Beyond consent: A framework for unified privacy management

Kevin L. Miller
Labyrinth Research, Gainesville, Florida, USA, klmiller@labyrinth-research.com

Abstract

The presence of robotic devices in our environment gives rise to unique privacy problems unlike those in other domains. Despite rapid advancement in the perception, movement, and learning capabilities of robots, issues in robot privacy remain without an effective research program. This research advances the conversation by proposing technological solutions aimed at the nexus between privacy as a legal and sociological concept and robot control in multi-actor environments. The following explores the system architecture and characteristics of a technical framework for making available, fusing and reconciling the privacy preference data of multiple actors across every contextual level (cultural, societal, group, locational, individual, and situational) and transforming them into concrete instructions usable by the robot as behavioral controls.

To that end, a taxonomic schema is described that can be accessed by robotic device makers to inform sensor collection, data collection, storage parameters and constraints, and the permissible range of movements, motions, and activities of a robot based on individualized, context-sensitive, and role-sensitive privacy preference rules. A privacy preference enunciator device and associated transport mechanisms are introduced that allow individuals and the robots they encounter in ad hoc environments to exchange privacy preference data in accordance with the taxonomic schema. Privacy preference rule selection and comprehensive resolution protocols are developed that allow for the automated or interactive resolution of conflicts arising between individuals in multi-actor environments or ambiguous contexts. Accountability and audit mechanisms are also discussed to encourage responsible adoption.

Keywords: privacy; robotics; blockchain; smart contracts.

1 Introduction & Motivation

In the near future, robots will increasingly populate the physical space in which humans live and work—perceiving our presence, observing and interpreting what we say and do, recording video, audio, and other sensor data, and physically interacting with us. Methods are needed to govern the observation, movement, and recording activities of a robot in accordance with the privacy expectations of the humans with which it interacts. The goal of this interdisciplinary research is to investigate technological approaches that lay the groundwork for addressing the very real privacy management challenges arising when humans coexist with robots and other devices.

A privacy management scheme for robots has several aspects that make it unique. In the first place, privacy management concerns stemming from the use of robots are qualitatively different from those encountered in conventional website and mobile device apps. The current paradigm of website and mobile app data privacy is incentivized by the lack of a viable economic model to monetize most web services and content publication. Thus, website and mobile app privacy tends to be defined by "notice and consent" modalities that are primarily concerned with obtaining broad permissions from consumers to sell their personal information or behavioral data to third parties for marketing purposes. Participants in this system have allowed this notion of information privacy and its associated notice and consent modality to define most aspects of the data privacy conversation, from its regulatory motifs to the design of the privacy setting user interfaces for giving or denying consent. Compounding this issue is the fact that, in privacy jurisprudence, people tend to be protected against privacy violations only when the intrusion is unreasonable or unexpected. The interplay of the notice and consent modality with the amorphousness of the "reasonableness" doctrine means that, over time, our "reasonable expectation" of privacy becomes inexorably eroded as individuals give blanket permission for web service providers to freely use our personal information in return for "free" use of their services and apps.

Robot privacy is a much harder and more nuanced problem than web privacy. It includes some classic information privacy concerns like those in website data sharing, but it must also account for physical privacy. "Physical privacy," as understood here, includes concepts such as whether a robot may measure and record a person's physical characteristics with sensors (e.g., audio recording or heart rate monitoring); a robot's physical proximity when interacting with a person in certain contexts; and whether, and in what manner, a robot can touch a person. These kinds of physical privacy are much more closely related to those protected by classic privacy torts such as "intrusion upon seclusion" and battery. Traditional notice and consent mechanisms, considered by many to be largely ineffective even within their own purview (*see, e.g.,* [1] [2]), are likely to be completely insufficient when applied to robot privacy management, which requires granular and scenario-specific restrictions on the range of actions a robot can take in a wide variety of environments, from assisting an elderly man in the shower to handing out brochures at a shopping mall.

Second, robot privacy management is fundamentally dynamic and contextual. Unlike in web-based privacy models, people and robots are mobile—robots can move into different physical spaces inhabited by different people, and different people can enter or exit a robot's functional proximity at any time. Privacy expectations are also based on cultural norms, shared group values, and even on physical location. Sometimes, situational contexts such as an emergency will override all other concerns. Thus, any proposed solution must facilitate a common consistent standard that assists robots in acting in alignment with our contextually-informed values.

Further complicating matters, privacy management becomes exponentially more difficult in real-world scenarios where robots must select appropriate governance actions to accommodate the potentially conflicting privacy needs of multiple people simultaneously occupying a home, workplace, or public space. Many of these individuals may be encountering a particular robot for the first time. Robots will be required to dynamically navigate a matrix of complex privacy settings, customs, culture, and personal needs and, in some cases, the robot may need to ask people nearby for clarification or mediate compromise positions in order to take effective action. Notably, robots share many of these problems with other categories of devices that cohabitate with humans, such as drones and "Internet of Things" devices; therefore, solutions to robot-specific problems have wide applicability spanning many device types.

Despite rapid advancement in the perception, movement, and learning capabilities of robots, issues in robot privacy remain without an effective research program. We propose that technological solutions aimed at the nexus between privacy as a legal and sociological concept and robot control in multi-actor environments can advance the conversation beyond the normative posturing engendered by the information privacy milieu. In that light, this research explores the characteristics of a technical framework for sharing individualized privacy preference data, including a formal taxonomic schema that can be accessed by robotic device makers to inform sensor collection, data collection, storage parameters and constraints, and the permissible range of movements, motions, and activities of a device. Transport mechanisms for distribution are discussed that allow individuals to "publish" privacy preference data in accordance with the centralized schema and robots to "subscribe to" that data when they encounter the individual. Privacy preference rule selection protocols and merger techniques are developed that allow for the resolution of rule conflicts that arise between individuals in multi-actor environments or ambiguous contexts. Accountability and audit mechanisms are also introduced.

2 Research Scope & Methodology

This paper considers the basic robot privacy problem from an essentially cybernetic viewpoint: aligning command and control of robots with human expectations in an environmental context. We develop a technical architecture, or "framework," necessarily incomplete but arrayed as a multi-pronged research agenda, to define structural concerns and implementation options that can assist in meeting the privacy challenges entailed by this new robotic environment. While predominantly a technical framework, this work uses legal and sociological understandings to design a model that exposes systemic assumptions and neutrally adapts norms to account for cultural and contextual subtleties.

More specifically, the objective is to ensure that robot control functions—namely, sensor activation and recording, as well as movement and action—meet the contextually sensitive privacy expectations of individuals co-inhabiting the robot's zone of influence. In light of the unique issues involved in robot privacy management, the research agenda is guided by the following questions and design constraints:

Q1: How can individualized privacy expectations that are sensitive of cultural/contextual nuance be indicated and communicated to robots?

- How do humans indicate their privacy expectations in a way that robots understand?
- How do humans communicate their privacy expectations to robots dynamically and in real-time?

Q2: How do robots use privacy expectations to take an appropriate action in a specific setting?

- How much can be automated to minimize configuration and interaction burden?
- How do we ensure that robots always have a path forward--i.e., are always able to make some control decision even in difficult ad hoc cases?

Q3: How can robots understand and respect the privacy expectations of multiple people in complex scenarios?

- How do we resolve conflicts in privacy expectations when multiple people are involved?
- What design mechanisms permit negotiation and consensus when the expectations of multiple people conflict?

Q4: How do we ensure that robots (and their designers) are accountable for their adherence to any proposed model? How do we gather consent, when needed, for legal purposes?

3 Technical Architecture

In order for robots to be reactive to the privacy expectations of individuals or groups and use them to inform and control their sensor activity, data collection, storage parameters and constraints, and the robots' permissible range of movements, motions, and activities, a framework for privacy preference data exchange is needed. Such a framework has several components, each of which may have multiple design options. Figure 1 shows a basic system architectural model of such a framework.

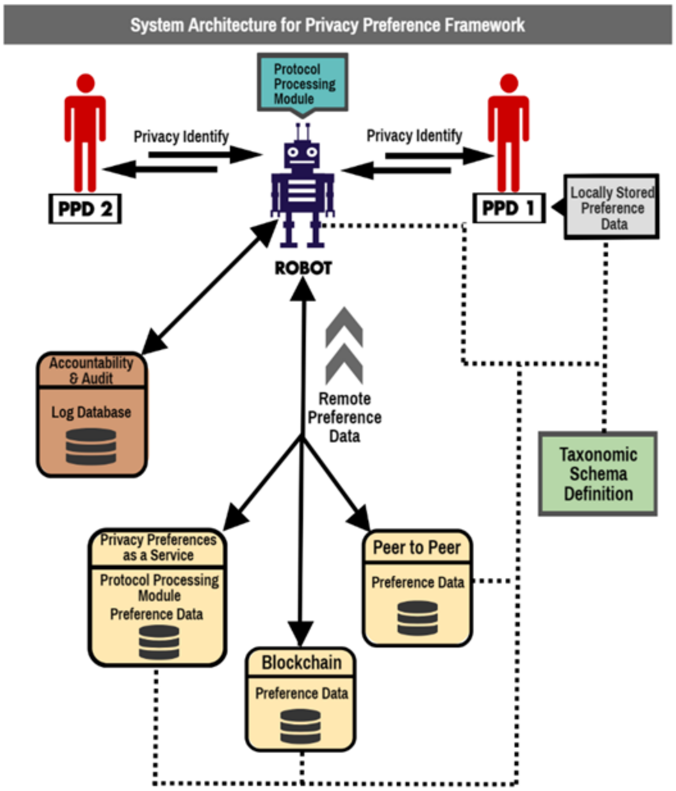
System Architecture for Privacy Preference Framework
Protocol Processing Module
Privacy Identify
Privacy Identify
PPD 2
PPD 1
Locally Stored Preference Data
ROBOT
Accountability & Audit
Log Database
Remote Preference Data
Taxonomic Schema Definition
Privacy Preferences as a Service
Protocol Processing Module
Preference Data
Peer to Peer
Preference Data
Blockchain
Preference Data

Figure 1: System Architecture

To allow robots to make ad hoc, real-time control adjustments to the privacy expectations of a dynamically-changing set of individuals, a mechanism is needed for recognizing the presence of individuals as they move in and out of a robot's sphere of influence. One way individuals may make their presence and privacy preferences known is by using a specialized device that serves as a personal enunciator or beacon, also called here a personal privacy device (PPD). A PPD may be embedded in a variety of hardware form factors, ranging from specialized wearable devices (e.g., a ring or other jewelry, clothing) to more generalized devices such as smartphones, smart watches, or fitness bands running enabling software. Signals sent by the enunciator are detectable by robots within a given range, or "detection zone," encompassing the robots' zone of activity or sensor capability. One potential mechanism for PPD communication uses Bluetooth Low Energy (BLE), a device-to-device networking technology supported on every major mobile and desktop operating system. In some cases "collective" enunciators may be used to indicate that the associated privacy preference data applies to every person present in a particular group, organization, or location. Collective enunciators may be valuable when the context governs the privacy environment, necessitating an override of individual preferences.

The actual privacy preference data remains unknown to robots in the detection zone until obtained from wherever it is being stored. In the most conceptually familiar model, robots use the privacy preference identity (e.g., BLE UUID) to retrieve the person or group's privacy preference data from a centralized cloud-based service. A disintermediated approach may be preferred when network performance is a concern, or to maintain data security or privacy isolation from centralized cloud service intermediaries. For example, the PPD may store privacy preference data locally and share it with nearby requestor robots using a local personal area networking technology, such as BLE, Zigbee, or DotDot. Disintermediation can also be achieved using peer-to-peer (P2P) networks, particularly the "blockchain." The blockchain is a public database or ledger of transactions consisting of "blocks" recorded by a network of nodes in chronological order. A distributed network of nodes can reach consensus on a particular state of affairs and record the consensus without any need for a controlling authority. The blockchain could be used by the privacy framework both to store privacy preference data and also to record the results of privacy-related decision-making for auditing.

3.1 Privacy Preference Taxonomic Schema

What comprises real-life privacy preference data? Individual privacy preference settings or rules that might inform an actual robot control instruction in a real scenario. For example, "For Alice, turn off audio recording when at a friend's house." This "privacy rule" defines a specific sensor state, for a specific person, applicable to all robots in a specific location-determined context. Let's examine the key design elements of a standardized taxonomic schema for preference data.

A recognizable structural form is apparent from the above example that may be generalized to the wider category of privacy preference settings relating to sensor control: SENSOR—SENSOR STATE—ROLE—CONTEXT. Although additional rule forms for robot movements/actions and other scaffolding will be necessary, this initial structural template for rule taxa provides a surprisingly useful generalization for robot sensor control with respect to privacy preferences.

A robot of any sophistication has dozens of sensors for a wide variety of purposes: to orient the robot in its environment, identify important objects or people, attenuate the force they apply when performing movements or other actions (e.g., grasping), determine the operating state of a machine or device they are controlling, and record the movements, actions, sounds, or other telemetry data of people or other entities for historical, accounting, behavioral analysis, and pure surveillance purposes. Depending on how and when a sensor is used and the duration for which it saves its sensor data, any sensor has the potential to violate a privacy preference. Each model of sensor being used in robotic devices collectively is a sensor taxon in the framework taxonomy. Each sensor model, in turn, has a set of specifications that describe its capabilities, operative ranges, and potential operating states. For representational efficiency, the sensor rule structure SENSOR may indicate an overall sensor capability such as "Video," "Audio," "Detect Motion," etc. A given robot will be able to process a rule in light of its own sensor capabilities, irrespective of the exact sensor model the robot is using.

The SENSOR STATE rule structure component represents an additional category of taxa for potential operating states of a sensor. The most simple and basic sensor state might be ON or OFF. More granular states for video sensors could include ORIENT-ONLY (camera is only used for the robot's own orientation purposes), RECORD PERMANENT (a permanent video record is stored on a robot's storage device), RECORD TRANSIENT [TIME] (the video record is kept for the designated time value, then disposed of), and RECORD OFF.

The CONTEXT component of rules reflects the reality that privacy choices are often context-dependent at a very granular level. Several conceptual layers of context are potentially important: **cultural** (cultural background of a region, religious affiliation, ethnic group), **societal** (economic system or political structure), **group** (voluntary or involuntary affiliation with a societal segment or group, such as charity, church, advocacy group, support group, political affiliation, formerly incarcerated), **locational** (home, place of employment, private meeting, friend's house, medical facility, region, country, state, city, country), **individual**, **situational** (ad hoc situations, emergencies or when security or safety of self/others is impacted), **legal** (constitutional, statutory, or regulatory constraint, e.g., compliance with a privacy law), and **trust relationship** (explicit or implicit interaction relationship with the robot entity, e.g., a robot one owns or that inhabits a place of employment; functional relationship, e.g., a personal care robot).

These layers may be organized hierarchically such that some layers are more generalized than others. Lower layers in the hierarchy can override the higher layers when rules are processed by robots to determine control behaviors. This enables the creation of powerful generalized default behaviors that can be used to apply sensible privacy controls to large groups with conceptual and representational efficiency. Meanwhile, the descending layers of hierarchy maintain the capability of the taxonomy to adapt to more granular choices at any arbitrary level of the context hierarchy. For instance, if culture is organized highest in the hierarchy of context layers, locational context rules, if present, override cultural rules.

A ROLE component allows rules to be representable in role-based forms, not merely identity-based forms. While a PPD and its associated rules represent the collective privacy expectations of a privacy identity, a rule itself pertains to the privacy expectations of a privacy identity's role in an identified context. This requires a taxonomic structure that can enumerate a privacy identity's roles, and the time, location, or other context in which those roles apply. For example, a possible role-qualified rule (defining recording states for a "doctor" role) might be: "FOR [ME] | VIDEO | RECORD PERMANENT | DOCTOR | IN HOSPITAL". In addition to accurately reflecting real-world scenarios, role-based rule capability has the additional advantage of enabling very generalized preference categories. This allows some rules to compactly represent privacy preferences applicable to large numbers of privacy identities. In fact, some rules may be generalized to such an extent that they encode valid "default" rules for certain contexts that largely apply to every privacy identity in that role-context.

The rule taxonomy has to be internally consistent and must be able resolve to a single outcome for every control state. In other words, the robotic control function outcome of any arbitrary collection of rules associated with the currently applicable context must be computable and decidable. "Decidability," as understood here, means that all control state processing paths in the privacy preference rule set applied to a single privacy identity lead to exactly one outcome, even if that outcome is a "default" rule at a very high level in the context hierarchy. In simple design terms, this means that robots can always take some action, and not freeze, crash, act erratically, or choose such a poor course of action that it disappoints the expectations of all concerned.

Additional action rule structures will be needed for a complete taxonomy, since some privacy expectations are met only when control can be exercised over a robots' proximity, motility, location, limb movements, functions, and conversation topics or interruptions that may be considered intrusive.

3.2 Merger and Resolution Protocol

A challenge and a design constraint of the proposed framework is that it be operative in environments where multiple privacy identities may exist. It is foreseeable that conflict states will occur frequently. How should rules from multiple privacy identities that direct the same robot control functions be merged, evaluated for consistency, and, if necessary, reconciled?

A robot must be able to take some action, even if the robot has to engage in a process of conflict reconciliation that includes dialogue, negotiation, and consensus-seeking with or between the relevant persons. The overall goal the merger and resolution protocol is to enable a robot to take appropriate action in spite of the inevitable situational ambiguities and conflicts between the privacy expectations of multiple actors that will arise in every day usage. The merger and resolution protocol's secondary design considerations include: (1) the protocol directs a robot's control behaviors toward contextually normative "best-case" outcomes whenever possible; (2) conflicting actors have the chance to discuss, clarify, and amend their privacy preference positions when necessary; (3) notice of sub-optimal control outcomes, as well as the opportunity to engage in mitigating actions, is provided to actors whose preferences were subjugated; and (4) the protocol acts autonomously, whenever possible, after considering the risk of harm.

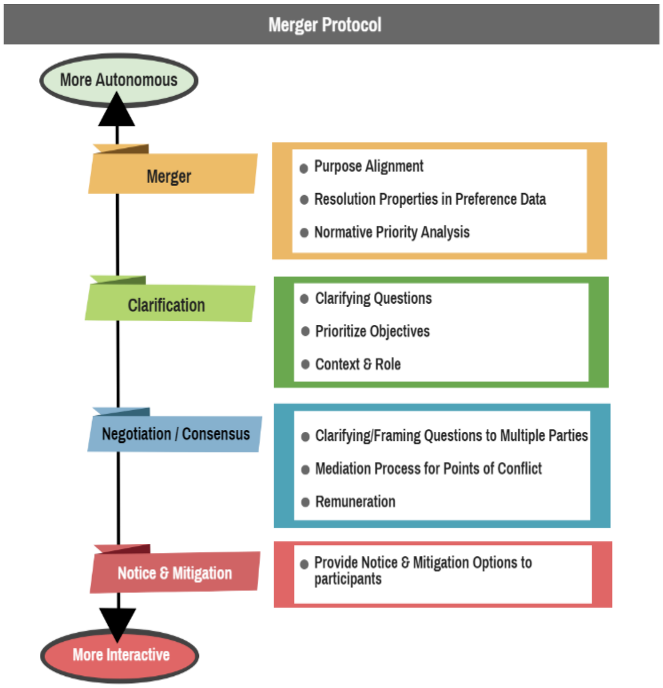
Merger Protocol
More Autonomous
Merger
Purpose Alignment
Resolution Properties in Preference Data
Normative Priority Analysis
Clarification
Clarifying Questions
Prioritize Objectives
Context & Role
Negotiation / Consensus
Clarifying/Framing Questions to Multiple Parties
Mediation Process for Points of Conflict
Remuneration
Notice & Mitigation
Provide Notice & Mitigation Options to participants
More Interactive

Figure 2: Merger and Resolution Protocol

In multi-actor environments, conflicts in compound rule sets, where two or more rules with matching sensor, role, and context content direct a robot to set conflicting or contradictory sensor states, are determined dynamically in real time. *Figure 2* shows a prototypical merger and resolution protocol that illustrates actions a robot can take to resolve the control paralysis induced by preference conflicts, grouped roughly according to the amount of guidance or intervention required of participating human actors. While *Figure 2* shows a broad default order, a robot's order of stepping through the actions of the resolution protocol may vary in each specific situation depending on such factors as the context, the actors' roles, the exact nature of the preference conflict, and the resolution option directives in an actor's privacy preference data.

Purpose alignment examines the robot's purpose in the environment, as well as the purpose of the specific interaction. First, a robot is presumptively in any given space for a discernable reason—what is the reason? For example, is the robot there to perform a small range of tasks (such as bringing tea or vacuuming), provide home health care, provide companionship, or greet the public? Second, does the control conflict block a specific robot activity that directly relates to the robot's reason to be there? How necessary is that specific activity to the robot's overall purpose? The robot can then use this background context to resolve conflicts by eliminating conflicting rules that clearly subvert its purpose.

Resolution properties are another conflict resolution stratagem that allow individual privacy identities to indicate their willingness to resolve certain kinds of control conflicts automatically. Resolution properties of the taxonomy can be used both to express the range and order of acceptable conflict resolution methodologies, as well as to indicate quantitatively the "importance" of a privacy preference rule to an actor. Resolution properties can be assigned to various taxonomic identifiers in the hierarchy—actor/identity, context, role, as well as to individual rules. They have an expressed hierarchy of priority in that properties attached to the higher levels are more general, and also more overridable, than properties assigned to lower levels of the hierarchy.

Some conflicts may be simplified by or completely automatable with the assistance of **normative priority analysis**. Normative statements generally take the form of a rule of conduct phrased as an imperative; for example, the "Golden Rule" ("Treat others as you would like to be treated.") is a normative statement. Normative statements claim how things "ought" to be and, by doing so, can serve as shortcuts to cut through ambiguities so that the robot can take action. Though a few normative statements may posit "universal" beliefs of humankind, most statements are likely to be based on cultural, national, or religious worldviews. Some normative rules may identify an implicit or explicit hierarchy of priority based on the privacy identity from which the conflicting rules originate. Encoding the priority rights as normative constructs allows them to be exposed and analyzed according to their cultural or legal context. The framework's taxonomic schema should support normative rules in automated resolution analysis that are applied from these perspectives in light of the robot's current operating context.

Autonomous methods may fail to completely resolve conflicts between actors' privacy preferences, either because additional clarifying information is needed from an actor about the situation, or because the robot needs to communicate the conflict and give the actors the option to negotiate a resolution.

An ability to ask **clarifying questions** allows the robot to request additional information about a human actor's goals or priority of objectives. A robot can ask further questions when it is having trouble discerning the current context from ambient conditions, when it is uncertain which taxonomic context maps to the current context, and when an actor's role in a given context may not be immediately apparent from preference data provided on the actor's PPD. A robot can also vocalize and receive assent to a putative control action when common knowledge assumptions seem to be violated in a given case. Clarifying communications can even enable an actor to spell out the modalities of resolution she is willing to undertake to resolve conflicts.

The basic strategy of **negotiation and consensus** protocols is to facilitate a conversation between parties whose privacy expectations have conflicts. Negotiation/consensus protocols can be considered to have two sub-stages. In the first sub-stage, a robot asks any clarifying/framing questions that require answers from multiple actors. These questions generally seek information that illuminates ambiguities in the operating context or conflicting knowledge assumptions. They also "frame" the conflict so that underlying presuppositions about the context or other matters are made apparent to the human actors. In the second sub-stage, the robot facilitates a mediation process between the conflicting parties. To do so, the robot may use mediation techniques common in alternative dispute resolution proceedings for legal matters or other fields of study.

Remuneration for consensus is a potential resolution option, as it is in many negotiation scenarios. Since the robot has arrived almost at the end of its available options to resolve conflicts, offering an actor the opportunity to yield for value is arguably more fair than the alternative, which is to subjugate the privacy preferences of one or more actors in order to be able to take some kind of action. Remuneration provides robots a powerful mechanism to impart a tangible benefit on actors who cannot otherwise agree. Whether individuals should be able to monetize their privacy in this way is debatable, but rather than engage in that discussion, we make two observations. First, people already monetize their privacy this way when they use free apps and web services in return for giving broad latitude to sell their personal information and behavioral data to advertisers. Second, our resolution rubric places monetization options at the end, not the beginning, of a number of options aimed at voluntary consensus. Indeed, an impetus for this framework is to move privacy preference selection back toward a more socially responsible model than the broad website-style notice and consent methods that will likely be used by robot manufacturers if no such framework is developed.

Who would pay? Robot manufacturers have an interest in their robots respecting the privacy of their customers and third parties, and may reward consumers to encourage consensus and avoid control conflicts. Alternatively, the payor might be the actor whose privacy preferences "win" the conflict. An actor-payor system allows those actors who highly value a given preference outcome to encourage consensus in their favor. To minimize unfairness resulting from wealth disparities, an actor-payor system could be built on a platform of micropayments wherein the exchange medium is not real money but "privacy credits." Actors could use privacy credits gained during negotiations where they were more flexible about the outcome to fund their desired outcomes during other negotiations. A decentralized value exchange platform based on blockchain smart contracts is an intriguing architectural foundation on which to implement a privacy credit micropayment system.

In some circumstances, a robot may have failed to resolve preference conflicts using autonomous or negotiated modalities. The robot's remaining resolution strategy is to determine the "best" final control state, during which it may revisit normative rules that were considered during earlier processing stages. When a control state is chosen, the robot provides **notice** and, if possible, **mitigation actions** (e.g., "leave the room") for those actors whose preferences were unavoidably subjugated.

3.3 Accountability: Logging and Audit

Commenters have increasingly come to believe that autonomous systems need mechanisms whereby their decision-making processes are more transparent to outside agents (*see* [3]). Robust accountability mechanisms are essential given the number of interacting inputs and the complexity of the decision-making processes as a robot moves through rule determination, merger protocols, and action determination. Moreover, any specific instance of rule conflict processing may involve negotiation and consensus between actors who make ad hoc decisions or compromises, give consent to override previously indicated preferences, or provide additional information the robot must accept as true. This ad hoc external information must be documented not only for liability purposes, but also because the legal constraints governing privacy demand actor notice and consent for some kinds of privacy intrusions.

Manufacturers participating in the privacy preference framework agree to comply with responsible use policies and to design their robots to act in compliance with its outcomes. In most cases, compliance auditing is confined largely to analyzing whether the control actions the robot took fit the parameters. The codified nature of the processing, as well as of the instrumentation, means that autonomous systems can perform most audits automatically. Manufacturers who fail audits are subject to incrementally escalating disincentives. Draconian as compliance auditing may seem, the dynamic and ad hoc nature of processing in real-life environments means that many times privacy outcomes will be conditional on actors' responses and other externalities. Therefore, heavy instrumentation of these externalities plays an important role in a robot device manufacturer's ability to show compliance for liability purposes.

4 Conclusion and Future Work

This paper began with the thesis that the omnipresence of robotic devices in our environment gives rise to unique privacy problems unlike those in other domains. To solve those problems, we delineated a technical framework that includes: the overall components of a system architecture, taxonomic data structures for mapping privacy preferences to robot control states, a processing protocol for enacting robot control states and identifying preference conflicts between multiple actors, automated and interactive methods of resolving those preference conflicts, and accountability and audit mechanisms.

Looking ahead, additional research and coordinated action is needed. Newly emerging architectures, such as the blockchain, can be used as a platform backbone for disintermediated privacy preference data storage. Privacy preference data exchange could then be conducted using "smart contracts" on the blockchain backbone, both to ensure robot control functions are implemented as expected and to provide an architectural basis for the negotiation and remuneration features of this framework. Although our taxonomic structure and processing architecture are built hierarchically so that a functioning rule set can be developed from very few rules, an approach to the problem of populating the data structures of the framework—particularly the taxa for context hierarchies, roles, and the default control state rules that associate with them—must be selected from the available strategies or developed anew. Usable methods of enabling individuals to understand and self-select their own rules are also needed.

Many benefits will accrue from approaching the robot privacy problem from the coordinated perspective described here. First, for individual users, the proposed model presents a comprehensive technological structure that empowers people to make choices reflecting their cultural and personal values. The PPD-enunciator concept helps individuals maintain and configure consistent privacy preference settings. Instead of defaulting to lowest common denominator approaches, robots can use these personalized settings to achieve optimum privacy expectation alignment even in ad hoc and unforeseen privacy scenarios, and even when multiple actors are involved. A rich and interactive resolution protocol ensures that users are able to engage in participatory and adaptive notice, negotiation, and consent activities when conflicts do arise.

Second, manufacturers have much to gain from a standardized approach to meeting individuals' privacy expectations. Public comfort that robots are attempting to be respectful of their privacy choices in most situations will help stave off the perception that robots are surreptitiously watching and recording them and do much to drive marketplace adoption of these technologies. Furthermore, government agencies are likely to become increasingly active in regulating consumer privacy in robot devices, and consumers are increasingly likely to litigate against manufacturers for privacy harms. By actually approaching privacy design in a detailed and thorough manner—a design that respects people's choices, documents their consent, and in which the robot chronicles and is accountable for its decision-making—robot device-makers erect a powerful natural defense against regulators and other legal challenges.

Society and policy stand to gain from an approach enabling a renewed conception of privacy that re-centers itself on the individual. Instead of the privacy-eroding blanket "consent" paradigm that has grown up around current web monetization models, a new paradigm matures—one where individuals are able to make fine-grained choices about privacy expectations that are specifically and painstakingly enacted by robot devices. An individual's option to make a choice and have it respected, in itself, re-energizes the legal conversation about the "reasonable expectation" of privacy in almost any domain. In other words, one way of slowing privacy's seemingly inevitable erosion is to provide a method of contextually reactive granular control whereby a person can say "I expect to have privacy here, here, and here—I refuse to give it all away in one broad permission slip." Rather than making one big decision to permit everything, this gives us the chance to make numerous contextually-based decisions that keep some interactions private. Our approach creates another type of empowerment of the individual arising from the remuneration model for resolution of conflicts. Since our design attempts to be contextually-neutral about normative analysis, value judgments, and rule hierarchies, it encodes those evaluative and processing constructs in a way that exposes hidden assumptions and biases. These accountability attributes represent a significant improvement to the way that automated systems conventionally operate—as "black boxes" that provide no insight into their contrivances.

References

[1] Schaub, F., Balebako, R., Durity A. L., & Cranor, L. F. (2015). A Design Space for Effective Privacy Notices. In USENIX Assoc. Symposium on Usable Privacy and Security.

[2] Cranor, L. F. (2012). Necessary But Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice. J. on Telecomm. & High Tech. L., vol. 10, 273-308.

[3] Diakopoulos, N. (2016). Accountability in Algorithmic Decision Making. Comm. of the ACM, 59(2), 56-62.