Game theoretical approaches and trading bots

The theory of games is a principal tool for studying negotiation and bargaining [72, 76]. Game theory's dominant concern is with fully rational players and what each should optimally do. This approach is therefore called *symmetrically prescriptive* [63]. The focus is on either equilibrium strategies or protocols that can guarantee a good outcome for both players through mechanism design [76]. Agents have a reduced scope for self-directedness in such settings, as they are relatively simple and need to conform to certain strategies (e.g. to bid truthfully in an auction). Similarly, real-world trading bots mostly employ simple rule-based functions which have been hard-coded in advance. Examples of this type are among the most advanced autonomous negotiators in terms of self-sufficiency, such as high frequency trading agents for financial and advertising exchanges, and sniping agents used in eBay that place bids at the last possible second [39]. While these approaches are able to function without human intervention and can be highly self-sufficient, they are constrained in terms of freedom to direct the process.

# Self-directed: Negotiation analytical approaches

Negotiation analysis *prescribes* how players should act given a *description* of how others will act. That is, this field is concerned with an *asymmetrical prescriptive/descriptive* view of autonomous negotiation [63]. Much research on what are often dubbed simply 'negotiation agents' (or 'heuristics' in game theory literature) falls into this category; e.g. all negotiation agents from the annual automated negotiation competition [5]. A key feature of this approach is the agent's ability to make judgment calls without intervention (i.e. to construct beliefs based on partial information and act in best response to these beliefs, typically over opponent types or strategies), while the agent's preferences are often considered externally given. This locates the negotiation analytical approach around the self-directed axis.

#### Interdependent: Negotiation support systems

Negotiation support systems are designed to assist and train people in negotiation. Some of these systems, such as the Inspire system [45], have been widely employed in reallife. However, while negotiation support systems enable interdependence by design, humans predominately supervise and make decisions on the appropriate outcome, which results in low self-sufficiency and self-directedness.

As can be gleaned from the fields indicated above, autonomous negotiation has garnered attention from different research directions and has managed to advance in key aspects of autonomous behavior. As a result, we now have negotiators that exist independently of their owner in the real world, delegated with a gamut of available strategies to freely choose among, and with the ability to engage in supportive interdependence; *just not all at the same time*.

The varied set of requirements for adequately autonomous negotiation may explain why it has proven difficult to extend the progress made in this field to truly representative negotiating agents. Of course we acknowledge that to a lesser degree, combined work on all dimensions has been performed (as depicted by the three-colored cube in Figure 1); we simply argue that the main automated negotiation research lines have developed in parallel to one of the three autonomy directions. Research-wise, it is unquestionably a sound strategy to first explore the autonomy axes in separation. As Figure 1 suggests, we can make substantive progress in autonomous negotiation by continuing to advance along *the autonomy diagonal*, which has inspired the focal points of the challenges we present in the next section (as summarized in Table 1).

**Domain knowledge and preference elicitation** (Section 3.1) Addressing *self-sufficiency & interdependence* 

Building blocks	Solutions roadmap	Applications
Preference elicitation on-the-fly	Value of information indicators, ro-	Privacy and IoT
	bust performance estimates	
Domain modeling	Separate user/agent domain mod-	Smart grids
	els, expert mappings	

#### **Long-term perspective** (Section 3.2) Addressing *self-sufficiency & self-directedness*

Building blocks	Solutions roadmap	Applications	
Repeated interactions	Temporally integrative negotia- tions, reputation metrics	Communities, smart homes, autonomous driving	
Non-stationary preferences	Cost-efficient tracking, context- dependent models, preference dynamics	B2B, entertainment booking	

#### **User trust and adoption** (Section 3.3) Addressing *self-directedness* & *interdependence*

Building blocks	Solutions roadmap	Applications
Acceptability and participation	Co-creation, adjustable autonomy, transfer of control	Conflict resolution, customer retainment
Transparent consequences	Transparency and openness, worst- case bounds, risk measures	Sharing economy, decentralized mar- ketplaces

**Table 1.** Overview of major challenges in autonomous negotiation and the main dimensions of autonomy to which they relate. Each challenge is subdivided in building blocks along with a solution roadmap and illustrative example applications.

# **3** Major Challenges

The various aspects of autonomy drive three major open challenges for autonomous negotiation, of which the overall theme can be summarized as *trusted and sustained representation*. We describe the challenges and their building blocks below, together with a number of explicit opportunities in each case (see Table 1 for an overview).

Note that many of these challenges intersect and cannot be entirely untangled; for example, adequate user preference extraction will not only increase the user model accuracy, but may also boost user trust. Therefore, just like autonomy itself, each challenge outlined here is *multi-dimensional*; i.e., each challenge pertains to at least two dimensions of autonomy, thereby providing the impetus to further advance along the autonomy diagonal.

#### 3.1 Domain Knowledge and Preference Elicitation

A *negotiation domain* typically admits contracts that consist of multiple issues (e.g. price, amount, quality of service). The specific structure of a domain together with the *user preferences* associated with its outcomes (prescribed by e.g. a utility function or outcome orderings [4]) forms a *negotiation scenario*.

Individual preferences over specific scenarios provide the opportunity for joint improvement and trade-offs [19]. The co-dependence between user and agent requires that they synchronize their negotiation scenario model, which can be enhanced by imparting the agent with accurate and timely user preferences about the negotiation process and co-constructing the real-world intricacies of the domain.

**Preference elicitation on-the-fly** In order to faithfully represent the user, an autonomous negotiator needs to engage with the user to make sure it constructs an accurate preference model (see e.g. [40]). However, users are often unwilling or unable to engage with a negotiation system, and hence prudence needs to be exercised when interacting with the user to avoid elicitation fatigue. This is especially important in domains where people are notably reluctant to engage with the system at length, for instance in privacy negotiations.

As a consequence, automated negotiators of the future are required to not only strike deals with limited available user information, but also to assess which additional information should be elicited from the user, while minimizing user bother [6]. This challenge is still as relevant (and for the most part still unaddressed) as when it was raised by Jennings et al. in [42]. However, as a way forward, we believe future research should particularly emphasize *preference elicitation on-the-fly* [7]: that is, active preference extraction *during* negotiation(s) (see Figure 2). Potential benefits include a significantly reduced initial preference elicitation phase (which can otherwise be a nuisance in many negotiation support systems) and the ability to select the most informative query to pose to the user at the most relevant time. For example, while negotiating, the system could dynamically decide to ask the user to rate specific negotiation outcomes, or to compare two of them.

To facilitate this, new performance-based metrics are required that can assess how supplementary preference information influences negotiation performance. Adaptive



**Fig. 2.** A representative agent has a high uncertainty about the utility  $u_x$  of a negotiation outcome  $x \in \Omega$  (in purple, prior to posing a query). Preference extraction through a query q (e.g. "is  $u_x > 0.5$ ?") can reduce this uncertainty, against certain user bother cost, by distinguishing between bad outcomes ( $u_x \mid no$ , in red) and good ones ( $u_x \mid yes$ , in blue).

utility elicitation models provide a good starting point for representing probabilistic utility-based preferences that allow for incremental updating over time (e.g. by using Bayesian reasoning), in the vein of work by Chajewska et al. [16]. To continuously balance the expected negotiation payoff with the potential benefit of performing additional elicitation, the viability of a negotiation query can for instance be measured in terms of the expected value of information [12] in order to assess the marginal utility of altering belief states and to decide if a query is worth posing.

Another challenge is for a negotiation strategy to decide on actions effectively in light of its imprecise information state. Techniques for decision making under uncertainty could assist in this and could thereby give rise to novel negotiation strategy concepts, for instance by incorporating the notion of expected *expected* utility [13] to express the expected negotiation payoff over all possible instantiations of the user model.

Note that the above discussion largely follows the standard assumptions of rational choice theory: i.e. that people's preferences can be accurately elicited. Most approaches adopt the perspective that the value of possible agreements is expressible by a (relatively) time- and context-invariant utility function over material outcomes – we revisit this in Section 3.2. Unfortunately, several idiosyncrasies of human psychology complicate these assumptions. While preference-elicitation methods can often extract coherent utility functions that capture people's rankings over possible agreements, people often have difficulty explicitly expressing their preferences. Further, a person's willingness to accept an agreement is also only partially determined by how they feel about the final

agreement; these feelings are also highly sensitive to contextual factors, such as how the deal is reached and how it is described. Research on human negotiation emphasizes that people attend to many factors besides the final outcome, as identified, for example, by Curhan's Subjective Value Inventory [20].

Research also illustrates that elicited utility functions are highly sensitive to subtle contextual factors. For example, *framing effects* emphasize that preferences between outcomes can reverse depending on whether they are seen as losses or gains with respect to some reference point [70]. In a negotiation, the reference point is often the perceived value that the other party receives, even though this knowledge does not change the individual's objective outcome. As a result, outcomes can be readily manipulated simply by changing the form and nature of information conveyed [31]. More broadly, valuations in a negotiation are shaped by emotion, including emotions that arise from the process, but also beliefs about what other parties feel (see, e.g., [10]). Given the highly context-sensitive nature of on-the-fly preference elicitation, such considerations will have to be taken into account in its design and implementation.

**Domain modeling** The quality of the negotiation outcome depends not only on the faithfulness of the preference model of an autonomous negotiator, but also on the accuracy of the domain model. The old 'garbage in, garbage out' truism applies here, as the quality of the offered solution depends so heavily on a correct domain description.

However, domain modeling, and certainly formal modeling, is an expertise that cannot be expected from an arbitrary user. Therefore, users require either expert guidance, or explicit domain modeling support. Modeling in close cooperation with a domain expert runs the risk of perpetuating people's uncertainty about the model, thereby limiting their ability to make necessary adjustments. When modeling support is provided by the system, the knowledge representation language used will be inherently simple as it has to be understood by arbitrary negotiators. This is especially important in domains where users can employ automated negotiation without any expertise, such as in the smart grid, which can result in the wrong evaluation of bids. Highly accurate models, on the other hand, also have their disadvantages: they can display complex non-linearities [41, 50], in which case even assessing the utility of a proposal can prove NP-hard [21].

This inspires the following open research question: what is the impact of simplifying the domain and preference models to facilitate layman understanding? An answer might come from using two models, as suggested in [37]: an accurate, but complex model that serves as a reference for the agent, and a more comprehensive one for interacting with the user. Proper clarification and explanation could then be elicited from a process of co-creation [62] or participatory design [65] between modeling experts and domain experts. Ideally, a reflection phase should be included during and after negotiations, in which the human (and perhaps eventually the agent) can provide feedback to allow for long-term co-evolution.

The above points also apply to the appropriateness and understandability of the *ne-gotiation protocol*, which governs the rules of the negotiation. A pre-negotiation phase provides the opportunity for the negotiation parties to engage in a debate about what protocol to employ and how to enforce the rules. To reduce the chance of parties exploiting loopholes in the rules, horizontal governance [68] approaches can be employed.

Such techniques are applied in border customs regulations, where the responsibility for fairness is carried by all participants. A corresponding challenge is to construct a best practice repository for negotiation techniques [42]. This has been tackled at least partially through recent efforts in creating a negotiation handbook for negotiation protocols [52].

Whatever approach is chosen, experts in formal modeling will be needed to instantiate a domain model that sufficiently captures all salient features. Those experts are pivotal to the negotiation agent business model and will be responsible for mapping user-understandable interests to the negotiation issues within complex domains. These are likely to become future jobs; i.e., real estate agents informing procurement agents of the future. Relevant research areas, and courses for training these experts, will be on collaborative and supportive modeling.

### 3.2 Long-term Perspective

Given the effort involved in domain modeling and preference elicitation, the opportunities for automated negotiation are even clearer in long-term scenarios where an agent frequently faces similar negotiation situations. Most research on negotiation agents, however, has focused on single encounters. The different challenges and opportunities for such long-term negotiations hinge on the volatility of both the opponent pool and the user's preferences.

**Repeated encounters** Given the efforts required to obtain an accurate user model, the benefits of autonomous agents become especially apparent when repeatedly dealing with similar situations, as is the case when negotiating multiple times with the same set of opponents. Indeed, there are many promising opportunities for applying negotiation in such repeated encounters. For example, in community energy exchange [2], agents can trade energy from storage and local renewable sources between neighboring homes and businesses to reduce peaks, carbon emissions and the load on the local network. These interactions would occur on a daily basis or even more frequently. Another example is the smart home, where different occupants have different needs and preferences and have to reach mutual agreements, e.g. about the trade-off between comfort and energy cost [59] and the use of IoT devices [57]. Other settings, in which the agent faces many different opponents, include self-driving vehicles, where vehicle-to-vehicle and vehicle-to-infrastructure negotiation will play an important role by, e.g., negotiating priority at intersections [71].

Negotiation opportunities for isolated encounters can be very limited, since often a resource (e.g. electricity or giving way) is needed without necessarily offering anything immediately in return (except possibly money or virtual currencies). In a single negotiation, the only truly interesting interactions revolve around multi-issue negotiation in which trade-offs can be made between the parties varying interests. However, explicitly considering the *temporal dimension* allows agents to receive or concede something now in return for conceding or receiving the same resource later. In other words, sequential, distributive negotiations can be turned into richer, multi-issue, integrative negotiations, with more scope to achieve win-win solutions (as presented in [53]; see Figure 3). This

is entirely analogous to how in single negotiations, package deals are more efficient than settling the issues independently due to the possibility of making fair trade-offs across issues [27]. Or likewise, how economic efficiency is enhanced when combinations of assets rather than individual items are considered [23]. By carefully bundling interdependent issues together, the exponential complexity of the resulting outcome space can in principle be mitigated (see e.g. [32]).



**Fig. 3.** By considering the Cartesian product of the outcome spaces, two sequential, distributive negotiations (one with 2 outcomes, in red; one with 3 outcomes, in blue) can be transformed to produce one integrative negotiation (with 6 outcomes, in purple).

Another significant challenge for long-term reciprocal encounters is that future needs are often uncertain, and so it is difficult to commit to giving up or requesting specific future resources. Possible solutions involve money or virtual currencies which can be redeemed at a later stage and can undergo temporal discounting if necessary, but they do not take advantage of the distributive nature of multi-issue negotiation. They also introduce additional challenges: using actual money requires an exchange rate with the resources involved (and other dimensions such as the scarcity at the time they are requested), while it may not be desirable to introduce money in certain settings; e.g. when they rely, to some degree, on unincentivized cooperation and altruistic behavior. Virtual currencies (including distributed ledger approaches) can be traded bilaterally in a "like for like" manner, addressing the exchange problem, but then other issues arise, e.g. how much of a currency each agent receives to begin with, and what happens if an agent runs out.

Another possible solution is to rely on altruism and using trust ratings and reputation metrics to provide the desired incentives (e.g. using favors and ledgers [53]). In such cases, 'altruism' can be a self-interested strategy if this is reciprocated at a later state,

possibly involving a different opponent. While reputation mechanisms are well-known to incentivize cooperation in the prisoner's dilemma, more research on this is needed in the context of (repeated) automated negotiation.

Unfortunately, negotiation methods that seek to establish a value-creating relationship by identifying efficient and fair (e.g., envy-free) agreements face, in addition to the above, a number of *psychological* challenges. People adopt a variety of interpretations as to what is fair and negotiations often involve disputes over which principle to apply [73]. For example, in the context of organ donation, the *equality principle* would allocate resources on the basis of ability, effort or merit, the *equality rule* would treat individuals the same, whereas the *principle of need* is achieved by allocating according to individuals medical condition, socio-economical status or other relevant needs. Even people's willingness to engage in negotiation is shaped by their views toward these principles [55]. Other complications involve moral constraints on certain exchanges. For example, it is considered morally repugnant to exchange money for bodily organs, so an agreement that combines material interests with sacred values may be seen as substantially worse than an independent evaluation of these elements would suggest [24].

Although these challenges might seem insurmountable, there are several ways to incorporate these biases into conventional computational methods. One approach is to incorporate psychological factors into the utility function, which can be done without violating the basic tenets of utility theory [28]. Some of the challenges with fairness can be addressed by making the process more transparent (Section 3.3). Another approach is to incorporate modest psychological extensions to rational methods. For example, framing effects can be handled through the use of prospect theory (e.g., [75]).

**Non-stationary preferences** While short-lived instantiations of representational agents may assume that there are some true and stationary preferences to be elicited from the user, in long-term negotiations, these very preferences may evolve over the course of weeks or months according to certain *preference dynamics*. For example, exposure to the view of wind generators can impact the preference profile over various renewable generation types [47]. This is related to the machine learning notion of *concept drift*, which expresses that statistical properties of a target variable (e.g. in a data stream) may change over time, possibly *abruptly* or *gradually*, which makes modeling challenging [69].

If an autonomous negotiator acts on elicited information for an extended period of time without accounting for existing drift in preferences, it may erroneously fulfill outdated design objectives. This leads to a plunge in user trust and adoption, or a de-facto shortened time of deployment. This is a typical example of opacity that can result from an excess of unchecked autonomy [56]. As a result, long-term negotiation requires increased co-dependence at the cost of throttled-down self-directedness; e.g., by repeated assessment of the preference representation quality, with intermittent elicitation actions whenever their anticipated benefits exceed their costs.

It is important to distinguish *complex preference models* from *dynamic preferences*, albeit the two provide complementary views on the same process. In representational negotiation, the target is to accurately model the users preferences, which may be

elicited at a cost, and which the user applies to judge the agent's performance. The user may perceive an exogenous evolution in her preferences (e.g., in risk tolerance or fairness attitudes), or update her preferences actively based on experience – she thus maintains dynamic preferences, that may themselves be *learned* over time. In contrast, the agent may employ a complex preference model, describing user preferences dependent on possibly uncertain user state variables. Changes in user preference may thus be ascribed to updates in the belief over the user's state, based on dynamics or observed information. The crucial deviation from previous approaches is an acknowledgment of the possibility that preferences may not settle, but remain in a state of flux. Tracking can mitigate the effects of evolving negotiation preferences in order to facilitate sustained representation [15].

Assuming non-stationary preferences reframes the challenge posed in Section 3.1 of preference elicitation to *cost-efficient tracking of non-stationary preferences* in long-term negotiation. Possible applications range from secretary tasks (e.g., 'book a restaurant/hotel/holiday') to representational business-to-business (B2B) negotiations [58]. Inspiration for tackling this challenge may come from the area of news recommender systems, which has embraced context-dependent models [1] and preference dynamics [48] in response to the inherent need to capture fast-paced preference evolution. Such models have promising merit for being transfered to negotiation strategies that balance the preciseness of preference representation with relevant and timely but costly elicitation, extending preliminary work in that area [7]. Reinforcement learning techniques could provide another possible route way by which to deal with this challenge by extending research into negotiation with non-stationary opponents [36, 54].

Beyond the passive modeling of dynamic user preferences, an opportunity for *managing user preferences* arises once a model of the user preference dynamics is available, e.g. through nudges and manipulation of cognitive biases [11, 30]. If preferences are learned, then the agent can choose to guide the user's experience (with intermediate results of negotiated outcomes) to promote certain preference profiles for which higher utility can be achieved in the long run.

### 3.3 User Trust and Adoption

While the agent depends on the user for knowledge and guidance (as described in Section 3.1), the user relies on a self-directed agent for a good outcome. To alleviate unwillingness to relinquish control and to guarantee user satisfaction with and adherence to the final outcome, the user needs to trust the system through co-participation, transparency, and proper representation.

**User participation** Lessons learned from collaborative human-robot teams indicate that it is important to be able to escalate to the meta-level (i.e. have humans participate) when necessary [43]. The need for escalating to a higher authority applies whenever a negotiator represents a group or a company (e.g., a union, or stakeholder organizations in general). In such cases, the negotiator can only make deals that fall within certain margins. Take, for example, a helpdesk operator with a telecom provider, authorized to offer new deals on a contract renewal. She has only limited freedom in terms of the

bounded range of possible deals she can sign off on; in fact, she does not even really possess the freedom to decide *whether* to negotiate. In case of doubt, the decision is escalated to a different authority level.

Similar to preference elicitation on-the-fly, user escalation should only occur through a minimum number of timely and pertinent questions (cf. [17], [46], [60]). As automated negotiators become more general and domain-independent [49], the need for a co-active design increases; i.e. one that requires the automated negotiator to be aware of the strengths and weaknesses of itself and that of the user, together with the ability to enhance the team model with domain-specific knowledge, preferred strategies, and interpretation of incoming bids.

The idea of collaborative control, or mixed-initiative control (see e.g. [29, 43]) might become essential to achieve the best outcome in complex, real-life negotiations. In this envisioned line of research, each negotiation party consists of at least one human and one negotiation agent. The agent should do the brunt of the negotiation work to find possible agreements with the other negotiation parties and which can presented to their human partners for feedback and new input, which necessitates an understanding of their behavior [33], attitudes [77], and preferred interaction method [61]. The research challenge is to determine when, how, and how often to switch the initiative from human to agent and vice versa.

**Transparent consequences** There exists an inherent tension between increased selfdirectedness and trust, which dampens the adoption of increasingly autonomous negotiators: on the one hand, an autonomous negotiatior's relevance is directly proportional to its ability to impact the user independently in meaningful ways (e.g. fiscal, wellbeing, reputation, and so on); but, in turn, the user's trust and willingness to relinquish control is conditional on understanding the agent's reasoning and consequences of its actions. The two can be reconciled by making the outcome space more *transparent* to the user, and by enabling the user to specify the permissible means in the form of *principles*. The challenge is that the negotiation agent's reasoning abilities may very well exceed the domain insights of a nonspecialist user, thus requiring a translation from stochastic performance models of self-directed expert reasoning into laymen terms that adequately convey expectations and risks.

Note that we suggest transparency as the key concept here, which subsumes Jennings' notion of predictability [42]. Predictability is essential towards the user to instill trust, but can be disastrous towards the opponent because of the potential for exploitability. Unpredictable behavior is in fact desirable as a negotiation tactic as a confusing and randomization device, as long as the consequences are transparently explained to the user.

The uncertainty inherent in negotiation can be captured in performance models and risk metrics, where the complexity should be scaled to the criticality of the consequences for the user. If the performance intervals are sub-critical, then simple guarantees on the range of possible outcomes may suffice (such as price bounds provided by Uber for individual rides), leaving it up to the user to build and judge the average performance model; otherwise, measures of risk are required, such as Conditional Value at Risk (CVar) [64].

# 4 Concluding Observations

Autonomous systems that are capable of negotiating on our behalf are among society's key technological challenges for the near future, and their uptake is important for many critical economical application areas. In this paper, we present a roadmap to arrive at representative and trusted negotiators that are endowed with a long-term perspective. By continuing along this trajectory, negotiation research can address perhaps the biggest challenge of all: a co-active approach that can propel swift adoption of computerized negotiation by simultaneously advancing the autonomy of a negotiation agent in all its aspects.

On the other hand, the sensitive nature of negotiation requires keeping a watchful eye for potentially adverse effects of increased computer autonomy. For instance, autonomous negotiators need to encapsulate information locally as part of their decentralized nature, which entails an inherent privacy risk. Privacy concerns can be important enough to restrict negotiation information flow to representative entities (as illustrated, for instance, by the leaked memo controversy of Brexit negotiators [74]). This also means that even when it is possible from a user's comfort level, it is not necessarily desired to extract the maximum user knowledge possible. This reiterates the need for improved measures for the value of information in negotiation; e.g. a metric that encapsulates the increased value of accurate preference estimates to reach win-win outcomes. For example, if a desire for maternity leave is inconsequential to the outcome (e.g. when it is not part of the contract), it does not need to be known by the agent.

Another consideration is when negotiation (in the sense of exchanging offers and counter offers) is the appropriate choice, and when other mechanisms can or should be used. Often alternatives are available that could be more efficient or simpler to use. These include auctions, which are especially fitting in cases of competing settings and resource allocation problems, and voting protocols, which are a good way of reaching agreements when a consensus is needed. These approaches are often used in settings with only one issue to decide on (in case of voting), or where the main issue is a monetary payment for a commodity (in case of auctions). Another, much simpler, alternative to negotiation is to offer an exhaustive menu of choices and allow the other player to pick one of these (as in the case of e.g. insurance policies). The advantage of negotiation is the ability to strategize over information revelation and to personalize deals by reaching differentiated agreements over illiquid, customized contracts. If these features of negotiation are not taken advantage of, then other approaches might be preferable.

On a societal level, negotiation has potential benefits and costs. In terms of benefits, negotiation allows for much more efficient allocation of resources than fixed pricing schemes as it optimizes value with respect to an individual's willingness-to-pay [38] rather than appealing to the "average customer". This helps to reach win-win outcomes and to get closer to what economists call first-degree price discrimination. Negotiation has fallen out of favor in wealthier societies because the time and anxiety associated with it, except for very costly transactions such as salary negotiations and home purchases, yet automated negotiation agents can mitigate these concerns and generate value for society as a whole. Widespread adopting of negotiation technology could benefit members of society that are reluctant or unable to negotiate effectively and could potentially address friction costs as well as inequities across society. For example, women are

especially averse to negotiating their salaries and this is a major contributor to gender pay inequality [9]. But negotiation technology could be abused in ways that outweigh or even undermine these potential benefits. For example, a benefit of fixed-pricing schemes is that they are transparent and applied uniformly, regardless of a persons gender, race, income, or negotiation skills. Indeed, Amazon's attempts at price discrimination have provoked lawsuits from consumer protection groups for exactly this concern [38]. More broadly, if negotiation agents are only available to the wealthy, they could exacerbate existing societal injustice. Regulations and ethical guidelines are needed to balance these benefits and costs.

In the end, the potency of autonomous negotiators is as much contingent on the acceptance by their users as by their counter-parties. Possible sources of resistance to adoption include established business models based on human inefficiencies (e.g., phone and media contracts) or anti-competitive practices (e.g., proprietary lock-in), which could become invalidated by autonomous (re-)negotiation. The most promising incubators of autonomous negotiators are ecosystems in which autonomous agents provide a unique source of societal value that is distributed over all stakeholders, as in the application of demand response for smart grids. Open platforms for value distribution have recently seen increased attention in flagship applications such as the cryptocurrency *bitcoin* and the decentralized world wide web *Blockstack* [3]. The digital API of these systems offers fertile grounds for a level playing field for competition and may soon provide a common interface for automated negotiators.

Finally, looking even further forward, it is worth noting that people negotiate differently through intermediaries than they would face-to-face. The literature on representation effects suggests that people may show less regard for fairness and ethical behavior when negotiating through a third (human) party [18]. Indeed, human lawyers are ethically permitted and, to some extent, expected to lie on behalf of their clients [31]. This raises the question as to whether agents should similarly lie on behalf of a user, e.g. by using argumentation and persuasion technology [25]. Analogous to recent research on ethical dilemmas in self-driving cars, people may claim that negotiation agents should be ethical, but sacrifice these ideals if it maximizes their profits. The natural dichotomy between recognizing the agent's autonomy and taking responsibility for its actions is best resolved by acknowledging user responsibility for the agent's design objectives (what should be achieved) and principles (how it should be achieved, as discussed in Section 3.3). This also illustrates an additional impetus for having humans understand the agent: feeling responsibility for the agent's actions implies an understanding what the agent is doing. Fortunately, some recent research on agent negotiators suggests that people may act more ethically when negotiating via computer agents [22], but far more research is needed to understand how *artificial representation effects* arise.

# Acknowledgments

This research has received funding through the ERA-Net Smart Grids Plus project Grid-Friends, with support from the European Union's Horizon 2020 research and innovation programme.

# Bibliography

- [1] Gediminas Adomavicius and Alexander Tuzhilin. Context-aware recommender systems. In *Recommender systems handbook*, pages 191–226. Springer, 2015.
- [2] Muddasser Alam, Enrico H. Gerding, Alex Rogers, and Sarvapali D Ramchurn. A scalable interdependent multi-issue negotiation protocol for energy exchange. In 24th International Joint Conference on AI (IJCAI), pages 1098–1104, 2015.
- [3] Muneeb Ali, Jude Nelson, Ryan Shea, and Michael J Freedman. Blockstack: A global naming and storage system secured by blockchains. In 2016 USENIX Annual Technical Conference, pages 181–194. USENIX Association, 2016.
- [4] Reyhan Aydoğan, Tim Baarslag, Koen V. Hindriks, Catholijn M. Jonker, and Pinar Yolum. Heuristics for using CP-nets in utility-based negotiation without knowing utilities. *Knowledge and Information Systems*, 45(2):357–388, November 2015.
- [5] Tim Baarslag, Reyhan Aydoğan, Koen V. Hindriks, Katsuhide Fuijita, Takayuki Ito, and Catholijn M. Jonker. The automated negotiating agents competition, 2010-2015. AI Magazine, 36(4):115–118, 12/2015 2015.
- [6] Tim Baarslag and Enrico H. Gerding. Optimal incremental preference elicitation during negotiation. In *Proceedings of the Twenty-fourth International Joint Conference on Artificial Intelligence*, IJCAI'15, pages 3–9. AAAI Press, 2015.
- [7] Tim Baarslag and Michael Kaisers. The value of information in automated negotiation: A decision model for eliciting user preferences. In *Proceedings of the* 16th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2017. IFAAMAS, 2017.
- [8] Tim Baarslag, Michael Kaisers, Enrico H. Gerding, Catholijn M. Jonker, and Jonathan Gratch. When will negotiation agents be able to represent us? the challenges and opportunities for autonomous negotiators. In *Proceedings of the Twenty-sixth International Joint Conference on Artificial Intelligence*, IJCAI'17, 2017.
- [9] Linda Babcock and Sara Laschever. *Women don't ask: Negotiation and the gender divide*. Princeton University Press, 2009.
- [10] Bruce Barry, Ingrid Smithey Fulmer, and Gerben A. Van Kleef. I laughed, I cried, I settled: The role of emotion in negotiation. *The handbook of negotiation and culture*, pages 71–94, 2004.
- [11] Max Bazerman and Margaret Neale. Nonrational escalation of commitment in negotiation. *European Management Journal*, 10(2):163–168, 1992.
- [12] Craig Boutilier. A POMDP formulation of preference elicitation problems. In *Eighteenth National Conference on Artificial Intelligence*, pages 239–246, Menlo Park, CA, USA, 2002. American Association for Artificial Intelligence.
- [13] Craig Boutilier. On the foundations of expected expected utility. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, IJCAI'03, pages 285–290, San Francisco, CA, USA, 2003. Morgan Kaufmann Publishers Inc.
- [14] Jeffrey M Bradshaw, Paul J Feltovich, Hyuckchul Jung, Shriniwas Kulkarni, William Taysom, and Andrzej Uszok. Dimensions of adjustable autonomy and

mixed-initiative interaction. In International Workshop on Computational Autonomy, pages 17–39. Springer, 2003.

- [15] Anais Cadilhac, Nicholas Asher, Farah Benamara, and Alex Lascarides. Grounding strategic conversation: Using negotiation dialogues to predict trades in a winlose game. In *Proceedings of Empirical Methods in Natural Language Processing* (*EMNLP*), pages 357–368, 2013.
- [16] Urszula Chajewska, Daphne Koller, and Ronald Parr. Making rational decisions using adaptive utility elicitation. In *In Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 363–369, 2000.
- [17] W. Chu and Z. Ghahramani. Preference learning with gaussian processes. In *ICML*, pages 137–144, 2005.
- [18] Dolly Chugh, Max H Bazerman, and Mahzarin R Banaji. Bounded ethicality as a psychological barrier to recognizing conflicts of interest. *Conflicts of interest: Challenges and solutions in business, law, medicine, and public policy*, pages 74– 95, 2005.
- [19] Robert M. Coehoorn and Nicholas R. Jennings. Learning on opponent's preferences to make effective multi-issue negotiation trade-offs. In *Proceedings of the* 6th International Conference on Electronic Commerce, ICEC '04, pages 59–68, New York, NY, USA, 2004. ACM.
- [20] Jared R Curhan, Hillary Anger Elfenbein, and Heng Xu. What do people value when they negotiate? Mapping the domain of subjective value in negotiation. *Journal of personality and social psychology*, 91(3):493, 2006.
- [21] Dave de Jonge, Carles Sierra, and Jordi González Sabaté. Negotiations over large agreement spaces. PhD thesis, 2015.
- [22] Celso M de Melo, Stacy Marsella, and Jonathan Gratch. Do as I say, not as I do: Challenges in delegating decisions to automated agents. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 949–956. IFAAMAS, 2016.
- [23] Sven de Vries and Rakesh V. Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, 15(3):284–309, 2003.
- [24] Morteza Dehghani, Scott Atran, Rumen Iliev, Sonya Sachdeva, Douglas Medin, and Jeremy Ginges. Sacred values and conflict over Iran's nuclear program. Judgment and Decision Making, 5(7):540, 2010.
- [25] Yannis Dimopoulos and Pavlos Moraitis. Advances in argumentation based negotiation. Negotiation and Argumentation in Multi-agent Systems: Fundamentals, Theories, Systems and Applications, pages 82–125, 2014.
- [26] Theodore Eisenberg and Charlotte Lanvers. What is the settlement rate and why should we care? *Journal of Empirical Legal Studies*, 6(1):111–146, 2009.
- [27] Shaheen S. Fatima, Michael J. Wooldridge, and Nicholas R. Jennings. On efficient procedures for multi-issue negotiation. *Agent-Mediated Electronic Commerce*, 4452:31–45 249, 2007.
- [28] Ernst Fehr and Klaus M Schmidt. The economics of fairness, reciprocity and altruism–experimental evidence and new theories. *Handbook of the economics of giving, altruism and reciprocity*, 1:615–691, 2006.
- [29] Terrence Fong, Charles Thorpe, and Charles Baur. Collaborative control: A robotcentric model for vehicle teleoperation, volume 1. Carnegie Mellon University, The Robotics Institute, 2001.

- [30] Michele J Gelfand and Jeanne M Brett. *The handbook of negotiation and culture*. Stanford University Press, 2004.
- [31] Jonathan Gratch, Zahra Nazari, and Emmanuel Johnson. The misrepresentation game: How to win at negotiation while seeming like a nice guy. In *Proceedings of* the 2016 International Conference on Autonomous Agents & Multiagent Systems, pages 728–737. IFAAMAS, 2016.
- [32] Rafik Hadfi and Takayuki Ito. Modeling complex nonlinear utility spaces using utility hyper-graphs. In *Modeling Decisions for Artificial Intelligence*, pages 14– 25. Springer, 2014.
- [33] Galit Haim, Ya'akov (Kobi) Gal, Michele Gelfand, and Sarit Kraus. A cultural sensitive agent for human-computer negotiation. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '12, pages 451–458, Richland, SC, 2012. IFAAMAS.
- [34] Richard L Hall. Measuring legislative influence. Legislative Studies Quarterly, pages 205–231, 1992.
- [35] Karen L. Harris. Content analysis in negotiation research: A review and guide. Behavior Research Methods, Instruments, & Computers, 28(3):458–467, 1996.
- [36] Pablo Hernandez-Leal, Yusen Zhan, Matthew E. Taylor, L. Enrique Sucar, and Enrique Munoz de Cote. An exploration strategy for non-stationary opponents. *Autonomous Agents and Multi-Agent Systems*, 31(5):971–1002, 2017.
- [37] Koen Hindriks, Catholijn Jonker, and Dmytro Tykhonov. Avoiding approximation errors in multi-issue negotiation with issue dependencies. In Proc. of The 1st International Workshop on Agent-based Complex Automated Negotiations (ACAN 2008), pages 1347–1352, 2008.
- [38] Oliver Hinz, II-Horn Hann, and Martin Spann. Price discrimination in ecommerce? An examination of dynamic pricing in name-your-own price markets. *MIS Quarterly*, 35(1):81–98, 2011.
- [39] Wenyan Hu and Alvaro Bolivar. Online auctions efficiency: a survey of eBay auctions. In *WWW*, pages 925–934, 2008.
- [40] Anthony Hunter. Modelling the persuadee in asymmetric argumentation dialogues for persuasion. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, pages 3055–3061. AAAI Press, 2015.
- [41] Takayuki Ito, Mark Klein, and Hiromitsu Hattori. A multi-issue negotiation protocol among agents with nonlinear utility functions. *Multiagent and Grid Systems*, 4(1):67–83, Jan 2008.
- [42] Nicholas R. Jennings, Peyman Faratin, Alessio R. Lomuscio, Simon Parsons, Michael J. Wooldridge, and Carles Sierra. Automated negotiation: Prospects, methods and challenges. *Group Decision and Negotiation*, 10(2):199–215, 2001.
- [43] Matthew Johnson, Jeffrey M. Bradshaw, Paul J. Feltovich, Catholijn Jonker, Birna van Riemsdijk, and Maarten Sierhuis. Autonomy and interdependence in humanagent-robot teams. *IEEE Intelligent Systems*, 27(2):43–51, March 2012.
- [44] Matthew Johnson, Jeffrey M. Bradshaw, Paul J. Feltovich, Catholijn M. Jonker, Birna van Riemsdijk, and Maarten Sierhuis. *The Fundamental Principle of Coactive Design: Interdependence Must Shape Autonomy*, pages 172–191. Springer Berlin Heidelberg, 2011.

- [45] Gregory E. Kersten and Gordon Lo. Negotiation support systems and software agents in e-business negotiations. In *The First International Conference on Electronic Business, Hong Kong, December*, pages 19–21, 2001.
- [46] D. C. Kingsley. Preference uncertainty, preference refinement and paired comparison choice experiments. University of Colorado, 2006.
- [47] Jacob Ladenburg. Dynamic properties of the preferences for renewable energy sources a wind power experience-based approach. *Energy*, 76:542 551, 2014.
- [48] Lei Li, Li Zheng, Fan Yang, and Tao Li. Modeling and broadening temporal user interest in personalized news recommendation. *Expert Systems with Applications*, 41(7):3168–3177, 2014.
- [49] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen V. Hindriks, and Catholijn M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014.
- [50] Miguel A. Lopez-Carmona, Ivan Marsa-Maestre, Mark Klein, and Takayuki Ito. Addressing stability issues in mediated complex contract negotiations for constraint-based, non-monotonic utility spaces. *Autonomous Agents and Multi-Agent Systems*, 24(3):485–535, 2012.
- [51] Michael Luck, Mark D'Inverno, and Steve Munroe. Autonomy: Variable and Generative, pages 11–28. Springer US, Boston, MA, 2003.
- [52] Ivan Marsa-Maestre, Mark Klein, Catholijn M. Jonker, and Reyhan Aydoğan. From problems to protocols: Towards a negotiation handbook. *Decision Support Systems*, 60(0):39–54, 2014.
- [53] Johnathan Mell, Gale Lucas, and Jonathan Gratch. An effective conversation tactic for creating value over repeated negotiations. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 1567–1576, Richland, SC, 2015. IFAAMAS.
- [54] Vidya Narayanan and Nicholas Jennings. Learning to negotiate optimally in nonstationary environments. *Cooperative information agents X*, pages 288–300, 2006.
- [55] James Neuberger and Peter A Ubel. Finding a place for public preferences in liver allocation decisions. *Transplantation*, 70(10):1411–1413, 2000.
- [56] Donald A Norman. The 'problem' with automation: inappropriate feedback and interaction, not 'over-automation'. *Philosophical Transactions of the Royal Soci*ety of London B: Biological Sciences, 327(1241):585–593, 1990.
- [57] Charith Perera, Susan Y. L. Wakenshaw, Tim Baarslag, Hamed Haddadi, Arosha K. Bandara, Richard Mortier, Andy Crabtree, Irene C. L. Ng, Derek McAuley, and Jon Crowcroft. Valorising the IoT databox: creating value for everyone. *Transactions on Emerging Telecommunications Technologies*, 28(1):1– 17, 2017.
- [58] Serban Radu. A negotiation framework with strategies based on agent preferences. In 2017 21st International Conference on Control Systems and Computer Science (CSCS), pages 529–535, May 2017.
- [59] Sarvapali D Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R Jennings. Putting the 'smarts' into the smart grid: a grand challenge for artificial intelligence. *Communications of the ACM*, 55(4):86–97, 2012.
- [60] Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.

- [61] Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, and Sarit Kraus. Negochat: A chat-based negotiation agent. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 525–532, Richland, SC, 2014. IFAAMAS.
- [62] Elizabeth B-N Sanders and Pieter Jan Stappers. Co-creation and the new land-scapes of design. *Co-design*, 4(1):5–18, 2008.
- [63] James K Sebenius. Negotiation analysis: A characterization and review. Management Science, 38(1):18–38, 1992.
- [64] M. Shafie-khah, D.Z. Fitiwi, J.P.S. Catalão, E. Heydarian-Forushani, and M.E.H. Golshan. Simultaneous participation of demand response aggregators in ancillary services and demand response exchange markets. In *Transmission and Distribution Conference and Exposition (T&D), 2016 IEEE/PES*, pages 1–5. IEEE, May 2016.
- [65] Jesper Simonsen and Toni Robertson. *Routledge international handbook of participatory design.* Routledge, 2012.
- [66] Reid G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, 29(12):1104– 1113, Dec 1980.
- [67] Katia P. Sycara. Arguments of persuasion in labour mediation. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, volume 1, pages 294–296, San Francisco, CA, USA, 1985.
- [68] Jacob Torfing, B. Guy Peters, Jon Pierre, and Eva Sørensen. Interactive governance: Advancing the paradigm. *Public Administration*, 91(4):1071–1082, 2013.
- [69] Alexey Tsymbal. The problem of concept drift: definitions and related work. Technical report, Trinity College Dublin, Department of Computer Science., 2004.
- [70] Amos Tversky and Daniel Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics*, 106(4):1039– 1061, 1991.
- [71] F. Visintainer, L. Altomare, A. Toffetti, A. Kovacs, and A. Amditis. Towards manoeuver negotiation: Autonet2030 project from a car maker perspective. *Transportation Research Procedia*, 14:2237–2244, 2016.
- [72] John Von Neumann and Oskar Morgenstern. *Theory of games and economic be-havior*. Princeton University Press, Commemorative Ed edition (1 May 2007), 1953.
- [73] Nancy A Welsh. Perceptions of fairness in negotiation. Marq. L. Rev., 87:753, 2003.
- [74] Ben Worthy. Brexit and open government in the UK. *Open Government Partnership*, 2016.
- [75] Rong Yang, Christopher Kiekintveld, Fernando Ordonez, Milind Tambe, and Richard John. Improving resource allocation strategies against human adversaries in security games: An extended study. *Artificial Intelligence*, 195:440–469, 2013.
- [76] H Peyton Young. Negotiation analysis. University of Michigan Press, 1991.
- [77] Jieyu Zhan, Xudong Luo, Cong Feng, and Minghua He. A multi-demand negotiation model based on fuzzy rules elicited via psychological experiments. *Applied Soft Computing*, 2017.