

Social Robotics and Societies of Robots

Antonio Bicchi, Centro di Ricerca E. Piaggio and Dipartimento di Ingegneria dell'Informazione (DII), Università di Pisa, Pisa, Italy, and Department of Advanced Robotics, Istituto Italiano di Tecnologia, Genova, Italy

Guglielmo Tamburrini, Dipartimento di Ingegneria Elettrica e Tecnologie dell'Informazione (DIETI), Università di Napoli Federico II, Napoli, Italy.

Contact Information:

Antonio Bicchi, Centro di Ricerca E. Piaggio & Dipartimento di Ingegneria dell'Informazione (DII), Università di Pisa, Via Diotisalvi 2, 56100 Pisa, Italy.

Abstract

The sustainability of social robotics, like other ambitious research programs, depends on the identification of lines of inquiry that are coherent with its visionary goals while satisfying more stringent constraints of feasibility and near-term pay-offs. Within these constraints, this paper outlines one line of inquiry that seems especially viable: development of a society of robots operating within the physical environments of everyday human life, developing there rich robot-robot social exchanges, and yet refraining from any physical contact with human beings. To pursue effectively this line of inquiry, sustained interactions between specialized research communities in robotics are needed. Notably, suitable robotic hand design and control principles must be adopted to achieve proper robotic manipulation of objects designed for human hands that one finds in human habitats. The *Pisa-IIT SoftHand* project promises to meet these manipulation needs by a principled combination of sensory-motor synergies and soft robotics actuation, which aim at capturing how the biomechanical structure and neural control strategies of the human hand interact so as to simplify and solve both control and sensing problems.

Keywords: society of robots, neighbor collision avoidance, hand synergies, soft robotics, multidisciplinary research in social robotics.

1. The role of visionary goals in robotics

Visions of social robotics point to futuristic scenarios in which robots fit flawlessly into homes, offices, workshops, hospitals, and entertainment settings, fulfilling their roles of service providing

helpers and tutors, trustworthy care-givers, dexterous assistants or even enjoyable robotic companions. Turning these scenarios into technological achievements presupposes far-reaching advances in sensorimotor and cognitive skills of robotic systems, to such an extent that one may sensibly doubt whether the research efforts of a few generations of committed scientists and engineers will suffice to bridge the gap between vision and reality.

In spite of their remote and possibly unattainable character, visionary goals may play a variety of useful roles in robotics research. Consider from this perspective RoboCup's visionary goal of putting together a robot soccer team that eventually beats the human world champion team. The RoboCup manifesto forthrightly states that accomplishing this goal "will take decades of efforts, if not centuries" (<http://www.robocup.org/about-robocup/objective/>). At the same time, however, it points out that the more down-to-earth objective of robotic soccer tournaments is to bolster research activities in the near term. Here the project to develop a robotic world champion team plays at least two important roles in the framework of this research program. On the one hand, it makes a unifying horizon available to scientists working on a wide variety of problems – notably including sensor fusion and perception, learning, reactive navigation, contextual awareness and strategic decision-making in multi-agent environments. On the other hand, it suggests more feasible objectives and lines of research that promise to nurture near-term advances in robotics, enabling one to achieve better robotic models, novel technologies, and valuable industrial applications.¹ The visionary scenarios of other ambitious research programs play similar roles. For example, the ideal horizon of fully autonomous robotic systems pursued by the DARPA Robotic Challenge (http://www.darpa.mil/Our_Work/TTO/Programs/DARPA_Robotics_Challenge.aspx) or its predecessors in military robotic research, like the DARPA Grand Challenge, accommodates periodic competitions suggesting more feasible objectives in the near term.

The sustainability of social robotics as a research program depends likewise on the identification of lines of research that are coherent with its visionary goals while satisfying at the same time more stringent requirements of feasibility and near-term pay-offs. The question then is: how does one step from visionary scenarios to more feasible research objectives?

The current efforts to reduce robot interactions with unexpected behaviors point to fruitful lines of research for social robotics. In industrial robotics, factory floors have been typically modified so as to resemble closed worlds characterized by well-known and precisely modelled dynamic evolutions. For example, human workers and robots have been often confined to separate workspaces, so as to exclude major sources of unexpected perturbations arising from intentional or unintentional human behaviours. Clearly, one cannot apply similar robot-in-a-cage policies in social robotics, for a social robot is by definition a robot that shares the same physical environment with human beings and other intelligent agents. This circumstance poses a unique challenge for social robotics research: one has to model and successfully predict the behaviors of intelligent agents – most notably of

¹ Fruitful lines of research satisfying these demands are called there *well-directed subgoals*: "Needless to say, the accomplishment of the ultimate goal will take decades of efforts, if not centuries. It is not feasible, with the current technologies, to accomplish this goal in any near term. However, this goal can easily create a series of well-directed subgoals. Such an approach is common in any ambitious, or overly ambitious, project." (see <http://www.robocup.org/about-robocup/objective/>)

intentional agents driven by beliefs and desires – dwelling in the robot’s operating environment. However, sharing of the same physical space does not mean that physical contacts have to occur between robots and other intelligent agents. Are there robotic tasks which involve no physical contact between humans and robots, and yet preserve sufficient interest for social robotics? The task assignments that we turn now to examine satisfy this constraint and also provide a unifying horizon for social robotics. Equally important, research work on these task assignments promises to produce near-term pay-offs in terms of both scientific results and industrial applications.

2. Societies of robots embedded in human societies

To begin with, consider an indoor environment inhabited by humans, like a museum or a similar exhibition space, where a number of robots are given the task of serving as information providers in daytime, and as supervisors at night.

During the night shift, robots would simply have to detect the presence of animated agents in the area and maintain a map of where these agents are. Each robot having sensors suited for detection, they should be deployed in the museum rooms so that every corner is detected by at least one robot (e.g., a central supervisor assigns these locations beforehand). If the number of robots is insufficient to obtain a complete static coverage, team members should move around, providing coverage over time. For each robot to be able to construct a map of the whole museum area, team members should pass information among themselves, through a multi-hop communication network, where neighboring agents can exchange information with each other, as well as pass along information from neighbors.

While building a general consensus among robot guards on what the intruder map is in the patrolled area is pretty straightforward to achieve, things get much more complicated when one takes into account that some of the agents may simply not supply veridical information to other agents in the team. This situation may arise because of a simple malfunctioning of the sensors or of the network, causing random, false or corrupted data to be passed on. And it may occur if a malignant attacker tampers with one or more of the robots, making them report intentionally falsified information. Clearly, the false information supplied by a malfunctioning robot could easily be propagated through the network by means of the very same consensus-building mechanism, which can easily be turned into a lie-propagating gossiping machine. The question then is, how can a system of many autonomous robots be built so that these faults can be detected and possibly resisted by the team, i.e. make so that all properly working agents have a correct view of the overall intruder map, and that faulty robots are identified and flagged for repair? The problem is akin to the classical Byzantine Generals conundrum: how many trusted robots are enough to fend off an attack by unfaithful robots (because of faults or malignant reprogramming)? Intuitively, keys to success are a sufficient degree of first-hand information in the system (in our example, overlap of coverage areas guaranteeing that a sufficiently large number of faithful agents directly see all events), and a sufficiently rich social connectivity, guaranteeing that majority voting rules can prevail everywhere in the system.

Although physical interactions between humans and robots are missing in the night-shift scenario, areas ripe for social robotics research are identifiable there in the rich social connectivity between robots and the typically human environment they operate in. In the day-shift scenario, the team of surveillance robots must be endowed with additional collision avoidance skills and cognitive-level

interaction capabilities. Indeed, the robotic team must be capable of recognizing calls by museum visitors, listening to their questions and answering them with proper information. However, for safety reasons one typically would have the robots move around their operational space strictly avoiding physical interactions, e.g. collisions with humans and other robots. While collision avoidance in slowly changing, simple environments is rather easy, it rapidly becomes a challenging task in dynamic and cluttered environments. Humans avoid collisions in crowded spaces through application of a number of social rules that are often unconscious but very effective: dynamics estimation, trajectory prediction, intention detection, and social hierarchies all play a role in negotiating passages through a corridor or narrow space, in very much the same way –albeit far less explicitly codified – as car drivers use rules of the road to drive their vehicles in the urban or highway traffic. In order to build teams of robots that can coexist and move safely, not only their individual behaviors have thus to be regulated, but their social behaviors as well. A robot needs to have a *model* of what humans or other robots can be expected to do. The behavior of a neighbor – be it a human or robot – of course depends also on what the neighbor’s goals are and what are his surroundings, or rather on what *does the agent know* about its own surroundings.

In the ignorance of this information, the only safe policy is to use extra caution. Such is the case, e.g. when a person turns a corner, without seeing whether or not another person is coming across her way. Building a suitable structure of information exchange between robots in a social environment is thus crucial. The question here is, what and how much should each robot know about other moving agents in the environment?

Clearly, omniscience of intentions and information available to all neighbors would in principle afford each robot with the potential to avoid every collision, and optimize its performance in e.g. minimizing travel time to destination (e.g., reaching a caller in the museum example). However, optimality comes at the cost of computational complexity, which is increasing exponentially with the connectedness of the network of social relations between robots. Accordingly, one has to come up with an alternative to omniscience in the way of navigation control strategy. Intuitively, a key to success in this context is the observation that accomplishing dynamic modeling, trajectory prediction, and collision avoidance is feasible by taking into account few neighbors at a time, and disregarding those that are farther away (where possibly “far” is to be understood in a different topology than usual metric distance). After all, this is what we do every day in driving our cars through traffic: look at the immediate neighbors, and disregard the rest – although in the rest are our neighbor’s neighbors (Bicchi, Fagiolini and Pallottino 2010).

Similar neighborhood-based strategies to trajectory modeling and collision avoidance might be fruitfully applied to deal with incrementally more challenging tasks for groups of robots in cluttered and dynamic environments. An admittedly more distant scenario of this sort involves robotic butlers running errands in a shopping mall, which is variously populated by human users and salespersons, by other robots and software agents embedded into its smart environments.

The shopping mall scenario spurs from the consideration that a wide diffusion of personal robots is expected to occur in the near term – there are already millions of personal robots in houses, and in few years from now more will be available to help people also in more complex chores than vacuum cleaning. The user or the household smart appliances compile the shopping list for the household weekly needs and provides the robot-butler with this information. User brings robot to

the local mall; the mall supervisor authority system authenticates the robot and accepts its presence. The robot obtains information on goods and their location, and clearance to perform transactions within the mall shops. Robot navigates the mall avoiding collisions with people, goods, and other carts. It fills the cart and waits the queues, while user is free to get involved in “higher-level,” more rewarding tasks.

The robots envisaged in the museum night shift, day shift, and shopping mall scenarios illustrate incrementally more challenging instances of the idea of groups of robots embedded into human habitats and societies. These groups are capable of rich social exchanges between robots and with other artificial agents, but look pretty unsociable to human bystanders: actively avoiding physical contacts with them, they negotiate a typically human environment and manage those social dealings with humans that are strictly necessary to mind their own business. These characteristic traits suggest a simile between each group of such robots to a subculture (or co-culture) embedded within the larger context of human culture: their members share distinctive interests and communication styles, engaging into more extensive social transactions within, rather than without, their own subculture. One may envision many other robotic subcultures resulting from different instantiations of the general strategy of limiting and streamlining – let alone avoiding – physical contacts between humans and robots. Collectively, these robotic subcultures may give rise to a composite society of robots embedded within human society.

3. Coordinating social robotics with other research communities

The subculture metaphor has another meaningful role to play in connection with social robotics and similarly ambitious research programs in robotics. Indeed, many scientific and technological fields of inquiry have been informatively compared to a mosaic of subcultures: the partly autonomous practices of instrumentation, experimentation, and theory exemplify this state of affairs in physics (Galison 1997). Interactions within each research subculture are typically more extensive than interactions between subcultures. Members of a particular research subculture are more strongly committed to the distinctive research objectives of their own subculture. However, the objectives of other subcultures are not completely ignored, insofar as a sufficient mutual coordination between subcultures enables them to enter fruitful scientific and technological transactions.

To identify what are the mechanisms enabling one to establish scientific and technological transactions within a markedly varied landscape of research interests, one may profitably learn, as Peter Galison suggests, “... from the anthropologists who regularly study unlike cultures that do interact, most notably by trade. Two groups can agree on rules of exchange even if they ascribe utterly different significance to the objects being exchanged; they may even disagree on the meaning of the exchange process itself. Nonetheless, the trading partners can hammer out a *local* coordination despite vast *global* differences.” (Galison 1997, p. 783).

Social robotics and other ambitious research programs play a crucial role in the processes which enable one to achieve local coordination within a disunified constellation of robotic subcultures – each one of them working asynchronously on regional objectives that are not endorsed with the same level of commitment by other robotic subcultures. In particular, both visionary goals and the more realistic objectives of social robotics require coordinated work of and exchanges between different robotic subcultures. Consider, for example, the shopping mall scenario discussed above. The development of robotic butlers requires one to solve significant problems of, e. g., navigation

and collision avoidance, perceptual recognition, modeling of multi-agent systems involving human-robot and robot-robot interactions. In addition to this, robotic butlers must be able to load goods and manipulate dexterously objects that are specifically designed for human hand manipulation. Thus, local coordination with research communities working on robotic hand design and control is needed, at least insofar as social robotics must address challenging manipulation tasks in the unstructured environments of human daily life. The question then is whether state-of-art research on artificial hands can satisfactorily respond to the robotic manipulation needs that are emerging in social robotics.

Researchers have been interested in the design and control of robot hands since the very early years of robotics, and therefore well before the birth of social robotics. Indeed, the history of sustained investigations in artificial hands spans at least 30 years and millions of euros in research funding worldwide. Yet, most researchers would frankly acknowledge that the state-of-the-art is not anywhere near to where many research objectives in social robotics need it to be: so far no device has been demonstrated that achieves robust and adaptive grasping in unstructured environments; concerning dexterous manipulation the goal is even farther from being attained. In particular, although many advances have been made in the mechatronics and computational hardware of artificial hands, the state of the art appears to be only marginally closer to a satisfactorily robust and usable approximation of the human hand than it was twenty years ago. A plausible explanation is that the main reasons of the gap are not merely technical, but invest some fundamental issues in the understanding of the organization and control of hands. Ultimately, the main problem appears to be the lack of a principled approach, that is, of a theory guiding scientists in their effort to taming the complexity of hands – meant here as the physical embodiments of the sense of active touch, and comprised of the sensorimotor apparatus that creates the link between perception and action.

In a recent project, the University of Pisa and the Italian Institute of Technology (IIT) teamed up to try and break through the state-of-art of artificial hands by a principled combination of two crucial and innovative concepts: *sensory-motor synergies* and *soft robotics actuation*. It is noteworthy that both concepts were identified through the close interaction of researchers in different areas of engineering, neuroscience, and movement science. This project may bring artificial hands to meet a variety of distinctive manipulation needs that are emerging in social robotics. Let's see.

The human hand is an enormously complex system, with a largely redundant number of receptors, muscles and articular joints. The central nervous system's capacity to control such a complex system in such an extremely simple and effective way is an astonishing fact, considering that neural communication of afferent and efferent signals is much slower than the time constants of the physical phenomena under control. To explain this observation, a principle of *dimensionality reduction* has been often invoked. Neuroscientists have proposed that complexity must be constrained and organized in structures, which are sometimes referred to as *synergies* (Santello, Flanders, & Soechting, 1998). Interestingly, in the mirror neuron system of humans and other primates (Cattaneo and Rizzolatti 2009), motor synergies have been invoked as organizing structures which play a significant role in both action control and perceptual recognition processes. Indeed, mirror neurons have been hypothesized to code simplified motor information based on hand motor synergies; and this information is then made available to achieve computationally feasible and effective action control and perceptual recognition processes (Tessitore et al. 2010).

The *Pisa-IIT SoftHand* project aims at transferring in the sciences of the artificial this principled, synergy-based approach to human hand organization, so as to exploit there its enormous potential for more effective modelling of robotic hands and their practical applications in various areas of robotics. This project hinges on two systems of enabling and interacting synergies, in the hand motor system and in the tactile and kinaesthetic sensory system, respectively. The basic idea is to replicate in a robot hand an organized set of synergies, ordered by increasing complexity, so that a correspondence can be made between any specified task set (in terms of a number of different grasps, explorative actions and manipulations) and the least number of synergies whose aggregation make the task set feasible. Thus, the prime theoretical enabler is an approach to the description of the organization of the hand sensorimotor system in terms of geometric constraints: those are correlations in redundant hand mobility (motor synergies), correlations in redundant cutaneous and kinaesthetic receptor readings (multi-cue integration), and overall sensorimotor control synergies. This sensorimotor organization will be replicated in artificial hands which have to perform various sorts of grasps, explorative actions and manipulations. For instance, a hand whose goal is to realize basic grasps only could use the first two or three synergies in the basis, thus reducing drastically the number and complexity of the actuation and sensory system in most manipulative tasks. One should be careful to note, however, that in some special manipulation tasks this approach may fail to be equally fruitful: a manipulative hand with fine motion control of single joints (such as a piano player's hand) may require coordination of many synergies – perhaps all of them, which in the human hand are around 20.

Let's go on. The hand posture must adapt to task requirements and object properties soon after contact is detected and established, so as to capture the task and object geometry, but need not perfectly match either of these. This approximation can be driven by searching not only in the space of feasible hand configurations but also – and maybe primarily so – in an ideally reduced space of task-specific constraints (feasible set of forces and torques, etc.). Humans are very quick and efficient in learning how to choose a suitable mapping between hand configurations, points of force application, and forces, in an effortless and effective manner. The second key innovation of the *Pisa-IIT SoftHand* is the possibility of controlling forces through tuning the variable compliance of muscle-like, “soft” actuators.

The combination of these ideas leads to the notion of “soft synergies”, which consists in regarding synergy eigenspaces as equilibrium manifolds for the hand (Gabbicini & Bicchi, 2010). Notably, the implementation in the hand of *variable compliance, muscle-like actuators* will allow one to shape this potential field in a suitable way, so as to control contact forces according to the task and the constraints (e.g., slippage avoidance). This project goal builds on previous extensive work concerning the modeling and implementation of *variable stiffness actuators* (Tonietti, Schiavi, & Bicchi, 2005), (Catalano, Schiavi, & Bicchi, 2010). This is necessary to avoid the shortcomings of interpreting synergies as mere mechanical shape primitives, and to introduce the possibility of implementing the idea of soft synergies equilibrium manifolds for the hand, towards which the hand is attracted by a potential field, while being repelled by the obstacle physical boundary.

The overall idea underlying the proposed approach is that the hand embodied speaks a language whose words are the sensorimotor synergies, and that only the understanding of this language will enable us to build artificial systems that bear a resemblance to the human counterpart at a deeper level than mere appearance. Interestingly, by embodying into robotic hands the language of

sensorimotor synergies that one finds in human hands, one would achieve, without additional efforts, a shared action code which may serve the purpose of facilitating human-robot interactions in the context of, e. g., service or social robotics (Prevete et al. 2008).

The thrust of the *Pisa-IIT SoftHand* research (Catalano, Grioli, Farnioli, Serio, Piazza, & Bicchi, 2014) is to capture the fundamental principles of the organization of the hand embodied not by trying to copy the complexity of the biological processes, but rather by capturing how the biomechanical structure and neural control strategies interact to simplify control and sensing problems. One of the *Pisa-IIT SoftHand* results so far is a hand with 19 rolling joints that close in an anthropomorphic way under the control of a single motor (see figure 1). The hand is very robust, and can adapt its grasp to a wide variety of object shapes, by virtue of its implementation of the soft synergy concept. The soft-hand approach to simplification, inspired to the principles of synergistic organization, is expected to lead to designing the mechanics and low-level control of a new hand that exactly match the specifications given, thus enabling practical applications of such devices in industrial, service, and social robotics.



Figure 1 A preliminary prototype of synergy-inspired hand developed

4. Concluding remarks

In spite of their seemingly remote and possibly unattainable character, visionary goals may come to play very useful roles in scientific and technological inquiry. In classical physics, for example, the grand objective of reducing to the laws of mechanics every kind of phenomena studied in the natural sciences extended its influence and was fruitfully pursued throughout the XVIII and XIX centuries (Nagel 1979). In mathematics, the discovery of Gödel's incompleteness theorems dashed Hilbert's ambitious goal of establishing in a mathematically conclusive way that abstract mathematical concepts and theories were free from internal contradictions. Nevertheless, Hilbert's foundational program was productive of significant advances in various areas of mathematics, let alone of entirely new mathematical disciplines (Sieg 2013).

One should be careful to note that visionary goals do not invariably play similarly useful roles in science and technology. In the early days of artificial intelligence, for example, Alan Turing advanced a daring vision of computers possessing rich natural language processing capabilities and

passing what is now known as the Turing test (Turing 1950). This visionary scenario, however, raised many methodological controversies, and arguably played a relatively minor role in orienting the development of AI inquiries towards research objectives which were both rewarding and feasible (Cordeschi 2002, 2007).

In social robotics, limiting and streamlining interactions between robots and human beings appears to be a sensible heuristic strategy enabling one to move from visionary scenarios towards more feasible and rewarding research goals. Enforcing these constraints naturally suggests the idea of a society of robots performing within human societies a variety of useful tasks that require rich social interchanges with other artificial agents, but limited forms only of human-robot interaction, if any.

In addition to providing a unifying horizon for more feasible research objectives, the visionary goals of social robotics and other similarly broad research programs play crucial roles by facilitating local coordination within the relatively disunified landscape of research communities both within and without robotics. These visionary goals establish a trading zone (Galison 1997) involving research communities which seek mutual advantages from the exchange of models, technologies, and systems – even though their research agendas are usually heterogeneous and the pursuit of their regional objectives usually requires no mutual coordination. Accordingly, the need arising in social robotics for dexterous manipulation of objects that are primarily conceived for human use paves the way to potentially rewarding exchanges with the community of researchers working on artificial hands and hand control mechanisms that are based on interacting principles of sensori-motor synergy.

References

Bicchi, A., Fagiolini, A., and Pallottino, L. (2010) “Towards a Society of Robots: Behaviors, Misbehaviors, and Security”, IEEE Robotics and Automation Magazine, vol. 17, p. 26 - 36.

Catalano, M., Schiavi, R., & Bicchi, A. (2010). Design of Variable Stiffness Actuators Mechanisms based on Enumeration and Analysis of Performance. In *Robotics and Automation (ICRA 2010)*. New York: IEEE.

Catalano, M., Grioli, G., Farnioli, E. Serio, A., Piazza, C., and A. Bicchi, (2014) “Adaptive Synergies for the Design and Control of the Pisa/IIT SoftHand”, *International Journal of Robotics Research* 33(5): 768-782.

Cattaneo L, Rizzolatti G (2009). The mirror neuron system. *Archives of Neurology*. 66, 557–560

Cordeschi, R. (2002). *The discovery of the artificial. Behavior, mind and machines before and beyond cybernetics*. Dordrecht: Kluwer.

Cordeschi, R. (2007). AI turns fifty: revisiting its origins. *Applied Artificial Intelligence*, 21, 259-279.

M. Gabbicini, Bicchi, A., Prattichizzo, D., Malvezzi, M. (2011). On the Role of Hand Synergies in the Optimal Choice of Grasping Forces. *Autonomous Robots*. 31, 235 – 252.

Galison, P. (1997). *Image and logic. A material culture of microphysics*. Chicago: University of Chicago Press.

Nagel, E. (1979). *The structure of science*. 2nd edition, Indianapolis: Hackett.

Prevete R, Tessitore G, Catanzariti E., Tamburrini G (2008). Towards a Shared Action Code for Human-Robot Interaction. In *ICKS08, International Conference on Informatics Education and Research for Knowledge-Circulating Society*, pp. 6-12. New York: IEEE.

Sieg, W. (2013). *Hilbert's programs and beyond*. New York: Oxford University Press.

Tessitore G., Prevete R., Catanzariti E., Tamburrini G. (2010). From motor to sensory processing in mirror neuron computational modelling. *Biological Cybernetics* 103(6): 471-485.

Tonietti, G., Schiavi, R., & Bicchi, A. (2005). Design and Control of a Variable Stiffness Actuator for Safe and Fast Physical Human/Robot Interaction. *Robotics and Automation (ICRA 2005)*, pp. 528-533. New York: IEEE.

Turing, A.M. (1950). Computing Machinery and Intelligence. *Mind* 59, 433–460. Reprinted in D.C. Ince, ed. (1992). *Mechanical Intelligence*. Vol. III of *Collected Works of A. M. Turing*, Amsterdam: North-Holland.