Formalizing the Design Space and Product Development Cycle for Socially Interactive Robots

Eric C. Deng
University of Southern California
denge@usc.edu

Bilge Mutlu
University of Wisconsin–Madison
bilge@cs.wisc.edu

Maja J. Matarić
University of Southern California
mataric@usc.edu

ABSTRACT

As socially interactive robots move out of the lab and "into the wild," the systems being developed are shifting from being general-use, research-focused platforms to application-specific products. In product development, designers systematically and iteratively explore the *design space* of a product to arrive at a design that best fulfills business, engineering, and design requirements. Socially interactive robots are no different; they are envisioned for specific applications, contexts, and scenarios of use and must be designed appropriately. Although robot designers can take cues from related fields such as industrial design and animation, there has yet to be a formalized approach to the development of these systems. We outline a characterization of the design and application spaces of socially interactive robots based on a survey of 65 publications. We also present a product development cycle for this new category given our design-space characterization and findings of the survey.

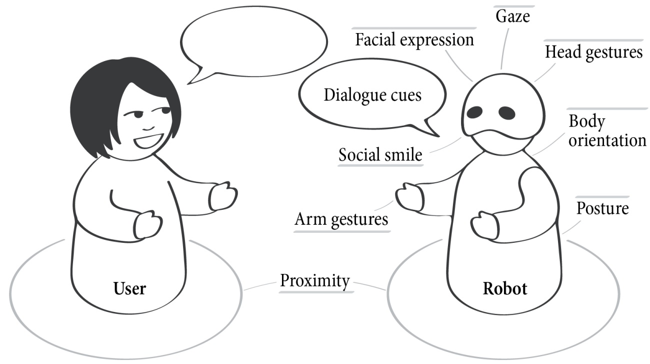

Figure 1: An example formulation of the *design space* of social, expressive cues for socially interactive robots

1 PRODUCT DEVELOPMENT PROCESS

Our environments consist of products designed to provide value to their users in specific ways. Products, when designed properly, are intuitive, appealing, and effective. These desirable traits result from systematic explorations of the design space through many iterations of product prototypes and testing that eventually converge to a singular product with a combination of features that best fits the business, engineering, and design requirements. This systematic evaluation is well-studied and has produced many frameworks, philosophies, and tools [2, 14]. These tools are designed to be combined and adapted to fit various product categories with differing requirements. While there are many dimensions to product development, we focus specifically on the development cycle pertaining to the engineering and usability of socially interactive robots as products. In this section, we introduce existing *processes, models,* and *frameworks* in product development that we will use as a foundation for a proposed model for product development for socially interactive robots.

1.1 Product Development Process Models

Using a formal, process-based approach is a commonly-used method for product development [4]. A wide variety of such approaches have been proposed for the varying needs of different organizations and teams, some focusing on an "ideal processes," some on the makeup of the team, and some on the customer. In this work, we focus on two of the most popular models: (1) *the Stage-Gate model* and (2) *the IDEO Process.*

The *Stage-Gate model* is an approach to managing and organizing the overall product development process through a set of hierarchical requirements and defined ordering of work [3, 4]. Complex products such as socially interactive robots have many components that can be developed in parallel. For example, in the early development process, infrastructure for mechanical, electrical, and software subsystems may be developed concurrently without blocking. As the product becomes more mature, and therefore more integrated, more dependencies arise, producing a network of blockers and "gates."

The *IDEO process* was developed in the global design-consulting firm IDEO; it emphasizes *situated observation and evaluation* of users to achieve *Form-Fit-Function (FFF)* in design solutions [13]. Unlike the Stage-Gate model, the IDEO process is explicitly human-centered and is more focused on the "what" than "how" in product development and consists of five core steps [13] Figure 4:

(1) *Understand and Observe:* observe people in real-world contexts of use to gain first-hand knowledge in order to appropriately scope the project;
(2) *Synthesize:* organize and interpret research data and insights into potential design features;
(3) *Visualize:* create representations of synthesis results using visible and tangible media;
(4) *Prototype, Evaluate, and Refine:* iteratively improve design ideas through prototypes of different quality and evaluate those with users in pseudo-realistic settings;
(5) *Implement:* finalize refined product features that have been evaluated and grounded in research and experimentation.

In this paper, we propose that the IDEO approach can serve as a blueprint for a design process for socially interactive robots. We focus on adapting iterative design methods to make them suitable for developing robotic products. Thus, we assume that the designer has a good understanding of the interactive task(s) for the robot. The IDEO process emphasizes systematic evaluation of designs such that, if a part of a design is not needed for *form, fit,* or *function*, then it is removed from the next iteration. In order to effectively execute this in practice, the IDEO method leverages the foundational concept of *design spaces* that can facilitate the design of complex systems such as socially interactive robots. Figure 1 illustrates a formulation of the design space for social, expressive cues for socially interactive robots.

2 DESIGN SPACE FOR SOCIALLY INTERACTIVE ROBOTS

The *design space* comprises the elements that designers can manipulate to create variations in the *appearance*, *behavior*, and *overall structure* of a product. Formally defining the design space can scaffold systematic experimentation and facilitate constructive discussion and analysis [10]. Defining a design space for socially interactive robots can facilitate utilization of prior research, guide the trajectory of future experimentation, and critically scaffold future work and discussion both in academia and in industry. Toward that end, we reviewed 65 studies published between 2002 and 2017, and analyzed the experimental methods, robot designs, and findings from those studies in order to identify patterns that can guide the development of future robots.

Our characterization of the design space for socially interactive robots captures three factors that also serve as facets of the industrial design, animation, and interaction-design of a robot: (1) *social role*, (2) *embodiment*, and (3) *communicative behaviors*. These factors allow for variations in physical construction, motion and behavior, and interactive capabilities of robot systems, respectively. While these factors are not independent, they are useful constructs for characterizing the complex design space.

The *social role* of a robot is a facet of its interaction design; is is defined relative to a user, and can be represented on a continuous scale ranging from *subordinate* to *superior*. These roles are closely tied to a robot's ability and approach to achieving its goals. For instance, a superior robot may be more effective as an instructor or coach, or an enforcer for rule-following in performance tasks, based on its perceived authority [1] and reliability, while a peer-like robot may be more effective in maintaining an appropriate levels of challenge for the user engaged in competitive tasks [8].

The robot's social role informs the design of its *embodiment*, a facet of the industrial design, and the robot's *communicative behaviors*, a facet of its animation. As the robot's embodiment is inextricably tied to its behavioral capabilities (e.g., a robot without arms cannot perform pointing gestures), we represent these two dimensions at the same level of abstraction in the design space.

The robot's embodiment can be represented by two features: the *design metaphor* and the *level of abstraction* at which the metaphor is followed in the design. The design metaphor for a robot's embodiment makes up a discrete, non-linear space that consists of familiar entities that afford certain expectations for the robot's interaction partners and scaffold its social interaction. For example, a *human* design metaphor can set expectations about the robot's intellect, reliability, and interaction modalities based on people's prior experience with other people [1, 5]. In such cases, people are more likely to adhere to human-human communication norms, such as maintaining eye contact [9], respecting personal space [12], and using multimodal iconic or metaphoric gestures to augment speech [7]. However, robots are meant to follow design metaphors at some level of abstraction rather than attempt to be perfect replicas of the source of the metaphor, which informs the designer on how to map affordances from the original design metaphor to the specific robot design Figure 6.

3 PRODUCT DEVELOPMENT IN SOCIALLY INTERACTIVE ROBOTS

Given the above characterization of the design space for socially interactive robots, we now turn to the problem of efficiently exploring that space to identify the *social role*, *embodiment*, and *communicative behaviors* for a specific robotic product that support a particular application or use. Because socially interactive robotic products, by definition, emphasize social interaction as a core pillar to their function, not only do these robots need to be effective in achieving their primary tasks, but they need to achieve these tasks while also maintaining social acceptance and interactivity.

3.1 Iterative Design in Socially Interactive Robots

Due to the holistic nature of measuring performance of socially interactive robots, testing design features in isolation is not a viable approach. The only way to find sets of features that work well together is to evaluate those features in sets and iteratively refine them. We propose a systematic way to iterate *efficiently* through such feature sets.

Figure 5 shows our proposed process for designing socially interactive robots. We begin with the task, which can be classified into one of eight task types based on the "task circumplex" proposed by Mcgrath [11], as seen in Figure 3. For each type of task, some social roles have been shown to be more effective than others, and based on prior data, intuition, and testing, designers can identify the best social role for a robot in the context of a given task. The social role is implemented through a combination of the robot's embodiment, behaviors, and interaction strategies. Because the appearance and behaviors of robots are inter-dependent and must be evaluated together, exhaustively exploring this design subspace can be challenging, if not impossible. By leveraging existing work as well as using our characterization of this design subspace as an analytical tool, designers may effectively contextualize and test their design ideas.

We propose a development process that involves the integration of the "gated" structure of the Phase-Gate model and the IDEO process presented above. We consider Modules 2 and 3 of our iterative process (Figure 2) as a socially interactive robot-specific approach to Step 4 of the IDEO process (Figure 4). By implementing "fuzzy" gates in the pipeline (i.e., setting soft requirements on getting each module *mostly finalized* before moving on to the next module), we can leverage the abstracted structure of the design space of socially interactive robots to most efficiently explore that space.

Because social roles are a higher-level construct, their iterative exploration requires more time and effort; not only does the process require conducting more experiments, but iteration may involve the development of different robots, a costly process. Furthermore, the applicability of previous work on robots with different social roles is unknown. We therefore recommend relying on existing work to

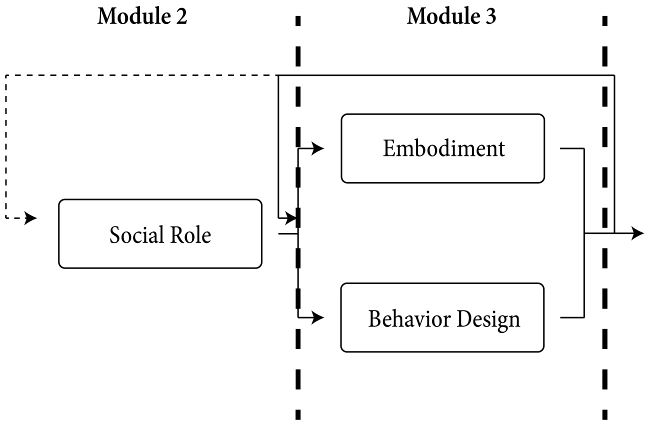

Figure 2: Design Metaphor and Abstraction

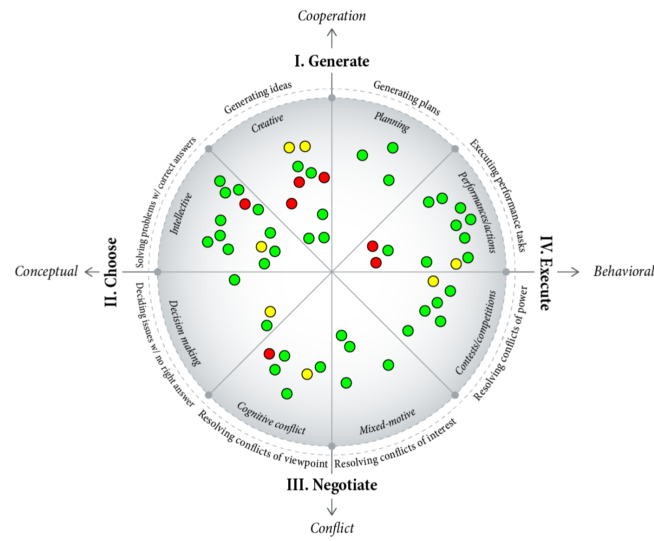

Figure 3: McGrath's Task Circumplex

guide decisions on the social role and focusing on exploring the robot's design metaphor and abstraction through different robot designs. In Figure 2, this approach is illustrated with the dotted line connecting the output of Module 3 to the input of social role as opposed to the solid line to the input of Module 3.

In the following sections, we discuss the results of our analysis of the past studies and illustrate how the aggregated data may guide the design process of future socially interactive robots when used with our characterization of the design space and product development process.

3.1.1 Reviewed Studies. The reviewed studies were published between 2002 and 2017, inclusive, and were comparisons between a physical robot and a baseline system, e.g., a non-physical embodiment of the same agent used in the same task. We classified results into two categories: *task performance* and *agent perception*. Task performance is a measure of how the participant's evaluated performance differed between the robot and baseline system. Agent perception, typically obtained through self-report [6], is a measure of the participant's feelings toward the robot or baseline system. The studies usually relied on task performance as a measure of the benefits of a physical robot. Agent perception served as an indicator of longer-term benefits of these systems, since the perceived desirability of an agent can increase length of engagement, trust of the agent, and compliance to robot instruction or suggestion.

In order to identify patterns in the results from the reviewed studies, we plotted each study relative to the task and the social role it involved. We represented each experiment with a point on McGrath's [11] circumplex with the color of the dot representing the combined results of the study with green representing positive results, yellow representing neutral results, and red representing negative results (Figure 3). This visualization reveals patterns that enrich our formulation of the design space as a decision-making tool in the development process as discussed in the following sections.

3.1.2 Navigating Social Roles. In our product development process, the module that designers must first evaluate is social role. The robot's role in the selected task shapes perceptions of it and scaffolds people's interactions with it. For each type of task, our review has shown that robots playing certain social roles tend to outperform others and that, for both task performance and for the person's perceptions of the agent, appropriately selecting a robot's social role is critical to its overall success. Two critical patterns can be seen in the plots and are especially visible when we "unravel" the circumplex to isolate each octant (Figure 7). The first pattern we observe is the distribution of the points over the space, which represents the social roles that past researchers had hypothesized to be effective for each role type. The second pattern is the clustering of the colored dots along each task dimension, representing the experimental results of the social roles in practice.

Given these results, we aim to facilitate mapping a given task to potentially effective social roles. In our review, we assigned values to the social role and performance of robots for each study. In Figure 9, we plotted these values for each task category and performed a regression on the points in order to create a plot that represents the expected performance of a robotic agent given its social role and task. We then overlaid a density distribution of the reviewed studies to represent where the experiments fall along the social roles so as to indicate the reliability of that estimate (the darker the plot, the more reliable the estimate is for that social role) (Figure 9).

We hypothesize that task and "optimal" social roles have a one-to-one relationship–for every task there is a *singular* "best" social role for a robot (Figure 9). By using Figure 9, or future revisions of Figure 9, designers can develop more grounded hypotheses about effective social roles for their tasks and researchers can more efficiently study previously unexplored social roles for each type of task.

3.1.3 Navigating Embodiment Design. In order to test the social role of robots for a given task, the designer must first realize a robot design to assume the selected social role. As discussed above, the two other key components of the design of a socially interactive robot are its physical embodiment and its communicative behavior. Although both are critical in the context of social roles, we chose to focus on the design of a robot's physical embodiment, because it is sufficiently constrained to explore thoroughly and is often the main determinant of the robot's behavior. We represent a robot's embodiment with a *design metaphor* and the *level of abstraction* at which this metaphor is implemented (Figure 6). The design metaphor "jump-starts" mental models of the robot's capabilities for task competence, reasoning, and interaction. Because these design metaphors are by definition familiar entities, such as everyday objects and biological systems, people have expectations of how they would interact with an actual instantiation of the design metaphor (e.g., people will likely interact with a real dog by petting it). The level of abstraction represents how closely those affordances will be mapped to robot designs inspired by the metaphors.

Our analysis of the reviewed studies classified robot systems by design metaphor and assigned values to each embodiment for its level of abstraction, ranging from *metaphoric* to *literal*. By combining these values with the social role values for the experiments in which each robot was used, we plot the social role over level of abstraction for each design metaphor (Figure 9). Similar to Figure 8, this figure is overlaid with a density distribution representing the reliability of the plot at different levels of abstraction for each metaphor based on where the robot designs are on the plot. When a designer is exploring possible designs for a robot-design task with an idea of the social role that the robot is expected to play, this figure can guide design decisions regarding what level of abstraction may be an effective implementation of the social role for the design metaphor. For example, if the designer is inclined to build a cat-like robot, and the task calls for a peer-like robot, Figure 9 suggests that a robot following a more literal cat metaphor is likely to be effective and that confidence in this estimate is high based on prior evidence that has shown slightly-literal, cat-like robots to be effective peer-like agents.

The described approach is particularly useful in designing multipurpose robotic products that have varied applications, fall into different task categories, and span a range of desired social roles. For example, if a designer wants a robot to be able to engage in *competitive* tasks as well as *creative* tasks, Figure 8 suggests that the robot is most effective in a superior role or a subordinate role, respectively. Because of the wide range of social roles that the robot needs to play, the designer can use Figure 9 to determine that a car-like design metaphor with a literal implementation may be a good choice, because prior research has shown this combination of metaphor and level of abstraction to be effective for both subordinate and superior roles. This example demonstrates the close coupling between social roles and robot embodiment. Unlike the relationship between task and "optimal" social role, which we see as one-to-one, robot embodiment can be seen as setting "bounds" on the social roles for a particular robot. For instance, based on the reviewed experiments, a metaphoric, human-inspired robot will likely be effective as a subordinate/peer-like agent and a superior agent, but it may not be effective as a peer/superior agent or a purely subordinate agent.

4 CONCLUSIONS

In this work, we aimed to develop an evidence-based decision-making tool to aid the design of socially interactive robotic products. Application-specific products are fundamentally different from general-use research platforms in their development processes, and by studying existing product development practices, processes, and models as well as by reviewing existing research in human interaction with social robots, we first presented a characterization of the design space for these products and then outlined a development process specifically for socially interactive robots. This pipeline can be more efficiently traversed using existing research results. We presented aggregated results from 65 studies in a variety of formats tailored to different modules within that process and demonstrated how these results can be used to guide design exploration.

The analyses and findings presented here serve as preliminary work for a more in-depth analysis of the reviewed results in a more general context beyond product development. We acknowledge that there are many intricacies involved in designing and developing products as complex as socially interactive robots that extend beyond the discussion in this paper. For example, the subtleties of behavior design, cultural context, and individual differences all require further consideration and analysis that we plan to explore in future work. We believe that, even in its preliminary form, our analysis and discussion offer informative guidelines for the design of socially interactive robots and establish a shared language and process for researchers and developers for future work on design frameworks for robotic products.

REFERENCES

[1] Wilma A Bainbridge, Justin W Hart, Elizabeth S Kim, and Brian Scassellati. 2011. The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics* 3, 1 (2011), 41–52.
[2] Peter H Bloch. 1995. Seeking the ideal form: Product design and consumer response. *The Journal of Marketing* (1995), 16–29.
[3] Robert G Cooper. 1990. Stage-gate systems: a new tool for managing new products. *Business horizons* 33, 3 (1990), 44–54.
[4] Robert G Cooper and Scott J Edgett. 2012. Best practices in the idea-to-launch process and its governance. *Research-Technology Management* 55, 2 (2012), 43–54.
[5] Juan Fasola and Maja Mataric. 2013. A socially assistive robot exercise coach for the elderly. *Journal of Human-Robot Interaction* 2, 2 (2013), 3–32.
[6] Maia Garau, Mel Slater, David-Paul Pertaub, and Sharif Razzaque. 2005. The responses of people to virtual humans in an immersive virtual environment. *Presence: Teleoperators and Virtual Environments* 14, 1 (2005), 104–116.
[7] Chien-Ming Huang and Bilge Mutlu. 2014. Learning-based modeling of multimodal behaviors for humanlike robots. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*. ACM, 57–64.
[8] Céline Jost, Vanessa André, Brigitte Le Pévédic, Alban Lemasson, Martine Hausberger, and Dominique Duhaut. 2012. Ethological evaluation of Human-Robot Interaction: are children more efficient and motivated with computer, virtual agent or robots?. In *Robotics and Biomimetics (ROBIO), 2012 IEEE International Conference on*. IEEE, 1368–1373.
[9] Katrin Solveig Lohan, Sebastian Gieselmann, Anna-Lisa Vollmer, Katharina Rohlfing, and Britta Wrede. 2010. Does embodiment affect tutoring behavior. In *IEEE international conference on development and learning (ICDL) conference*.
[10] Allan MacLean, Richard M Young, Victoria ME Bellotti, and Thomas P Moran. 1991. Questions, options, and criteria: Elements of design space analysis. *Human–computer interaction* 6, 3-4 (1991), 201–250.
[11] E Mcgrath. 1995. Methodology matters: Doing research in the behavioral and social sciences. In *Readings in Human-Computer Interaction: Toward the Year 2000 (2nd ed.* Citeseer.
[12] Ross Mead and Maja J. Matarić. 2016. Autonomous human-robot proxemics: socially aware navigation based on interaction potential. *Autonomous Robots* (Jun 2016), 1–13. https://doi.org/10.1007/s10514-016-9572-2
[13] Ron Moen. 2001. A Review of the IDEO Process. *Robert Wood Johnson Foundation, Toolkit* (2001).
[14] Timothy W Simpson, JD Poplinski, Patrick N Koch, and Janet K Allen. 2001. Metamodels for computer-based engineering design: survey and recommendations. *Engineering with computers* 17, 2 (2001), 129–150.

A APPENDIX

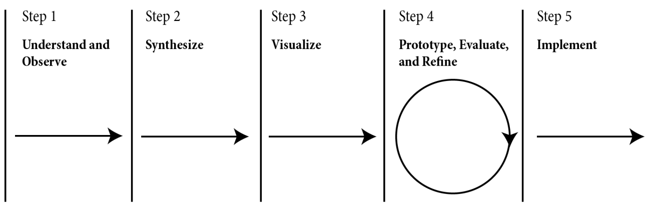

Figure 4: The IDEO Process [13]

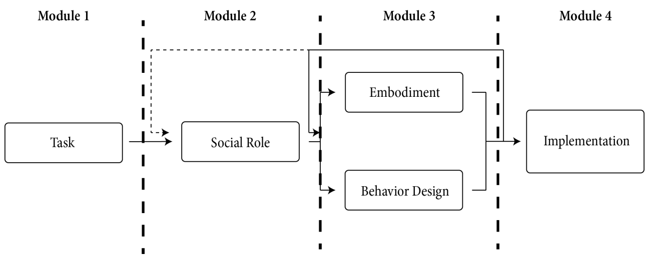

Figure 5: Product development pipeline specific to Socially Interactive Robots

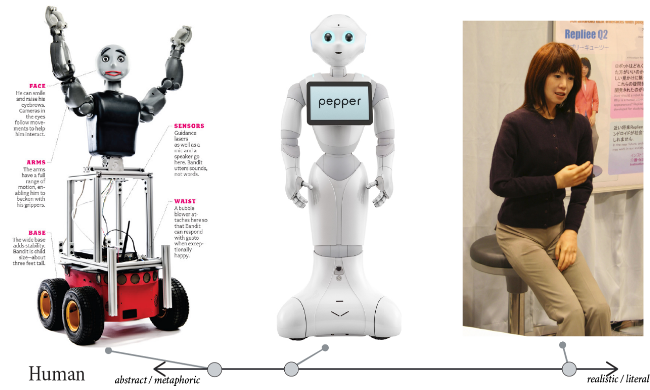

Figure 6: Design Metaphors and Abstraction

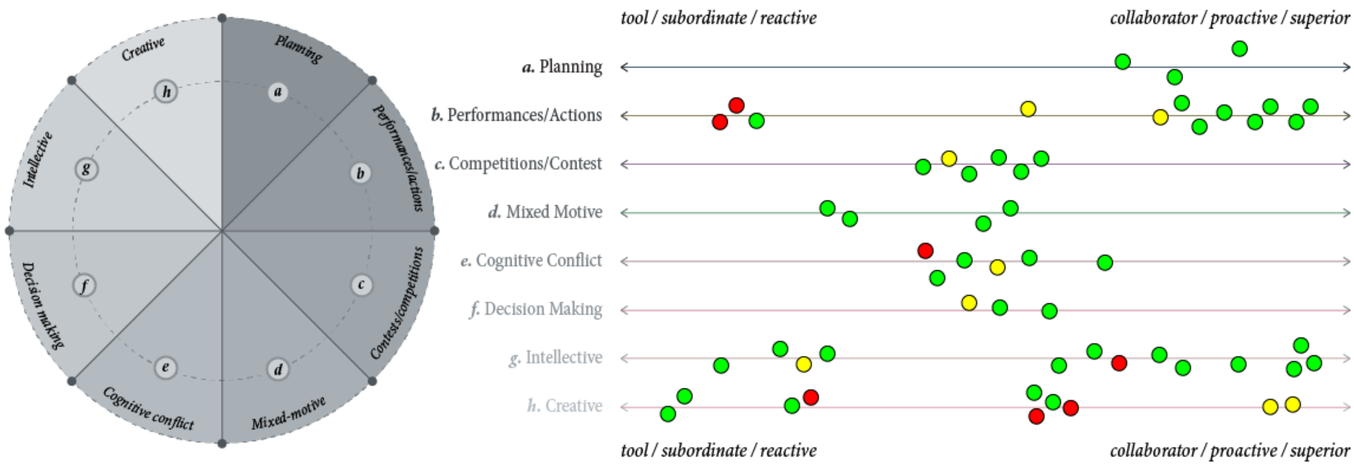

Figure 7: "Unraveled" Version of McGrath's Task Circumplex with the overlaid results

Performance of Social Roles by Task

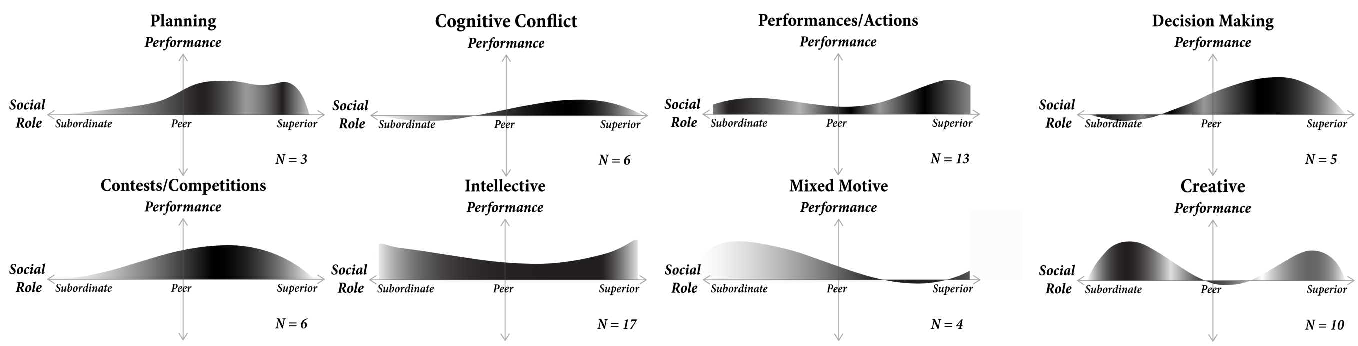

Figure 8: Performance of Robot Social Roles by Task Category

Design of Social Roles by Embodiment

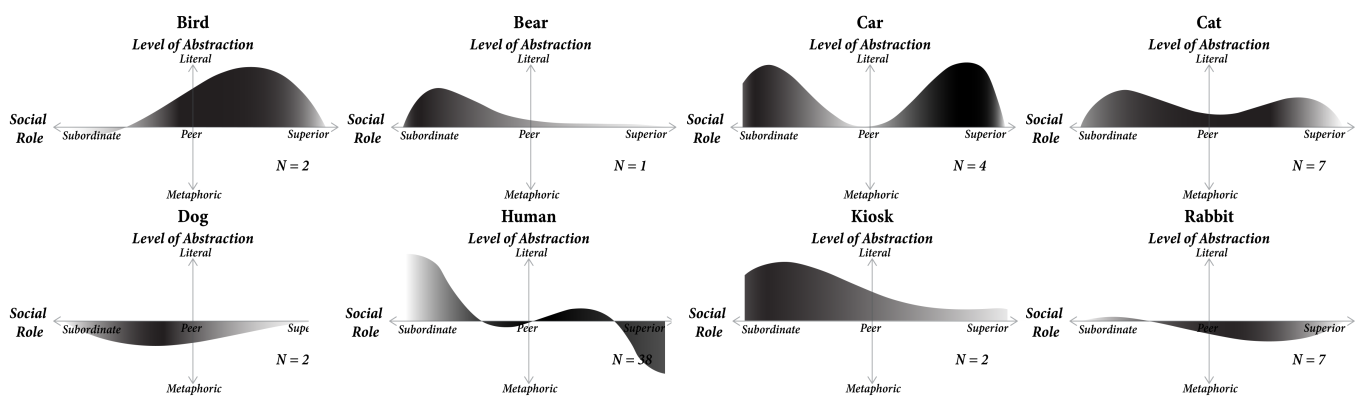

Figure 9: Embodiment Design and Metaphors by Social Roles