Decision-Theoretic Human-Robot Interaction: Designing Reasonable and Rational Robot Behavior

Mary-Anne Williams(✉)

QCIS, University of Technology Sydney and Codex, Stanford University, Stanford, USA
Mary-Anne@TheMagicLab.org

Abstract. Autonomous robots are moving out of research labs and factory cages into public spaces; people's homes, workplaces, and lives. A key design challenge in this migration is how to build autonomous robots that people want to use and can safely collaborate with in undertaking complex tasks. In order for people to work closely and productively with robots, robots must behave in way that people can predict and anticipate. Robots *chose* their next action using the classical sense-think-act processing cycle. Robotists design actions and action choice mechanisms for robots. This design process determines robot behaviors, and how well people are able to interact with the robot. Crafting how a robot will *choose* its next action is critical in designing social robots for interaction and collaboration. This paper identifies *reasonableness* and *rationality*, two key concepts that are well known in Choice Theory, that can be used to guide the robot design process so that the resulting robot behaviors are easier for humans to predict, and as a result it is more enjoyable for humans to interact and collaborate. Designers can use the notions of *reasonableness* and *rationality* to design action selection mechanisms to achieve better robot designs for human-robot interaction. We show how Choice Theory can be used to *prove* that specific robot behaviors are reasonable and/or rational, thus providing a formal, useful and powerful design guide for developing robot behaviors that people find more intuitive, predictable and fun, resulting in more reliable and safe human-robot interaction and collaboration.

Keywords: Human-robot interaction · Designing robot behavior · Legible robot behavior · Predictable robot behavior · Choice theory

1 Introduction

There is a quiet revolution taking place. Robots have been moving from research labs and cages on factory floors into spaces inhabited by people over the last decade. As the service robot industry continues to seek new opportunities that generate value in people's lives, homes and workplaces there is a pressing need to design robots whose behavior people find intuitive so that their interaction and collaboration are enjoyable.

By *enjoyable* we mean legible [4], predictable [6, 8], safe, fluent, and effective. A robot that behaves unintuitive, unpredictably, unsafely, awkwardly or ineffectively is not enjoyable to be near. To be enjoyable robot interactions need to be easy, seamless and natural for humans. We should expect that some training might be needed but that

A. Agah et al. (Eds.): ICSR 2016, LNAI 9979, pp. 72–82, 2016.
DOI: 10.1007/978-3-319-47437-3_8

after being trained a human can work with a robot safely without being surprised, irritated or frustrated.

The current trend in robotics sees robots taking on increasingly expansive roles in society as slave, enabler, protector, companion, entertainer, collaborator and partner. Robot behaviors and social intelligence are fast becoming hot research topics in the field of robotics as researchers frantically seek to develop new scientific methods to: (i) assist in the design of intuitive robot behaviours for complex human-robot interactions and (ii) support synchronized and flexible human-robot real-time joint actions for collaboration.

The challenge for social robot design is to create robot behaviors that people can anticipate because when people work closely with robots they must be able to predict what a robot will do next in order to know how best to respond while undertaking cooperative action with a robot in real-time.

Robots are distributed computer systems that chose what actions they will enact in real-time numerous times a second. Robot behaviour designers craft robot actions and specify the action choice mechanism that determine how a robot will choose its next action: autonomous robots spend their entire life gathering sensor data and creating perceptions of their own body (proprioception) and the external environment (exteroception) that they then use to execute the *action choice* mechanisms.

There are no widely accepted principles that can be used to guide designers on how best to build action choice mechanisms, and as a result robot systems can be difficult to work with, unsafe and unintuitive because their behaviors are hard to interpret, predict, anticipate and explain. It is time to explore how we can introduce more rigorous methods into the design of robot action choice mechanisms, particularly in human-robot interaction and social robot applications.

This paper uses ideas from behavioral Choice Theory [7] to develop a new approach to robot behavior design that leads to *reasonable* and *rational* behaviors that people can understand more easily, and importantly, predict. As a result human-robot interaction and collaboration can be designed to be more enjoyable and productive. Section 2 describes state-of-the-art in robot design. Section 3 discusses the importance of people being able to interpret and predict robot behavior in human robot interaction and collaboration. Section 4 introduces Choice Theory as a formal tool for describing and analysing the robot action choice problem, and explores the idea of what it would mean for a robot to act *reasonably* or *rationally*. Section 5 shows how robot designers can develop reasonable and rational behaviors for robots, thus creating robots that people find easier and more enjoyable to work with during interactions and collaboration. Section 6 summarizes the contribution.

2 State-of-the-Art Robot Behavior Design and Interpretation

A robot is a *real-time distributed computer system* that essentially executes the classical robot *sense-think-act processing cycle*. During this cycle a robot gathers and interprets sensor data from its body and the environment to build high-level perceptions. Using perceptions and knowledge about the available resources, like body parts, the robot "thinks" and *chooses an action* to perform next. As an example, if a robot wants to *lift* a coffee cup and its left arm is busy undertaking another action, the robot could chose to *wait* or use its idle right arm to *lift* the cup. The final stage of the cycle is the *enaction* or performance of the chosen action.

During the time it takes a robot to gather sensor data, select the next action and execute it, there are typically changes within the robot, and in environment. The robot will gather and analyze new data to choose the next action to execute, and so it continues. Robots typically complete the sense-think-act cycle many times a second.

Robots are real-time distributed systems, coordination of their body parts and actuators is complex. A robot's behavior arises from its action choices, the set of available actions and a choice mechanism.

Robot designers play a critical role. They design and develop the set of possible robot actions and craft the robot's *action choice mechanism* that robots use to select the next action to execute in real-time. Once deployed autonomous robots essentially spend their life making sense of their perceptions and choosing their next action.

Designing action choice mechanisms so that the robot makes autonomous and appropriate next action selections is a serious challenge. It is important to note that many robots are reactive and do not have explicit goals, plans or intentions. Even robot soccer players do not typically have intentions, however, *people regularly attribute intention to robots*. For example, they say the robot is "going after the ball", "trying to kick a goal", "looking for a team member to pass to". Studies of human understanding show that if a person fully understands a system they are able to explain it in terms of *underlying mechanisms* e.g. "the robot can calculate how far away the ball is with only one camera using the size of the ball in the image because it knows the size of the ball". If a person is not sure how a system works they often provide a *functional description* e.g. "robots have cameras to see where the ball is". If a person has little idea how a system works then they tend to attribute intention as a means to explain behavior, e.g. "the robot goalie dived because it wanted to stop the other team from scoring a goal". There is a tendency for people to anthropomorphise robots as a means to explain and predict their behaviour as intelligent machines.

The action choice mechanism ultimately determines robot behavior. Robotists develop system architectures and designs that enable robots to make the critical decision of what action to perform next. Robot decisions are complex and always involve uncertainty and risk because a robot's sensor data is noisy; its knowledge and ability to reason is limited; its understanding of its environment and the real world is superficial and often flawed; its perceptions are crude and not always faithful to reality; the robot may not have a clear understanding of its goals, roles, objectives or specific deployment tasks; situations become even more complex for robots when interacting with people.

The field of social robot design is full of *ad hoc* procedures, folk philosophy and folk psychology, and as a result robot behaviors do not follow any guidelines or principles. A designer simply develops behaviors based on their experience regarding what they know "works". It is difficult to scientifically analyse and compare robot behaviors and action selection mechanisms because they are typically idiosyncratic and/or incredibly complex as they attempt to mimic the human brain.

3 Importance of Legible and Predictable Robot Behavior

In human social settings and during collaborative activities it is critically important to be able to interpret, explain, predict, and anticipate other people's behavior. Of course, it is impossible for people to explain or predict every aspect of another person's behavior. However, since people have similar morphologies and use similar communicative signals, and tend to act in roughly "rational" ways it is certainly possible to predict other people's behavior to a large extent such that society can function reasonably effectively.

People do not behave entirely randomly, instead they exhibit predictable patterns of behavior that other people take into account when they plan and execute joint/collaborative actions. People are also able to develop strategies to mitigate the risk of failing to predict other people's behavior, and to respond in real-time when collaborative action goes awry. People tend to mostly behave reasonably and rationally, thus making working together easier and more enjoyable than if they acted unreasonably and irrationally. By contrast, unreasonable or irrational people are hard to understand, difficult to predict and typically not enjoyable to work with.

As robots become increasingly prevalent in society and their tasks require increasingly complex human-robot interactions there is a pressing need to design and develop robots that are easy for people to understand and predict, so that interactions and collaborations are more legible, predictable, safe, fluent, and effective, i.e. enjoyable. Sharing the same physiology helps people interpret and predict each other's behavior because similar sensory stimuli have similar effects on human brains. For example, we all know that a flash of light or loud noise will typically attract a person's attention when it is in their sensor range, we can imagine and explain other people's behavior by introspection and a study of ourselves. Most people make an audible sound when they experience sharp strong pain, and when someone falls over and cries, we know why. Some people are easier to predict than others, and people can adopt deceptive behaviors to mask their action choices and intentions.

Robots do not share the same morphology as people, and our bodily experiences are entirely different, and yet, they can still exhibit behaviors that people can understand and predict. By way of comparison, people are able to predict and control certain aspects of other biological species behavior and can work with some animals in highly productive ways. Not all animals can be tamed and trained. Consider, horse riding where a human controls much of the behavior of a horse: people and horses can work seamlessly together. In contrast, zebras are difficult to harness and work with. Horse riding comes with risk: no matter how skilled a rider, if a horse is surprised or afraid it can react in unpredictable ways. Just as people are able to "predict" animal behavior, there is a need to deploy robots with behaviors that people can predict to an appropriate degree and interact with in an enjoyable way.

It is easier for people to predict certain animals like horses and dogs, than it is for them to predict robots of today: partly because people have little experience with robots, they are not sure what to expect. However, there is a critical difference between animals and robots: robots are designed, and the quality of the design can have a massive impact on how well and how easily humans can predict them.

People should not expect to predict robot behavior all the time, however, they should absolutely expect to be able to predict robot behavior most of the time [5] particularly in circumstances when deviations from expectations are dangerous. Determining when the behavior of a robot is unpredictable is crucial; as this is when it is time to give a robot more physical space to undertake its maneuvers.

Having to deal with unpredictable robots is *not enjoyable*, and it will be costly for society in the same way dealing with unpredictable people can be time consuming and exhausting. Not only are unpredictable robots unsafe, difficult and unpleasant to be around, the lack of predictability is a major obstacle to technological innovation adoption and the expansion of the robot market.

In addition, to the need to be predictable, robots should be able to help people understand some of their actions and to explain their action choices. Choice designers in disciplines like marketing have been able to improve prediction by learning more about how people perceive and make choices. Robotists and robot users will also benefit if they can learn robot choice patterns and interrogate robots to discover their preferences as an explanation for their action choices and subsequent behavior.

We define a robot to be *unpredictable if* its action choices do not have a predictable pattern from a human perspective. Apart from being annoying and irritating, an unpredictable robot may threaten people's well being and cause all kinds of *havoc*, and so there is a pressing need to develop robot systems that people can predict.

Unfortunately, designing and developing predictable robots has proved to be a major challenge and has led to the design of highly deterministic robot designs with limited scripted robot behaviors, which are predictable but hopelessly inflexible, not adaptive and not scalable, thus restricting the range of tasks that robots can be deployed to undertake. The real challenge though lies in building robots that can work closely with people in enjoyable ways. On one hand, people must be able to anticipate robot behaviors, and on the other hand, robots must be able to interpret people's behavior and anticipate them as well.

4 Rational and Reasonable Action Choices

Choice Theory provides a sound approach to reasonable and rational decision-making. It turns out that all rational choices are reasonable, but there are some reasonable choices that are not rational. So rational choices are a subclass of reasonable choices. Rational and reasonable action choices can be used to design more predictable and legible robot behaviors. In this section we describe how robot action choice mechanisms can be described in a Choice Theory framework. In the following section we show how this allows robot behaviors to be designed so that they are *reasonable* or *rational.*

Robot behavior can be specified as a combination of desires, intentions, perceptions, beliefs, skills, actions and action choices: ***goals/desires*** are explicit representations of what a robot is aiming to do; ***plans/intentions*** are series of actions that can be performed to achieve a goal; ***perceptions*** are created from interpreting and combining sensor data; ***actions*** are processes that the robot can execute and/or enact; ***beliefs*** include facts and rules; ***skills*** involve information about when an action can be undertaken; ***action choice***

mechanisms determine the action choices for the robot to select the next action to execute and enact.

Choice Theory focuses on the set of actions that a decision maker, in this case a robot, can choose to enact. It formalizes the use of a preference relation/ordering to encapsulate goals and plans, and drive action choices. Choice Theory explores the selections that underlie patterns of choice and it can be used to prove that robot action choices are *reasonable* or *rational*.

A *robot action choice model* comprises a set of all possible *actions*, **A**, that the robot can perform. A *robot action choice function,* **c** is used to determine the set of actions that a robot could execute **c(A)** at a given time under certain circumstances. For example, a robot might be able to execute any of the actions in its set of possible actions: **A** = {*rotate_head_left, evaluate(x*6), rotate_head_right}* but not all of them simultaneous. At any given time when in a specific state the robot must chose an applicable set of actions, called the *choice set,* that it can actually execute: **c(A)** = {*rotate_head_left), evaluate(x*6)*}.

A *robot action choice function* for a binary relation > and a set of actions **A** is a function **c**(A, >) defined by {x∈A: for all y∈A and y not > x} where the ordering, > , is a *preference relation. a* > *b* is read as *a* is *at least as good as b,* and if it were the case then the robot in a particular state essentially prefers action *a* over action *b*.

It turns out that if the preference relation, > , over actions is acyclic then the robot action function **c**(**A**, >) gives rise to a simple choice function **c**(**A**). Preference relations can be designed and used by robots to prefer action *a* over action *b*, or vice versa, or to be indifferent. Choice Theory provides a number of basic conditions that allow us to classify different kinds of choice functions that are useful in robot design.

In order to define what it would mean for a robot to make reasonable or rational choices we introduce three key conditions. They govern how choices are made across subsets and supersets of choices and impose forms of consistency across these choices: Given a set of robot actions **A**:

i. Choice function **c** satisfies the *contraction condition* if for any choice, **c**(**A**), then **c**(**A**) is chosen if **c**(**A**) is available.
ii. Choice function **c** satisfies the *expansion condition* if actions a, b ∈ **c**(**A**) **A**⊆ **B** and **b** ∈**c**(**B**), then **a**∈ **c**(**B**).
iii. Choice function **c** satisfies the *revelation condition* if actions **a**, **b** are in **A** and **a** ∈ **c**(**A**) then for all **A'**⊆ **A** whenever **b**∈**c(A')** we have **a** ∈ **c(A').**

Robot choices satisfy the contraction condition if whenever the robot chooses a particular action, say **a**, from a set of possible actions, if the possible actions were fewer and action ***a*** is still available, then the robot should choose action ***a*** again. It turns out that the contraction and expansion conditions are consistent and independent, and revelation entails both contraction and expansion, but not conversely. These three properties are used in the next section to show how to construct rational and reasonable robot choices.

5 Designing Reasonable and Rational Robot Behaviors

In this section we consider several important conditions that the action choice mechanism can be designed to satisfy in order to make robot behavior rational and/or reasonable.

Decision makers' appetite for risk often influences the choices selected. Choice Theory uses a notion of "rationality" to mean that an individual acts *as if* balancing costs against benefits to arrive at an action that maximizes personal advantage [24] Applying Choice Theory to robots raises the question of what constitutes "personal advantage" for robots. But for robot designers it is clear, we want robots to achieve their specific deployment tasks.

Proposition 1: Let A_S denote the set of all actions available to a robot in state S. Robot action choices satisfy the *contraction condition* iff the robot action choice $\mathbf{c}(A_S) \subseteq A'$ and $\mathbf{c}(A_S) \subseteq \mathbf{c}(A')$ whenever $A_S \subseteq A'$.

In other words, if a robot's action choices satisfy the contraction condition then reducing the size of the possible set of actions in state S does not change the robot's choice if the selected actions are still available, and conversely.

Proposition 2: Let a robot be in state S and let A_S denote the set of actions available to the robot in state S from the set of all actions A. If there is an additional set of actions B, then robot action choices satisfy the *expansion condition* iff the robot chooses $\mathbf{c}(A_S)$ among $A_S \cup b$ for each action $b \in B$.

Expansion says that if a robot chooses the same set of actions, say. $\mathbf{c}(A_S)$, from an expanded set of actions from B that includes $\mathbf{c}(A_S)$ and any $b \in B$ then it will chose $\mathbf{c}(A_S)$ from the expanded set.

If a robot's action choices satisfy the contraction and expansion conditions, then Choice Theory says its behavior is defined to be ***reasonable***.

The following simple proposition relating choices to preferences is immediate from standard results in Choice Theory, however it is a striking claim in robotics. The notion that a robot's behavior could be classified as "reasonable" is novel in robotics.

Proposition 3: If a *robot exhibits reasonable behavior then it has a preference relation* over its set of actions, A, for every state S.

If a robot is in state S and A_S is the set of actions available to the robot in state S. Robot action choices satisfy the *revelation condition* iff the robot chooses A_i when $A_i \subseteq A$, and whenever the robot chooses A_i it also choses $A_j \subseteq A$. In other words, if the robot chooses action A_i over a second action A_j, then whenever it chooses A_j, then it also chooses A_i whenever it is available.

If a robot's choices satisfy the revelation condition, then Choice Theory says its actions are defined to be ***rational***.

Reasonable action choices are weaker than rational action choices, i.e. rational choices are stricter than reasonable choices as they must satisfy the much stronger revelation condition. As noted earlier rational choices are a special case of reasonable choices.

If a robot acts reasonably people could predict its actions some of the time, but if it acts rationally then it would be possible to predict the robot all the time. In order to achieve this level of perspicuity a robot's preferences would need to be known.

A choice function based on a preference ordering is utility maximizing if for some assignment of utilities the actions chosen are precisely those whose utility is at least the utility of every action.

A robot action choice is *rational* if and only if it can be explained by a preference ordering; an action choice is *rational* if and only if it is utility maximizing. Choice gives rise to utility, and utility is a measure of preference [1].

There is an important difference between using choice models to describe behavior as reasonable or rational, and choice models that can be used to make predictions about *actual* behavior. Since robots are designed decision makers, it is possible to use preference relations to describe and explain robot behavior.

Proposition 4: If a robot exhibits *reasonable behavior* then it has a *preference ordering* over the power set of its actions.

Proposition 5: If a robot's actions are rational, then they are reasonable.

Propositions 1–5, above, show that in order for robots to exhibit reasonable or rational behavior their choices must be disciplined. This will not happen without proactive design steps.

Value-based action selection naturally aligns with Choice Theory because the robot action selection can be described using ordinal or cardinal ranking of actions based on a set of criteria [25]. Hoffman and Breazel [6] aggregate values of actions from several sources to drive robot behavior using a variety of explicit and implicit feedback mechanisms: (i) the strength of the sensory input, (ii) the strength of the motivation, (iii) level of interest to model boredom or behavior-specific fatigue, and (iv) various forms of inhibition. Value-based approaches have also been used to guide action selection in robot teams. Stroupe and Balch [13] used probabilistic values to direct next-step movements of robot teams as they map objects in their environment. It turns out that these methods resulted in robot paths that found vantage points that maximized information gain by reducing the uncertainty of each robot team member's next observation.

6 Improving Human-Robot Interaction

Henzinger and Sifakis (2006) and many others have identified a major chasm between analytical and computational models, and the gap between safety critical and best effort engineering practices. This chasm is particularly disturbing in the robot-human interaction space where people increasingly work in close proximity with robotic technologies, e.g. manufacturing robots, robotic surgery, exoskeletons, and underwater robots. Unless robots are safe and easy to work with, their utility and adoption as a technology will be limited. Unfortunately, the prevailing approach to developing robots that are safe and easy to work with has delivered robot designs that are not adaptable or suitable for open, complex or dynamic environments. This typically means that robots can only achieve structured tasks in predictable and scripted ways; their ability to adapt to new circumstances or achieve complex tasks in dynamic open environments is severely limited.

Robot design delivers a set of actions and a mechanism that allows a robot to choose the next action to execute and enact. Actions are computational processes that robots can execute and enact. Actions can be general computation processes, e.g. pause/wait, arithmetic, database manipulation, or control programs that involve physical actuation such as actuator control. Actions involving actuation require appropriate access to relevant actuators, and in order to be deemed successful they may require certain expectations to be filled, e.g. at the end of the *lift_cup* action, the robot should have lifted a cup.

There are several basic action choice mechanisms widely used in robot systems, which include *reactive mechanisms* that rely on look-up tables in which each stimuli is linked with an explicit response action. Reactive mechanisms are highly deterministic and generate inflexible behaviors: they encapsulate skills with a fixed set of stimulus-response relationships that govern robots behavior. Reactive mechanisms can be implemented as finite state machines. *Behavior-based mechanisms* build on Brook's idea of subsumption [22], which is essentially a layered reactive model. Other kinds of action selection mechanisms include *rule-based selection* [21]; *blackboard architectures* [23]; and *value based selection* using ordinal and cardinal measures of value like cost and risk [6] and concepts of attention competitions [9–11].

7 Discussion

As robots become increasingly prevalent in society and their tasks involve more complex human-robot interactions there is a pressing need to design and develop robot behaviours that are easy for people to understand and predict, so that interactions and collaborations are more legible, predictable, safe, fluent, and effective, i.e. enjoyable.

There are no widely accepted design principles that can be used to guide action selection for social robots that engage in human interaction and collaboration. We addressed this gap by approaching the robot design as *a problem of designing an action choice mechanism*: robots spend their entire life interpreting their sensor data and using it to choose their next action to execute. Action choice mechanisms are fundamental to robot capability and behaviours. Robots behaviours need to be legible and predictable in order for humans to find working with robots to enjoyable. In this paper we used a decision-theoretic approach to argue that the Choice Theory concepts of reasonable and rational choices can be used to show that designed robot behavior is more predictable and legible. The robots that will be the most successful working with people will be the ones that people find enjoyable to work with, and that means those that people can understand and anticipate.

Future work will explore three key research questions (i) how to extend the use of Choice Theory for robots in changing and uncertain circumstances in complex social settings and human-robot interaction scenarios, (ii) how to incorporate theory of mind reasoning mechanisms to enrich robot choices of action in social settings and human-robot interaction scenarios, and (iii) explore the tension between rationality and insanity, where insanity is defined as making the same choices but expecting a different outcome.

References

1. Allingham, M.: Choice Theory. Oxford Press, Oxford (2002)
2. Arrow, K.: Essays on the Theory of Risk Bearing (1971)
3. Bem, D.J., Allen, A.: On predicting some of the people some of the time: the search for cross-situational consistencies in behavior. Psychol. Rev. **81**, 506–520 (1974)
4. Dragan, A.D., Lee, K.C., Srinivasa, S.S.: Legibility and predictability of robot motion. In: Human-Robot Interaction, 2013 8th ACM/IEEE International Conference, pp. 301–308. IEEE, March 2013
5. Epstein, S.: The stability of behavior: I. On predicting most of the people much of the time. J. Pers. Soc. Psychol. **37**(7), 1097–1126 (1979)
6. Hoffman, G., Breazeal, C.: Cost-based anticipatory action selection for human-robot fluency. IEEE Trans. Robot. (T-RO) **23**, 952–961 (2007)
7. Kreps, D.M.: Notes on the Theory of Choice. Westview Press, Boulder (1988)
8. Kruse, T., Basili, P., Glasauer, S., Kirsch, A.: Legible robot navigation in the proximity of moving humans. In: May 2012 IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO), pp. 83–88. IEEE (2012)
9. Novianto, R., Johnston, B., Williams, M-A.: Attention in the ASMO cognitive architecture. In: Proceedings of the First Annual Meeting of the BICA Society, vol. 221 of Frontiers in Artificial Intelligence and Applications, pp. 98–105 (2012)
10. Novianto, R., Johnston, B., Williams, M.-A.: Habituation and sensitisation learning in ASMO cognitive architecture. In: Herrmann, G., Pearson, M.J., Lenz, A., Bremner, P., Spiers, A., Leonards, U. (eds.) ICSR 2012. LNCS, vol. 8239, pp. 249–259. Springer, Heidelberg (2013)
11. Novianto, R., Williams, M.-A.: The role of attention in robot self-awareness. In: 18th IEEE International Symposium Robot and Human Interactive Communication, pp. 1047–1053 (2009)
12. Simon, H.: A behavioral model of rational choice. Q. J. Econ. **69**, 99–188 (1955)
13. Stroupe, A., Balch, T.: Value-based action selection for observation with robot teams using probabilistic techniques. J. Robot. Auton. Syst. **50**, 85–97 (2004)
14. Takayama, L., Dooley, D., Ju, W.: Expressing thought: improving robot readability with animation principles. In: Proceedings 6th International Conference Human-Robot Interaction, pp. 69–76. ACM, March 2011
15. Williams, M.-A.: Robot Social ıntelligence. In: Ge, S.S., Khatib, O., Cabibihan, J.-J., Simmons, R., Williams, M.-A. (eds.) ICSR 2012. LNCS, vol. 7621, pp. 45–55. Springer, Heidelberg (2012)
16. The Fugitive Robot Video, https://www.youtue.com/watch?v=rF_-TmrTan8. IJCAI Best Video Award 2013
17. Mutlu, B.: Nonverbal leakage in robots: communication of intentions through seemingly unintentional behavior. In: 4thACM/IEEE Int'l Conference on Human robot interaction (2009)
18. Terada, K., Shamoto, T., Mei, H., Ito, A.: Reactive movements of non-humanoid robots cause intention attribution in humans. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2007, pp. 3715–3720. IEEE, October 2007
19. Goldman, A.: Desire, ıntention and the simulation theory. In: Malle, B., Moses, L., Baldwin, D. (eds.) Intentions and Intentionality, pp. 207–224. MIT Press, Cambridge (2001)
20. Kögler, H.H., Stueber, K. (eds.): Empathy and Agency: The Problem of Understanding in the Human Sciences. Westview Press, Bolder (2000)
21. Salehie, M., Tahvildari, L.: Self-adaptive software: landscape and research challenges. ACM Trans. Auton. Adapt. Syst., 1–40 (2009)

22. Brooks, R.A.: Intelligence without reason. In: Proceedings of 12th International Joint Conferenceon Artificial Intelligence, Sydney, Australia, August 1991, pp. 569–595 (1991)
23. Engelmore, R.S., Morgan, T. (eds.): Blackboard Systems. Addison-Wesley, Reading (1998)
24. Friedman, M.: Essays in Positive Economics, Chicago (1953)
25. Blumberg, B., Galyean, T.: Multi-level direction of autonomous creatures for real-time virtual environments. In: Proceedings of SIGGRAPH 95 (1995)
26. Henzinger, A., Sifakis, J.: The embedded systems design challenge. In: Proceedings of the 14th International Conference on Formal Methods, pp. 1–15