17 Cognitive Control for Decision and Human-Robot Collaboration

Erwin Jose Lopez Pulgarin, Ute Leonards, and Guido Herrmann

This chapter focuses on the concept of cognitive control in robotics and how it is linked to decision, control, and human-robot interaction (HRI). Achieving a control paradigm that enables robust, flexible goal-driven performance in a myriad of scenarios involving unstructured changing environments and interaction between robots and other agents such as humans has been pursued during the last decade (e.g., Avery, Kelley, and Davani 2006; Baud-Bovy et al. 2014; Herrmann and Leonards 2018). In order to achieve this, inspiration has been taken from nature, with a focus on the way humans and other animals undertake their decision and control processes (see chapter 1). Indeed, by creating controllers inspired by human flexibility and adaptability, some or all of the qualities found in human cognitive processes can be pursued (i.e., adaptability, robustness, goal-driven behavior with sensor and subtask prioritization) in artificial programmable systems.

First, this chapter includes an introduction to the concept of control in the context of industrial processes and expands it to robotics in general; challenges behind robot control will be raised, highlighting the need for novel decision and control architectures for modern robotics such as those involved in closely interacting with humans, dealing with unstructured environments, and learning to better perform a task—hence cognitive control.

Second, the word "cognitive" in the context of control will be defined after an overview about how "cognition" has been used in the literature; the definition of what a cognitive controller is will include aspects about both its architecture and inputs, highlighting how it relates to the term originally used in human behavioral studies and cognitive neuroscience.

Finally, a modeling approach for cognitive control, which integrates the principles of multiagent interaction into a decision-making (i.e., discrete and probabilistic) and control action (i.e., continuous and dynamic) framework will be proposed. This will be followed by a discussion around the framework's elements and their wider impact in different areas of application, such as autonomous driving, teleoperation, and human-humanoid interaction.

17.1 Control in Robotics

When considering any system that interacts with the environment and manipulates it by any means or in any way, the concept of control needs to be considered. Starting from industrial control or process control (Ogata 2010), the main objective of "control" science is to be able to manipulate one or several variables of a system and make them behave as one desires. Control problems can be described generically as either a trajectory-following problem (i.e., make a system's variable follow a set of values) or a regulation problem (i.e., keep a system's variable at a fixed value). Most modern control problems deal with closed-loop control architectures, using sensor or estimation inputs from the system to feed back to the controller; this feedback allows a comparison of expected system outputs with real outputs, which is a prerequisite to modify control outputs based on the state of the system (i.e., outputs). When considering the controller in a system description (Ogata 2010; Maciejowski 2002), it can be described based on its inputs (i.e., single input [SI] or multiple input [MI]) and its outputs (i.e., single output [SO] or multiple output [MO]), with its subsequent combinations (e.g., single input, single output [SISO] or single input, multiple output [SIMO] and so on). This relates to the system's complexity and the control goals—that is, the amount of inputs being how many sensor inputs or control goals the system requires and the amount of outputs as control signals or controlled variables. When considering the controller's inner workings, an explicit understanding of the system to be controlled is used and most of the time is needed in the form of a mathematical description of its dynamics. This understanding and the requirements for control determine how the controller inputs relate to the desired outputs. Based on the level of detail these models require, they could be described using any sort of mathematical description, such as linear operators, nonlinear equations, and probability distributions, usually in a dynamic framework. Performance criteria are imposed on the controller in order to have a complete description of how each variable is controlled (e.g., time to reach the desired value, percentage of error when reaching the desired value, maximum error if the controller overshoots). Finally, controllers can be designed to deal with uncertainty from the system model and to be adaptable to changes in the environment or changes in the model itself. Figure 17.1 shows a general description of a control architecture, considering its required input (i.e., the system's demanded output), the controller that looks to achieve this input, the system, plant, or environment to be controlled, and the sensory input that comes from the system itself.

Bringing these concepts from industrial machinery to the realm of robotics was a straightforward task in the early stages of robotics, as industrial robots had similar physical shapes and objectives compared to industrial machinery (i.e., industrial manipulators were dealing with repetitive tasks with high precision at high speeds). Indeed, most industrial controller designs focused on dealing with low-level control for each link or motor, while

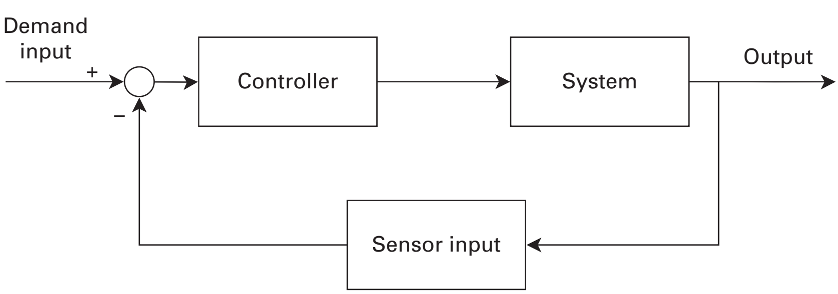

Figure 17.1
General control architecture.

high-level path and trajectory planning was dealt with through solutions based on the robot's geometric properties (e.g., a kinematic description using a Jacobian for end-effector positioning or forward kinematics and motion planning or inverse kinematics). These approaches were highly successful for a wide variety of industrial applications (see LaValle 2006; Scassellati 2002; Visioli and Legnani 2002).

However, as the field of robotics expands, the desire to move robots from industrial setups to more general environments brings challenges beyond what previous approaches can solve. First of all, the goals of a robot outside an industrial setup are potentially more generic and difficult to define completely in advance. Robots thus need to be able to change/adapt over time. For example, taking care of an elderly person could start with only checking their temperature and helping with mobility inside a room but might then evolve into reaching for objects, general social companionship, administering medicine, and more. In addition, using robots outside an industrial setup involves dealing with unstructured, complex, and changing environments that could be difficult to assess or predict at all times. Finally, some applications, such as robots for retailing, teaching, and medical care, would require interaction and/or cooperation with other autonomous agents—be it other robots or human beings (i.e., human-robot interaction). These are all challenges that go well beyond what traditional frameworks focusing on motion control would be able to deal with. Going back to the closed-loop controller description, any such system requires many multiple-input, multiple-output (MIMO) controllers with potentially nonlinear models, configured for both trajectory following and regulation just to focus on general movement alone—for example, to move the robot body to a known location, traverse unknown terrain, or mediate closeness to interacting robots or humans while maintaining safety. Additional components such as high-level decision-making and multimodal communication, supported by specialized hardware such as sensing, actuating, and communication devices, would be necessary to complement the proposed controller (Whitsell and Artemiadis 2017).

The goal here is to find an architecture or methodological approach that can help solve such problems in a complete and integrated manner. To achieve this, inspiration has been drawn from nature and, particularly, from human cognitive processes to better replicate and improve robots in "humanlike capabilities" such as dealing with unstructured and uncertain environments or prioritizing between subtasks and sensory input while maintaining a goal-driven task execution that is adaptable and changes over time. Indeed, human beings are the best-known system to date for adapting to new environments, performing robustly, and prioritizing while reaching a goal. In addition, it has been suggested that a robot that tries to copy or mimic human capabilities by relying on similar mechanisms as the person it is interacting with might be the easiest to understand intuitively (e.g., nonverbally) when interaction between artificial agents and humans is needed (Eder, Harper, and Leonards 2014).

17.2 Cognition in Control and Robotics

The use of the word "cognition" for control has been suggested because it takes inspiration from human cognitive processes. Cognition in humans covers mental processes and their role in thinking, feeling, and behaving, as defined by Kellogg (2015). Cognition includes *perception*—that is, the processing and understanding of the outside world by sensory inputs (Fischer and Demiris 2019); *memory*, or how information is stored, manipulated, and used (Baddeley 2012); *decision-making*, or how to decide on the best action to reach a certain goal (Haefner, Berkes, and Fiser 2016); and *acquisition of knowledge and expertise*, including abstracting high-level understanding and learning from its interaction with the world (Moulin-Frier et al. 2018), among other factors, such as creativity and reasoning, as aspects of human capabilities.

17.2.1 Cognitive Architectures

In robotics, the concept of cognitive architecture comes from research in the field of artificial intelligence to describe a list of components, organizational structures, information flows, representations, and computational procedures that enable some intelligent behavior (Kotseruba and Tsotsos 2020; see also chapter 10); these mechanisms mimic ways the brain is thought to deal with and manipulate information. Such architectures tend to work as blueprints, with no consideration or explanation of how to be implemented in any specific agent. This means that they can be software based only or embodied in the form of a robot body (Kawamura et al. 2008; Wei and Hindriks 2013; see chapter 11). They focus on describing different "cognitive" modules that enable the mimicking of certain intelligent capabilities such as short- and long-term memory modules for better decision-making (Ratanaswasd, Gordon, and Dodd 2005). Such modular descriptions tend to focus on the modules' interconnections, their interaction with the outside world (i.e., environment) in the form of sensor inputs (i.e., stimuli), and their possible control outputs (i.e., action).

A wide range of cognitive architectures have been proposed over the past forty years, each author tackling the problem of representing humanlike intelligence or capabilities in their own way (see Kotseruba and Tsotsos [2020] for a recent review and chapter 10). A possible general classification for these architectures lies in the way information is processed and represented, either by using a handcrafted symbolic representation (i.e., symbolic or cognitivist systems), a sensor and data-based representation (i.e., emergent or connectionist systems), or a mix of both (i.e., hybrid systems; Kotseruba and Tsotsos 2020). Symbolic systems tend to have a long design process because they require a large initial knowledge base including rules, conditions, label descriptions, or possible scenario descriptions. They achieve great predictability and reproducibility, although at the expense of flexibility and robustness to changing environments. In contrast, emergent systems are highly adaptable, suited for learning from the environment and easier to design, but they require potentially long training processes, losing transparency in their results and traceability due to these learning processes. It thus becomes difficult to know what to learn, what exactly is being learned, and when to stop learning in order to achieve optimal performance.

The above classification serves as a parallel to one often employed in control science to describe the mathematics used to design and create the controller itself (Lopez Pulgarin et al. 2018): model-based controllers are designed using a mathematical representation of the system (i.e., plant model) that describes the dynamics surrounding the system. Such controllers are in stark contrast to a data-driven controller that uses available environment measurements to construct a relation between how a system is manipulated (i.e., actions) and the system itself (i.e., states) based on rewarding or punishing certain behaviors and limited to no knowledge of the system itself (e.g., Al-Tamimi, Lewis, and Abu-Khalaf 2007; Na et al. 2012; Lewis, Vrabie, and Vamvoudakis 2012).

17.2.2 Cognitive Controllers

Cognitive controllers then are those that allow the creation of a controller by either implementing or taking inspiration from cognitive architectures (Haykin et al. 2012; Fatemi and Haykin 2014; Kawamura and Gordon 2006). Note that some authors define cognitive control as an addition to other low-level adaptive controllers (Haykin et al. 2012) or as a supplementary way to deal with high sensor input in parallel in a data-driven fashion while ignoring noncritical information (Kawamura and Gordon 2006). Yet, even for such alternative uses of the word "cognitive," authors generally agree on the idea of drawing inspiration from mental models or brain-inspired cognitive architectures. As many cognitive architectures exist, however, there is no single standard of how the components should look (i.e., submodules, types of inputs or outputs, functionality implemented) and thus how these intelligent/mental capabilities are achieved. Figure 17.2 shows an adaptive controller, an extension of the architecture shown in figure 17.1, that allows the model to learn from the environment and inform the controller of some previously unknown parameters in the system to allow it to adapt (Khan et al. 2012; Na et al. 2015). In cognitive architectures, these capabilities are embedded in a cognitive action module, where information derived from perception inform the system how to learn and adapt to the changing and unknown environment.

The main difference between a modern or smart controller (Kawamura et al. 2008) and a cognitive controller is their flexibility in goal description. Although both include interaction with the environment via sensory input and actuation output, having some kind of memory of the environment and the interaction of the controller with it, the cognitive controller is not restricted to one particular task; it has the capability to translate information to other tasks and thus goes beyond initial requirements. In other words, cognitive controllers have the ability to go beyond an initial task definition in order to

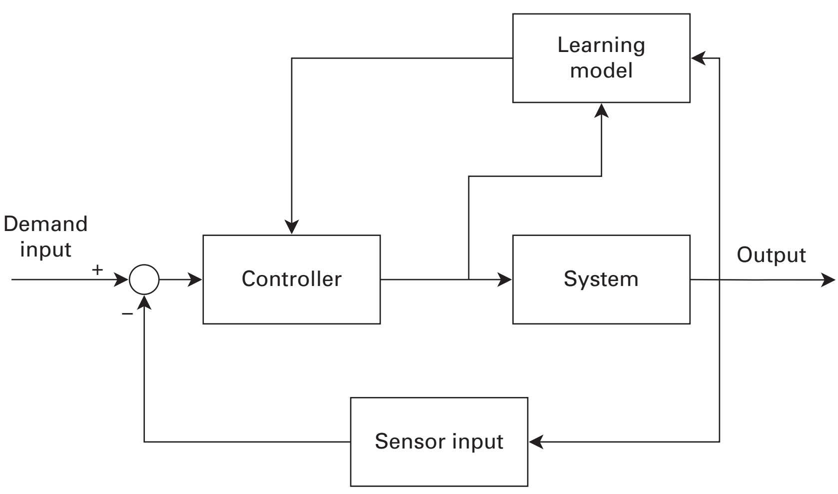

Figure 17.2
Adaptive controller architecture.

achieve an overarching goal through generalization and flexibility (Kotseruba and Tsotsos 2020).

Considering the goal of allowing high-level decision-making and control by a cognitive controller (Kotseruba and Tsotsos 2020), a more detailed cognitive architecture can be formulated by reviewing the specifics of human cognitive processes (Kellogg 2015; e.g., perception, memory, learning). Figure 17.3 introduces the general information loop used in many cognition-inspired applications (Kawamura and Gordon 2006; Ratanaswasd, Gordon, and Dodd 2005), expanding the previously introduced perception and action modules. Sensing and actuation are separated, suggesting they deal only with how sensory information is transformed into useful knowledge and information (i.e., perception) and how the selected decision or sets of actions are performed (i.e., actuation and low-level control), respectively (Haefner, Berkes, and Fiser 2016). A module is added that deals with both the regulation and control of how perception outcomes are used (Gold and Heekeren 2013) and how they can relate to a specific goal such as executive functions or more general goal-related information. An additional module (Ratanaswasd, Gordon, and Dodd 2005) is added that considers how all remaining modules can generate relevant information that could be stored and used to improve their functioning over time and how this process is performed (i.e., learning and memory); the inner workings of this module tend to take inspiration from working-memory models in humans (e.g., Baddeley 2000, 2012).

The information loop of decision-making and control in figure 17.3 implies that for a certain scenario the best possible decision is selected from any set of possibilities by cycling through them and performing any necessary motor control (e.g., limb movement, gaze control, speech). This loop resembles the problem faced in nonlinear control when

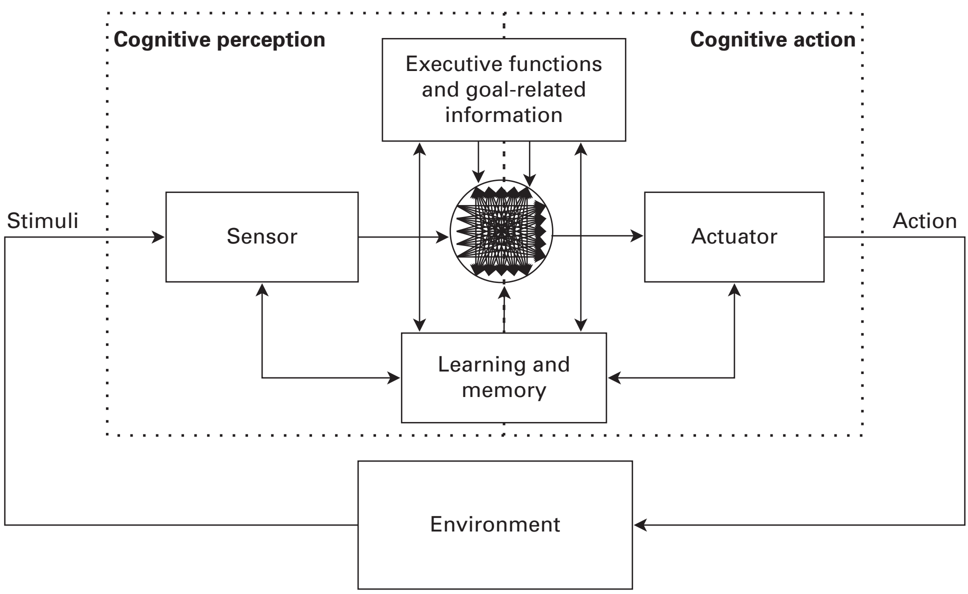

Figure 17.3
Cognitive control architecture with general functional blocks. *Source:* Inspired by Kazahiko Kawamura and Gordon 2006; Ratanaswasd, Gordon, and Dodd 2005.

dealing with uncertain or highly dynamic environments for which a certain controller has been specifically designed or tuned for optimal performance within a specific range of the dynamics, called *gain scheduling* (Yang et al. 2010). The challenges faced in gain scheduling could be seen as a reduced set of those arising in cognitive control: in the former, the cases for which a set of controllers is designed and the controllers themselves are known in advance, and the challenge is to tune the controllers and change from one to the other to maintain performance and stability; in the latter, an additional challenge is to select from an only vaguely defined set of uncertain possibilities and to perform control over them with little to no prior knowledge.

17.2.3 Control in Cognitive Robotics and HRI

Cognitive robotics (Levesque and Lakemeyer 2008; see also chapter 1) arises with the use of cognitive architectures or concepts inspired by these architectures in order to tackle challenges faced in robotics at both task (e.g., object manipulation, exploration) and application levels (e.g., autonomous operation, teleoperation, HRI), respectively. Tasks that have been performed in cognitive robotics range from command following for object manipulation (e.g., Ratanaswasd, Gordon, and Dodd 2005; Dodd and Gutierrez 2005; Kawamura and Gordon 2006; Kawamura et al. 2008) to autonomous navigation (e.g., Avery, Kelley, and Davani 2006; Wei and Hindriks 2013) to reaching a goal by changing tasks (e.g., Khamassi et al. 2011).

Building on such achieved robotic capabilities (e.g., object reaching and navigation), applications that go beyond following an explicit human command have been proposed that tend to involve humans in some aspect or another (e.g., medical aid; Neerincx et al. 2019); hence, Human-robot interaction (HRI) is involved. HRI is the term used to include all the tools and studies around the actuation and interaction of robots with human beings in any possible way (see also chapter 19). Cognitive robotics has proposed a range of methodologies to better interact with humans, such as knowledge and skill transfer from human to robot (e.g., Tan and Liang 2011), knowledge acquisition and learning through interaction (e.g., Moulin-Frier et al. 2018; Nakamura, Nagai, and Taniguchi 2018), and perspective taking (Fischer and Demiris 2019), to name but a few. However, robots with full autonomy have not yet been achieved.

Building from the definition of HRI, a special category focused on scenarios in which robot and human work together to reach a common goal is called human-robot collaboration (HRC). Two key methodological aspects of HRC highlighted by Bauer, Wollherr, and Buss (2008) in their review of the most challenging aspects of HRC are *intention* and *action*; the former considers an initial agreement of the common/joint goal either by explicit (e.g., speech and haptic commands) or implicit (e.g., hand gestures, eye gaze, estimation from physiological signals) means, and the latter considers planning and replanning capabilities to deal with unstructured dynamic environments and a potential joint action (e.g., carrying and sharing a moving load).

HRI brings challenges beyond those previously stated. Even if cognitive processes could be mimicked to better deal with an unstructured and uncertain environment following a certain goal, the challenge of interacting with an autonomous agent who deals with a similar cognitive architecture that requires dynamic change and adaptation is a daunting task. As human beings can perform many different tasks and actions with no guarantee that they will do what the interacting robot expects, robots need to be equipped with the ability to both predict human actions effectively and to clearly communicate their intentions to the interacting human (e.g., Scassellati 2002; Grigore et al. 2013; Eder, Harper, and Leonards 2014; Herrmann and Leonards 2018).

17.3 A Multiagent-Inspired Approach to Control in Cognitive Robots

After having introduced cognitive robotics and its challenges, particularly for advanced HRI applications, we now move on to a decision and control action scheme (DCAS) that provides a clear application framework in which we try to tackle some of the issues raised above. This framework is focused on applications in which spatially close interaction or cooperation between human and robot is either a necessity or would at least improve overall task performance (e.g., semiautonomous vehicles or robotic care). The main challenge in these applications is to achieve safe, cooperative, human-centered, and human-predictive decision-making between a technological robotic device and a goal-oriented human through intelligent control and decision-making.

17.3.1 Paradigm Proposal for a Multiagent-Inspired Dynamic Decision and Action Framework for Human-Robot Interaction

Current state-of-the-art HRI sees the human as "in the loop" and thus as an unpredictable part of the robot's cognitive control system (see, e.g., Eder, Harper, and Leonards 2014). The addition of the human inside a control loop means trying to model the human's requirements, needs, or general behavior in order to minimize any negative effect on task performance or any risk of harming the human in close proximity to the robot while the robot navigates an environment (Dondrup et al. 2015). The uncertainty that arises from the "unpredictable" human can be dealt with safely and reliably as long as the environment in which such interactions happen is well controlled (Eder, Harper, and Leonards 2014). However, problems arise as soon as the environment itself becomes unpredictable. For most everyday environments, this is the case because they often include both other humans and animals (i.e., autonomous agents), making the environment unpredictable and demanding the system to interact or coordinate not only with one unpredictable partner but, potentially, with a whole range of external agents at the same time. Moreover, many physical environments themselves are too complex to be predicted in their entirety, thus leaving further risk of unpredictability. This means that we have an unpredictable part within the system itself as well as an unpredictable, continuously changing environment, a problem that is very hard to solve.

One way to solve this issue is by changing how one understands the directly collaborating partner and their role relative to the robot. If we understand the robot as an autonomous yet collaborative agent in its own right and take the human out of its direct loop by understanding them as an autonomous partner in the robot's environment, then we have to solve only one issue—namely, the dynamic environmental uncertainty or unpredictability. As a partner, the human has built an internal model of the autonomous agent (e.g., robot or another human), as much as the autonomous agent has an internal model regarding the human colead/any other human in the environment. In cognitive psychology terms, such an internal model of an interaction partner's mind would be based on a concept known as *theory of mind* (Baron-Cohen et al. 1985). Theory of mind refers to the attribution of mental states (e.g., intentions, beliefs, and desires) to living beings; for an interaction scenario between two people, an understanding of the other agent's intentions and decision-making process is essential for seamless interaction. Translated to HRI, there is thus only an "intensity" proximity difference or connectivity between the human and other autonomous agents in the environment, comparable to human-human interaction in close proximity or further away (i.e., personal space or extrapersonal space; Curioni, Knoblich, and Sebanz 2017). Hence, we suggest a scenario in which an autonomous system and a human each act as independent autonomous agents. As in human-human interaction, the two interacting partners can then have substantially different abilities as long as their internal representation of each other is sufficiently accurate.

This creates a redundant, safe, and interchangeable cooperative dynamic partnership between the "lead" and "colead" in which both robot and human can take on either role (Curioni, Knoblich, and Sebanz 2017). Communication and cooperation between the autonomous system and the human are a necessity not only for safety reasons but also for the accomplishment of common objectives as determined by the human. The joint action process between an artificial agent and a human being can only realize the optimal outcome of safe and efficient cooperation (i.e., shared control) if the autonomous system is able to synthesize, evaluate, and predict the human colead's intentions and communicate its own possibly limited aims and capabilities to the interacting partner and the environment more generally (figure 17.4). This can be achieved as a cooperative decision and a subsequent dynamic

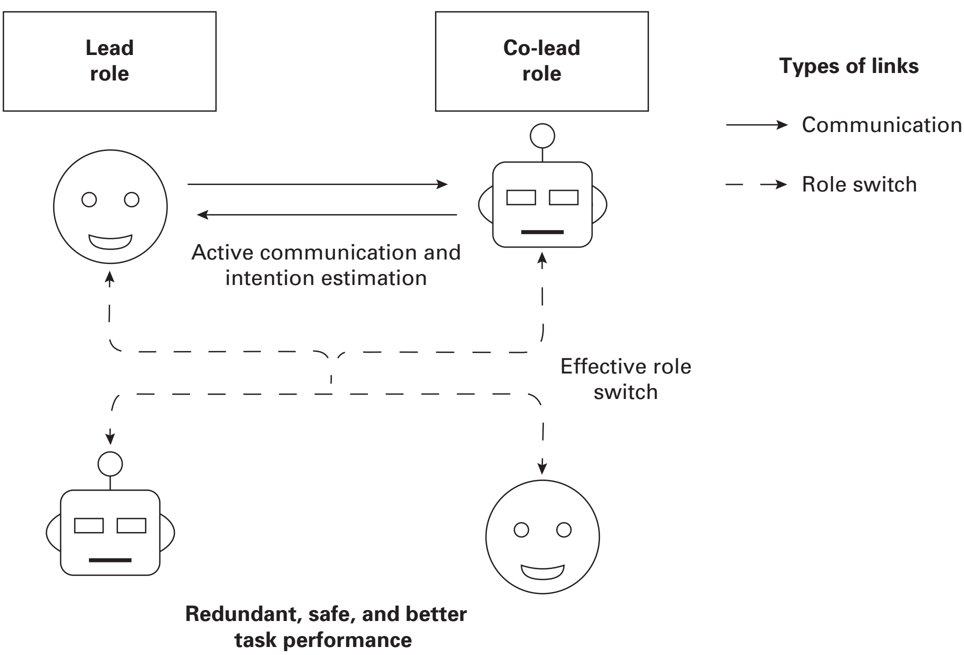

Figure 17.4
Autonomous robot: roles and their interchange between a human and an autonomous system.

action. In fact, such communication and cooperation are key, yet highly problematic, for HRI in general (Herrmann and Leonards 2018). The following suggestion of a decision-and-action framework provides a possible basis for a technical, dynamic HRI control paradigm to deal with interaction issues.

The above proposed paradigm shift in the way we think about fast dynamic interactions between people and artificial autonomous systems (i.e., robots) looks at the interaction and cooperation of two cooperative autonomous agents (figure 17.5) who operate as an interchangeable lead and colead (figure 17.4). Both agents are engaged in the task, and any inattention or objective track loss can be detected immediately. We propose a fluent change between who leads and who follows in joint actions in line with what is known for human-human interaction (Curioni, Knoblich, and Sebanz 2017). Indeed, coordination with others is implicit in many of our human behaviors. Such principles of cooperation can be nicely framed in a theoretical cooperative hybrid decision and dynamic control framework, the technical instantiation of the paradigm shift in dynamic HRI.

We propose that the solution to any human-artificial agent interaction lies in creating an intelligent cooperative decision and actuation framework in which decision-making-relevant information can be seamlessly merged with the human's goals and interests through theory of mind, to the extent necessary and possible. Similar to human-human joint action, the autonomous agent becomes a partner in its own right that is jointly involved in the decision-making process.

Within this cooperative framework, it is important for each agent to be aware that there are other, possibly less capable, autonomous agents in the environment. In this development context, the autonomous system-human relationship can be seen as the pupil (robot)-teacher (human) relationship in a learning stage, with a relationship of a close set of trusted partners as the end goal. The willing and supportive autonomous agent learns how to better interpret and interact (i.e., the autonomous agent learns from and adapts to the human agent). There is also the need for a "human-agent-detection" method to pick up on "error signals" induced during a task (e.g., inattentiveness within the teacher) so corrective actions can be made.

The successful interaction between human and autonomous agent would have to be fluid. This requires both cooperative decisions and cooperative dynamic actions to guarantee a safe and trusted cooperative process during the decisive changeover of leader and follower. For such a technical mechanism of cooperative interaction between two autonomous agents to work, the guiding principle that underlies this cooperation needs to be based on optimality, a principle well known in engineering (Turnbull et al. 2016) and robotics (Mombaur, Truong, and Laumond 2010; Khan et al. 2012) as well as an underlying concept to cognitive science (Berkes et al. 2011; Fiser et al. 2010), where it has been shown that under most circumstances humans decide and dynamically act in an optimal sense (e.g., Spiers, Khan, and Herrmann 2016; Haefner, Berkes, and Fiser 2016).

Putting the different concepts together, a hybrid optimal, yet adaptive, cooperative agent-based decision and control action scheme (i.e., DCAS) must provide the "intelligence" as an active negotiation scheme between autonomous agent and human. This scheme must resolve both the dynamic, the physical, and the behavioral event-driven interaction between human and autonomous system. To date, this is still an important unresolved step.

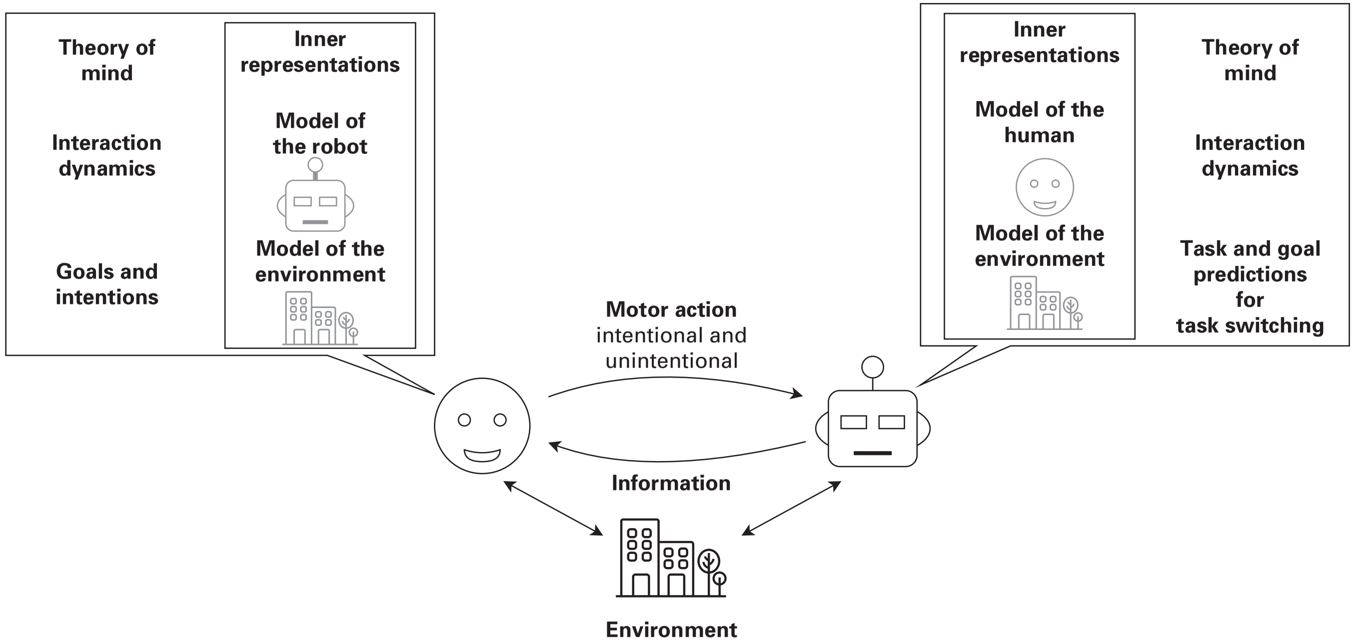

Figure 17.5
Human-autonomous agent interaction principles, appropriated from human-human interaction.

17.3.2 Principles and Characteristics of the Dynamic Decision and Action Framework

Based on predictions of the possible decisions a human agent could make (Lopez Pulgarin, Herrmann, and Leonards 2018), any DCAS should look at making decisions within a fraction of those prediction windows (e.g., one second) to then dynamically cue and actively influence the decisions of the human partner. Hence, human and autonomous agent would be able to cooperatively act within the time period of the predicted decisions and actions thereafter. The following axioms would lead to the DCAS:

1) The realization that we can treat the human as an "external source" or independent collaboration partner in relation to the autonomous agent instead of "in the loop."

2) Learning from and adapting to the human and the signals people send in joint action situations, proxemics, and so on; learning to understand and predict adaptability within the human and their trust of the autonomous system as an autonomous, collaborative agent or partner.

3) Identification and subsequent learning from the "error" signals when situations go wrong. The principle of optimality of decisions and actions in human agents and control technology is an exploited commonality.

4) Identification and enabling of verbal and nonverbal communication channels in a human to indicate changes in "who is the leader, who is the follower" in joint action.

5) Subsequent "joint" cooperative, agent-based decision-making and dynamic action taking.

Overall, a coherent modeling methodology for decisions and actions would have to be developed that is deeply rooted in complementary research on human decision-making, cognition, communication, dynamic actions, dynamic decisions, and action theories in control and computer science.

This requires that agent models and their uncertainties involved in the joint decision-making process be predetermined. This includes both the human and the autonomous system. The more the autonomous system relies on principles that underlie successful human-human interaction, the easier it will be for the human to develop a theory of mind of the robot. Only an approach that allows the human to intuitively understand the "mind" of the robot and that takes into account that an agent's own actions influence other agents' actions and vice versa will make joint actions among intelligent autonomous systems and humans possible (King, Rowe, and Leonards 2011).

Autonomous artificial agent models take inspiration from the fact that human decision models (e.g., Bellet et al. 2009; Berkes et al. 2011) have strong similarity to discrete hybrid stochastic automata (DHSA; Bemporad and Di Cairano 2005). There is a decision-making level that is responsible for the decisions, resulting in subsequent dynamic actions at the automatic level. Hence, the decision-making level may imply a set of discrete yet uncertain decisions, each followed by an uncertain dynamic action. Decisions are carried out within a fraction of a second, while dynamic actions can extend over intervals of several seconds.

The probabilistic approach for the analysis of human decision-making based on Fiser's sampling-based probabilistic representational framework (Haefner, Berkes, and Fiser 2016; Fiser et al. 2010) is a possible guidance for the development of such agent models. In

Fiser's framework, both the human's internal representation of visual, aural, and tactile events during acting as a colead and the decision-making process in lead situations must be assessed. For the sequential character of decisions and dynamic actions, it therefore becomes necessary to explore how decisions in the present moment depend on the series of decisions made in the recent past. This leads to an assessment process of cues given to the human and the decisions made. For modeling the human decision-making process, the optimality principle following a Bayesian method can be used, such as the "cognitive tomography" method of Houlsby et al. (2013). Applied to behavioral tasks, this allows for a quantitative description of an internal representation of a human based on discrete test choices (figure 17.6). Alternatively, a machine-learning-based understanding of the decision-making model (Lopez Pulgarin, Herrmann, and Leonards 2017) could be deployed and the synergies explored in which decision probabilities determine decision costs. Though such methods resemble emergent methods in cognitive architectures, they aim at presenting their results in a clearer and more predictable manner than traditional data-driven methods.

For the lower automatic dynamic action level—that is, the dynamic action following the decision—learning-based, regressive models based on data-driven methods might be preferable to strongly physical model-based methods; they may provide a continuous integral or summative optimal cost function that the human follows. Optimal cost function models allow for a more flexible prediction of the human's actions. This is, for example, used in inverse optimal reinforcement learning (Mombaur, Truong, and Laumond 2010). Both levels are joint via the DHSA (Bemporad and Di Cairano 2005) and exploit mechanisms

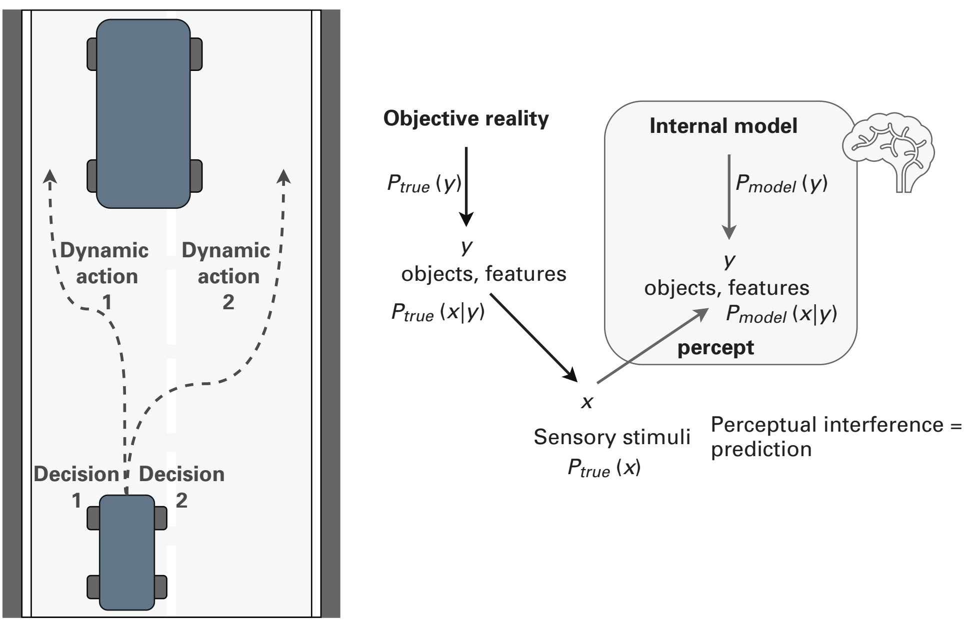

Figure 17.6
Probabilistic internal decision model of a driver attempting to pass a car in front of them. *Source:* Adapted from Berkes et al. 2011.

like model predictive control (Morari and Lee 1999; Di Cairano et al. 2014; Rosolia, Zhang, and Borrelli 2018).

As mentioned earlier, a joining principle in human decision, human action dynamics, and many artificially designed technological processes is *optimality*. Each decision and action can be quantitatively associated with a cost. For a robot, a part of the cost in a dynamic action can be characterized, for instance, by its distance, either to a target for target tracking or to the distance from a human for safety. For both humans and robots, their energy consumption could be included in the cost and would be expected to increase over time while remaining limited in order for it to be optimized. The optimization of energy consumption underlies many human functions, such as locomotion (e.g., Warren 2006).

In terms of decision-making, the synergetic power of cognitive-science-founded models (e.g., Fiser et al. 2010; Berkes et al. 2011; figure 17.6) and machine-learning models (e.g., Lopez Pulgarin, Herrmann, and Leonards 2017, 2018) has to be exploited. Humans develop an internal model for each perceptive decision that guarantees that the decision regarding an intended task is carried out with the highest probability of success (Fiser et al. 2010; Berkes et al. 2011) considering the uncertainty of the environment (figure 17.7). Hence, decision costs are inversely related to the probability of the decision made. Identifying not only the models and their uncertainty sources but the optimal criteria for joint action between agents is key (Fiser et al. 2010).

The cooperative decision-making process can use the set of aforementioned DHSAs within a cooperative agent-based process, using model predictive control principles, to speed up the decision process and to allow fast computation of dynamic control actions from the multiagent framework. A probabilistic decision framework would possibly enhance such a process (Turnbull et al. 2016). For this, a virtual autonomous agent (figure 17.7) can

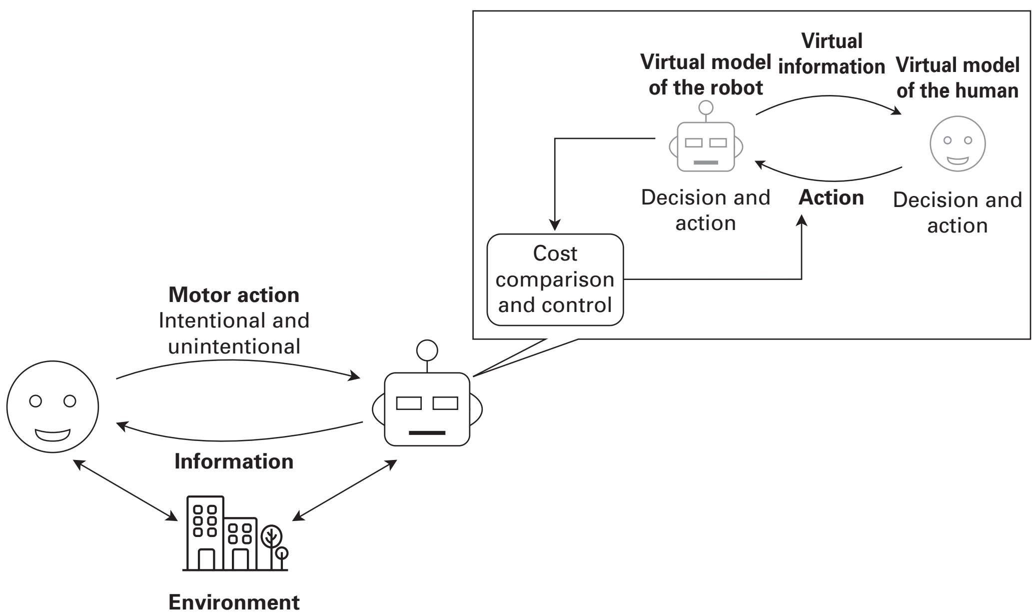

Figure 17.7
DCAS overview.

be developed by applying the principles behind DCAS. This virtual agent represents the nominal action computed from a joint optimal criterion for safety and a nominal understanding of the human, the internal model of the virtual leader, including individual differences between humans and in their intentions. The virtual model will act as an agent to be compared with the human characteristics and its short-term predictions using the unconstrained human model. Hence, the human and the virtual agent are assessed for their cost function, which evaluates whether the human cooperative partner is in line with the virtual human model. Subcomponents for safety are prioritized in decision-making, together with representations of human intention to decide to what extent the human or the virtual agent lead within the collaboration within a network of decentralized agents. Principles of game-theoretic approaches and agent synchronization can be used for a control policy in the vital time frame of dynamic actions following a decision, thus leading to an action of the cooperative decision-making process. To minimize conflicts, the autonomous system will take the human's desired actions as long as they do not compromise safety.

17.3.3 Impact on Autonomous Systems and HRI

Below we will analyze cognitive control for HRI with different types of robots and human-robot collaboration scenarios.

Humanoid robots

Humanoid robots (Oh, Kim, and Kim 2005) and interaction with them (i.e., human-humanoid interaction (HHI; Herrmann and Leonards 2018) could be enabled or improved by implementing a DCAS similar to that described above. Robots performing tasks that benefit from understanding the interacting human(s) while aiming at a final goal, such as to jointly move an object, to keep the human safe, or to maintain a human's vital signs inside a desired threshold, are the key benefits of the DCAS. Similar to existing cognitive architectures that aim at achieving humanlike capabilities, DCASs would allow many robotics applications to improve human life.

Understanding the interacting human and having the ability to share certain goals would be a big step toward safe, trustworthy HRI. For example, applications in medical assistive robotics could range from robots serving as partial nurses or assistants to medical professionals to shared physical cooperative work (e.g., object carrying; Parker and Croft 2012) or object manipulation (Sheng, Thobbi, and Gu 2015; Whitsell and Artemiadis 2017). By understanding the final goal that both the robot and medical professional share, meaning patient care, auxiliary actions could be performed by the robot across the whole care experience.

For cases in which the human is the recipient of the robot's actions and not the cooperative leader or companion, substantial benefits would be derived from understanding the human recipient's mindset in order to take the appropriate decisions at the best time possible.

Although the DCAS's main goal is not restricted to better understanding a robot's surrounding environment, it is one of its planned capabilities. Hence, the DCAS should improve the robot's autonomy during its sensing and decision-making processes by means of a collaborative learning strategy (e.g., supervised learning). By learning from the sensed environment while keeping a preset goal, long-term goals can be achieved autonomously and cooperatively as decision-making is improved across task iterations.

Teleoperated robots

Robotic teleoperation, understood to be the operation of a robot at a distance that allows one or many operators to interact with an environment (Li, Xia, and Su 2015), can benefit from the use of DCAS. As the scope of both operation and distance in teleoperation can be very wide (e.g., operation being by direct control or control by commands and distance understood as either a physical distance or difference in scale), many applications include a teleoperation setup (e.g., robotic surgeon, robotic manipulator for maintenance).

As in other HRI examples, DCAS would improve interaction to achieve a shared goal. Even if teleoperated robots are not considered autonomous or able to make decisions, the robot could possess intelligent mechanisms to help improve overall task performance—for example, to deal with potential delays in communication channels or complications introduced by control means or interfaces. By considering the robotic teleoperation device as a cooperative agent that understands and predicts the human operator's actions, the impact of delays could be minimized, as shared control would be made possible. This has been proposed before (e.g., Corredor, Sofrony, and Peer 2017), but here the idea is applied to a multitask and multidimensional space. Following a paradigm of a shared control, the level of autonomy in teleoperation devices could increase with improved understanding of the teleoperation task and increased safety.

In particular, higher autonomy of the system could speed up the operator's learning curve to use the device. Learning curve theory started empirically in the 1930s as cost reduction due to repetitive procedures in production plants was observed (see Anzanello and Fogliatto [2011] for the full reference); its goal is to exemplify and track how proficiency in performing a task or in the use of a device is improved via repetition (i.e., experience). Learning curves have been applied in teleoperation (e.g., Anvari 2007) to evaluate how much training is needed with using a device to achieve proficiency (Doumerc et al. 2010). Learning curves have been used in the field of medicine, particularly to evaluate both manual surgical procedures (e.g., Hopper, Jamison, and Lewis 2007; de Oliveira Filho 2002) and robotically assisted surgical procedures (e.g., Kaul, Shah, and Menon 2006; Chen et al. 2017) and to compare the two types of procedures with each other.

Building on the results around learning curves for robotic teleoperated devices, particularly in medicine (e.g., Yamaguchi et al. 2015; Samadi et al. 2007), a general learning curve can be proposed. Figure 17.8 shows the potential shape of the learning process behind a robotic device when plotting performance against experience. Three different phases can be identified: 1) an initial slow learning phase in which the operator gets used to the device until it reaches some minimal proficiency pg_1 after certain experience tp_1, 2) a second practicing phase in which an acceptable proficiency pg_2 is achieved after continuous training tp_2, and 3) a mastery phase in which optimal performance pg_3 is reached with continuous training and repetition.

DCAS could reduce training times tp_1 and tp_2 by making the teleoperation device both more intuitive and more responsive to the operator's needs. In addition, the gap between pg_1 and pg_2 could be reduced following the principle previously explained, ultimately leading to improvement in overall performance (i.e., push pg_3 higher).

The training of operators is an important task of teleoperation devices when autonomy levels of the teleoperation system are low. However, as the autonomy of a teleoperated robot increases, following autonomy levels similar to those declared by the Society of

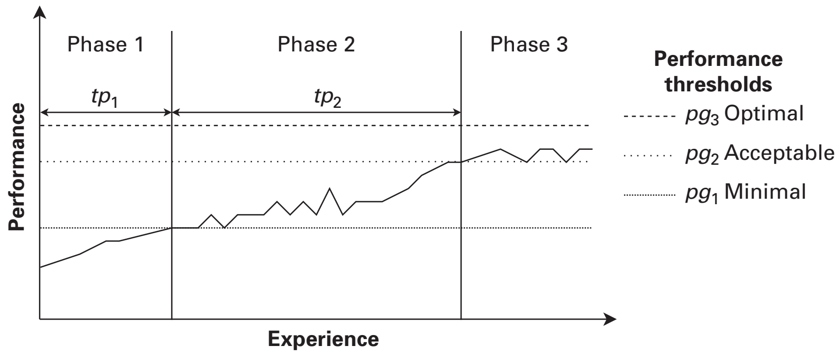

Figure 17.8
Potential learning curve for the teleoperation of robotic devices.

Automotive Engineers (SAE; SAE International 2016), the DCAS could be an important enabler of improved teleoperation. Indeed, in many respects a teleoperated task is similar to a vehicular driving task in which increasing autonomy is introduced for improved performance, decreased human operator workload, and, ultimately, higher levels of safety.

Autonomous vehicles

Autonomous vehicles are a key target of many companies, as they could potentially bring significant economic and societal benefits (Fagnant and Kockelman 2015). Enormous structural efforts have been undertaken in terms of legislation and technology to enable autonomous driving. This includes the introduction of high-bandwidth G5 communication technology as an important enabler of autonomous driving through connectivity between cars or for high-precision maps. At the same time, the diverse and historically grown character of cities poses a challenge in its own right, with partially outdated infrastructure, differences in road regulations, and a highly dynamic environment due to other road users.

Albeit error prone, humans are fully capable of steering around a city's complexities. They can interpret complex situations, make decisions, resolve problems, and even reinterpret rules and road regulations within new contexts. Autonomous vehicles fail in such situations (see, e.g., fatal accidents with regard to Uber and Tesla; Banks, Plant, and Stanton 2018), meaning the human needs to remain included in the driving process. In addition, a significant number of countries, especially within Europe, demand a human-focused approach that requires the driver to be able to retake control at any moment—something that is not possible if a person has been occupied with a different task.

However, not only autonomous cars make mistakes. Drivers can be expected to make mistakes commensurate with the cognitive load they have to deal with or when they lack situational awareness through distraction or mind wandering (de Winter et al. 2014). While some advanced driver assistance systems (ADAS) and semiautonomous driving technologies try to account for human inattentiveness (e.g., Fagnant and Kockelman 2015), the majority work independently. Yet a more alert and experienced driving partner and copilot would be able to help steer a driver out of a temporary problem by direct communication or by providing supportive and intuitive cues for the driver.

Whether to allow human passengers to interact with autonomous cars remains an unresolved problem affecting cockpit design (Fagnant and Kockelman 2015), in addition to the aforementioned uncertainty of understanding the human within the vehicle as part of the car's system (i.e., the human in the loop) and separating it conceptually from the external environment.

The DCAS suggested in this chapter interacts with the human in the car (i.e., the driver) in a cooperative way (figure 17.5), like a human pilot would with their human copilot. As pilot and copilot swap roles, so do the artificial agent (i.e., autonomous car) and the human driver, considering the requirements at hand and allowing the human agent to retake control of the driving process if desired.

17.4 Conclusion

A DCAS was introduced as a response to some of the challenges faced in modern robotics, such as goal-driven task performance and flexible and robust interaction with autonomous agents and the environment, as well as learning and knowledge acquisition. This decision and control framework was inspired by cognitive architectures and is expected to benefit many fields of application inside and beyond robotics. A list of a DCAS's major capabilities would be to

1. enable the robot's interaction with humans by understanding the human's goals and current state,
2. provide an agent-based description for both human and robot in order to enable joint action or cooperative work,
3. deal with partial or incomplete representations of the environment and the interacting agents using learning, and
4. exploit commonalities of recent research in human decisions and actions and existing predictive decision and action methodologies in control and decision theory.

However, many aspects of such a DCAS remain open questions, specifically of how to implement a cohesive mathematical framework around each of the scheme's components or capabilities. Going back to the previous list, some of its key challenges are as follows:

1. Human state and intention estimation and prediction
   1.1. What measurements can we use to help estimate or predict a human intention related to a certain task?
   1.2. How do we generate estimations or predictions of a human before, during, or after a task is being performed?
   1.3. How do we keep track of these estimations or predictions and update them as a task is being performed?
2. Task performance and coordination
   2.1. How do we make the robot perform a certain task or part of it?
   2.2. How do we let the robot know when to stop performing the task?
   2.3. How do we make the robot stop performing the task and release partial or complete control over a task?
   2.4. How do we let the robot know when to take back partial or full control of the task?
   2.5. How do we make the robot take back control of the task?

3. Decision-making and action with incomplete models
   3.1. How do we integrate a learning process in a decision-making and control application?
   3.2. How do we learn from performing a task and interacting with a human?
   3.3. How do we convert sensed data and the learning process into knowledge useful for task completion and goal reaching?

Some technical insight has been given into how to answer these questions. A data-driven approach, taking advantage of both machine-learning (e.g., Lopez Pulgarin, Herrmann, and Leonards 2017, 2018; Khamassi et al. 2011) and probabilistic-sampling techniques (e.g., Nakamura, Nagai, and Taniguchi 2018; Haefner, Berkes, and Fiser 2016; Fiser et al. 2010), has been proposed as a feasible solution to improve understanding of the environment and to create knowledge, acknowledging challenges around modeling and validating and integrating the proposed methods into a more general cognitive control framework. Discrete hybrid automata (e.g., Bemporad and Di Cairano 2005) and model predictive control (e.g., Morari and Lee 1999) have been proposed as solutions for handling several action paths simultaneously (i.e., decision-making) and implementing controllers, with some others using reinforcement learning (i.e., data-driven methods) to deal with both situations (e.g., Lopez Pulgarin et al. 2018; Haykin et al. 2012; Khan et al. 2012; Khamassi et al. 2011). Hence, a suggested major joint guiding principle of these methods is optimality in discrete decisions and dynamic actions for dynamic autonomous agent-based cooperation. Some authors have managed to integrate data-driven methods with dynamical systems for control (e.g., Warren 2006), which again keeps the discussion going about how to better achieve a cognitive controller that takes advantage of symbolic (i.e., model-based) and emergent (i.e., data-driven) representations in cognitive architectures for control.

After introducing the concept of cognitive control and cognitive robotics, including its benefits and challenges, we hope to have sparked more interest in this promising research field while sharing some ideas and concepts developed over the past few years.

Acknowledgments

We would like to acknowledge the enormous contributions given by the following people in the form of discussions and idea sharing, which shaped the concepts described in this document. We would like to thank, in alphabetical order, Murad Abu-Khalaf, Eric Armengaud, Phil Barber, Gabriel Baud-Bovy, József Fiser, Tobias Kessler, Alois Knoll, Weiru Liu, Majid Mirmehdi, Henrik J. Putzer, Francesco Rea, Arthur Richards, Markus Rickert, Giulio Sandini, Alessandra Sciutti, and Robert Wragge-Morley.

Additional Reading and Resources

- An interesting book with applied examples of controllers for robotic arms movement: Spiers, Adam, Said Ghani Khan, and Guido Herrmann. 2016. *Biologically Inspired Control of Humanoid Robot Arms*. Cham, Switzerland: Springer.
- A comprehensive overview of some of the challenges in human-humanoid interaction inspiring work in cognitive robotics: Eder, Kerstin, Chris Harper, and Ute Leonards. 2014.

"Towards the Safety of Human-in-the-Loop Robotics: Challenges and Opportunities for Safety Assurance of Robotic Co-workers." In *23rd IEEE International Symposium on Robot and Human Interactive Communication*, 660–665. New York: IEEE.

- A specific overview on optimal control and reinforcement learning, some of the techniques used in advanced control applications: Khan, Said G., Guido Herrmann, Frank L. Lewis, Tony Pipe, and Chris Melhuish. 2012. "Reinforcement Learning and Optimal Adaptive Control: An Overview and Implementation Examples." *Annual Reviews in Control* 36 (1): 42–59. https://doi.org/10.1016/j.arcontrol.2012.03.004.
- ROS packages for symbolic planning and robot task planning: https://moveit.ros.org/, http://wiki.ros.org/smach, http://wiki.ros.org/flexbe.
- Software packages to get started with data-driven control (RL):
  - MATLAB (proprietary but with better documentation): https://uk.mathworks.com /products/reinforcement-learning.html.
  - PYTHON (free and more popular) for algorithms: https://github.com/openai/baselines; testing environments: https://github.com/openai/gym; use with robotic simulators http: //wiki.ros.org/openai_ros.
- Software packages to get started with traditional control and model-based control (MPC):
  - Optimization solver: https://osqp.org/.
  - MATLAB (proprietary but with better documentation) control toolbox: https://uk .mathworks.com/products/control.html; MPC toolbox: https://uk.mathworks.com/products /mpc.html; modeling language wrapper: https://yalmip.github.io/.
  - PYTHON (free) control library: https://python-control.readthedocs.io/en/latest/; free modeling language wrapper: https://www.cvxpy.org /.

References

Al-Tamimi, Asma, Frank L. Lewis, and Murad Abu-Khalaf. 2007. "Model-Free Q-Learning Designs for Linear Discrete-Time Zero-Sum Games with Application to H-Infinity Control." *Automatica* 43 (3): 473–481. https://doi .org/10.1016/j.automatica.2006.09.019.

Anvari, M. 2007. "Remote Telepresence Surgery: The Canadian Experience." *Surgical Endoscopy and Other Interventional Techniques*. Berlin: Springer. https://doi.org/10.1007/s00464-006-9040-8.

Anzanello, Michel Jose, and Flavio Sanson Fogliatto. 2011. "Learning Curve Models and Applications: Literature Review and Research Directions." *International Journal of Industrial Ergonomics* 41 (5): 573–583. https://doi .org/10.1016/j.ergon.2011.05.001.

Avery, Eric, Troy Kelley, and Darush Davani. 2006. "Using Cognitive Architectures to Improve Robot Control: Integrating Production Systems, Semantic Networks, and Sub-Symbolic Processing." In *Simulation Interoperability Standards Organization : 15th Conference on Behavior Representation in Modeling and Simulation*, 190–198.

Baddeley, Alan. 2000. "The Episodic Buffer: A New Component of Working Memory?" *Trends in Cognitive Sciences* 4 (11): 417–423. https://doi.org/10.1016/S1364-6613(00)01538-2.

Baddeley, Alan. 2012. "Working Memory: Theories, Models, and Controversies." *Annual Review of Psychology* 63 (1): 1–29. https://doi.org/10.1146/annurev-psych-120710-100422.

Banks, Victoria A., Katherine L. Plant, and Neville A. Stanton. 2018. "Driver Error or Designer Error: Using the Perceptual Cycle Model to Explore the Circumstances Surrounding the Fatal Tesla Crash on 7th May 2016." *Safety Science* 108:278–285. https://doi.org/10.1016/j.ssci.2017.12.023.

Baron-Cohen, Simon, Alan M. Leslie, and Uta Frith. 1985. "Does the Autistic Child Have a 'Theory of Mind'?" *Cognition* 21 (1): 37–46.

Baud-Bovy, Gabriel, Pietro Morasso, Francesco Nori, Giulio Sandini, and Alessandra Sciutti. 2014. "Human Machine Interaction and Communication in Cooperative Actions." In *Bioinspired Approaches for Human-Centric Technologies*, 241–268. Dordrecht: Springer. https://doi.org/10.1007/978-3-319-04924-3_8.

Bauer, Andrea, Dirk Wollherr, and Martin Buss. 2008. “Human-Robot Collaboration: A Survey.” *International Journal of Humanoid Robotics* 5 (1): 47–66. https://doi.org/10.1142/S0219843608001303.

Bellet, Thierry, Béatrice Bailly-Asuni, Pierre Mayenobe, and Aurélie Banet. 2009. “A Theoretical and Methodological Framework for Studying and Modelling Drivers’ Mental Representations.” *Safety Science* 47 (9): 1205–1221. https://doi.org/10.1016/j.ssci.2009.03.014.

Bemporad, Alberto, and Stefano Di Cairano. 2005. “Optimal Control of Discrete Hybrid Stochastic Automata.” In *Lecture Notes in Computer Science*, 151–167. 3414. Berlin: Springer. https://doi.org/10.1007/978-3-540-31954-2_10.

Berkes, Pietro, Gergo Orbán, Máté Lengyel, and József Fiser. 2011. “Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment.” *Science* 331 (6013): 83–87. https://doi.org/10.1126/science.1195870.

Breazeal, Cynthia. 2004. “Social Interactions in HRI: The Robot View.” *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews* 34 (2): 181–186. https://doi.org/10.1109/TSMCC.2004.826268.

Chen, Po Da, Chao Yin Wu, Rey Heng Hu, Chiung Nien Chen, Ray Hwang Yuan, Jin Tung Liang, Hong Shiee Lai, and Yao Ming Wu. 2017. “Robotic Major Hepatectomy: Is There a Learning Curve?” *Surgery* (United States) 161 (3): 642–649. https://doi.org/10.1016/j.surg.2016.09.025.

Corredor, Javier, Jorge Sofrony, and Angelika Peer. 2017. “Decision-Making Model for Adaptive Impedance Control of Teleoperation Systems.” *IEEE Transactions on Haptics* 10 (1): 5–16. https://doi.org/10.1109/TOH.2016.2581807.

Curioni, Arianna, Gunther Knoblich, and Natalie Sebanz. 2017. “Joint Action in Humans: A Model for Human-Robot Interactions.” In *Humanoid Robotics: A Reference*, edited by A. Goswami and P. Vadakkepat, 1–19. Dordrecht, Switzerland: Springer. https://doi.org/10.1007/978-94-007-7194-9_126-1.

de Oliveira Filho, Getúlio Rodrigues. 2002. “The Construction of Learning Curves for Basic Skills in Anesthetic Procedures: An Application for the Cumulative Sum Method.” *Anesthesia and Analgesia* 95 (2): 411–416. https://doi.org/10.1213/00000539-200208000-00033.

de Winter, Joost C. F., Riender Happee, Marieke H. Martens, and Neville A. Stanton. 2014. “Effects of Adaptive Cruise Control and Highly Automated Driving on Workload and Situation Awareness: A Review of the Empirical Evidence.” *Transportation Research Part F: Traffic Psychology and Behavior* 27:196–217. https://doi.org/10.1016/j.trf.2014.06.016.

Di Cairano, Stefano, Daniele Bernardini, Alberto Bemporad, and Ilya V. Kolmanovsky. 2014. “Stochastic MPC with Learning for Driver-Predictive Vehicle Control and Its Application to HEV Energy Management.” *IEEE Transactions on Control Systems Technology* 22 (3): 1018–1031. https://doi.org/10.1109/tcst.2013.2272179.

Dodd, Will, and Ridelto Gutierrez. 2005. “The Role of Episodic Memory and Emotion in a Cognitive Robot.” In *IEEE International Workshop on Robot and Human Interactive Communication, 2005*, 692–697. New York: IEEE. https://doi.org/10.1109/ROMAN.2005.1513860.

Dondrup, Christian, Nicola Bellotto, Marc Hanheide, Kerstin Eder, and Ute Leonards. 2015. “A Computational Model of Human-Robot Spatial Interactions Based on a Qualitative Trajectory Calculus.” *Robotics* 4 (1): 63–102. https://doi.org/10.3390/robotics4010063.

Doumerc, Nicolas, Carlo Yuen, Richard Savdie, M. Bayzidur Rahman, Kris K. Rasiah, Ruth Pe Benito, Warick Delprado, Jayne Matthews, Anne Maree Haynes, and Phillip D. Stricker. 2010. “Should Experienced Open Prostatic Surgeons Convert to Robotic Surgery? The Real Learning Curve for One Surgeon over 3 Years.” *BJU International* 106 (3): 378–384. https://doi.org/10.1111/j.1464-410X.2009.09158.x.

Eder, Kerstin, Chris Harper, and Ute Leonards. 2014. “Towards the Safety of Human-in-the-Loop Robotics: Challenges and Opportunities for Safety Assurance of Robotic Co-workers.” In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, 660–665. New York: IEEE. https://doi.org/10.1109/ROMAN.2014.6926328.

Fagnant, Daniel J., and Kara Kockelman. 2015. “Preparing a Nation for Autonomous Vehicles: Opportunities, Barriers and Policy Recommendations.” *Transportation Research Part A: Policy and Practice* 77:167–181. https://doi.org/10.1016/j.tra.2015.04.003.

Fatemi, Mehdi, and Simon Haykin. 2014. “Cognitive Control: Theory and Application.” *IEEE Access* 2:698–710. https://doi.org/10.1109/ACCESS.2014.2332333.

Fischer, Tobias, and Yiannis Demiris. 2019. “Computational Modelling of Embodied Visual Perspective-Taking.” *IEEE Transactions on Cognitive and Developmental Systems* 12 (4): 723–732. https://doi.org/10.1109/TCDS.2019.2949861.

Fiser, József, Pietro Berkes, Gergo Orbán, and Máté Lengyel. 2010. “Statistically Optimal Perception and Learning: From Behavior to Neural Representations.” *Trends in Cognitive Sciences* 14 (3): 119–130. https://doi.org/10.1016/j.tics.2010.01.003.

Gold, Joshua I., and Hauke R. Heekeren. 2013. “Neural Mechanisms for Perceptual Decision Making.” In *Neuroeconomics: Decision Making and the Brain*, edited by P. Glimcher and E. Fehr, 355–372. 2nd ed. San Diego: Elsevier. https://doi.org/10.1016/B978-0-12-416008-8.00019-X.

Grigore, Elena Corina, Kerstin Eder, Anthony G. Pipe, Chris Melhuish, and Ute Leonards. 2013. "Joint Action Understanding Improves Robot-to-Human Object Handover." In *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4622–4629. New York: IEEE. https://doi.org/10.1109/IROS.2013.6697021.

Haefner, Ralf M., Pietro Berkes, and József Fiser. 2016. "Perceptual Decision-Making as Probabilistic Inference by Neural Sampling." *Neuron* 90 (3): 649–660. https://doi.org/10.1016/j.neuron.2016.03.020.

Haykin, Simon, Mehdi Fatemi, Peyman Setoodeh, and Yanbo Xue. 2012. "Cognitive Control." *Proceedings of the IEEE* 100 (12): 3156–3169. https://doi.org/10.1109/jproc.2012.2215773.

Herrmann, Guido, and Ute Leonards. 2018. "Human-Humanoid Interaction: Overview." In *Humanoid Robotics: A Reference*, edited by A. Goswami and P. Vadakkepat, 1–16. Dordrecht: Springer. https://doi.org/10.1007/978-94-007-7194-9_146-1.

Hopper, A. N., M. H. Jamison, and W. G. Lewis. 2007. "Learning Curves in Surgical Practice." *Postgraduate Medical Journal* 83 (986): 777–779. https://doi.org/10.1136/pgmj.2007.057190.

Houlsby, Neil M. T., Ferenc Huszár, Mohammad M. Ghassemi, Gergő Orbán, Daniel M. Wolpert, and Máté Lengyel. 2013. "Cognitive Tomography Reveals Complex, Task-Independent Mental Representations." *Current Biology* 23 (21): 2169–2175. https://doi.org/10.1016/j.cub.2013.09.012.

Kaul, Sanjeev, Nikhil L. Shah, and Mani Menon. 2006. "Learning Curve Using Robotic Surgery." *Current Urology Reports*. Berlin: Springer. https://doi.org/10.1007/s11934-006-0071-4.

Kawamura, Kazuhiko, and Stephen M. Gordon. 2006. "From Intelligent Control to Cognitive Control." In *2006 World Automation Congress, WAC'06*. New York: IEEE. https://doi.org/10.1109/wac.2006.376003.

Kawamura, Kazuhiko, Stephen M. Gordon, Palis Ratanaswasd, Erdem Erdemir, and Joseph F. Hall. 2008. "Implementation of Cognitive Control for a Humanoid Robot." *International Journal of Humanoid Robotics* 5 (4): 547–586. https://doi.org/10.1142/S0219843608001558.

Kellogg, Ronald Thomas. 2015. *Fundamentals of Cognitive Psychology*. 3rd ed. Thousand Oaks, CA: Sage.

Khamassi, Mehdi, Stéphane Lallée, Pierre Enel, Emmanuel Procyk, and Peter F. Dominey. 2011. "Robot Cognitive Control with a Neurophysiologically Inspired Reinforcement Learning Model." *Frontiers in Neurorobotics* 5:1. https://doi.org/10.3389/fnbot.2011.00001.

Khan, Said G., Guido Herrmann, Frank L. Lewis, Tony Pipe, and Chris Melhuish. 2012. "Reinforcement Learning and Optimal Adaptive Control: An Overview and Implementation Examples." *Annual Reviews in Control* 36 (1): 42–59. https://doi.org/10.1016/j.arcontrol.2012.03.004.

King, Dorothy, Angela Rowe, and Ute Leonards. 2011. "I Trust You; Hence I like the Things You Look At: Gaze Cueing and Sender Trustworthiness Influence Object Evaluation." *Social Cognition* 29 (4): 476–485. https://doi.org/10.1521/soco.2011.29.4.476.

Kotseruba, Iuliia, and John K. Tsotsos. 2020. "40 Years of Cognitive Architectures: Core Cognitive Abilities and Practical Applications." *Artificial Intelligence Review* 53 (1): 17–94. https://doi.org/10.1007/s10462-018-9646-y.

LaValle, Steven M. 2006. *Planning Algorithms*. Cambridge: Cambridge University Press.

Levesque, Hector, and Gerhard Lakemeyer. 2008. "Cognitive Robotics." In *Foundations of Artificial Intelligence*, chap. 23. San Diego: Elsevier. https://doi.org/10.1016/S1574-6526(07)03023-4.

Lewis, Frank L., Draguna Vrabie, and Kyriakos G. Vamvoudakis. 2012. "Reinforcement Learning and Feedback Control: Using Natural Decision Methods to Design Optimal Adaptive Controllers." *IEEE Control Systems* 32 (6): 76–105. https://doi.org/10.1109/mcs.2012.2214134.

Li, Zhijun, Yuanqing Xia, and Chun Yi Su. 2015. *Intelligent Networked Teleoperation Control*. Berlin: Springer. https://doi.org/10.1007/978-3-662-46898-2.

Lopez Pulgarin, Erwin Jose, Guido Herrmann, and Ute Leonards. 2017. "Drivers' Manoeuvre Classification for Safe HRI." In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 475–483. 10454 LNAI. Berlin: Springer. https://doi.org/10.1007/978-3-319-64107-2_37.

Lopez Pulgarin, Erwin Jose, Guido Herrmann, and Ute Leonards. 2018. "Drivers' Manoeuvre Prediction for Safe HRI." In *IEEE International Conference on Intelligent Robots and Systems*, 8609–8614. New York: IEEE. https://doi.org/10.1109/iros.2018.8593957.

Lopez Pulgarin, Erwin Jose, Tugrul Irmak, Joel Variath Paul, Arisara Meekul, Guido Herrmann, and Ute Leonards. 2018. "Comparing Model-Based and Data-Driven Controllers for an Autonomous Vehicle Task." In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 170–182. 10965 LNAI. Berlin: Springer. https://doi.org/10.1007/978-3-319-96728-8_15.

Maciejowski, J. M. 2002. *Predictive Control: With Constraints*. Essex, UK: Pearson Education.

Mombaur, Katja, Anh Truong, and Jean Paul Laumond. 2010. "From Human to Humanoid Locomotion—an Inverse Optimal Control Approach." *Autonomous Robots* 28 (3): 369–383. https://doi.org/10.1007/s10514-009-9170-7.

Morari, Manfred, and Jay H. Lee. 1999. "Model Predictive Control: Past, Present and Future." *Computers and Chemical Engineering* 23:667–682. https://doi.org/10.1016/S0098-1354(98)00301-9.

Moulin-Frier, Clement, Tobias Fischer, Maxime Petit, Gregoire Pointeau, Jordi Ysard Puigbo, Ugo Pattacini, Sock Ching Low, et al. 2018. "DAC-H3: A Proactive Robot Cognitive Architecture to Acquire and Express Knowledge about the World and the Self." *IEEE Transactions on Cognitive and Developmental Systems* 10 (4): 1005–1022. https://doi.org/10.1109/tcds.2017.2754143.

Na, Jing, Muhammad Nasiruddin Mahyuddin, Guido Herrmann, Xuemei Ren, and Phil Barber. 2015. "Robust Adaptive Finite-Time Parameter Estimation and Control for Robotic Systems." *International Journal of Robust and Nonlinear Control* 25 (16): 3045–3071. https://doi.org/10.1002/rnc.3247.

Na, Jing, Xuemei Ren, Cong Shang, and Yu Guo. 2012. "Adaptive Neural Network Predictive Control for Nonlinear Pure Feedback Systems with Input Delay." *Journal of Process Control* 22:194–206. https://doi.org/10.1016/j.jprocont.2011.09.003.

Nakamura, Tomoaki, Takayuki Nagai, and Tadahiro Taniguchi. 2018. "SERKET: An Architecture for Connecting Stochastic Models to Realize a Large-Scale Cognitive Model." *Frontiers in Neurorobotics* 12 (12): 25. https://doi.org/10.3389/fnbot.2018.00025.

Neerincx, Mark A., Willeke van Vught, Olivier Blanson Henkemans, Elettra Oleari, Joost Broekens, Rifca Peters, Frank Kaptein, et al. 2019. "Socio-cognitive Engineering of a Robotic Partner for Child's Diabetes Self-Management." *Frontiers in Robotics and AI* 6:118. https://doi.org/10.3389/frobt.2019.00118.

Ogata, Katsuhiko. 2010. *Modern Control Engineering*. 5th ed. London: Pearson. https://doi.org/10.1201/9781315214573.

Oh, Kwang-Myung, Ji-Hoon Kim, and Myung-Suk Kim. 2005. "Development of Humanoid Robot Design Process-Focused on the Concurrent Engineering Based Humanoid Robot Design." In *IDC International Design Congress 2005*, 1–13. International Design Congress. Yunlin, Taiwan: National Yunlin University of Science and Technology.

Parker, Chris A. C., and Elizabeth A. Croft. 2012. "Design and Personalization of a Cooperative Carrying Robot Controller." In *Proceedings—IEEE International Conference on Robotics and Automation*, 3916–3921. New York: IEEE. https://doi.org/10.1109/icra.2012.6225120.

Ratanaswasd, Palis, Stephen Gordon, and Will Dodd. 2005. "Cognitive Control for Robot Task Execution." In *IEEE International Workshop on Robot and Human Interactive Communication, 2005*, 440–445. New York: IEEE. https://doi.org/10.1109/roman.2005.1513818.

Rosolia, Ugo, Xiaojing Zhang, and Francesco Borrelli. 2018. "Data-Driven Predictive Control for Autonomous Systems." *Annual Review of Control, Robotics, and Autonomous Systems* 1 (1): 259–286. https://doi.org/10.1146/annurev-control-060117-105215.

SAE International. 2016. *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. https://doi.org/10.4271/J3016_201609.

Samadi, David, Adam Levinson, Ari Hakimi, Ridwan Shabsigh, and Mitchell C. Benson. 2007. "From Proficiency to Expert, When Does the Learning Curve for Robotic-Assisted Prostatectomies Plateau? The Columbia University Experience." *World Journal of Urology* 25 (1): 105–110. https://doi.org/10.1007/s00345-006-0137-4.

Scassellati, Brian. 2002. "Theory of Mind for a Humanoid Robot." *Autonomous Robots* 12 (1): 13–24. https://doi.org/10.1023/A:1013298507114.

Sheng, Weihua, Anand Thobbi, and Ye Gu. 2015. "An Integrated Framework for Human-Robot Collaborative Manipulation." *IEEE Transactions on Cybernetics* 45 (10): 2030–2041. https://doi.org/10.1109/tcyb.2014.2363664.

Spiers, Adam, Said Ghani Khan, and Guido Herrmann. 2016. *Biologically Inspired Control of Humanoid Robot Arms*. Cham, Switzerland: Springer.

Tan, Huan, and Chen Liang. 2011. "A Conceptual Cognitive Architecture for Robots to Learn Behaviors from Demonstrations in Robotic Aid Area." In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 1249–1252. New York: IEEE. https://doi.org/10.1109/IEMBS.2011.6090294.

Turnbull, Oliver, Jonathan Lawry, Mark Lowenberg, and Arthur Richards. 2016. "A Cloned Linguistic Decision Tree Controller for Real-Time Path Planning in Hostile Environments." *Fuzzy Sets and Systems* 293:1–29. https://doi.org/10.1016/j.fss.2015.08.017.

Visioli, Antonio, and Giovanni Legnani. 2002. "On the Trajectory Tracking Control of Industrial SCARA Robot Manipulators." *IEEE Transactions on Industrial Electronics* 49 (1): 224–232. https://doi.org/10.1109/41.982266.

Warren, William. 2006. "The Dynamics of Perception and Action." *Psychological Review* 113 (2): 358–389. http://search.proquest.com/docview/214221535/.

Wei, Changyun, and Koen V. Hindriks. 2013. "An Agent-Based Cognitive Robot Architecture." In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in*

Bioinformatics), edited by M. Dastani, J. F. Hübner, and B. Logan, 54–71. 7837 LNAI. Berlin: Springer. https://doi.org/10.1007/978-3-642-38700-5_4.

Whitsell, Bryan, and Panagiotis Artemiadis. 2017. "Physical Human-Robot Interaction (PHRI) in 6 DOF with Asymmetric Cooperation." *IEEE Access* 5:10834–10845. https://doi.org/10.1109/ACCESS.2017.2708658.

Yamaguchi, Tomohiro, Yusuke Kinugasa, Akio Shiomi, Sumito Sato, Yushi Yamakawa, Hiroyasu Kagawa, Hiroyuki Tomioka, and Keita Mori. 2015. "Learning Curve for Robotic-Assisted Surgery for Rectal Cancer: Use of the Cumulative Sum Method." *Surgical Endoscopy* 29 (7): 1679–1685. https://doi.org/10.1007/s00464-014-3855-5.

Yang, Weiwei, Guido Herrmann, Mark Lowenberg, and Xiaoqian Chen. 2010. "Dynamic Gain Scheduled Control in a Multi-variable Control Framework." In *Proceedings of the IEEE Conference on Decision and Control*, 7081–7086. New York: IEEE. https://doi.org/10.1109/cdc.2010.5717054.