Designing a minimalist socially aware robotic agent for the home

Matthew R. Francisco¹, Ian Wood¹, Selma Šabanović^{1,2}, and Luis M. Rocha^{1,2,3}

¹School of Informatics and Computing, Indiana University, Bloomington, IN 47406, USA

²Cognitive Science Program, Indiana University, Bloomington, IN 47406, USA

³Instituto Gulbenkian de Ciencia, Oeiras, Portugal

francm@indiana.edu

Abstract

We present a minimalist social robot that relies on long timeseries of low resolution data such as mechanical vibration, temperature, lighting, sounds and collisions. Our goal is to develop an experimental system for growing socially situated robotic agents whose behavioral repertoire is subsumed by the social order of the space. To get there we are designing robots that use their simple sensors and motion feedback routines to recognize different classes of human activity and then associate to each class a range of appropriate behaviors. We use the Katie Family of robots, built on the iRobot Create platform, an Arduino Uno, and a Raspberry Pi. We describe its sensor abilities and exploratory tests that allow us to develop hypotheses about what objects (sensor data) correspond to something known and observable by a human subject. We use machine learning methods to classify three social scenarios from over a hundred experiments, demonstrating that it is possible to detect social situations with high accuracy, using the low-resolution sensors from our minimalist robot.

Introduction

In 2003, Rodney Brooks suggested that "by 2020 robots will be pervasive in our lives" (Brooks, 2002, p. 113). As an example of this trend, he conceptualized an autonomous robotic vacuum cleaner for the home with a bottom-up design which allowed the robot's behaviors to emerge in interaction with its physical environment. The robot used the amount of light it sensed as a measure of the dirtiness of the floor, readings from its bump sensors as a signal to change direction, and its cliff sensors to know when to stop so as not to fall down stairs. Without full knowledge of the physical environment, the robot could randomly cover yet fully clean a wide variety of floors. The *iRobot Roomba* robotic vacuum, commercialized in 2002, is the materialization of Brooks' idea; more than 10 million Roombas have so far been sold worldwide.¹

As a robust, commercially available robotic product, the Roomba was one of the first robots to be used naturalistic and long-term studies of human-robot interaction in the home. These pioneering studies ascertained that, along with the physical environment, the *social context* also had an effect on the cleaning robot's ability to function successfully. Researchers described that domestic Roombas were given names (Sung et al., 2007) and treated as "social agents" (Forlizzi, 2007); the use of Roombas also had a reciprocal effect on the social organization and practices in the home, as men and teenagers participated more in domestic cleaning chores. Such findings call attention to the importance of understanding the social as well as the physical dynamics of the context of use for robotic products.

Inspired by the Roomba as a commercial and social product, we propose that future robotic technologies that can coexist and collaborate with people in everyday environments should have a sense of the social as well as physical contexts in which they operate. Contemporary robots are largely ignorant of the social significance of their actions and of the bustle of human life around them. As robots spend more time around humans, they will profit from being able to take advantage of the social as well as the physical characteristics of the environment to support their successful functioning (Dautenhahn et al., 2002).

Artificial Life has contributed greatly to the development of situated robots whose behavior emerges from the nonlinear interaction between machine and environment (Almeida e Costa and Rocha, 2005). But just like human cognition and social intelligence is extended into the environment (Clark, 1998), robots can use the bottom-up principles of artificial life to develop social competency. Uexkull's concept of *umwelt* (the self-centered sensorial world of animals) has served as a guiding principle to generate robot behavior that is grounded on their own perception-action interaction with an environment (Hoffmeyer, 1997). While the concept of an *umwelt* makes the case for personal sensory experience, Uexkull (2001) notes that it gives us a way to understand sociality as an *intersubjective* process rather than a subject-object dualism:

... the idea of an objective universe, that embraces all living things, is undeniably very useful for ordinary life. The conventional universe, where all our relationships to our fellow human beings are enacted, has

¹http://www.irobot.com/en/us/Company/About.aspx?pageid=79

brought all personal Umwelt spaces under a common denominator, and this has become indispensable for civilized human beings. Without it, we cannot draw the simplest map, because it is impossible to combine all subjective points of view in a single picture (p. 109).

Similarly, we now need to develop social umwelten for robots. By this we mean that, rather than designing robotic social behavior in a top-down manner, we need to develop it from the bottom-up in a manner that is most consistent with the robots own sensors and its interaction with the social environment (including designers). This also means that the robot and humans must have more ways of recognizing, remembering, and building upon intersubjective experiences. Within this multi-agency environment the idea is to develop umwelt overlap between robotic and human agents that scaffolds cooperative action (Ferreira and Caldas, 2013).

We describe an initial approach to developing social awareness and presence for a domestic robot using minimalist robots, a combination of simple sensors, and bio-inspired computational techniques to develop a social unwelt for domestic robots. We seek to develop socially situated robotic agents whose behavioral repertoire is subsumed by the social order of the space. The goal is to recognize different classes of human activity and then associate to each class a range of appropriate behaviors. While human-robot interaction research has largely focused on developing algorithms that use high resolution data such as audio and video, our system relies on long time-series of low resolution data such as mechanical vibration, temperature, lighting, sounds and collisions. We rely on low-resolution data because that is the reality of the sensors in the robot platform we use (see below). This allows us precisely to test if such cheap sensors, which are immediately and widely available, are capable of developing minimal social awareness.

We begin the paper with a discussion of the Roomba and its relationship to the social and cultural models of the home. We then introduce the robot and test how it experiences the world through its sensors. These exploratory tests allow us to develop hypotheses about what objects (sensor data) correspond to something known and observable by a human subject. In the third section we use machine learning methods to classify three scenarios from over a hundred experiments involving human interaction. We conclude with some future directions for our research.

Navigating social spaces

The Roomba is one of the first instantiations of robots that work in everyday human environments with untrained users. Technology corporations and governments around the world expect that such technologies will proliferate and provide a new era of technological and economic production. This future direction for robot development is highlighted by the US National Science Foundation's National Robotics Initiative, which funds the development of co-robots that "work beside, or cooperatively with, people... acting in direct support of and in a symbiotic relationship with human partners" (National Science Foundation, 2013). We use the Roomba as a model case for studying how robotic technologies might be socially integrated into human environments.

Roomba's design, inspired by Brooks' subsumption approach to artificial intelligence, provides a robust and workable solution to issues posed by diverse and constantly changing human environments. Studies of Roomba's use in actual homes, however, have pointed out that there are important challenges and resources in the environment that the Roomba's design does not take into account. For example, while the Roomba is designed to function in rooms of different shape, size, and organization, it is limited in the kinds of terrain it can cover. Therefore owners need to adapt the home to their Roombas by moving furniture, objects, and moving them between different levels of a house (Forlizzi, 2007). More relevant to our goals, Roomba's current limitations in terms of social awareness also limit its use; users may find the Roomba's random coverage of space may disrupt their activities, and therefore turn the Roomba on only when they are away from home or at specific times.

Existing research on robots in the home has shown that even functional robots like the Roomba are interpreted in social ways when they are situated in social environments (Forlizzi, 2007). The robot's emergent interactions with the social organization, cultural norms, political dynamics, and people's interpretations within the social environment can have a significant effect on whether the robot is accepted or rejected by users (Mutlu and Forlizzi, 2008). We have also shown that users envision and evaluate robots and other technologies in the context of the social hierarchies and relationships they regularly inhabit (Lee and Sabanovic, 2013). Initial research on service robots further suggests that personalization is an important component of robotic functions in open-ended human environments (Forlizzi and DiSalvo, 2006)

The idea that human interpretations of a robot's functioning can support its behavioral repertoire was suggested decades ago in Braitenberg (1984) description of robotic "vehicles" whose simple behaviors in relation to the environment evoke ascriptions of affective and cognitive meaning by human observers. Several social robotic projects, such as Keepon (Kozima et al., 2009), Muu (Matsumoto et al., 2005), and PARO (Shibata, 2012), have embraced the possibility of communicating social presence and agency through simple relational cues. Alač et al. (2011) has shown that the social agency of robots is not constructed solely, or even primarily, through their functional capabilities, but is scaffolded through the ways in which human actors orient themselves towards the robot.

This research in human-robot interaction corroborates prior theories in human-computer interaction that describe everyday contexts as not only physical, but social and cul-



Figure 1: Katie in operation.



Figure 2: Sensor locations on a Katie robot.

tural spaces that hold both personal and shared meanings for their human inhabitants, which are constructed through embodied interaction (Dourish, 2001). We extend this understanding of how space relates to the design of robots and human-robot interaction by developing minimalist robots that are becoming aware of the social and cultural characteristics of their environment.

The Katie family social robot

In this section we give a description of our robots and the kind of sensorial world in which they operate. The graphs of the sensor data in this section are meant to demonstrate the flow of data from the robots that can be interpreted by most people. Can the robot tell the difference if it is next to a wall versus in the middle of a room? If a person greets the robot, which sensors exhibit change and how much? The graphs of this data are a resource of the robot's social umwelt as it is the first place where both humans (researchers) and our robots establish a common denominator together with an object in the world. Most importantly, graphs such as these help us to understand the limits of real time sensor responses, graphic visualization of data, and humans linking those to meaningful environmental changes. To translate further between robot and human will require more sophisticated techniques, which we discuss in the next section.

Robot design

The robots we are using, called the *Katie family*, are built on the *iRobot Create* platform, an *Arduino Uno*, and a *Raspberry Pi*. The Create has several native sensors and we have included more sensors to the Arduino (see Figure 2). The Raspberry Pi collects images which are used for coding data and verifying our classes. It also relays data to and commands from our server. Our research group can begin and start experiments remotely and download data from each experiment. These basic parts comprise the embodiment of the family and gives each their unique umwelt.

We chose this minimalist design for Katie to keep all components lightweight, low power, and housed inside the cargo bay of the Create. This allows the robot to reside in a location for longer time spans and to explore (or seek shelter) under objects such as chairs, couches and tables. Retaining size and weight of the Create also maximizes mobility of the platform in physically tight or socially constrained indoor spaces, allowing the robot to do things that humans don't normally do in a space, such as getting on the ground to look under objects or to look closely at the ground and baseboards of a space.

The sensorial world of Katie

The results from multiple hour tests demonstrate the resolution of the sensor data, behavior of the sensors, and the starting range of patterns that comprise the robot's unwelt. For designers, tests such as these help to decide where to place sensors and when to focus on a given data stream or not. Because there is so much complexity in even some of our simplest spaces we give the robot a basic scanning behavior, which is a 30 degree pivot every 20 seconds.

For all of our tests in this section and experiments in the next section, the following sensors are examined: two infrared (IR) range sensor aimed rearwards (integer)²; a photovoltaic cell for sensing light (integer); two IR thermometers (floats)³; and Create platform internal sensors, namely 4 integer IR sensors around the bumper to detect cliffs, a wall IR sensor, and 5 boolean wheel and bumper sensors. This results in a total of 15 sensor variables. An observation is recorded approximately once every tenth of a second.

²The Sharp GP2Y0A02YK0F has a range of 15 to 150 cm and Sharp GP2Y0A41SK0F has a range of 4 to 30cm.

³This sensor, the Melexis MLX90614, report hundredth's of a degree resolution, with an accuracy of \pm .5°C for most ranges in room temperature and \pm .1°C in human body temperature range.



Figure 3: Infrared distance sensor data from two positions. The medium and long range sensors (p3 and p4) disagree little. In future tests the long range sensor p4 will face upward to detect when the robot is under furniture and possibly detect when in a door frame.

The robots have multiple infrared distance sensors. There are three outward facing sensors on the robot each with a different range. Figure 3 shows two 90 minute runs of data collected from a robot positioned in an open floor with no objects nearby and one positioned next to a wall. On the Create there are an additional four distance sensors that are pointed toward the ground for cliff detection. The short range on these allow for detection of small position changes if, for instance, the robot is moved or if the floor moves.

There are two infrared thermometers positioned on the robot at an upward angle of 45 degrees. Each face in opposite directions. Data from a four hour test in Figure 4 shows two cyclical patterns. The first cycle is from the rotation of robot in its scan behavior where we see a $1^{\circ}F$ difference in the sensors at one point in each rotation.

The light sensor has a slight angle towards the front of the robot. An angle on this sensor give some direction of where light is coming into a room. Figure 5 shows data when the robot is in a sunny location in a house. Depending on the intensity of light in a specific direction, fast changes in the light can be caused by humans or animals casting a shadow on the sensor. Very fast changes, like an activation of an electric lamp, cause sharp and consistent patterns from the light sensor reading while changes in the light coming through windows caused by clouds are smooth and stable.



Figure 4: Infrared thermometers pick up changes as the robot faces a new direction every 20 seconds and when the furnace turns on in the house.

While the robot is stationary, the bump sensors are our the most reliable detector of presence of people and animals. The bump sensor is a large plastic bumper that covers the entire front end of the robot with a switch on each side. When an object depresses the right side the right switch is triggered. Contact with the center of the bumper triggers both sensors. All three states (contact with left side is bump = 1, right side is bump = 2, and center is bump = 3). The wheel drop sensors are also switches. These trigger when the wheels extend all the way down (each wheel has spring to force them down when they are off the floor) are are therefore reliable for detecting when the robot is lifted.

Finally each Katie has a high resolution camera. The main purpose of the camera is to annotate the data with socially meaningful categories and to verify if the robot's classifications make sense. In the current form the camera is angled slightly upward. Robot mounted cameras can collect potentially sensitive or embarrassing information and this increases within a private space like a home. Since Americans are becoming increasingly aware and concerned about privacy we design into the robots some deference to privacy by mounting the camera with an angle toward the ground. Lowering the gaze is an embodied signal that shows deference to broader cultural concerns. The images offer another affordance that will be crucial as the robot learns more about the household and begins to move easily within it. Images can be used by people in the home to identify meaningful objects and places and such annotations can used to respond to a richer understanding of the environment.

Classifying social situations

As a first step in developing the social capabilities of the system, we would like to know if a Katie can detect differences in simple human-robot interactions. The scenarios we wish to discriminate between are: (0) an empty room, (1) someone walking across a room and (2) someone walking around the robot. Scenarios 1 and 2 are represented in figure 6. The



Figure 5: Light can provide some directional context if the photo sensor has a slight angle and the light source is intense. The picture was taken from the robot while it was collecting data in the graph.

robot is placed in the center of a room with measurements of distance from the robot taped to the floor. Each scenario is run at 5 different proximities: contact with the robot, 1-20 cm, 21-40 cm, 41-60 cm, and 61-80 cm. Each scenarioproximity condition is run 10 times, for a total of 150 labeled experiments. For each experiment, two doors were randomly selected (by computer) from the three entrances to the room to be the starting and ending doors. During all experiments, the robot performed its scanning behavior, turning 30 degrees every 20 seconds. Certain sensors like the photo-sensor and the thermometers are correlated with time due to natural ambient variations. In order to create independence between the scenarios and the observations, the order in which the 150 experiments were conducted was random. The robot recorded an *observation*, a set of readings from its sensors, about once every 0.1 to 0.2 seconds. The mean (and standard deviation) of the number of observations per scenarios 0, 1, and 2 are: 134.92 (57.18), 38.54 (15.59) and 70.46 (21.19), respectively.

We are also interested in understanding how important various sensors are to the performance of the system, in order to avoid building new versions of the robot with useless sensors, and to help decide what new sensors should be added to the robot.



Figure 6: Scenarios 1 (left) and 2 (right).

Methods

To classify the labeled data we used three well-known classifiers implemented in the Python sckit-learn library (Pedregosa et al., 2011): *random forest, boosting*, and *logistic regression*. The first two are decision-tree classifiers, the last is a maximum-likelihood method that separates data based on linear relationships between variables. Ten-fold crossvalidation was performed for each classifier as follows: the validation set of each fold contains a single, randomlyselected experiment from every scenario-proximity condition, for a total of 15 out of 150 experiments; the remaining 135 experiments (9 from each condition) comprise the training set of each fold. Sensor data is normalized by subtracting its mean and standard deviation calculated from the training set, which is especially helpful for logistic regression.

Decision tree learning is useful for classifying data with nonlinear relations. A decision tree partitions the data into regions through recursive binary splits, choosing the best predictor for the split at each step according to an impurity measure. We choose the *Gini index* for training, since it is more sensitive than misclassification error and more interpretable than cross-entropy. This impurity measure can be interpreted as the training error rate at the split. However, to evaluate the performance on validation data, we use standard *misclassification error* (Hastie et al., 2009).

The growth of decision trees is highly sensitive to noise in the data. Any errors in the first splits are propagated down to all splits below it. To reduce this variance, ensemble methods like bagging, boosting, and random forest can be used. These ensemble methods produce a forest of trees, with the final classification determined by a majority vote among them (Hastie et al., 2009). We use two such ensemble methods: *random forest*, which creates trees trained on bootstrapped data with limited access to variables at each split; and the *SAMME boosting* algorithm, which iteratively trains trees while weighting data points by their difficulty of classification, and weighting trees by their training accuracy. These methods should be able to perform well even in the presence of highly nonlinear patterns in the data, and deal well with ambiguous data points.

Finally, *logistic regression* classifies data according to linear relationships between predictors. The coefficients of this relationship are estimated by maximizing the log-likelihood



Figure 7: Mean cross-validation error of Random Forest classifier over observations, for number of trees and number of variables at each split.

of the conditional probability distribution of the class given an observation, modeled as a linear function of the variables. This classifier should do well if the data can be linearly separated, and has the additional benefit of returning a probability of class membership for future analysis. Two regularization penalties to the size of the regression coefficients are investigated: L1 (linear) and L2 (squared). These penalties subtract from the objective function of the regression, respectively, the sum of the absolute values of the coefficients and the sum of the squares of the coefficients. When the data are normalized to similar ranges, this can be used for variable selection (Hastie et al., 2009).

Results

We analyzed classifier performance for both observations and experiments. The average validation misclassification error of the random forest across all ten cross-validation folds is depicted in figures 7 and 8 for observations and experiments, respectively. Since we ultimately want the robot to classify scenario-proximity conditions, rather than single sensor observations, each experiment is classified according to a majority vote among the labels predicted for observations taken during that experiment. This also weighs each scenario-proximity condition equally, whereas per-observation error favors conditions with the most data. The performance is fairly robust to the number of trees, although more trees, as expected, tends to produce a smoother curve (more robust to changes in number of variables tried at each tree split). We can see that performance tends to be best when the classifier has access to 3 or 4 randomly selected variables at each split.

The best classifier parameters were selected according to the average validation misclassification error across all ten cross-validation folds. The corresponding errors and their



Figure 8: Mean cross-validation error of Random Forest classifier over experiments, for number of trees and number of variables at each split.

Classifier	Obs. Error	Exp. Error
Random Forest	$0.182\pm.041$	$0.273 \pm .094$
SAMME Boosting	$0.204\pm.055$	$0.293 \pm .106$
Log. Reg.	$0.209\pm.065$	$0.320\pm.086$
Trivial	$0.711\pm.019$	$0.667 \pm .000$
Random	$0.586\pm.007$	$0.664 \pm .001$

Table 1: Observation and experiment errors on 3-Scenario classification. Mean and 95% confidence interval.

95% confidence interval are shown in Table 1. The confidence intervals were calculated using the errors on each fold, assuming a t-distribution. The random forest classifier achieves observation and experiment mean error rates as low as 0.182 and 0.273, respectively. The boosting classifier achieved similar performance (figures with mean crossvalidation error not shown). The performance of the logistic regression classifiers per experiment is shown in figure 9. Performance improves with smaller regularization penalties, and for sufficiently small penalties, the performance is not significantly worse than random forest. Overall, the three classifiers can classify correctly about 80% of the time in which of the three scenarios an observation was taken, and the experiments 70% of the time.

We also computed the performance of two null-model classifiers. The *trivial* classifier labels every observation in the test set with the most frequent class label in the training set of each fold. The *random* classifier randomly labels observations in the test set with the same frequency that classes appear in the training set. The three (non-null) classifiers perform quite well, given the low-resolution sensors. They significantly outperformed the null models, although not significantly different from each other. This suggests that any nonlinear patterns in the data are not significant for performance of the same frequency of the sensors.



Figure 9: Mean cross-validation error of Logistic Regression classifier over experiments, for both regularization penalties

mance on this minimal social classification task—since logistic regression was not significantly worse than decisiontree classifiers.

While the performance of the classifiers is significantly better than null models, it could still be the case that most of the error would be between the two walking scenarios 1 and 2. These scenarios may be harder to distinguish, since both begin the same way. It is important that a social robot be able to determine an empty room from an occupied one. However, it is more important for social awareness that Katie can distinguish between an indifferent person and one that displays an interest in interacting, situations whose proxies here are scenarios 1 and 2 respectively.

In order to investigate whether these two scenarios are discriminated well by the same classifiers, we performed binary classifications, where scenario 2 is "positive" and scenario 1 is "negative". These values can be interpreted as answers to the question: does a person want to interact with the Katie? The average ten-fold cross-validation performance of the classifiers is presented in Tables 2 and 3 for the best parameters. Performance is reported for accuracy, balanced F1 measure (harmonic mean of precision and sensitivity) and the Matthew's correlation coefficient (MCC)⁴.

The performance of the classifiers is again quite good and significantly better than the null models, though not significantly different from one another. It is clear that the two social scenarios can be distinguished by Katie's low-resolution sensors most of the time; with accuracy reaching 90% of the time with random forest. The MCC measures the correlation between observed and predicted labels. It is zero for random prediction (as is the case of our random null model).

Classifier	Accuracy	F1	MCC
R. Forest	$0.874 \pm .077$	$0.893 \pm .079$	$0.765 \pm .118$
SAMME	$0.852\pm.057$	$0.882 \pm .053$	$0.703 \pm .104$
Log. Reg.	$0.845\pm.078$	$0.861 \pm .083$	$0.734 \pm .112$
Trivial	$0.647 \pm .029$	$0.785\pm.021$	4
Random	$0.542\pm.000$	$0.645\pm.000$	$-0.001 \pm .001$

Table 2: 2-Scenario classification for observations

Classifier	Accuracy	F1	MCC
R. Forest	$0.900\pm.075$	$0.892\pm.094$	$0.826\pm.128$
SAMME	$0.860\pm.077$	$0.856 \pm .092$	$0.748 \pm .139$
Log. Reg.	$0.880\pm.066$	$0.862\pm.090$	$0.789\pm.110$
Trivial	0.500 ± 0	$0.667\pm.000$	4
Random	$0.519\pm.001$	$0.673\pm.001$	4

Table 3: 2-Scenario performance for experiments

For experiments, it reaches 0.826 which is a very high correlation between observation and prediction. It is also worth noticing that in the 2-scenario classification, contrary to the 3-scenario case, the performance was slightly higher for experiments than observations. This is likely due the larger number of observations gathered by scenario 0 experiments.

The relative importance of each sensor to classification performance is calculated as the expected fraction of observations that each sensor variable contributes to in the classification. This is depicted in figure 10 for the random forest classifiers. In this case, the photo sensor is the most useful, followed by the IR thermometers, the cliff sensors, and the rear-facing IR range sensors. The bumps, wheel, and wall sensor are not useful for discriminating these social scenarios. Results are similar for the boosting classifier, but with the photo sensor greatly emphasized (figure not shown).

As a proxy for variable importance, in the case of logistic regression, we can examine the average coefficients produced by the classifier for variables across folds (figure not shown). However, different sets of coefficients correspond to different scenarios. As in the case of the random forest and boosting classifiers, the photo sensor and IR thermometers are found to be important for all scenarios, but the bump sensors are also relevant in distinguishing scenario 0, since they are not activated during any scenario 0 experiment.

Future work

We see robotic development as moving from an understanding of robots as techno-scientific artifacts to sociotechnical ones. The social umwelt is woven together through technical infrastructure and we have illustrated this infrastructure with a simple robot. Furthermore, we showed that such a robot endowed with low-resolution sensors is capable of distinguishing minimal social scenarios with high accuracy. We plan on designing more sophisticated communication

⁴ The trivial classifier for observations and experiments, and the random classifier for experiments label all data into a single class, which results in a division by 0 in the calculations of the MCC



Figure 10: Random Forest variable importance

loops for humans and robots to come to a shared experience of their environment. These will extend beyond the social world of our household lab and incorporate ways for the robot to select different forms of motion. To move in this direction of more complicated social robot unwelts, we will design robots capable of interacting directly with people running common household scenarios such as arriving home from work, reading a book, watching tv, etc. We will use this information for further classification tasks and hope to refine the kinds of social states the robot can detect using simple sensors. At the same time we will have people interact with the backend of the system, looking at the world from the perspective of the robot, showing graphs and associated images, and classifying sensor data. Once the robot has a basic repertoire of meaningful classifiers it can begin making guesses about interesting/anomalous social events and then eliciting humans for annotations.

Acknowledgements

This work was supported by an Indiana University Collaborative Research Grant.

References

- Alač, M., Movellan, J., and Tanaka, F. (2011). When a robot is social: Spatial arrangements and multimodal semiotic engagement in the practice of social robotics. *Social Studies of Science*, 41(6):893–926.
- Almeida e Costa, F. and Rocha, L. M. (2005). Embodied and situated cognition. Artif. Life, 11(1-2):5–12.
- Braitenberg, V. (1984). Vehicles: Experiments in synthetic psychology. MIT Press, Cambridge, MA.
- Brooks, R. A. (2002). *Flesh and machines: How robots will change us*. Pantheon Books, New York.
- Clark, A. (1998). Being there: Putting brain, body, and world together again. MIT press.

- Dautenhahn, K., Ogden, B., and Quick, T. (2002). From embodied to socially embedded agents - implications for interactionaware robots. *Cognitive Systems Research*, 3:397–428.
- Dourish, P. (2001). Where the action is: The foundations of embodied interaction. MIT Press.
- Ferreira, M. I. A. and Caldas, M. G. (2013). The concept of Umwelt overlap and its application to cooperative action in multi-agent systems. *Biosemiotics*, 6(3):497–514.
- Forlizzi, J. (2007). How robotic products become social products: An ethnographic study of cleaning in the home. In ACM/IEEE International Conference on Human-Robot Interaction (HRI), volume 2, pages 129–137. ACM.
- Forlizzi, J. and DiSalvo, C. (2006). Service robots in the domestic environment: A study of the Roomba vacuum in the home. In ACM SIGCHI/SIGART Conference on Human-Robot Interaction (HRI), HRI '06, pages 258–265, New York, NY, USA. ACM.
- Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Friedman, J., and Tibshirani, R. (2009). *The elements of statistical learning*, volume 2. Springer.
- Hoffmeyer, J. (1997). *Signs of meaning in the universe*. Indiana University Press.
- Kozima, H., Michalowski, M. P., and Nakagawa, C. (2009). Keepon: A playful robot for research, therapy, and entertainment. *International Journal of Social Robotics*, 1(1):3–18.
- Lee, H. R. and Sabanovic, S. (2013). Culturally variable preferences for robot design and use in South Korea, Turkey, and the United States. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 17–24.
- Matsumoto, N., Fujii, H., Goan, M., and Okada, M. (2005). Minimal design strategy for embodied communication agents. In *IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN)*, volume 14, pages 335–340.
- Mutlu, B. and Forlizzi, J. (2008). Robots in organizations: The role of workflow, social, and environmental factors in humanrobot interaction. In ACM/IEEE International Conference on Human-Robot Interaction (HRI), volume 3, pages 287–294.
- National Science Foundation (2013). National robotics initiative.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikitlearn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Shibata, T. (2012). Therapeutic seal robot as biofeedback medical device: Qualitative and quantitative evaluations of robot therapy in dementia care. *Proceedings of the IEEE*, 100(8):2527– 2538.
- Sung, J.-Y., Guo, L., Grinter, R. E., and Christensen, H. I. (2007). "My roomba is rambo": Intimate home appliances. In *International conference on Ubiquitous computing (UbiComp)*, volume 9 of *UbiComp*, pages 145–162. ACM.
- Uexkull, J. (2001). An introduction to umwelt. Semiotica, 134(1):107–110.