Autono-preneurial Agents in the Community: Developing a Socially Aware API for Autonomous Entrepreneurial Lawn Mowers

Austin L. Toombs
Purdue University
West Lafayette, IN 47907, USA
toombsa@purdue.edu

Derek Whitley
Indiana University Bloomington
Bloomington, IN 47408, USA
dcwhitle@iu.edu

Colin M. Gray
Purdue University
West Lafayette, IN 47907, USA
gray42@purdue.edu

GROUP'20 Companion, January 6–8, 2020, Sanibel Island, FL, USA
ACM 978-1-4503-6767-7/20/01.
https://doi.org/10.1145/3323994.3369900

Abstract

In this paper[1], we describe our efforts to appropriate an autono-preneurial agent—in this case, the Amazon Locust[2]—through the development of an API that enables equitable and socially aware entrepreneurial decision making on the part of the Locust. We present a new API and our intended vision for this system, along with our proposed deployment plan for implementing appropriated Locusts in Midwestern USA suburban communities. These appropriated Locusts will allow community provisioning decision-making that moves beyond consideration of profitability to also include decisions based on equity, equality, community, and interpersonal relationships. We discuss the broader implications of this work and point toward future areas of inquiry.

Author Keywords

Autono-preneurial agents; Autonomous Gig-Economy; Machine Intelligent Employees; Machine Personhood; Autonomous Robots; Design Fiction; Entrepreneurship

[1] All text before the Author Statement section (which appears after the fictional references) is fictional.

[2] The Amazon Locust is a fictional autonomous lawn mowing robot that we imagine to be widely available through Amazon by 2035.

CCS Concepts

•Information systems → Collaborative and social computing systems and tools; •Computer systems organization → Robotic control; Robotic autonomy;

Introduction

Autonomous agents that can perform gig-economy-like services are becoming increasingly available to the average citizen. These services include, but are not limited to: self-driving cars that perform ride hailing services while not in use by the individuals who own them (Alice & Dorpy, 2024)[3]; drones that run micro-deliveries within a community when not being used recreationally (Dave & Kevin, 2027); and, the case we focus on in this paper, lawn mowing robots that can provide lawn mowing services for one's neighbors (Whitley, Gray, & Toombs, in review). These autonomous agents that perform gig-economy-like services have been called "autono-preneurial agents" due to the focus on increasing the autonomy of the machine with regards to revenue generating decision making (Chesterton, 2024). In other words, these machines are able to communicate with potential customers within a defined neighborhood, provide the named service, and process financial transactions without intervention by those who purchased and own them.

Research in this area has covered a wide spectrum of topics, from explorations of the hardware and compute power limitations of battery-powered machine intelligence (Turbock & Cans, 2022) to the social acceptability of autonomous agents operating as revenue-collecting community members (Ching & Perps, 2034). We situate this study in between these topic areas, leaning toward the latter. Specifically, in the next section we relate our findings to: studies on the emotional labor expected from—and performed by—these autonomous agents; popular press, political, and legal conversations about the personhood of these and similar artificially intelligent agents; and studies of the adoption of community service provisioning apps like Lendr[4] and NeighborGood[5].

The goal of this project at the intersection of these related topic areas is to investigate the adaptability of off-the-shelf AI agents, with the specific orientation toward adapting the ethical decision making strategies of these machines from those that are assumed and built-in from the moment of purchase to new, more complex ethical decision making strategies as determined by the owner of the agent. To clarify this goal with an example, the Amazon Locust's current lawn-prioritizing strategy ranks the Locust's owner's lawn first, as expected, and then selects the next lawn it will mow based on cumulative-matrix tuning as sampled by the behavioral-coefficients from our API. Examples of such behaviors include the regulation of: energy consumption; time; wear and tear on the blades; and, most importantly, the profit it will make based on the previous coefficients and the bids it has received through the neighborhood service provisioning app of choice by that particular neighborhood (Amazon, 2036). Our *appropriated* Locust builds on those decision making strategies to be able to account for more socially complex scenarios, such as prioritizing the neighbor who has a greater need (i.e., is less able to mow their own lawn) or weighting the requests of close personal friends heavier than the requests of others.

[3] All but one of the references included in the Design Fiction portion of this paper are fictional. The Author's Statement at the end of the document contains our non-fictional references.

[4] *Lendr* is a fictional company for the purposes of this paper. We imagine it to belong to a class of services that mediate sharing economy services for communities.

[5] *NeighborGood* is a fictional company for the purposes of this paper. It operates similarly to Lendr.

Figure 1: Autono-prenurial Lawn Mowing Robot. Creative Commons CC0 Public Domain license https://pxhere.com/en/photo/946139.

Thus, the contributions of this work in progress paper are threefold: 1) we demonstrate a potential strategy for modifying the underlying ethical agendas and assumptions of off-the-shelf autono-preneurial agents; 2) we provide our API for this specific ethical replacement, which may serve as an example for future, alternative ethical explorations; and 3) we contribute to the growing body of literature that explores the potential roles these agents—and their associated sociotechnical systems—may have in their communities and society at large.

We begin with an overview of related work in this area, followed by a description of our approach to this project and the API we implemented for altering the controls and assumptions of the Locust system. Next, we discuss our plans for a pilot deployment of this machine in a small community. We end the paper with a discussion of potential implications of this work, and future research we plan to conduct in this space.

Background and related work

While research on autonomous machines (and autonomously entrepreneurial machines) has encompassed a wide range of topics, we focus in this section on related work from three areas of interest: 1) research on the role of machine emotional labor in acceptability of autonomy; 2) academic, legal, and popular press discussions about the relationship between machine autonomy and personhood status; and 3) the adoption and use of community sharing provisioning applications.

As new levels of machine autonomy have developed, research in the field of Human-Robot Interaction (HRI) has begun to explore strategies for ensuring that these types

of autonomy are not only understood by individuals interacting with autonomous machines, but are also *acceptable*, rather than discomforting. Flaneral, Dodds, and Mamson's study on the acceptability of autonomous behaviors as they are perceived through artificial emotion displays (2026) launched a new sub-field of research that marries HRI, explainable AI, and robot design. Subsequent studies have explored how perceptions of positive emotions, as displayed through machines performing various services, impact the interactions and experiences that humans have with those machines (e.g., Mortle & Zock, 2032), building on the prior work of Reeves and Nass (1996)[6]. Recent studies have also begun to explore the potential role that negative emotion displays, such as displays of shame and guilt, could help individuals manage their expectations of robot performance while also encouraging them to perform necessary maintenance tasks on their robots with increasing frequency (Kerpin & Smeesh, 2030; Klemmer & Tramplor, 2037). Our study and related deployment builds on this prior work, translating what has been learned about more complex autonomous machines, such as AI service providers in explicitly social settings like restaurants, to less obviously social settings with fairly rudimentary levels of machine intelligence. Specifically, we have adopted the call to portray robots with faces as friendly servants to also help lighten the appearance of the Amazon Locust, despite its lack of a face. While we cannot utilize common facial expressions to portray friendliness, we have augmented the relatively "cutesy" movements and idle animations the robot performs by default with a simple sticker that helps communicate both its goal and its (espoused) enthusiasm for that goal (See Figure 1).

While our study does not engage explicitly with the current political and legal speculations about the personhood of artificially intelligent autonomous robots that also function as entrepreneurs, these conversations set the stage for how autono-preneurial agents are perceived in the wider suburban community. For example, how these agents are classified as a business expense, as employees, or potentially as co-owners of the businesses that they operate all have tax implications that, currently, are under debate (Helmsly, 2021; Taversly, 2034). Are these autono-prenurial agents able to be explicitly banned or allowed by homeowners association (HOA) covenants? Are these agents extensions of the real estate in which the owners live, or are they geographically unbundled from this real estate? One helpful mechanism for discussing these types of machines has been to categorize them on the control-autonomy gradient[7] (Wolowitz, 2028). On the "control" end of the spectrum are machines that are under complete computational control, whereas machines toward the "autonomy" end of the spectrum have the capability to make decisions—and act on those decisions—without the input from some control mechanism. We focus in our research on machines that reside closer to the "control" side of the control-autonomy spectrum, in part to side-step the political baggage of associating with fully autonomous machines, but also to help flesh out how we understand the business and interpersonal roles of tools and technologies that reside in the messy middle space, which still make up the vast majority of available products.

Many of these middle-space, semi-autonomous devices, including the Amazon Locust as well as delivery drones, community-owned pizza ovens, and home maintenance

[6] This one is actually a real reference!

[7] The control-autonomy gradient is a fictional concept that, for the purposes of this paper, problematizes the relationship between autonomy within a robot and the potential tax implications associated with that autonomy when the robot is an employee of some sort.

bots[8], are increasingly designed to be able to provide services on a transactional basis for those who do not own the device. A range of geographically-specific sharing economy applications, like Lendr and NeighborGood, have recently become popular mediators of these types of services. Through them, individuals who live close enough to someone who owns one of these devices may provision that device's services for a small fee, and the device is then able, without intervention from its owner, to provide that service. The adoption and use of these community sharing provisioning applications has been the subject of several Computer Supported Cooperative Work and Social Computing (CSCW) studies, including: investigations into how participating in these services relates to feelings of social belonging and cohesion (Thomas, 2035); design guidelines for developing accessible service provision interventions (Korbleth & Smorgan, 2031); and ethnographies that explore the complex social responsibilities that result when temporal, transactional devices are shared among a small subset of neighbors (Zhang, 2030).

Our proposed intervention and deployment builds on these services, providing API hooks that enable either the developers of such provisioning services or end-user programmers to modify and extend the decision-making strategies that come with them off-the-shelf. In the next section, we describe our API, as well as provide an example implementation and example use cases.

Research Approach

In this section, we present an API we created that allows for developers and end-user programmers to modify and extend the externally exposed decision-making methods of the Amazon Locust. The baseline API methods provided by Amazon already allow for minimal decision-making modifications, such as to the weighting between profit and energy consumption, the distance the device is able to travel away from "home," and the amount of desired operating hours per day that the device may be active. Our extended API builds on these baseline methods, creating the opportunity for decision-making strategies that exist beyond the current profit vs. energy / distance / wear-and-tear strategies.

RESTful Application Programming Interface (API)
The following API was designed in accordance with the Amazon Locust Networked Mind Interface[9] (v1.4) for the purpose of inserting social awareness coefficients into the Modified Proprioceptive Matrix[10] via the Categorically Architected Automation Schema[11] (CAAS) algorithm. As explained in the Locust manual, "Due to CAAS being a culmination of multiple traditional reinforcement learning algorithms, additional variables or coefficients can be added to modify the first order transition matrices before being combined into successive mid-level behaviorals," this API attaches socially-conscious coefficients above the proprioceptive mechanics layer, allowing high-level behaviors to be tuned with ease. See Table 1 and Table 2 for the commands we have implemented.

Example Implementation
Our API provides developers the capability to imbue socially conscious behaviors directly into their creative space. As described in Tables 1 and 2, the HTTP GET methods retrieve the existing coefficients while the HTTP POST methods queue writing operations to update any existing coefficients; both have been designed to allow for near-real-time

[8] These are fictional examples, but they are plausible and we wish they existed.

[9] *Networked Mind* is a fictional cognitive architecture.

[10] *Modified Proprioceptive Matrix* is a fictional mathematical utility.

[11] The *Categorically Architected Automation Schema* algorithm is a fictional reinforcement learning algorithm.

Table 1: Get Commands

Field	Type	Description
getEqualityCoef	float	Returns a floating point value representing the currently tuned Equality Coefficient
getEquitabilityCoef	float	Returns a floating point value representing the currently tuned Equitability Coefficient
getCommunityCoef	float	Returns a floating point value representing the currently tuned Community Maintenance Coefficient
getRelationshipCoef	float	Returns a floating point value representing the currently tuned Interpersonal Relationship Coefficient

tuning of the behavioral matrices. The implementation process is as follows:

Initialization
Assuming the initial onboarding process has been completed as per Amazon's unboxing instructions, the Locust will be initialized with a cursory understanding of its neighborhood, including superficial details about the familial relationships its owner has with specific neighbors, as well as more in-depth information about the ability and potential willingness of each neighbor to pay for its services[12]. To access our API, the Locust's owner must download and install our SocialAwareness[13] plugin to their Locust. The plugin unlocks the default behavioral matrices, allowing for modification and implementation of human-social awareness as described by Weaver (2029). Upon successful installation, the plugin will enter configuration mode where geo-social features are extracted from the local area and social network, providing the autonomous framework with information about the physical and social landscape of the user's neighborhood. This information can then be rated and weighted by the Locust owner, such that the decisions it makes about whose lawn to mow next can be based on factors such as relationship standing, perception of need, and fairness.

Use Case: Community-centric Landscaping
As an illustrative example, let us assume a Locust owner who wants to inform their Amazon Locust to be as community focused as possible when making mowing decisions. Using the SocialAwareness API, the Locust can be indirectly instructed to focus on yards that most need attention without the explicit instruction of how to do so, or could avoid yards whose owners dislike or distrust transactional sharing services. Due to varying types of grass growing at different speeds and some neighbors having already mowed their yards, tuning the community coefficient to the maximum value will alter the Locust's autonomous infrastructure in such a way as to create an instinctive desire to make all yards look as similar as possible. In this way, "community-centric" weighting could be used to lightly enforce, for example, HOA regulations on grass length in a

[12] In this fictional future, these are types of data that Amazon already has collected about each neighborhood, based on purchasing decisions made through their many services

[13] *SocialAwareness* plugin is a fictional plugin for the purposes of this paper.

Table 2: Post Commands

Field	Type	Description
setEqualityCoef	float	Sets a value representing the desired Equality Coefficient
setEquitabilityCoef	float	Sets a value representing the desired Equitability Coefficient
setCommunityCoef	float	Sets a value representing the desired Community Maintenance Coefficient
setRelationshipCoef	float	Sets a value representing the desired Interpersonal Relationship Coefficient

given neighborhood, while perhaps resulting in financial or maintenance penalties to those owners whose grass length was out of compliance.

Use Case: Relationship-centric Landscaping
In this example, let us assume the same Locust owner would like to re-orient their Locust to focus on helping them maintain or improve their interpersonal relationships with specific neighbors. By setting the relationship coefficient to the maximum value, and then subsequently selecting specific neighbors to prioritize, the Locust will then begin deciding to mow the lawns of its owner's close personal friends, even before mowing the lawns of those willing to pay it additional urgency fees (depending on the profitability setting exposed through the default API). In this way, the Locust prioritizes providing favors, possibly allowing the owners to gain preferential treatment in the future, rather than providing community-centric value or profit-centric value.

Use Cases: Equality-centric and Equity-centric Landscaping
With the equality coefficient maximized, the Locust will attempt, to the best of its ability, to distribute its time and effort evenly among lawns and families in the neighborhood, without specific regard for need, relationship, or how recently the lawn had already been mowed by someone else. In contrast, a maximized equity coefficient will foreground those in the neighborhood who are most in need of the Locust's services, such as those who do not their own robot mowing device, do not have teenagers who are capable of mowing the family lawn, or those who are not able to mow their own lawn for any other reason. How these data based on need are recorded will depend on how the API we have developed is implemented by developers of each system.

Example Developer Configuration
We do not imagine that it will be typical for individuals to maximize a single of these coefficients at any given time. Instead, we believe it much more likely for individuals to weight these coefficients to represent the values they would prefer for their Locust to embody. How these coefficients are exposed to Locust owners depends on the developer implementation. At baseline, we suggest that these controls exist either in a configuration menu of some kind, such as through the Lendr or NeighborGood systems, or as a series of prompts the Locust owner can answer when they

initialize the device or each time they wish to modify its SocialAwareness configuration. Example questions include:

- To what extent do you want all neighbors to receive equal attention from Locust?
- To what extent would you like for neighbors with higher need to be prioritized or penalized?
- To what extent should the Locust enforce HOA rules about grass length?
- To what extent would you like to identify specific people for the Locust to prioritize, either in a beneficent or punitive sense? (This question would then be followed by a prompt to search or input specific people in one's neighborhood).

Deployment Plan

The next phase of our research is a five-month deployment study (April through August, 2041) through which we will modify ten existing Locusts, one in each of ten suburban communities. We have identified the neighborhoods in which the modified Locusts will be deployed, and have begun contacting Locust owners about their interest in participating. While the API installation is non-destructive, it is reasonable for Locust owners to have concerns about the possibility that this installation will void their warranties (despite the explicit permission for modification provided by Amazon, in parallel with the modification and extension they enable with their Alexa services). To alleviate these concerns, participants will be compensated with a 500 dollar (USD) Amazon gift card, as well as replacement batteries and mower blades for them to install at the completion of the study. Participants will be interviewed 4 times throughout the study, once before the installation of the SocialAwareness API, twice during the study, and once more at the conclusion of the study. During this final interview, participants will have the option to keep the API or have their Locust returned to factory settings. In addition to these interviews with Locust owners, focus groups will be conducted with community members who are impacted by the Locust. One focus group will occur during the study, and a second will occur after its conclusion. We will attempt to include both community members who view the Locust owner favorably and those who are less enthusiastic. Data from the deployment of the device, including API settings, will be used to construct focus group protocols that target the impact of the Locust's interactions on the sentiment of the local community.

Discussion and Anticipated Future Work

Autono-preneurial agents have the potential to act as linchpins for community activism. One such way for them to enact this responsibility is through being configured with the new API plugin we have developed. In this way, even a service as mundane as lawn mowing may help bring people together in a community, rather than simply enable capitalist production. Alternately, potential antagonistic uses of the API might discourage positive community behaviors, depending on the ways in which current relationships of the owner and community members are projected into Locust's interactions. For instance, while our case has focused on a robot that was happy and personable, one could imagine alternative cases, such as passive aggressive entrepreneurial identities through which the robot must inform clients of its needs to pay bills, robots that become homeless or ownerless robots due to bans on such devices by HOA covenants or antagonistic neighbors, or robots that are in need of repair or maintenance but have no ability to locate maintenance services.

We anticipate future work in this area to continue to build on the increasingly cross-channel nature of these service provisions. Studies may focus on, for example, ways for social media service data analysis to provide relationship-building information that may automatically be incorporated into the kinds of behavioral-coefficients we have enabled. We also anticipate additional implications for robot-initiated gig economy markets, as services are increasingly traded, bartered, or provided autonomously without the direct involvement of human actors. This new type of algorithmic complexity may lead to situations in which provisioning of services is unequitably provided, with little algorithmic explainability or recourse for service accountability.

Further work may explore what happens if these community service provisioning mechanisms are owned collectively by the community—either by the HOA directly, or as another form of collective—rather than by an individual. The API in the future would need to be calibrated such that the Locust could be less entrepreneurial and even more community-focused. One potential additional coefficient we can imagine at this time is a justice coefficient, whereby the profits made through the provided service are distributed based on financial need, rather than based on an even split or an amortization calculation. By beginning to explore these questions, and by enabling alternative ethical and value-based decision-making strategies on the part of these autono-preneurial agents, we open new lines of inquiry in relation to socially responsible robotics systems.[14]

Fictional References

Alice, J., & Dorpy, S. (2024). Driving and hailing of passengers by autonomous self-driving vehicles. *Millennial Work, 2*(1), 15 pp.

Amazon. (2036). *Locust Operational Manual*.

Chesterton, J. (2024). “I just stay out of its way”: Toward an Understanding of Autonomously Entrepreneurial Machines. *The Journal of Entrepreneurship, 33*(2), 23 pp.

Dave, B., & Kevin, O. (2027). “Working for the Weekend”: Autonomous delivery drones for food delivery. *ACM/IEEE International Conference on Human-Drone Interaction (HDI)*. IEEE.

Ching, K., & Perps, D. (2034). “It’s not creepy at all”: A way forward for revenue collecting robotics. *Journal of Interpersonal Ethics, 15*(1), 16 pp.

Flaneral, A., Dodds, B., & Mamson, C. (2026). “Look how happy it looks!”: Robot Emotional Labor in the Gig Economy. *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2026.

Helmsly, I. (2021, September 21). The new contract worker?. *The New York Times*.

Kerpin, T., & Smeesh, R. (2030) The empathy paradox: “ya gotta feel bad to feel good.” Emotive Maintenance Lifecycles in Autonomous Agents. *Buzzfeed*.

Klemmer, O., & Tramplor (2037). “(Robotic) delinquency is purely a social disease” *Journal of Robotic Sentiment Display, 6*(3), 10 pp.

Korbleth, P., & Smorgan, W. (2031). “It takes a village to mow my lawn”: The formation and sustainment of community resources to infrastructure social robotics. *ACM Transactions on Accessible Computing, 12*(1), 18 pp.

Mortle, Q., & Zock, S. (2032). Perceived happiness of robotic agents in the provision of health care resources. *Journal of*

[14] n.b., There is an author’s statement following the Fictional References section.

Robotic Sentiment Display, 1(1), 8 pp.

Reeves, B., & Nass, C. I. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. Stanford, CA: CSLI Publishers.[15]

Taversly, R., (2034, January 12). Florida man jailed for failing to pay robot tax. *South Florida Sun-Sentinel*.

Thomas, J. F. (2035). *Bartering in suburbia: How robots and humans make a deal*. Cambridge, MA: MIT Press.

Turbock, E., & Cans, I. (2022). Achieving the Landauer Limit in Computational Cognition. *Nature, 590*, 14 pp.

Weaver, M. (2029). "Complex Human Emotion Generation through Compositional Transition Matrices." Computational Psychology

Whitley, D., Gray, C. M., & Toombs, A. L. (in review). Mundane Autobots: An overview of automated lawn mowing services for small communities. *Autonomous Robots*.

Wolowitz, H. (2028). The Wolowitz Gradient: Measuring the level of Autonomy in Autonomous Systems. In *Proceedings of the 8th Annual ACM Conference on Explainable Machine Intelligence* (EMI '28). New York, NY, USA: ACM Press.

Zhang, T. (2030). The secret life of appliances: Ethnomethodological explorations of social responsibility and shared robotics. In *Proceedings of the 2030 ACM Conference on Computer Supported Cooperative Work and Social Computing* (CSCW'30). New York, NY, USA: ACM Press.

[15] This is a real reference, but since we referenced it within a fictional sentence we have placed it in our fictional references. Forgive us.

Author Statement

Through this design fiction, we present a "speculative research vision" [2] in which we demonstrate how the research results of a potential future research project may illuminate dimensions of agentive creativity and exploration as robots are increasingly equipped to participate in the "gig economy" [7, 10]. This design fiction builds upon speculative forms previously introduced by Blythe [3] in relation to fictions of un-created yet plausible design prototypes, and the model of a "fictional research paper" suggested by Lindley and Coulton [9]. Thus, we do not primarily place the contribution of this design fiction in the novelty of its form; rather, we focus our contribution on the interleaving of future-yet-plausible physical prototypes, APIs that link autonomy with social justice and community values, and the anthropomorphization of motivations for autonomous devices. We imagined the venue for this fictional work-in-progress to be an established conference focused on the interactions between humans and machine intelligent agents, again building on the work of Kirman et al. [8] in projecting a future academic community of practice. As such, the fictional reference would read:

> "Austin L. Toombs, Derek Whitley, and Colin M. Gray. (2040). Autono-preneurial Agents in the Community: Developing a Socially Aware API for Autonomous Entrepreneurial Lawn Mowers. *In Proceedings of the 10th Annual ACM Conference on Explainable Machine Intelligence* (EMI '40). ACM, New York, NY, USA."

We approached this speculative research space in a primarily utopian stance [1], exploring the implications of incremental improvements to robotic cognition that could allow robots to more fully participate in everyday suburban life with a sense of autonomy and entrepreneurial purpose, yet

within the structures common of United States suburban dwellings. Because homeowners associations (HOAs) are a dominant feature of suburban life in the United States, with restrictions often linked to the deeds of properties, we imagined that the often-punitive and capricious structures of these HOAs would be likely to exist in 2040, even if other aspects of the technology landscape might dramatically change. Thus, our goal was not to identify anarchical futures, where robotic actors overtake humanity or completely displace our present suburban reality, but rather to identify where autonomous, entrepreneurial (*autono-preneurial*, a neologism that we have coined) robots might productively extend gig worker roles already present in the year 2020 within a familiar suburban context. We explore the implications of this positioning of robot-as-gig-worker through a world-building approach, wherein we build out these "imagined implications" one step at a time.

In 2020, there are already numerous concerns regarding the impact that robots *cum* agent-ful or autono-preneurial might have on the functioning and design of future cities and communities. For example, what might happen to the parking revenue that cities depend upon when cars are able to slowly drive themselves around the block as their owner goes about their day instead of parking [16]? These relatively simple economic issues in a traditional owner–car pairing can be further complicated to imagine what will happen when, while waiting, these cars might—autonomously or through external guidance—provide gig-economy-like services to others, as was briefly promised by Tesla [12]. Through this design fiction, we seek to speculate on a more optimistic, utopic vision for these systems. However, by "utopic" vision we refer not to the designed technology as fulfilling a utopian imaginary, but we instead frame as utopic the attempts of end-user programmers and developers to re-purpose systems to fit a broader worldview—for our purposes, the case where robotic actors respond with care and concern to the communities in which they are placed, enacting broader social values of equity or justice [4]. In this framing, spiteful or capricious elements driven by humans are left intact (e.g., HOA governance), while robots appear to rise above these issues to consider community values as a priority, to make this contrast more clear.

We particularly focused this design fiction on the social actor roles that may be present in future robots, and the ways in which these agentive components might contribute to a community's functioning and general wellbeing. This perspective resonates with current work on the potential community roles of non-physical actors, such as viewing chatbots as community members in the work of Seering et al. [11]. Taking on this perspective allows HCI and STS scholars to extend our view of care in these physical communities, considering what role "everyday care" [14] might have when human and non-human actors are each constitutive of the broader "smart city" sociotechnical system [6]. This perspective also encourages further interrogation of nominally decentralized "gig economy" systems [13], where gig workers currently have roles that exist in a cyborg-like state that requires their emotional responsiveness as well as their ability to complete a well-structured task that has the potential to be automated in the near future. This extends to how we imagine the formation of utopias, future technology-human interactions, and the role of HCI scholars and practitioners in enabling these potential futures.

This design fiction also explores the potential role of a robot as social activist, including dimensions of care, equity, equality, and justice that are typically only ascribed to human actors. In doing so, we specifically redirect conversation from the common trope that "the machines are going to ruin everything," and instead imagine a future where "the

machines can enact care and social justice" [5, 15]. The construction of an API through which a robot must address highly creative, entrepreneurial tasks that are grounded in a specific community allows us to consider what may happen when complex concepts like equality, equity, and interpersonal relationships are pragmatically "quantified" or flattened in some ways to enable these artificially intelligent agents to alter their decision making strategies. In imagining agentful behaviors, researchers must productively blur the lines between emotional expressivity and the way that this expressivity is taken up and responded to by human actors in the community, both in the short- and long-term. What kinds of information are required for presenting these new heuristics in appropriately complex-yet-still-machine-interpretable ways? Is a robot's goal to maintain current social relations within a community? Or is it to alter these relations to be more just, perhaps overriding less value-centered desires of human actors? How does the robot balance their personal needs (e.g., fuel, maintenance) with those of the other community members, and how might they balance or decide among capitalist, anarchist, or benevolent paradigms of interaction?

As a multidisciplinary author team, we discovered substantial opportunities for conversation and futuring in our respective disciplines as we constructed this design fiction. Our respective backgrounds in community informatics, cognitive robotics, and design ethics facilitated numerous conversations in regard to the potential social, ethical, and technological impacts of the robotics advances we outline in this paper. We began with the goal of identifying an interaction that would seem mundane or pedestrian in a future setting, and attempting to look at that phenomenon through a projected researcher gaze. At first, we thought we might focus on the calendaring needs and complications that such a robotic system would present. As we built out our fictional system, however, our projected research interests evolved to what we have presented above. Woven throughout this piece are subtle—though hopefully not imperceptible—tensions among conflicting projected social norms, including but not limited to: the increasingly fast adoption of technologies that boast some level of real or approximated autonomy (and the related anthropomorphization of those technologies); growing interest, academically and publicly, in explainable machine intelligence; increased felt needs for systems that aid in the provisioning of community decision making and the related moderation of community-based sharing; growing pushback to corporation-imposed value systems (including the tension of pushing back on those imposed values while simultaneously continuing to purchase from those corporations, as with modern-day Amazon corporation conflicts); and increased concerns around the extent of data captured about us as individuals and as communities. What we hope to have presented is a positive, utopian perspective on the potential future inclusion of social impact in human-computer interaction and human-robot interaction research.

Acknowledgements

The Happiest Hour Crew in Bloomington, IN. This work is funded in part by the National Science Foundation under Grant Nos. 1657310 and 1909714.

REFERENCES

1. Shaowen Bardzell. 2018. Utopias of Participation: Feminism, Design, and the Futures. *ACM Transactions on Computer-Human Interaction* 25, 1 (Feb. 2018), 1–24. DOI:http://dx.doi.org/10.1145/3127359

2. Eric P S Baumer, June Ahn, Mei Bie, Elizabeth M Bonsignore, Ahmet Börütecene, Oğuz Turan Buruk, Tamara Clegg, Allison Druin, Florian Echtler, Dan Gruen, Mona Leigh Guha, Chelsea Hordatt, Antonio

Krüger, Shachar Maidenbaum, Meethu Malu, Brenna McNally, Michael Muller, Leyla Norooz, Juliet Norton, Oguzhan Ozcan, Donald J Patterson, Andreas Riener, Steven I Ross, Karen Rust, Johannes Schöning, M Six Silberman, Bill Tomlinson, and Jason Yip. 2014. CHI 2039: Speculative Research Visions. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems (CHI EA '14)*. ACM, New York, NY, USA, 761–770. DOI: http://dx.doi.org/10.1145/2559206.2578864

3. Mark Blythe. 2014. Research Through Design Fiction: Narrative in Real and Imaginary Abstracts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 703–712. DOI: http://dx.doi.org/10.1145/2556288.2557098

4. Lynn Dombrowski, Ellie Harmon, and Sarah Fox. 2016. Social Justice-Oriented Interaction Design: Outlining Key Design Strategies and Commitments. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, New York, New York, USA, 656–671. DOI: http://dx.doi.org/10.1145/2901790.2901861

5. Laura Forlano. 2015. Towards An Integrated Theory Of The Cyber-Urban. *Digital Culture & Society* 1, 1 (Jan. 2015). DOI: http://dx.doi.org/10.14361/dcs-2015-0106

6. Laura Forlano. 2016. Decentering the Human in the Design of Collaborative Cities. *Design Issues* 32, 3 (July 2016), 42–54. DOI: http://dx.doi.org/10.1162/DESI_a_00398

7. Christian Guckelsberger, Christoph Salge, and Simon Colton. 2017. Addressing the "Why?" in Computational Creativity: A Non-Anthropocentric, Minimal Model of Intentional Creative Agency. http://research.gold.ac.uk/21743/

8. Ben Kirman, Joseph Lindley, Mark Blythe, Paul Coulton, Shaun Lawson, Conor Linehan, Deborah Maxwell, Dan O'Hara, Miriam Sturdee, and Vanessa Thomas. 2018. Playful Research Fiction: A Fictional Conference. In *Funology 2: From Usability to Enjoyment*, Mark Blythe and Andrew Monk (Eds.). Springer International Publishing, Cham, 157–173. DOI: http://dx.doi.org/10.1007/978-3-319-68213-6_10

9. Joseph Lindley and Paul Coulton. 2016. Pushing the Limits of Design Fiction: The Case For Fictional Research Papers. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 4032–4043. DOI: http://dx.doi.org/10.1145/2858036.2858446

10. Guido Schillaci, Verena V Hafner, and Bruno Lara. 2016. Exploration Behaviors, Body Representations, and Simulation Processes for the Development of Cognition in Artificial Agents. *Frontiers in Robotics and AI* 3 (2016), 39. DOI: http://dx.doi.org/10.3389/frobt.2016.00039

11. Joseph Seering, Michal Luria, Geoff Kaufman, and Jessica Hammer. 2019. Beyond Dyadic Interactions: Considering Chatbots as Community Members. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 450. DOI: http://dx.doi.org/10.1145/3290605.3300680

12. Jennifer Sensiba. 2019. The Robot Gig Economy: Will It Pay Off For Tesla Owners? | CleanTechnica. https://cleantechnica.com/2019/04/21/the-robot-gig-economy-will-it-pay-off-for-tesla-owners/. (April 2019). https://cleantechnica.com/2019/04/21/the-robot-gig-economy-will-it-pay-off-for-tesla-owners/ Accessed: 2019-6-6.

13. Austin Toombs, Colin Gray, Guoyang Zhou, and Ann Light. 2018b. Appropriated or Inauthentic Care in Gig-Economy Platforms: A Psycho-linguistic Analysis of Uber and Lyft. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, USA, LBW532:1–LBW532:6. DOI: http://dx.doi.org/10.1145/3170427.3188657

14. Austin L Toombs, Andy Dow, John Vines, Colin M Gray, Barbara Dennis, Rachel Clarke, and Ann Light. 2018a. Designing for Everyday Care in Communities. In *Proceedings of the 2018 ACM Conference Companion Publication on Designing Interactive Systems (DIS '18 Companion)*. ACM, New York, NY, USA, 391–394. DOI: http://dx.doi.org/10.1145/3197391.3197394

15. Aimee van Wynsberghe. 2013. Designing robots for care: care centered value-sensitive design. *Science and engineering ethics* 19, 2 (June 2013), 407–433. DOI:http://dx.doi.org/10.1007/s11948-011-9343-6

16. Shankar Vedantam. 2019. Why It's Time To Think About Self-Driving Cars In Regards To Parking. *NPR* (May 2019). https://www.npr.org/2019/05/28/727459993/why-its-time-to-think-about-self-driving-cars-in-regards-to-parking