

# When Exceptions are the Norm

Exploring the Role of Consent in HRI

VASANTH SARATHY, Tufts University, USA

THOMAS ARNOLD, Tufts University, USA

MATTHIAS SCHEUTZ, Tufts University, USA

HRI researchers have made major strides in developing robotic architectures that are capable of reading a limited set of social cues and producing behaviors that enhance their likeability and feeling of comfort amongst humans. However, the cues in these models are fairly direct and the interactions largely dyadic. To capture the normative qualities of interaction more robustly, we propose “consent” as a distinct, critical area for HRI research. Convening important insights in existing HRI work around topics like touch, proxemics, gaze, and moral norms, the notion of consent reveals key expectations that can shape how a robot acts in social space. Consent need not be limited to just an explicit permission given in ethically-charged or normatively-risky scenarios. Instead, it is a richer notion, one that covers even implicit acquiescence in scenarios that otherwise seem normatively neutral. By sorting various kinds of consent through social and legal doctrine, we delineate empirical and technical questions to meet consent challenges faced in major application domains and robotic roles. Attention to consent could show, for example, how extraordinary, norm-violating actions can be justified by agents and accepted by those around them. We argue that operationalizing ideas from legal scholarship can better guide how robotic systems might cultivate and sustain proper forms of consent.

CCS Concepts: • **Human-centered computing** → **Interaction design**.

Additional Key Words and Phrases: implicit and explicit consent, normative HRI, architectural requirements

## ACM Reference Format:

Vasanth Sarathy, Thomas Arnold, and Matthias Scheutz. 2010. When Exceptions are the Norm: Exploring the Role of Consent in HRI. *ACM Comput. Entertain.* 9, 4, Article 39 (March 2010), 22 pages. <https://doi.org/0000001.0000001>

## 1 INTRODUCTION AND MOTIVATION

Social robots are designed to engage people in an interpersonal manner, often as partners, in order to achieve positive outcomes in domains such as education, therapy, and healthcare. Their sociality also aids task-related goals in areas such as coordinated teamwork for manufacturing, search and rescue, domestic chores, and more [Breazeal et al. 2016]. HRI research has made tremendous strides in exploring and discovering the social competencies that allow for an “efficient, enjoyable, natural and meaningful relationship” [Breazeal et al. 2016]. Understanding what sorts of conduct and behavioral choices enhance or detract from positive social interactions is crucial.

---

Authors’ addresses: Vasanth Sarathy, Tufts University, 200 Boston Ave. Medford, MA, 02155, USA, [vasanth.sarathy@tufts.edu](mailto:vasanth.sarathy@tufts.edu); Thomas Arnold, Tufts University, 200 Boston Ave. Medford, MA, 02155, USA, [thomasarnold@alumni.stanford.edu](mailto:thomasarnold@alumni.stanford.edu); Matthias Scheutz, Tufts University, 200 Boston Ave. Medford, MA, 02155, USA, [matthias.scheutz@tufts.edu](mailto:matthias.scheutz@tufts.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](mailto:permissions@acm.org).

© 2010 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

1

53 In pursuit of this understanding, HRI research has organized itself along different lines. There are topics represented as  
54 types of interactions (e.g., touch based, long-term), and there are HRI subfields mapped as general aspects of interactions:  
55 proxemics [Mead and Matarić 2017; Mumm and Mutlu 2011]; approach [Walters et al. 2007]; eye-gaze and joint attention  
56 [Admoni and Scassellati 2017]; interruption [Thomaz et al. 2016]; touch [Arnold and Scheutz 2018a; Van Erp and  
57 Toet 2013]; level of politeness [Banerjee and Chernova 2017; Briggs and Scheutz 2016; Palinko et al. 2018]. Technical  
58 approaches to equipping robotic systems to interact along such lines has generally involved computationalizing an  
59 existing social or anthropological theory of interaction applied to the aspect being studied. HRI’s empirical research  
60 helps evaluate how robots so-equipped interact in experimental settings.  
61

62  
63 Even as the field has pushed itself to more robust kinds of interactions and more sophisticated systems of planning,  
64 the prevalent notion of “social interaction” still focuses largely on dyadic exchanges. These one-on-one engagements  
65 usually rely on fairly direct cues and straightforward measurements of user preferences. For example, in approach  
66 research, it has been found that people generally prefer robots to approach them from the front and do not prefer robots  
67 to approach them from the back [Walters et al. 2007]. But there is a major “elephant in the room” in such findings, and  
68 in this paper we propose that it deserves distinct recognition as a way to organize HRI research: the notion of *consent*.  
69 Our argument is that HRI’s progress in investigating interactive dynamics will not be generalizable enough without  
70 tackling the issue of consent head-on.  
71

72  
73 Consider “approach” again. There are instances when an approach from behind is acceptable and maybe even  
74 preferred, for example if a robot was pushing a person out of harm’s way from an oncoming car. Wouldn’t the human  
75 have consented to this act if asked and provided time? Should the robot presume consent? At first glance, it seems  
76 that that sort of reasoning is highly circumstantial and potentially too ungainly to be a cogent research problem. Like  
77 other subtleties of interaction, it could be chalked up as “context” and tackled at a future date when technology has  
78 evolved further. We argue here instead that the tools already exist to directly explore the question further and that  
79 understanding consent is central to HRI’s future research.  
80

81 The word “consent” typically brings to mind either notions of legality around highly deviant sexual behavior or  
82 around the appropriateness of information given to a patient about a medical procedure. However, consent’s rich  
83 nuances and intricacies permeate our daily interactions in even normatively neutral environments. For example, when  
84 should a robot waiter take your plate at a restaurant? Consent can not only influence a robot’s behavioral choice, but  
85 also the normative valence of a situation. Consent can turn a normatively neutral situation into a charged one, and vice  
86 versa. While consent is one factor in the larger normative fabric of society, it is its normative power and pervasiveness  
87 that provides the motivation for HRI researchers to take it seriously.  
88

89  
90 More generally, social interactions among humans in societies are based on social and moral norms [Arnold and  
91 Scheutz 2017]. These norms are deeply ingrained in us and we expect others to abide by them. Now, failing to abide by  
92 these norms causes social reactions in humans, from blaming and reprimanding to potentially harsher legal consequences  
93 [Scheutz 2017]. Given that we have a propensity to anthropomorphize robots, especially human-like ones, we are likely to  
94 extend these expectations to robots and perceive them as moral agents as well [Malle and Scheutz 2015]. Consent is a  
95 key component of this normative setting, continually influencing our expectations of one another. So, it may be natural  
96 to expect that ethical competence in robots cannot be achieved without a careful consideration of consent.  
97

98  
99 Consent is often thought of as a permission (either general or specific) granted to someone to make something  
100 happen or perform certain behavior. To consent often means giving this permission, either implicitly or explicitly.  
101 While this intuitive definition provides a good starting point, it does not quite capture the intricacies and nuances of  
102 consent. How is the permission to be manifested? What does it mean to give this permission? What exactly am I giving  
103

104 Manuscript submitted to ACM

105 permission to? To whom am I giving permission and for how long? Should I be giving permission? What if I would  
106 have given permission had I known it was being sought? Would someone else in my shoes have given permission?

107 As we explicate below, the dimensions of consent can inform and flesh out existing HRI topics and findings. Returning  
108 to the subfields listed above, work on touch has already broached questions of trust, social cues, and the delicate terms  
109 on which robots should presume to touch a person [Chen et al. 2014; Van Erp and Toet 2013]. Even simple forms  
110 of robot-initiated touch evidence gender effects and workplace norms, and are evaluated accordingly [Arnold and  
111 Scheutz 2018b; Wullenkord et al. 2016]. Proxemics has broached many consent-related themes, including how people  
112 compensate for robots intruding into their personal space [Sardar et al. 2012]. Spatial preferences, for example how a  
113 robot should approach someone, have received sustained attention [Koay et al. 2014; Syrdal et al. 2007], including how  
114 people adjust those preferences to aid in a robot’s understanding of social signals [Mead and Mataric 2015]. The gaze of  
115 a robot can exert similar force on an interactant’s social “space” and shape how a robot’s role is perceived [Fiore et al.  
116 2013; Stanton and Stevens 2017].

117 Sex robots have directly raised various concerns about consent, including the threat of being “ever-consenting”  
118 models of submission to abusive sexuality [Gutiú 2012; Scheutz and Arnold 2016, 2017]. The risk of neglecting consent  
119 as a component of interaction is to risk robots being submissive by default, allowing them to be held up as ideals against  
120 people who refuse to consent to abuse. Indeed, some wonder what consent is for such an intimate interaction [Frank  
121 and Nyholm 2017]. But it is not just a headline-grabbing issue like sex robots that makes consent salient. The socially  
122 assistive role for robots in eldercare residences, especially when those served suffer from dementia, call up numerous  
123 questions about “informed consent” and privacy among all stakeholders involved [Ienca et al. 2016; Piatt et al. 2017].  
124 Consent may thus bind HRI even more strongly to issues for public deliberation around social institutions and their  
125 purposes.

126 In this paper, we introduce different kinds of consent with the help of legal theory. We then propose some concrete  
127 suggestions for how consent can be broken down and understood for settings of human-robot interaction, as well as  
128 how it can be used to drive both empirical and technical research forward. More specifically, in Section 2 we take a  
129 deep dive into five different forms of consent defined in common law, exploring, with examples, how they influence  
130 social behavior and normative expectation. In Section 3 we resurface and briefly discuss how consent issues can arise  
131 in robotic application scenarios that are relatively mundane, and not “normatively-weighty.” In Section 4, we discuss  
132 the implications for the HRI community and put forth four research directions – involving research questions and  
133 architectural considerations – that take into account this nuanced understanding of consent. In Section 5, we provide  
134 some practical recommendations for how HRI researchers can make their results more consent-aware in the near-term.

## 143 2 LEGAL LANDSCAPE OF CONSENT UNDER THE LAW OF INTENTIONAL TORTS

### 144 2.1 Conventions, Norms and Laws

145 The normative landscape of a society can be classified into conventions, norms, and laws. These govern individuals’  
146 behaviors, based on the degree of their normative strength, and the level or type of sanctioning for their violations. For  
147 example, a behavior or conduct that is a convention suggests a stable pattern of behavior [Haynes et al. 2017], a societal  
148 regularity like people choosing to walk on the right, even when indoors. There may or may not be consequences for  
149 violating a convention. Even if they are violated, the wrongdoer is typically presented with a mild push to conform or a  
150 small dose of scoffing. Certain types of wrongful conduct can rise to the level of a norm violation. We can think of norms  
151 as conventions with an added level of sanctioning, in a stricter sense than the scoffing associated with conventions.  
152  
153  
154  
155  
156

157 Sanctions for norm violations can be penalizing and are intended to change the behavior of the wrongdoer. For example,  
158 walking into a Hindu temple with shoes on is a serious norm violation that can incur social sanctioning and potential  
159 banishment from the corresponding social group. Behaviors achieve their highest normative level when they are codified  
160 into law. Laws deter wrongful behaviors that society has deemed to be harmful to the basic functioning of the society.  
161 Without laws, society cannot run. Because laws are codified and interpreted repeatedly over the course of centuries, the  
162 legal landscape has attained a degree of maturity and sophistication that can serve as a guide to situate less official  
163 normative elements like conventions and norms. That is, exploring and dissecting the legal landscape can provide us  
164 with some ammunition when we tackle less normatively charged situations, ones that are more common in everyday  
165 life. Below, we explore specifically the notion of *consent* from a legal standpoint, but note that the framing applies to  
166 conventions and norms as well.

169 In particular, we will primarily focus on the legal landscape of torts. In addition to tort law, the doctrine of consent  
170 heavily features in criminal law and contract law. In all three legal systems, there can be legal action for wrongful  
171 conduct, and consent can play a role in modulating the repercussions of the conduct. Contract law is concerned with  
172 cases where individuals can negotiate and craft their own rights and responsibilities towards each other. Most of our  
173 social interactions, however, are largely non-contractual, and tort law holds individuals legally accountable for the  
174 consequences of their actions in these cases. Tort law developed through judicial decisions in common-law courts, and  
175 later was codified into statutes and laws. Like tort law, criminal law is codified in statutes and laws and is aimed at  
176 addressing wrongful conduct. Still, criminal law is different in at least a couple of regards. First, tort law views certain  
177 acts as wrong in the private sense, one for which a defendant will owe damages to the victim. Criminal acts, on the  
178 other hand, are deemed to have wronged not just the victim, but the society as a whole, and the wrongdoer is deemed  
179 to be deserving of punishment. Second, while crimes can also be torts, criminal liability requires the defendant to  
180 have a more culpable state of mind than does tort liability.<sup>1</sup> Here, we are interested in exploring how consent can (or  
181 cannot) mitigate behavior that is deemed wrongful in the broadest sense – everything from the slightest deviations  
182 from customs to the most egregious violations of our moral code. To do so, we think that the law of intentional torts  
183 allows us to best begin this exploration, as it covers wrongful conduct broadly without also having to tackle issues of  
184 contractual meeting-of-the-minds, mens rea, and punishment.

189 As noted earlier, torts are a legal wrong committed upon person or property and, unlike contracts, require no prior  
190 agreement between parties. Every member of our society is obligated to pay damages if their actions harm others. One  
191 variety of torts, known as intentional torts, is particularly relevant and covers those situations in which intentional  
192 actions cause harm. There are several types of intentional torts, including harmful contact (battery), apprehension of  
193 harmful contact (assault), unlawful possession, use and alteration of personal property (trespass to chattels, conversion),  
194 restricting personal movement of others (false imprisonment), unlawful entry into the land of another (trespass to  
195 land), and intentional infliction of emotional distress. A crucial doctrine applicable in all types of intentional torts is the  
196 notion of consent, which can alter the morality of another's conduct [Hurd 1996] and potentially permit otherwise  
197 wrongful conduct. The doctrine is deeply rooted in history and stems from the Roman maxim *volenti non fit injuria*,  
198 latin for “to a willing person no injury is done.”

202 Consent can convert what would be “trespass into a dinner party, a battery into a boxing match, a theft into a gift,  
203 and rape into consensual sex” [Hurd 1996]. Consent can not only convert a wrongful behavior into right, but also can  
204 grant another the right to do wrong. In this latter sense, it is less about flipping the normative valence of an act to make  
205

206 <sup>1</sup>There are other differences between criminal law and tort law including who can bring a case - individuals or government, where can a case be brought -  
207 civil or criminal, and what are the possible penalties.

209 it right, and more about granting permission to perform what is ordinarily a morally wrong act. In both cases, consent  
 210 has a tremendous amount of normative power and is an important concept that is central to every interaction. There  
 211 are several forms the legal notion of consent can take and we explore each below.  
 212  
 213

## 214 2.2 Actual Consent

215 Actual consent is the “gold standard of consent” and is a person’s subjective willingness to permit tortious (wrongful)  
 216 conduct [Simons 2017]. It is generally considered to be the case that consent is limited to the person’s mental state  
 217 and does not also include the communicative act or signal. That a person demonstrates or evidences their consent via  
 218 language, gesture and other means is separate from their truest “consent,” which is their subjective acquiescence.  
 219  
 220

221 *A holds out her glass and asks B to take it and bring her some water. B takes the glass from A.* (1)  
 222  
 223

224 Here, A has consented to what would otherwise be trespass to chattels, or wrongful taking of property. Now, B’s  
 225 actions would need to cause some damages for them to be tortious, but we will set aside other such legal issues (damages,  
 226 causation, intent etc.). Instead, we will focus on conduct that could be wrongful, and explore how consent impacts it.  
 227  
 228

229 We can distinguish two forms of actual consent: *express* and *inferred*. Express consent refers to situations when the  
 230 consent is demonstrated explicitly either orally or in written form (i.e., linguistically). Inferred consent is when the  
 231 subjective willingness is inferred from the person’s conduct. (1) was an example of inferred consent, while (2) below is  
 232 express.  
 233

234 *A says to B “You may take my plate.” B takes A’s plate.* (2)  
 235  
 236

237 The statement “you may take my plate” by A literally grants permission to B to take away the plate. A slight variant  
 238 of this if A says “take my plate” leaving out the permission granting part. Note that while the statement “take my plate”  
 239 does not literally state “you are permitted to take my plate”, this permission can be directly inferred from the imperative  
 240 form (as it does not make sense to instruct someone to take the plate without allowing them to do so – this is the  
 241 crux of the old deontic principle ascribed to Immanuel Kant that “ought implies can”) Utterances where the utterance  
 242 semantics does not directly express the intended semantics are called “indirect speech acts” (ISAs) and are widely used  
 243 in social contexts for proper conduct, often to meet politeness norms (e.g., [Briggs and Scheutz 2013]). We will discuss  
 244 more about ISA’s below.  
 245  
 246

247 Returning to express consent, one distinction that is not often drawn in the legal literature is that between express  
 248 consent that is direct versus that which is indirect.  
 249

250 *A says to B, a bellhop, “Can you pick up my bag, please? It’s quite heavy.” B picks up A’s bag.* (3)  
 251  
 252

253 Here, A was clearly consenting to B’s actions, but it was certainly not direct in the sense of the semantics of the  
 254 utterance. B needed to determine that A’s utterance was in fact a command or at least a request for B to perform  
 255 a specific action. ISAs of this sort can be highly conventionalized, and thus it is possible B interpreted the specific  
 256 linguistic form just as he would the direct utterance “pick up my bag.” Crucially, because consent is a mental state, we  
 257 must consider the possibility that direct and indirect expressions are possible.  
 258  
 259  
 260

261 The degree to which a dialogue can be indirect is exemplified in (4), below.

262 *A arrives at an expensive restaurant and is greeted by B, the Maitre D'. B welcomes A and then walks over* (4)  
 263 *close to A and says "may I?." A extends his arms. B takes off A's coat.*

265

266 Here, A actually consents to B's contact and taking his property. Not only is consent signal indirect, even B's utterance  
 267 soliciting the consent is indirect. In fact, B's utterance does not even reference the particular act for which B is seeking  
 268 consent. Yet somehow, A has deciphered what act B wants to perform and consented to this act. The story would have  
 269 been quite different if B, after asking A's permission then proceeds to do a pat-down. This would be acceptable in  
 270 airport security, but not in a restaurant.  
 271  
 272  
 273

274

### 2.3 Apparent Consent

275

276 A second, and lesser, form of consent is one in which the person does not actually consent, but the actor reasonably  
 277 believes that the person actually consents. In these cases, the person has not formed the subjective willingness to permit  
 278 the conduct. Tort law requires not just that there be no actual consent, but also that there be no *reasonable belief* that  
 279 the person consents. Unlike actual consent, which transforms or negates the wrongfulness of the act, apparent consent  
 280 serves to grant permission to an act that still remains morally wrong.  
 281

282 *A (never been to a doctor) visits a doctor. B, the nurse, touches A to take her blood pressure. B does not reasonably* (5)  
 283 *know that A is averse to physical touching.*

285

286 *A gets up from her table at an expensive restaurant to go to the bathroom. She places her crumpled napkin on*  
 287 *the table by her plate. B, a waiter, comes over and folds her napkin. A has previously never been to expensive* (6)  
 288 *sit-down restaurants before and find this a bit strange that B touched her napkin.*

290

291 In (5) (a variant of [189 1891]) and (6) B cannot rely on actual consent, because A was not actually consenting, but B  
 292 might be able to use apparent consent (or possibly even presumed consent, discussed below) to preclude liability for  
 293 battery in (5) and normative sanctioning in (6). A may or may not have liked her napkin to have been folded in (6) but  
 294 whether she objected to it or even just did not consider it, we can argue that she did not actually consent to it on  
 295 either occasion.  
 296

297 The question of whether the actor can "reasonably believe" that the person consents can be based on evidence  
 298 obtained from the person's conduct, others' conduct, or customary norms at play.

299 *At the end of a successful business meeting, B reaches out and shakes A's hands before A can react. Unbeknown* (7)  
 300 *to B, A objects to touching and shaking hands when greeting people.*

302

303 *A and B are moving about in a busy restaurant. B bumps into A as he was walking to the bathroom. A objects* (8)  
 304 *to touching.*

306

307 *A is standing in front of B in line in a coffee shop. Believing A to be his wife based on her demeanor, pocket*  
 308 *book she was carrying and her outfit, B gives her a surprise hug from behind. A, however, turned out not to be* (9)  
 309 *B's wife and objected to touching.*

311

312

Manuscript submitted to ACM

313 In (5), (7) and (8) the evidence for B’s arguably reasonable belief that A was actually consenting may come from  
314 what is customary in a doctor’s office or in a business meeting or in a crowded space. However, if in either case B knew  
315 that A objects to being touched, then B cannot rely on apparent consent to preclude liability. In (9), the evidence for B’s  
316 reasonable belief came not from custom but directly from A’s conduct and appearance.  
317

## 319 2.4 Presumed Consent

320 We can next consider the case when the actor knows that the person did not actually consent, but reasonably believes  
321 that the person *would* actually consent, if asked. This is a counterfactual form of consent. This is different from apparent  
322 consent, where the actor reasonably believes that the person actually consents. This form of consent is typically, but  
323 not always, encountered in emergency situations when it is infeasible to obtain consent.  
324

325  
326 *A, a child, is chasing after his ball runs into oncoming traffic on a busy road. B, who’s nearby, notices this and* (10)  
327 *pushes A out of harm’s way, but in doing so A breaks his arm.*

328  
329  
330 *A is standing on the street corner. B walks up to A and taps her on the shoulder and asks her for directions.* (11)

331  
332  
333 In both (10) and (11) B is not liable to A because of presumed consent. In (10) the situation is emergent whereas in  
334 (11) it is not. Presumed consent has a high burden, additionally requiring a justification for the contact. For example,  
335 the actor must show that the contact was minor and was in fact customary in the community.  
336  
337

## 339 2.5 Constructive Consent

340 Often known as “implied-in-law” this form of consent is rare, but quite extreme. It considers conduct that causes socially  
341 justifiable minor acts that not only not require consent, but in fact can be performed *despite explicit objections* to the act  
342 and expressions of non-consent. Here, the standard is an objective one of whether a reasonable person would consent  
343 under the circumstances.  
344

345  
346 *A drives to work and arrives at the scene of an accident. B, a police officer, orders him to wait in the car while*  
347 *the scene is cleared. A demands that he be allowed to turn around and leave the scene. B refuses fearing safety* (12)  
348 *risks.*

349  
350 B is not liable to A for false imprisonment, a variety of intentional tort, because of constructive or implied-in-law  
351 consent. This form of consent is controversial particularly because of its blatant disregard for the person’s actual  
352 objections.  
353

354  
355 The concept of consent permeates our legal system beyond intentional torts. Consent is a critical defense to certain  
356 criminal conduct and plays a role in mutually beneficial contractual agreements. One interesting twist in criminal law  
357 is the existence of a category where consent is not a defense. That is, even if the person granted the actor consent, the  
358 conduct might still be deemed to be wrongful and the actor culpable. For example [198 1980; 201 2012]:  
359

360 *A and B engage in a sex act in which A consents to B using a dangerous weapon during sex. B uses the weapon* (13)  
361 *on A and A is harmed.*



365 Here, B might be criminally liable in the Commonwealth of Massachusetts as consent is no defense to a charge of  
366 battery and assault with a dangerous weapon.  
367

## 368 2.6 Reluctant consent

369 The framing of consent as a defense (to a tort claim or a criminal charge) does not consider how fervently one gives  
370 consent. If there is legal basis in evidence for one or more forms of the above consent it can be sufficient. However,  
371 exploring the amount of enthusiasm shown in regards to consent could have normative implications even if it does not  
372 have legal ones.  
373

374 *A enters a taxi driven by B. B takes an unusual route that involves driving through a part of town that A*  
375 *associates with past emotional trauma. A says nothing.* (14)  
376  
377

378  
379 By not objecting to this route, B might be able to rely on a subsequent defense of apparent consent when faced with  
380 an intentional infliction of emotional distress claim. It might even be the way most taxi drivers drive to the location that  
381 A intends to reach. However, A's consent is absent and, at best, reluctant. The level of A's consent might be evidenced by  
382 facial expressions suggesting it or by other gestures like looking away from the windows or closing eyes. Normatively,  
383 it seems B should take this reluctance into account and explore an alternative route. The norms at play include business  
384 aspects of ensuring that the passenger is kept happy, the human decency aspect of not having an interactant relive a  
385 past trauma and so on. Thus, the legal framework, while giving us a grounding for extreme behavioral violations, might  
386 be insufficient. We might need to consider lesser norm and convention violations that are especially evident in cases of  
387 reluctant consent.  
388

389 A thorny issue in (14) is whether B had the duty to present other alternatives. The legal system is not entirely clear  
390 on this issue and some scholars have compared reluctant consent with fullness (or partial-ness) of consent and their  
391 connections to the doctrine of the assumption of risk. What risk does A assume in taking B's route? If B presented A  
392 with a second less-traumatic route, and A still reluctantly chose the first route, then the doctrine of assumption of risk  
393 may play a role in clearing B. One takeaway here is that the consenting process is much more complicated than a single  
394 interaction and typically involves a complex mix of duties and responsibilities of the actor and the mental state and  
395 behaviors of the person who arguably is deemed to have consented.  
396  
397

## 400 2.7 Other Issues

401 We have only scratched the surface of the intentional tort landscape around the notion of consent. There are numerous  
402 other issues around whether a person is competent or has the capacity to even grant consent. The law has established  
403 that for consent to be valid the person must have capacity to consent. Capacity is typically gauged by age and mental  
404 competence. Children, those suffering from mental disorders, and intoxicated persons are deemed to not have the  
405 capacity to consent. This means tortious behavior will not be defensible even if consent is obtained from those without  
406 the capacity. Consent that is granted can also be revoked at any time. Consent can be conditional in that it is ineffective  
407 if there has been a failure to meet conditions that the consenter has imposed. Consent must also be voluntarily given  
408 and not coerced or obtained under duress.  
409

410 Beyond intentional torts, notions of consent appear as another related legal principle in torts of negligence known as  
411 "assumption of risk." Assumption of risk is often treated as consent to conduct that is merely negligent [Simons 1987,  
412 2006]. The doctrine of assumed risk connects consent with whether there might have been some uncertainty associated  
413  
414

415  
416 Manuscript submitted to ACM



417 with the consequence of the conduct. In the medical context, a related notion known as the doctrine of informed consent  
418 is used to receive a patient's acknowledgement that there are risks associated with a particular procedure even if the  
419 doctor performed it correctly.  
420

421 Consent is also a crucial doctrine in criminal law (as noted earlier) and in contract law. Unlike tort law, consent in  
422 contract law is centered not around internal subjective mental state, but external performance. Moreover, in contract  
423 law, consent serves as a proxy for evaluating whether the parties voluntarily entered into a contract.  
424

425 Finally, it is worth clarifying that the legal principles in the paper merely serve as a guide to unpacking an otherwise  
426 complex normative landscape in many social settings. We do not require or even suggest that robots be granted any  
427 form of legal status as a person or agent. On the other hand, we want to stress that it is insufficient to treat robots as any  
428 other form of technology where consent is equivalent to user preferences. Past research has shown that robots create  
429 expectations in people that they will be less like passive technologies (e.g., washing machines and cell phones) and  
430 more like persons. While we do not take a position on the question of whether robots have "agency" in a philosophical  
431 sense, we do encourage thinking about humans in social spaces not in terms of "users" but in terms of "interactants"  
432 [Calo 2010]. Robots, as part of their functions, might interact with other non-users as well. Moreover, it is not enough  
433 to express the limitations of the robot to users as users impute agency and capabilities beyond that which the robot has  
434 regardless of prior disclaimers.  
435  
436

437 Entangled with these foundational legal notions, the nuances of consent can be affected by other aspects of a social  
438 context like, for example, gender and culture. Gender and cultural variables associated with interactants and social  
439 groups can affect behavior and expectations associated with consent. For example, what is considered presumed consent  
440 in one culture might not be in another. Not only will robots need to account for the cultural considerations of the  
441 society in which they are deployed, but also potentially the cultural diversity of their interactants within that society.  
442 Many communal settings like restaurants and parks might require robots to carefully acknowledge cultural variations  
443 among their human interactants.  
444  
445

### 446 3 CONSENT ISSUES IN ROBOT APPLICATIONS 447

448 Consent issues can arise any time there is interaction between agents. As we have seen, these consent issues come in a  
449 variety of forms and are highly dependent on the social space in which the interaction is taking place. Thus, although  
450 the notion of consent might seem somewhat dyadic, what influences the dynamics of granting and revoking of consent  
451 is not merely dyadic, but in fact influenced by the normative expectations of the entire social space. This means consent  
452 issues can arise even in scenarios that on the face are not socially or morally charged, as the social setting affects the  
453 normative expectations of the interactants. In this section, we will consider two robot applications, in increasing order  
454 of "normative-charged-ness" of scenarios in which they are typically deployed: (1) robot vacuum cleaners and (2) robot  
455 waiters.  
456  
457  
458

#### 459 3.1 Robot Vacuum Cleaners 460

461 Some have estimated that robots account for over 20% of the world's vacuum cleaners. Even if this is a bit overstated,  
462 robot vacuum cleaners are ubiquitous. These robots have a range of actuation and sensing capabilities, albeit much  
463 more limited than a humanoid robot. They have cameras, IR and laser sensors and gyros to be able to map out a house  
464 and remember what parts of the house have been cleaned and what remains. They also have other features to avoid  
465 bumps, recognize no-go areas in the house and schedule cleaning sessions. Many are beginning to come equipped with  
466 WiFi capabilities and connections to Alexa and Google Home voice-activated personal assistants. These robots are not  
467  
468

Manuscript submitted to ACM

469 simply passive devices that provide a service, they are mobile robots that, numerous studies have shown, cause humans  
470 to impute agency and capabilities to, beyond even that which the robot actually possesses [de Graaf 2016; Scheutz 2011;  
471 Sung et al. 2007].  
472

473 The robot vacuum cleaner's primary designed function is to vacuum and clean the house, and not necessarily to  
474 interact socially with humans. Even so, before, during and after the robot performs its cleaning duties, it can and will  
475 interact with humans because it is deployed in an otherwise social space. It's owner might interact with it to decide  
476 appropriate cleaning schedules so as to take into account the presence or absence of other members of the household.  
477 The robot vacuum cleaner might inadvertently interact with non-owner humans who might be present in the space the  
478 robot is cleaning. Humans might direct the robot to go elsewhere or take an alternative cleaning route to avoid being  
479 disturbed. During its operation, the robot is also mapping the house and possibly recording video as it learns about  
480 the layout of the house and habits of the humans living there. Consent issues can arise quite easily in many of these  
481 interactions - robot getting in the way of humans moving about the house, cleaning parts of the house that are not  
482 necessarily no-go areas, but temporarily off-limits, cleaning near children's toys, some of which might be small parts,  
483 listening to private conversations in the name of listening for human commands, and recording video in private areas  
484 of the house. In many of these cases, the humans may not have actually consented, but can the robot assume there is  
485 apparent or presumed consent?  
486  
487  
488

### 489 3.2 Robot Waiters

490 We have seen in the examples presented in Section 2 that consent issues can arise quite frequently in restaurant  
491 scenarios. Robot waiters constantly interact with their assigned customers, other customers walking in the restaurant,  
492 chefs, and managers. During these interactions, the robots through their appearance and functional (actuation and  
493 sensing) capabilities generate a set of expectations in the interactant humans that together establishes the social space  
494 for that particular interaction [Riek and Robinson 2011].  
495

496 Distinct from robot vacuum cleaners, however, robot waiters are designed to provide a service that is inherently  
497 social, but not necessarily high-stakes or weighty. They direct customers to their seats, take orders, deliver food,  
498 check for comfort, and are meant to generally address grievances and comments issued by the customers. Each of  
499 these interactions involves entering personal spaces, touching humans, giving and taking possession of objects (some  
500 belonging to the customer). Most, if not all, of these interactions are consent-driven, in that some form of implicit or  
501 explicit consent was given or presumed, otherwise these behaviors would be problematic.  
502

503 For example, consider the situation where the robot waiter is tasked with removing plates from a customer's table  
504 once they have finished eating. Customers might explicitly call out to a robot and ask for their plate to be removed.  
505 They might actually let the robot waiter know that they have finished their meal, as an indirect way to signal their  
506 desire for the plate to be removed and a check to be delivered. More often than not, customers will not make this request  
507 explicitly or even indirectly. Instead, they will provide implicit cues to suggest that they are done with a particular part  
508 of the meal and that the waiter *can* take away their plate. They issue this consent by either arranging their forks and  
509 knives a particular way, or push their plate away a bit, or even lean back away from the table as the waiter approaches.  
510 It is also possible that the customer is not really thinking about the waiter taking away their plate, but would not mind  
511 if the waiter did so. This might be conveyed via a smile that is offered after the waiter has already leaned in and grabbed  
512 the plate. The waiter may have assumed the customer was done with their meal because they had not been eating  
513 anything from it for a while. Robot waiters need to be able to read these subtle cues. Robot waiters must also be able to  
514 read cues in which the waiter does not want their plate taken away from them, even if it seems as if they are done, or  
515  
516  
517  
518  
519  
520

Manuscript submitted to ACM

521 they have inadvertently placed their fork and knife in a “finished” position. If the robot waiter does not pick up on  
522 refusal hand-gestures or on the customer’s grabbing of a fork to show they are still eating, it can negatively impact the  
523 service and experience. It is worthwhile to note that regardless of the aesthetic appeal of the robot or of its ability to  
524 display pleasant emotions, taking away a customer’s plate while they are still eating without their consent (in one of  
525 the forms mentioned earlier) can ruin the interaction.

527 This is just one example of a scenario where a robot’s behavior, which would generally have been considered  
528 preferable, is suddenly deemed unwanted. Such a robot must be able to recognize the corresponding consent-cue to  
529 avoid the entire interaction from being impacted. Other such examples of behaviors that might require a consent-cue  
530 include: friendly touches of the customer’s shoulder, filling or not filling water or wine, approaching the table to check  
531 on the customer’s experience, picking up an item the customer may have dropped on the floor, clearing a path between  
532 tables by moving customer’s garments, bags and chairs or requesting customers to make way, helping a customer take  
533 off their coat, and many more.

#### 534 4 RESEARCH DIRECTIONS FOR HRI

535 In this section, we explore possible research directions for HRI as the legal structure presented in Section 2 allows us to  
536 think about consent more deeply. But, how can this legal structure help with designing better human-robot interactions?  
537 We specifically consider four research directions or themes that the HRI research community is best positioned to  
538 address: (1) when consent is applicable, (2) how consent can be detected, (3) what happens when snap judgments must  
539 be made, and (4) what the different roles for a robot are when it is fully integrated into the normative fabric of society.  
540 For each theme, we will raise some research questions and provide architectural considerations that the HRI community  
541 might consider.

##### 542 4.1 Applicability of Consent

543 When does a behavior need consent? Presumably, there are behaviors that robots can perform that do not need consent.  
544 For example, a vacuuming cleaning robot cleaning a kitchen floor might not need consent for choosing to follow one  
545 (random) cleaning path versus another (random) cleaning path. Even in more socially-charged settings, it would be  
546 unnecessary for say an elder-care service robot to have to secure approval for every single motor movement. This  
547 would be burdensome and possibly even negatively impact the interaction itself.

548 Current HRI research has provided some guidance into what behaviors are preferred versus those that are unpreferred.  
549 The angle of approach example presented in Section 1 is one such contribution. We can, at a first approximation, note  
550 that unpreferred behaviors will probably need to be consented or at least need some sort of justification (viz. apparent  
551 consent, presumed consent). Also, it is possible that in many cases even preferred behaviors could require consenting.  
552 Behaviors are situated in a dynamic environment and their normative valence and the applicability of consent can,  
553 therefore, shift in time. Consider this example:

554 *A and B are sitting at a shared library table that can seat two. While studying, B moves some textbooks around*  
555 *pushing several books over to A’s side of the table.* (15)

556 B’s behavior caused an intrusion into A’s space. B’s behavior, however, did not start out that way. It was when B’s  
557 books crossed the imaginary halfway line on the table between A and B that B needed consent. The implicit norm at

573 play was the idea of a fair partitioning of a shared work space. The norm became activated once B's book crossed this  
574 imaginary center line.

575 Or consider this service example:

576  
577 *A is a patient at a hospital and is recovering from a surgery in his assigned room. A receives guests and*  
578 *wellwishers who bring him cards, flowers and some of A's favorite take-out Chinese food, which they share* (16)  
579 *during the visit. B, a janitor, is tasked with cleaning A's room and throws out not only the Chinese food*  
580 *containers, but also some of the cards.*

581  
582  
583 Once allowed into the room, B's behavior of clearing the dirty food containers does not need consenting from A, but  
584 clearing the adjacent greeting card might. The underlying norm hinges on the agent being able to detect and consider  
585 questions of ownership, possession, and the use of space, in addition to the normative affordances offered by the objects  
586 in the environment. The greeting card serves the purpose of lifting A's spirits and therefore is an important component  
587 of the environment. B's action of clearing it might eliminate this positive utility. Note, it might be entirely inappropriate  
588 to even ask for consent to clear the greeting card as that would suggest B's normative proclivities.

589  
590  
591 **4.1.1 Research Questions.** When does a behavior turn from being consent-free to one that requires it? Are there some  
592 behaviors that require robot consent, but not human consent? Are social interaction cues of proximity, eye gaze, and  
593 so on also cues suggesting that consent may be needed? Who should be consenting in a particular situation and if  
594 that person is unavailable, should the behavior be postponed? Do humans expect robots to have the same level of  
595 understanding of consent as they do with other humans? Do they hold robots to higher standard, i.e., more conservative,  
596 when it comes to asking for consent? A popular maxim used in connection with consent in sexual acts is – “yes means  
597 yes, no means no.” – do we expect the same understanding from our robots?  
598  
599

600  
601 **4.1.2 Architectural Considerations.** One important shift in architecture design that is prompted by consent processing  
602 is that even what are otherwise considered to be automatic behaviors in robots (e.g., approach and obstacle avoidance  
603 behavior, eye gaze and blinking behavior, posture and gesturing, etc.) are subject to modulations based on consent. This,  
604 in turn, requires the architecture to have built-in suppression mechanisms that can be triggered by changing consent  
605 status and interrupt, suppress, or alter these reflex-like behaviors. In addition, the architecture requires a knowledge  
606 base with common sense knowledge about the different consenting conditions and the appropriate modulations. For  
607 example, it will require general knowledge about possession of property, what actions are and are not allowed with  
608 another person's property, and under what circumstances the robot can dispense with certain rules and norms because  
609 there are role-based expectations that trump them. In the above example, even though the food container was technically  
610 owned by patient A, a janitor robot B would have had to make the inference that based on its role to clean up used  
611 items – and an empty food container was a used item – the ownership of the container was no longer an issue and it  
612 was, in fact, expected from B to take the container away from A. This is a case where two norms clash, the norm of  
613 following ownership rules versus the norm of fulfilling one's duties based on one's roles. The challenge then is to detect  
614 that the role-based norm of cleaning up the container trumps the ownership-based norm of not touching it precisely  
615 because it implies consent by the owner to do so. Recently, research progress in explicit norm representations [Sarathy  
616 et al. 2017b,a], norms relating to object affordances [Sarathy and Scheutz 2016], norms relating to ownership [X. Tan  
617 2019] and other prerequisites for consent reference what social norms might be present and how they legitimize behavior,  
618 which in turn might suggest how they are triggered and modulated by consent.  
619  
620  
621  
622  
623

624 Manuscript submitted to ACM

625 Detecting shifting consent-dependencies of behaviors is particularly challenging from a technical standpoint. De-  
626 tecting consent applicability shares many similar technical aspects to detecting consent itself. The robot will need to  
627 detect a variety of cues from its interactant, others in the environment, and other animate and inanimate objects that  
628 constitute the social space. A technical designer must consider what level of symbolic abstraction does a behavior need  
629 to be represented as before it qualifies as one applicable for consent. Clearly, step-wise motor movements and joint  
630 angles might be too specific and not necessarily applicable. On the other hand, higher-level behaviors of swinging arms  
631 and even higher-level conduct of picking up and manipulating objects might be candidates for consent-dependency. It  
632 is also possible that the robot might need to consider the nature of objects it is handling (e.g., knives versus a bouquet  
633 of roses) when performing certain behaviors.  
634  
635  
636  
637  
638

#### 639 4.2 Detecting Consent

640 At the core of consent-related considerations that we have discussed in this paper is the question of how and what  
641 consent cues are detected. At first glance it seems like language, gestures, facial expressions, eye gaze, approach,  
642 proximity can all serve as valid vehicles to communicate one's subjective intent of actual consent or objection. But, how  
643 can a robot detect consent expressed through an ISA, or through the usage of objects in the environment? Consider this  
644 example:  
645  
646

647 *A was sitting at the bar getting a drink that he already paid for. A drinks most, but not all, of the drink and*  
648 *then leaves the bar. B a waiter did not see A leave. B takes away A's almost empty mug.* (17)  
649  
650

651 Presumably, B's conduct was normatively appropriate given that A might not have objected to the mug being taken  
652 away. Note, the mug actually belongs to the restaurant, so there is no legal consequence here. We nevertheless use the  
653 legal framework to inform our discussion about consent norms and consent conventions. So, even if A had just stepped  
654 out for a smoke or to take a call, B's actions might have been justified. Now, contrast this example with one where A  
655 leaves his bag on the seat and goes to the restroom. Here, B's actions of taking away the mug might not be justified. A  
656 might object that he had not finished his drink and his intentions were made clear by him leaving his bag on the chair.  
657  
658

659 There are many different types of consent cues. Some cues are *consent-granting*, in which case they have an  
660 illocutionary force of conveying permission to a behavior. Alternatively, some cues are *consent-denying*, in which case  
661 they have the force of conveying explicit non-consent to a behavior. In some instances, the person is not yet ready  
662 to signal, but they will eventually. These *consent-withholding* cues suggest that the actor return at a later specified or  
663 unspecified point to inquire about the consent. For example,  
664

665 *B, a waiter, is about refill A's wine glass. A gestures with a downward facing open palm and says, "not yet." B*  
666 *retracts.* (18)  
667  
668

669 Here, A did not say to never fill his glass, instead withheld consent until some later point. Now, there is a subtlety  
670 here where A might have instead just said "no", which might have been a consent-denying cue, or if prior consent had  
671 been granted, a *consent-revoking* cue, taking away a prior consent. Finally, the consent itself might be conditioned and  
672 expressed via a *consent-conditional* cue of A saying "only if it is a 1991 Argentinian Malbec" suggesting that refilling is  
673 only acceptable for a particular variety of wine.  
674  
675  
676

677 These different types of cues not only convey whether or not consent is given, but also specify which behavior they  
678 refer to. In (18), when A says “not yet” they are not withholding eventual consent for B pouring the wine on them. The  
679 content of the consent is difficult question that poses many challenges.  
680

681 The legal landscape also suggests that for succeeding on a defense of consent, the actor must also prove that the  
682 person has the capacity to consent, is consenting voluntarily and has all the information necessary to decide. The latter  
683 of these elements is the basis for the doctrine of informed consent, common in medical settings.  
684

685  
686 *4.2.1 Research Questions.* How to detect social and environmental consent cues? Are there certain cues that transcend  
687 situations and universally signify consent or non-consent for any behavior? Is it enough for the robot to detect and  
688 then immediately act on a consent cue or should it ask for confirmation? What is the role of uncertainty in detection?  
689 Crucially, what if the cue was not only unclear in a noisy environment, but also ambiguous in its scope and content? If  
690 the robot is unable to glean actual consent, should it explicitly ask the human or should it be satisfied with inferring  
691 apparent consent?  
692  
693

694  
695 *4.2.2 Architectural Considerations.* To be able to detect cues, the robotic architecture must be able to first process  
696 and integrate percepts that can serve as cues obtained through various sensory modalities. Consider linguistic cues,  
697 which, although they seem like the clearest type of cue, can still be a technical challenge to process. Here, processing  
698 consent involves a degree of intent recognition and processing of potentially indirect cues (via indirect speech acts).  
699 Architecturally, it might not be enough to implement statistical natural language processing and speech recognition  
700 systems that do not have a way to represent contextual variations in the illocutionary and perlocutionary force of  
701 speech acts. A way around this issue would be to explicitly design an architecture that allows for performing cue-based  
702 inference. An approach might be to use a rule-based strategy where implication rules can represent consent-cues as  
703 antecedents and behaviors as consequents. Cues then trigger rules that determine appropriate behavior. The normative  
704 propriety of the behavior can then be computed via logical modus ponens based on the perception of the relevant consent  
705 cues. A variant of this rule-based approach can be seen in indirect speech research ([Briggs and Scheutz 2013]). This  
706 sort of architecture and knowledge representation has also been shown to be amenable to non-statistical, instruction  
707 based teaching and learning via natural language.  
708

709  
710 In addition to linguistic cues, the robot may have to perform inference from visual percepts as in (17). In such cases,  
711 the robotic architecture must be equipped to not only recognize objects, but also generate more holistic representations  
712 of a scene. This means object properties need to be inferred and, in dynamic environments, objects and their properties  
713 will need to be tracked. These object representations will then need to be integrated with other modalities including  
714 the linguistic cues noted above in order to generate a coherent *story* for whether or not the robot has deemed there to  
715 be consent in a particular situation.  
716  
717

718  
719 But that only covers cases of actual consent. The cues in actual consent are typically limited to those that result  
720 from the conduct of the person consenting. However, in apparent consent situations, the cues can come from not only  
721 from the person who is the target of the behavior in question, but also from others and from customary aspects of the  
722 situation.  
723

724 The robot must also be able to understand the interactants’ capacity to consent and ensure that they are provided  
725 with all the information they need to make their consent decision. Robots working with certain vulnerable populations  
726 – children, elderly, substance-abusers, mentally disordered – must account for these aspects.  
727



### 4.3 When Consent Does Not Matter

What the legal landscape has shown us is that there are situations where an otherwise unconsented behavior is appropriate (presumed consent and constructive consent) and consented behavior is still inappropriate (consent not a defense to certain criminal charges).

In situations like emergencies where consent is presumed, the law places a high burden on those looking to use this principle in an affirmative defense to an intentional tort. Presumed consent falls squarely below actual and apparent consent and has limited scope. But, these are situations in which our robots might find themselves, possibly even more frequently than imagined. Search-and-rescue robots, space robots, medical robots and robots in many emergent situations must take quick action in order to prevent a larger harm.

In certain cases, sex robots engaged in BDSM activities might have received consent for certain acts, but the law has held that performing these acts (particularly those with dangerous weapons) can still remain a criminal charge. On a less egregious note, sex robots that bring spontaneity and creativity into a sexual encounter (e.g., a surprise tickle) might need to rely on presumed consent to provide a positive social interaction.

*4.3.1 Research Questions.* Should a robot ever presume consent? If a robot does presume consent, should its scope be kept limited to a few specifically designed behaviors. Presumed consent is particularly controversial and it, combined with constructive consent, start tearing away at personal autonomy. If a person's consenting of a behavior is taken for granted or presumed given or worse, ignored, is there not a taking away of a moral right to autonomy? Should we give robots this privilege? What if it ends up helping us? Moreover, do we humans expect this from our rescue robots? Some in the HRI community have explored the role of supererogatory actions (like in (10)) and whether we expect a street cleaning robot to go out of its way to save a child from an oncoming car. The technical challenges for designing such a robot share many aspects of cue detection and norm representations noted earlier. In cases where consent is not a defense, should robots be designed to never perform these behaviors? How can we balance robot creativity, spontaneity and risk-taking behaviors (all of which might have positive social value) with normative propriety? Can the robot presume consent if it enhances a social experience?

*4.3.2 Architectural Considerations.* While the question of whether or not consent matters in certain special situations is largely an empirical and philosophical one, we can still evaluate what architectural capacities a robot must have. Uniquely, robots must be equipped with suitable architectural components to ensure that it does not become a passerby, or bystander, during critical situations. Thus, it is important for there to be coordination within the architecture for integrating cues suggesting an emergent situation and the ability to reason against norm expectations. Crucially, such reasoning is required as the robot needs to deduce that a norm violation might be necessary in order to achieve a greater good. The architecture will need mechanisms to reason with multiple such mutually inconsistent and conflicting norms, all in real-time under extreme time-pressures imposed by the emergent situation.

### 4.4 Robot Roles

Thus far, we have looked at several examples in which we examined the propriety of actions of B in interactions with A. While it is natural to think about the role of a robot as B, and the human as A, in reality, the role of the robot could be quite varied. Here, we can explore some of the alternative roles the robot plays in an interaction.

*4.4.1 Research Questions.* First, we can consider the robot as a consenter (or protester), i.e., A. Robots in this role raises many questions about robot rights and autonomy [Darling 2016]. Should the robot be able to consent at all? A narrow

Manuscript submitted to ACM



781 view of robot rights would suggest that robot consent does not matter. This view, in turn, might increase the scope of  
782 presumed consent or constructive consent in robot interactions. Can we then abuse a robot?  
783

784 *B verbally abuses A a personal assistant robot by shouting expletives.* (19)  
785

786  
787 Some might argue that A's consent towards this behavior is irrelevant as A is a machine. However, the counter  
788 argument to this view is that B's verbal abuse sets the wrong normative tone in the environment and therefore should  
789 be seen as wrongful. There may not be emotional distress for A, but there could be for another, say C, who is watching  
790 this interaction.  
791

792 We can bring in another interactant C who might be a third-party to the interaction. C could play the role of sanctioner  
793 and function to right normative wrongs in an environment. For example, C could be a manager of a restaurant in which  
794 B a waiter is serving A a customer. Alternatively, we can think of C as a sanctioner role that could be played by A  
795 themselves.  
796

797 In such a role, we must ask if it is okay for the robot to call-out and sanction normatively harmful behavior (see  
798 [Fessler 2018] for Alexa pushing back against abuse). In (19) could the robot defend itself and protest A's abuse? Should  
799 the robot do such a thing?  
800

801 Another role a robot could take is not that of a moral arbiter, but that of the elusive "reasonable person". In much of  
802 the legal scholarship one aspect that has been at the center of many debates is the the objective standard of a "reasonable  
803 person", one who exercises average skill and judgment and can objectively assess a situation and arrive at the correct  
804 solution. Many legal scholars have argued that such a person does not exist and is, in fact, by no means average [Gardner  
805 2015]. Nevertheless, robots might present themselves in the unique position to serve as such a reasonable person.  
806 Can these reasonable robots help resolve conflicts by performing an unbiased and objective analysis of the facts and  
807 situation?  
808  
809

810 **4.4.2 Architectural Considerations.** The architectural capabilities we have discussed thus far apply to the robots that  
811 serve as sanctioners or reasonable persons. Of particular note, these robots must be able to detect norm violations,  
812 and so also have architectural components to compare normative cues with norm-based expectations to detect norm  
813 violations. Robots in these varying roles also have the function of deciding if the justifications provided by parties is  
814 acceptable and as such must have mechanisms to infer consent by exploring the purposes of actions. Investigations  
815 into human-robot teamwork have even included their possible role as repairers of social conflict [Jung et al. 2015].  
816 Robots must also maintain honesty [Hoffman et al. 2015] and sometimes even keep secrets [Kahn Jr et al. 2015] in order  
817 to fulfill the responsibilities [Asaro 2007] imposed by their roles. This also means that robotic architectures serving  
818 any role must be able to represent various supporting aspects of the situations – ownership, possession, resulting  
819 permissions and obligations and whether or not agents have the capacity to consent and have done so voluntarily  
820 with the needed information. The architectures must also be able to determine precedence relationships to help enable  
821 consent inference.  
822  
823  
824  
825

## 826 **5 NEAR-TERM NEXT STEPS FOR HRI**

827  
828 In the previous section, we discussed four research directions (or themes) that are opened for the HRI community by  
829 the more nuanced consideration of consent. Critically, we proposed research questions and technical (architectural)  
830 design issues that will need to be addressed if we are to design fully consent-aware robots. In some sense, we presented  
831  
832

Manuscript submitted to ACM

833 a research vision for normative HRI that is focused on much-needed technical advancements that are likely to be  
834 longer-term. But, what can HRI researchers start doing today? We do not yet have all the technical capabilities to build  
835 the architectural components needed for fully consensual agents. However, as we have discussed in Section 3, robots  
836 are already deployed in many real world social settings and will need to be able to handle at least some partial set  
837 of consent cues. How can HRI researchers design and evaluate more consent-aware interactions with these current  
838 robots? In this section, we provide some recommendations for how current HRI researchers can adapt their design and  
839 evaluation methodologies to begin accounting for consent. We expect that by doing so, this will guide future technical  
840 developments in the right direction.  
841  
842

### 843 5.1 Early-Stage Interaction Design

844 HRI design is a truly collaborative endeavor and tackling some of the consent issues might need to go beyond  
845 architectural modifications of robots. That is, it is insufficient to only take a robot/technical-centric view and place  
846 all the burden of consent-management to the robot and its underlying A.I. Accounting for consent must begin at the  
847 earliest stages of interaction design. Interaction design involves considering not just behaviors in isolation, but the  
848 generation of sequences of behaviors and how such sequences influence and are influenced by particular application  
849 scenarios. Behaviors could include both dialogue and non-dialogue (body movements), and interaction forms the overall  
850 specification of the set of behaviors that technical researchers must build into their robots. Interaction designers can  
851 proactively consider consent when designing specific interactions by asking the research questions we presented above  
852 *early* in the design stages of storyboarding and prototyping. Asking when and how consent will apply to particular  
853 scenarios by designers at this stage will likely prove crucial when designing the specifics of the interaction itself and  
854 contain the set of robot capabilities. As behaviors are evaluated against a scenario and flows constructed, interaction  
855 designers can asking critically *when* and *why* certain behaviors are acceptable or not, clueing them into whether or  
856 not consent applies in a particular situation. For example, asking when and why it is okay for a robot waiter to fill  
857 an empty water glass raises many interesting questions about the situation and the nature of the social space that  
858 might need to be answered before a suitable interaction can be designed. The answers might suggest entirely different  
859 interactions for an expensive restaurant versus a fast food restaurant. A deeper inquiry at that point might shed some  
860 light on what factors might turn acceptable behaviors into unacceptable ones and vice versa. For example, customers  
861 in expensive restaurant might provide consent cues that are much more indirect and implicit than those at fast food  
862 restaurants. Similarly, the interaction designer could attempt to model the mental state of users in the scenario to  
863 consider what these users would consider reasonable deviations from acceptable behavior. If the scenario involves  
864 emergencies, interaction designers will be able to directly incorporate notions of presumed consent. For example, when  
865 a robot detects that a human is about to get run over, it can assume presumed consent for pushing the person out of  
866 harms way, even though pushing normally would violate a tort.  
867  
868  
869  
870  
871  
872  
873  
874

### 875 5.2 Experimental Evaluation and User Studies

876 When robotic systems are introduced into a real world application scenario, they can have a significant impact on  
877 usability, user experience, and social aspects like impact and acceptance. Prior to such an introduction, a strong  
878 evaluation methodology is needed to enable fair comparisons of competing HRI systems along various dimensions  
879 including efficiency, feasibility, safety, and social and psychological aspects. There have been many approaches proposed  
880 in the literature for the most suitable set of metrics for evaluating robots. A prominent one is by Steinfeld et al. [Steinfeld  
881 et al. 2006] who propose a set of task-specific as well as common metrics for assessing a robot's performance. In this  
882  
883  
884

Manuscript submitted to ACM

885 comprehensive work, they discuss some socially-driven metrics (persuasiveness, trust, engagement and compliance),  
886 but these are largely discussed in connection with robots whose primary function is to interact socially with people.  
887 As we have discussed, socially-relevant consent issues can arise even when the robot's primary function is not social  
888 interaction; consent issues can be evoked from the robot simply operating in a social space. The common metrics  
889 proposed by Steinfeld et al. do not consider such normative and socio-cultural aspects of a scenario in which the robot is  
890 deployed. Thus, urgent future work is needed to develop suitable metrics and measurement instruments that explicitly  
891 probe for consent.  
892

893  
894 In addition, we propose that interaction researchers studying behaviors, but not necessarily studying consent, can  
895 and should provide some more guidance about the ecological validity of their findings. They can do this by asking  
896 when preferred behaviors become dispreferred and vice-versa. User-study questionnaires already ask questions about  
897 user preferences and likeability of behaviors. These can be extended to include questions about dislikeability and what  
898 makes a certain behavior inappropriate, and when dispreferred behavior could become preferred. By exploring ways to  
899 extend these questionnaires we can study how consent influences existing interaction research findings and provide  
900 some new ideas for future research inquiries.  
901

902  
903 More generally, when robot behaviors are being evaluated for a particular use-case, HRI researchers must consider  
904 the nature of normative expectations held by humans present in that use-case, and ensure that their experimental  
905 subject-population is representative of those in the real world. Most HRI user studies obtain demographic information  
906 either to directly study gender or some other effect relating to the identity of participants, or indirectly to correct for  
907 individual differences across a population. We are suggesting here that consent-preferences are an equally important  
908 (an possibly biasing or confounding) factor to either study or at least correct for in user studies.  
909

910  
911 Finally, when drawing conclusions from user-studies HRI researchers must make explicit the types of consent-related  
912 aspects they have accounted for in their experiment. That is, their conclusions and claims about behavioral preferences  
913 must be accompanied by the scope and nature of the scenario(s) in which those behaviors are situated. During the  
914 peer-review process, we urge reviewers to extend this inquiry and consider how consent or lack thereof can impact  
915 these results and push the researchers to make explicit these limits.  
916

### 917 **5.3 Revisiting Past HRI Experiments**

918  
919 The future outlook for HRI research incorporating this rich notion of consent is quite exciting, and we hope the  
920 suggestions provided in this paper will facilitate a smoother transfer of future interaction research results into real  
921 world robotic application domains. That said, it is still unclear what we are to do with past and current research  
922 results like, for example, research on "robot approach" that we alluded to in Section 1. A formidable challenge lies in  
923 understanding the scope of the results in not only approach research, but also dialogue, eye gaze, proxemics, touch  
924 and other interactionally-relevant behaviors. We know that behavioral preferences can be highly contextualized and  
925 dictated by the task at hand. But, it is unclear how the results for particular studies can be generalized beyond the  
926 specific experimental setting in which they were designed.  
927

928  
929 For example, in eye gaze research, HRI researchers have suggested that high levels of mutual gaze express feelings of  
930 trust and extroversion, and gaze aversions express feelings of distrust and introversion [Admoni and Scassellati 2017].  
931 However, this is not universally true, as an unwelcome extended eye gaze might be suspect and can adversely impact  
932 trust. In conversational contexts, eye contact is seen as more acceptable in emotionally neutral topics but not so when  
933 the conversation is embarrassing [Admoni and Scassellati 2017], suggesting that the topic of conversation could be a  
934 modulating factor when evaluating robot perceptions. Although there could be other modulating factors in particular  
935

936 Manuscript submitted to ACM

scenarios, we propose that consent is a key modulating factor in all HRI settings, and whether or not certain eye gaze strategies are deemed socially acceptable (or any other socially relevant evaluation metric) is dependent on whether the interactants mutually consented to each others behavioral choices. In situations where there is no explicit consent, the question then becomes if the social setting provides a reasonable basis for establishing some implicit form of consent. A social space that places a higher bar on consent, will not allow for prolonged eye contact even if conversations are emotionally neutral, without a reasonable basis for it. In order to extend the applicability of many such HRI studies into *real* interactional contexts, what we need is a deeper understanding of the social spaces anticipated by these studies. We need a better sense for what sorts of factors need to be present or absent to reasonably establish that either consent can be apparent or presumed in the social space, or actual consent was indeed given.

As another example, consider research in politeness in natural language. HRI researchers first looked at speech modifications as a way to modulate politeness and thereby improve human perceptions of robot helpers. Torrey et al. [Torrey et al. 2013] argued that robot helpers create a positive impression when using hedges (“I guess” or “kind of” or “probably”) and discourse markers (“I mean” or “like you know”). Their study involved humans observing robots employing those dialogue strategies, not humans directly interacting with such robots. Strait et al. [Strait et al. 2014] subsequently noted that such speech modification strategies alone are insufficient and other aspects of the interactional scenario like participants’ personality, dialogue efficiency and task success also impacted the perceptions of the robot. They explored the relative effects of communication strategy as modulated by interaction modality and presence as well as robot appearance. Their results showed that some prior claims made by Torrey et al. do not generalize to real interactional scenarios where other modulatory factors might be at play. Their conclusions leave open the possibilities for further situational factors that might modulate interactions. We believe this is a step in the right direction, and we urge HRI researcher to make more explicit the spectrum of modulatory factors that might be present in a particular HRI scenario. Specifically though, we reiterate that consent is one such modulatory factor, and an extremely important one at that. We believe the effect of consent is critical to evaluating robot perceptions, and on par with other influential factors such as gender effects.

Regardless of where one stands on the relative importance of the role of consent over other situational factors, it is fair to say that since consent can significantly impact all HRI interactions, it must be explicitly accounted for in our HRI experiments. One direction might be to redo some of the more significant past experiments to include some explicit consent measures and possibly even using consent/no-consent as an independent variable in the design. So, the robot-approach results could be redesigned by taking the presupposed consent structure of the context in which they are performed into account. On the one hand this might sound like bad news, but we think it is not so and will actually substantially strengthen those results that survive this repetition, and provide new opportunities for those results that show some variation when corrected for consent. Importantly, such a revisiting of past experiments might be necessary in order to be able to understand what those interaction design results actually mean when those interaction designs are translated into real world applications deployed in social spaces requiring consent-awareness. We might need to go back to the drawing board on some experiments to get additional results for cases not previously considered or conflated due to the lack of a consent dimension in the experimental design.

## 6 CONCLUSION

In this paper we have presented consent as an crucial area for HRI research. Consent often defines human-robot interactions for their entire duration, rather than being just a precondition that is established or refused beforehand. We have outlined levels of consent to show how subtle social cues can change the nature of what a social robots roles

are and what actions they can do. From several exemplary scenarios, it is clear that the HRI community can extend insights of existing research into richer areas of human-robot interface. Our brief architectural considerations suggest how robotic systems, designed with architectures incorporating rules-based and context-aware elements, could best recognize how consent is sought, maintained, and respected.

As robots enter into the rough terrain of social spaces, consent will unavoidably take on cultural and political valences. Even seen as tools, robot will assume representations of broader values, for better or worse. Societal critiques have stressed the harm of gendered norms of silence and disregard for consent, and steering clear of consent still implicates robotic systems in how they are embodied. This reflects why for broader discussions of AI ethics, including law and policy, HRI research can be a strong voice for otherwise neglected risks and opportunities posed by embodied automation. It can be at the forefront of guiding how human society can more carefully consent to the technologies meant to serve it.

## REFERENCES

1891. O'Brien v. Cunard Steam Ship Co.. In *28 N.E. 266*.
1980. Commonwealth v. Appleby. In *402 N.E.2d 1051*. Mass: Supreme Judicial Court.
2012. Commonwealth v. Carey. In *463 Mass. 378*. Mass: Supreme Judicial Court.
- Henny Admoni and Brian Scassellati. 2017. Social eye gaze in human-robot interaction: a review. *Journal of Human-Robot Interaction* 6, 1 (2017), 25–63.
- Thomas Arnold and Matthias Scheutz. 2017. Beyond moral dilemmas: Exploring the ethical landscape in hri. In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 445–452.
- Thomas Arnold and Matthias Scheutz. 2018a. Observing Robot Touch in Context: How Does Touch and Attitude Affect Perception of a Robot's Social Qualities?. In *Proceedings of the 13th ACM/IEEE International Conference on Human-Robot Interaction*.
- Thomas Arnold and Matthias Scheutz. 2018b. Observing Robot Touch in Context: How Does Touch and Attitude Affect Perceptions of a Robot's Social Qualities?. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 352–360.
- Peter M Asaro. 2007. Robots and responsibility from a legal perspective. *Proc. IEEE* (2007), 20–24.
- Siddhartha Banerjee and Sonia Chernova. 2017. Temporal Models for Robot Classification of Human Interruptibility. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1350–1359.
- Cynthia Breazeal, Kerstin Dautenhahn, and Takayuki Kanda. 2016. Social robotics. In *Springer handbook of robotics*. Springer, 1935–1972.
- Gordon Briggs and Matthias Scheutz. 2016. The pragmatic social robot: Toward socially-sensitive utterance generation in human-robot interactions. In *Proceedings of the AAAI Fall Symposium Series on Artificial Intelligence for Human-Robot Interaction (AI-HRI)*. 12–15.
- Gordon Michael Briggs and Matthias Scheutz. 2013. A Hybrid Architectural Approach to Understanding and Appropriately Generating Indirect Speech Acts.. In *AAAI*.
- Ryan Calo. 2010. Robots and privacy. Robot Ethics: The Ethical and Social Implications of Robotics, Patrick Lin, George Bekey, and Keith Abney, eds.
- Tiffany L Chen, Chih-Hung Aaron King, Andrea L Thomaz, and Charles C Kemp. 2014. An investigation of responses to robot-initiated touch in a nursing context. *International Journal of Social Robotics* 6, 1 (2014), 141–161.
- Kate Darling. 2016. Extending legal protection to social robots: The effects of anthropomorphism, empathy, and violent behavior towards robotic objects. *Robot Law, Calo, Froomkin, Kerr eds., Edward Elgar* (2016).
- Maartje MA de Graaf. 2016. An ethical evaluation of human-robot relationships. *International journal of social robotics* 8, 4 (2016), 589–598.
- Leah Fessler. 2018. Amazon's Alexa is now a feminist, and she's sorry if that upsets you. <https://qz.com/work/1180607/amazons-alexa-is-now-a-feminist-and-shes-sorry-if-that-upsets-you/>
- Stephen M Fiore, Travis J Wiltshire, Emilio JC Lobato, Florian G Jentsch, Wesley H Huang, and Benjamin Axelrod. 2013. Toward understanding social cues and signals in human-robot interaction: effects of robot gaze and proxemic behavior. *Frontiers in psychology* 4 (2013), 859.
- Lily Frank and Sven Nyholm. 2017. Robot sex and consent: Is consent to sex between a robot and a human conceivable, possible, and desirable? *Artificial Intelligence and Law* 25, 3 (2017), 305–323.
- John Gardner. 2015. The many faces of the reasonable person. *Law Quarterly Review* 131, Oct (2015).
- Sinziana Gutiu. 2012. Sex robots and robotization of consent. In *We Robot 2012 conference, Coral Gables, Florida. Retrieved April*, Vol. 15. 2013.
- Chris Haynes, Michael Luck, Peter McBurney, Samhar Mahmoud, Tomáš Vitek, and Simon Miles. 2017. Engineering the emergence of norms: a review. *The Knowledge Engineering Review* 32 (2017).
- Guy Hoffman, Jodi Forlizzi, Shahar Ayal, Aaron Steinfeld, John Antanitis, Guy Hochman, Eric Hochendoner, and Justin Finkenaar. 2015. Robot presence and human honesty: Experimental evidence. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 181–188.
- Heidi M Hurd. 1996. The moral magic of consent. *Legal Theory* 2, 2 (1996), 121–146.
- Manuscript submitted to ACM

- 1041 Marcello Ienca, Fabrice Jotterand, Constantin Viçă, and Bernice Elger. 2016. Social and assistive robotics in dementia care: ethical recommendations for  
1042 research and practice. *International Journal of Social Robotics* 8, 4 (2016), 565–573.
- 1043 Malte F Jung, Nikolas Martelaro, and Pamela J Hinds. 2015. Using robots to moderate team conflict: the case of repairing violations. In *Proceedings of the*  
1044 *Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 229–236.
- 1045 Peter H Kahn Jr, Takayuki Kanda, Hiroshi Ishiguro, Brian T Gill, Solace Shen, Heather E Gary, and Jolina H Ruckert. 2015. Will people keep the secret of a  
1046 humanoid robot?: Psychological intimacy in HRI. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*.  
1047 ACM, 173–180.
- 1048 Kheng Lee Koay, Dag Sverre Syrdal, Mohammadreza Ashgari-Oskoei, Michael L Walters, and Kerstin Dautenhahn. 2014. Social roles and baseline  
1049 proxemic preferences for a domestic service robot. *International Journal of Social Robotics* 6, 4 (2014), 469–488.
- 1050 Bertram F Malle and Matthias Scheutz. 2015. When will people regard robots as morally competent social partners?. In *2015 24th IEEE International*  
1051 *Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 486–491.
- 1052 Ross Mead and Maja J Mataric. 2015. Robots have needs too: People adapt their proxemic preferences to improve autonomous robot recognition of human  
1053 social signals. *New Frontiers in Human-Robot Interaction* (2015), 100.
- 1054 Ross Mead and Maja J Mataric. 2017. Autonomous human–robot proxemics: socially aware navigation based on interaction potential. *Autonomous Robots*  
1055 41, 5 (2017), 1189–1201.
- 1056 Jonathan Mumm and Bilge Mutlu. 2011. Human-robot proxemics: physical and psychological distancing in human-robot interaction. In *Proceedings of the*  
1057 *6th international conference on Human-robot interaction*. ACM, 331–338.
- 1058 Oskar Palinko, Kohei Ogawa, Yuichiro Yoshikawa, and Hiroshi Ishiguro. 2018. How Should a Robot Interrupt a Conversation Between Multiple Humans.  
1059 In *International Conference on Social Robotics*. Springer, 149–159.
- 1060 Jennifer Piatt, Shinichi Nagata, Selma Šabanović, Wan-Ling Cheng, Casey Bennett, MS Hee Rin Lee, David Hakken, et al. 2017. Companionship with a  
1061 robot? Therapists’s perspectives on socially assistive robots as therapeutic interventions in community mental health for older adults. *American*  
1062 *Journal of Recreation Therapy* 15, 4 (2017), 29–39.
- 1063 Laurel D Riek and Peter Robinson. 2011. Challenges and opportunities in building socially intelligent machines [social sciences]. *IEEE Signal Processing*  
1064 *Magazine* 28, 3 (2011), 146–149.
- 1065 Vasanth Sarathy and Matthias Scheutz. 2016. A Logic-based Computational Framework for Inferring Cognitive Affordances. *IEEE Transactions on*  
1066 *Cognitive and Developmental Systems* 8, 3 (2016).
- 1067 Vasanth Sarathy, Matthias Scheutz, Joseph Austerweil, Yoed Kenett, Mowafak Allaham, and Bertram Malle. 2017b. Mental Representations and  
1068 Computational Modeling of Context-Specific Human Norm Systems. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*.
- 1069 Vasanth Sarathy, Matthias Scheutz, and Bertram Malle. 2017a. Learning Behavioral Norms in Uncertain and Changing Contexts. In *Proceedings of the 2017*  
1070 *8th IEEE International Conference on Cognitive Infocommunications (CogInfoCom)*.
- 1071 Aziez Sardar, Michiel Jousse, Astrid Weiss, and Vanessa Evers. 2012. Don’t stand so close to me: users’ attitudinal and behavioral responses to personal  
1072 space invasion by robots. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 229–230.
- 1073 Matthias Scheutz. 2011. 13 The Inherent Dangers of Unidirectional Emotional Bonds between Humans and Social Robots. *Robot ethics: The ethical and*  
1074 *social implications of robotics* (2011), 205.
- 1075 Matthias Scheutz. 2017. The Case for Explicit Ethical Agents. *AI Magazine* 38, 4 (2017), 57–64.
- 1076 Matthias Scheutz and Thomas Arnold. 2016. Are we ready for sex robots?. In *Proceedings of the 11th ACM/IEEE Conference on Human-Robot Interaction*.
- 1077 Matthias Scheutz and Thomas Arnold. 2017. Intimacy, Bonding, and Sex Robots: Examining Empirical Results and Exploring Ethical Ramifications. In  
1078 *Robot Sex: Social and Ethical Implications*, John Danaher and Neil McArthur (Eds.). MIT Press.
- 1079 Kenneth Simons. 2017. Actual, apparent and hypothetical consent in tort law. [https://www.law.uci.edu/centers/clp/images-pdfs/  
1080 simons-actual-apparent-and-hypothetical-consent.pdf](https://www.law.uci.edu/centers/clp/images-pdfs/simons-actual-apparent-and-hypothetical-consent.pdf).
- 1081 Kenneth W Simons. 1987. Assumption of risk and consent in the law of torts: A theory of full preference. *BUL Rev.* 67 (1987), 213.
- 1082 Kenneth W Simons. 2006. A Restatement (Third) of intentional torts. *Ariz. L. Rev.* 48 (2006), 1061.
- 1083 Christopher John Stanton and Catherine J Stevens. 2017. Don’t Stare at Me: The Impact of a Humanoid Robot’s Gaze upon Trust During a Cooperative  
1084 Human–Robot Visual Task. *International Journal of Social Robotics* 9, 5 (2017), 745–753.
- 1085 Aaron Steinfeld, Terrence Fong, David Kaber, Michael Lewis, Jean Scholtz, Alan Schultz, and Michael Goodrich. 2006. Common metrics for human-robot  
1086 interaction. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*. ACM, 33–40.
- 1087 Megan Strait, Cody Canning, and Matthias Scheutz. 2014. Let me tell you! investigating the effects of robot communication strategies in advice-giving  
1088 situations based on robot appearance, interaction modality and distance. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot*  
1089 *interaction*. ACM, 479–486.
- 1090 Ja-Young Sung, Lan Guo, Rebecca E Grinter, and Henrik I Christensen. 2007. “My Roomba is Rambo”: intimate home appliances. In *International*  
1091 *Conference on Ubiquitous Computing*. Springer, 145–162.
- 1092 Dag Sverre Syrdal, Kheng Lee Koay, Michael L Walters, and Kerstin Dautenhahn. 2007. A personalized robot companion?—The role of individual differences  
on spatial preferences in HRI scenarios. In *Robot and Human interactive Communication, 2007. RO-MAN 2007. The 16th IEEE International Symposium*  
on. IEEE, 1143–1148.
- Andrea Thomaz, Guy Hoffman, Maya Cakmak, et al. 2016. Computational human-robot interaction. *Foundations and Trends® in Robotics* 4, 2-3 (2016),  
105–223.



- 1093 Cristen Torrey, Susan Fussell, and Sara Kiesler. 2013. How a robot should give advice. In *Proceedings of the 8th ACM/IEEE international conference on*  
1094 *Human-robot interaction*. IEEE Press, 275–282.
- 1095 Jan BF Van Erp and Alexander Toet. 2013. How to touch humans: Guidelines for social agents and robots that can touch. In *Affective Computing and*  
1096 *Intelligent Interaction (ACII), 2013 Humaine Association Conference on*. IEEE, 780–785.
- 1097 Michael L Walters, Kerstin Dautenhahn, Sarah N Woods, and Kheng Lee Koay. 2007. Robotic etiquette: results from user studies involving a fetch and  
1098 carry task. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*. ACM, 317–324.
- 1099 Ricarda Wullenkord, Marlena R Fraune, Friederike Eyssel, and Selma Šabanović. 2016. Getting in Touch: How imagined, actual, and physical contact  
1100 affect evaluations of robots. In *Robot and Human Interactive Communication (RO-MAN), 2016 25th IEEE International Symposium on*. IEEE, 980–985.
- 1101 B. Scassellati X. Tan, J. Brawer. 2019. That’s Mine! Learning Ownership Relations and Norms for Robots. In *Proceedings of The Thirty-Third AAAI*  
1102 *Conference on Artificial Intelligence (AAAI)*.
- 1103
- 1104
- 1105
- 1106
- 1107
- 1108
- 1109
- 1110
- 1111
- 1112
- 1113
- 1114
- 1115
- 1116
- 1117
- 1118
- 1119
- 1120
- 1121
- 1122
- 1123
- 1124
- 1125
- 1126
- 1127
- 1128
- 1129
- 1130
- 1131
- 1132
- 1133
- 1134
- 1135
- 1136
- 1137
- 1138
- 1139
- 1140
- 1141
- 1142
- 1143
- 1144 Manuscript submitted to ACM