Cognitive Control

Martin Buss, Sandra Hirche, and Tariq Samad

Introduction

As the field of control engineering has evolved, its horizons have continually broadened. From regulation with simple proportional-integral-derivative (PID) loops, to model-based control and multivariable schemes, to explicit incorporation of uncertainty in robust and modern control theory, to hybrid and hierarchical architectures, and most recently, to control of and via networks, both theoretical foundations and application scope have seen dramatic advancement.

What's next, we might wonder? In this section we outline one prospective answer: "cognitive control." We believe that the incorporation of properties we usually associate with cognition—including reasoning, planning, and learning—within control systems holds the promise of greatly expanding the scope and impact of the field.

The scope and impact of control systems could be substantially increased with the incorporation of properties we usually associate with cognition, such as reasoning, planning, and learning.

We consider cognitive control to be an enabler for novel technologies in many diverse application areas. Field robotics, space and sea exploration systems, and next-generation unmanned aerial vehicles will achieve a higher degree of autonomy through cognitive function. Cognitive control systems for manufacturing plants will be partners to plant operators and engineers; less human intervention will be necessary even as the safety and performance of plants improve. Similar benefits can also be expected from cognitive systems assisting or ultimately replacing human operators in supervisory control applications (for example, in power generation/distribution, traffic control, and similar infrastructure-oriented domains). Search and rescue missions, especially in environments that are remote or inhospitable for humans, will also be an important application domain. Assistive technologies for the elderly are another target, and an increasingly important one given aging populations in many developed countries—cognitive control systems can help overcome both physical and cognitive impairments by enabling the elderly and infirm to live independently as well as by assisting human health workers in caring for them.

The behaviors, functions, and features required of envisioned cognitive control systems have always been part of the vision of control engineering—as articulated in motivating research in areas such as adaptive, robust, and intelligent control. This vision, however, is not much in evidence in the conferences and journals in the field. Specific research in control has focused on narrower—and better defined—problem formulations. Yet the relevance of control methodologies and tools to the broader vision is not in question. The rigor and "systems" orientation of control will be instrumental for realizing cognitive control systems in practice, and by virtue of both its intellectual depth and its record of success across all engineering fields, the controls community is ideally positioned to spearhead the development of cognitive control systems.

Below we first discuss what motivates cognitive control as a research field. We then explain in broad terms what we mean by cognitive control. Related work in other fields is outlined, and we highlight the crucial role of control science and engineering. We conclude with discussion of some challenge problems and associated research questions for cognitive control.

Motivation: Why Cognitive Control?

Current automated systems function well in environments they are designed for, that is, around their nominal operating conditions. They also function well in environments with "predictable" uncertainties as treated, for example, in the advanced adaptive and robust control frameworks—and as demonstrated in modern engineering systems such as unmanned aerial vehicles (UAVs) and process plants without on-site operators. However, control systems of today require substantial human intervention when faced with novel and unanticipated situations—situations that have not been considered at the controller's design stage. Such situations can arise from discrete changes in the environment, extreme disturbances, structural changes in the system (for example, as a result of damage), and the like. To illustrate, future autonomous robots in search and rescue operations, in mining, in the service domain, and in autonomous driving will regularly encounter novel situations that require perception, reasoning, decision making, fact generalization, and learning. Such cognitive control aspects will play a major role in future automated and autonomous systems and will advance "automation" to the next level.

But fully autonomous systems represent just one direction for cognitive control research. Today's control systems for applications such as aircraft, chemical factories, and building systems automate many operational functions while simultaneously aiding human operators in doing their jobs. More cognitive abilities in such control systems will enable safer and higher performance semiautonomous engineering systems.

This human-automation interaction aspect suggests another important focus for cognitive control: social and group environments. Multiagent coordination and control, cooperative execution of complex tasks, effective operation in competitive or mixed competitive-cooperative situations all require the participating agents, whether human or machine, to have cognitive capabilities. In this context, communication takes on added importance and complexity. Agents will need linguistic sophistication. Shared semantic models and ontologies will be necessary. Beyond semantics, just as people rely on pragmatics in their use of language—much of what we convey through speech or writing is not directly related to the literal meaning of our utterances—so will cognitive control systems.

Definition/Description of the Topic: What Is Cognitive Control?

Attempting to define the notions of "cognition" and "cognitive system" is a controversial endeavor, as shown dramatically by the 40-plus diverse definitions of cognition that were collected within the "euCognition" project funded by the European Commission [1]. Rather than attempt a necessary and sufficient definition, we describe several fundamental ingredients of cognitive control, without any claim of completeness.

A system under cognitive control

- exhibits goal-oriented behavior in sensing, reasoning, and action;
- flexibly changes its goals and behavior depending on situational context and experience;

- is able to act in unstructured environments without human intervention and robustly responds to surprise; and
- is able to interact with humans and other cognitive systems to jointly solve a complex task.

To achieve these properties, a system under cognitive control needs to

- understand the present situation (including awareness of itself, its environment, and other agents)—to this end, the cognitive control system must implement several functions, such as (active) sensing, the extraction and abstraction of relevant information, acquisition of semantic knowledge, comparison with previous experience, and knowledge updating;
- purposefully act to modify the current situation and react to any unpredicted changes in a reasonable (not necessarily optimal) way—components required include decision making, planning, reasoning, learning, and adaptation.

An important characteristic is that full information is rarely available to construct models. Hence, the mechanisms for estimating the current state as well as for purposeful modification of this state need to operate on partial/uncertain information.

In Fig. 1, a cognitive control system architecture is proposed showing the possible components of the system:

- Perception includes the acquisition of low-level sensor data, data fusion, information processing and abstraction, and the interpretation of the information for decision making. The question is, how can important (that is, task-relevant information) be reliably filtered from the vast amount of noisy and incomplete data. Major challenges are the inclusion of contextual/semantic knowledge for more robust signal processing and interpretation and the development of active (multi-modal) sensing and signal processing strategies.

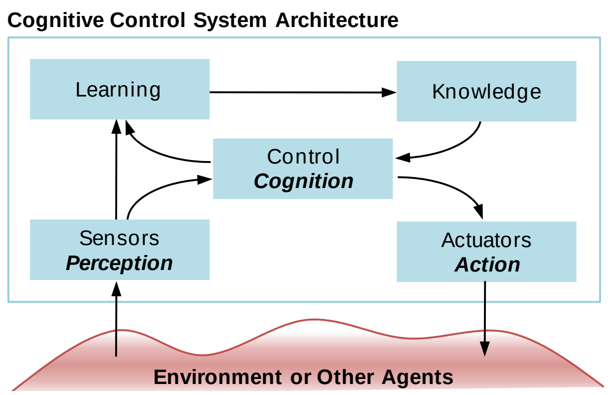

Figure 1. The perception-cognition-action loop—a proposal for a cognitive control system architecture (CoTeSys®) [2].

- Control maps percepts onto actions using existing knowledge/experience. One of the major challenges is to combine semantics with continuous and discrete signal-based representations and to produce a reasonable control decision in the presence of incomplete and/or uncertain information.
- Actions implement the output of the control element, thereby affecting the external environment of the cognitive control system. Both symbolic and continuous actions may be required, similar to the structure of the control output.
- Learning is essential to updating existing knowledge, resulting in the online adaptation of cognitive functionalities to changing environmental situations and contexts. Learning under

partial/incomplete information, hierarchical learning, and learning of symbolic temporal sequences, relations, and concepts present some of the major challenges in the area.

- Knowledge or memory/experience represents a fundamental feature of cognitive control systems. In contrast to classical approaches, this knowledge is continuously updated and modulates the task execution at runtime. An important aspect is the representational formalism for knowledge, such as the choice of representational primitives, compositions, and structure.

One limitation of Fig. 1 is that it does not show interagent interactions separately from the inputs and outputs associated with the environment. At some level of abstraction, other agents and the environment are both part of the external world of an agent, but an agent's ways of engaging will be very different with both. These differences need to be explicitly addressed in a more complete architectural design.

Relevant Neighboring Disciplines and Rationale for a Leadership Role for Controls

The area of engineered cognitive systems has so far been dominated by the artificial intelligence (AI) and computer science communities. These disciplines, together with areas of neuroscience, cognitive science, and psychology, represent the most relevant neighboring disciplines. Their contributions so far and their role in cognitive control are highlighted below. In addition, the contributions of operations research, embedded real-time systems, signal processing, and pattern recognition have been helping to advance the field and are expected to continue to do so in the future.

Artificial Intelligence and Computer Science

Within artificial intelligence and computer science research, advanced methods for reasoning, planning, decision making, and learning have been investigated over the past decades and successfully applied in information-based systems. However, their impact on systems interacting with the physical world has been limited. Such "cyber-physical systems" certainly require a deep understanding of dynamical systems (including hybrid systems that combine continuous and discrete dynamics) and feedback loops, concepts that are fundamental to control. Accordingly, existing theories need to be reformulated to include dynamical system properties. Relevant topics from AI for the area of cognitive control include

- theories of reasoning under uncertainty, sequential logic reasoning, rule-based systems, and inference machines;
- knowledge representation, reasoning about knowledge, and use of prior knowledge;
- machine learning, probabilistic learning methods, reinforcement learning, and statistical learning.

The state of the art is regularly demonstrated in benchmarking competitions such as the DARPA Grand Challenge (2005), Urban Challenge (2007), Grand Cooperative Driving Challenge (2011), and the RoboCup (yearly since 1997). (Control technologists have also been involved in, and in several cases have successfully led, entries in these competitions.)

Neuroscience, Cognitive Science, and Psychology

Neuroscience, cognitive science, and psychology can stimulate research in cognitive control by providing insights on fundamental mechanisms of natural (biological) cognition. Progress in technology for measuring brain activity has provided and will continue to provide results that are useful for engineering purposes concerning the function and architecture of the brain and their relationship to human behavior. These results are relevant for the area of cognitive control from two standpoints:

- Natural cognition as a role model for artificial cognition: The understanding of the principal mechanisms of decision making, learning, abstraction, and other functions may guide the development of artificial cognition.
- Joint human-machine cognition: To design machines with cognitive functionalities that help humans perform their tasks efficiently, the mechanisms of human perception, decision making, and action, as well as their fundamental limits, must be clearly understood. A major challenge is to obtain quantitative dynamical models suitable for cognitive control design.

The Role for Control

Given the contributions of the neighboring disciplines, what are the envisaged contributions of the controls community? As mentioned above, a fundamental ingredient of a cognitive system is goal-oriented behavior in unstructured environments. This is hardly a novel concept for control—achieving goal-oriented behavior is the basis of almost all control designs! Furthermore, control technology includes efficient and effective methods for addressing issues such as stability, optimality, and robustness. Formulations and solutions for modeling and control for uncertain, stochastic, and hybrid dynamical systems have been developed.

The controls community can contribute greatly to the area of cognitive control by using its strengths in the understanding of dynamical systems, advanced modeling concepts, feedback system analysis methods, and control synthesis tools. The methodical, system-oriented approaches and the mathematical rigor of control methods will be required for deriving provably correct results and for ensuring the safety and performance of engineering products. Even the critical importance of properties such as stability, controllability, and robustness are best appreciated, and the realization of these properties best assured, by experts in control. Without the rigor and analysis that are hallmarks of control science, we cannot expect to develop reliable, high-confidence cognitive control systems for complex applications. These arguments not only justify a role for control in cognitive control; they suggest that the controls community should adopt a leadership role in shaping the research agenda.

> *These arguments suggest that the controls community should take a leadership role in shaping the cognitive control research agenda.*

Challenge Problems for the Field

To provide a better and more specific sense of how a cognitive control system might bring novel capabilities to automation technology, and of the multidisciplinary aspects of such a system, we outline two broad challenge problems below.

Adaptive Management of Cognitive Resources in Real-Time Systems

In today's complex automation systems, human operators play the crucial roles of aggregating and consolidating information, balancing long-term and immediate priorities, and shifting attention dynamically as circumstances dictate. Such capabilities are especially important in large-scale systems, where hundreds, thousands, or more sensors and actuators must be managed. Examples include building automation, manufacturing or process control systems, and traffic management, but an everyday example can help make the point. We are all able to drive a car on a highway while carrying on a conversation with a passenger and listening off and on to the car radio. In the background, we know the route we are taking and effect appropriate actions. However, if another car suddenly cuts in front of us or some other emergency event occurs, we immediately divert our attention to focus on the urgent need of ensuring safety. Our cognitive resources are rescheduled flexibly and at a moment's notice. This flexible, robust behavior is in contrast to the scheduling of tasks in today's computational real-time systems, which is typically static and predefined.

The difference between biological cognition and computer-based attention management becomes more pronounced as the scale of the system under control increases. Learning becomes increasingly important with problem scale. Human operators learn over time what information is important to attend to and what (huge amount of) other information can be safely ignored. The performance improvement, in terms of the ability to monitor and control complex systems, that operators achieve as a result of experience is, in part, a consequence of improved attention management strategies that they have acquired over time.

As these examples illustrate, biological cognition suggests how much better our engineered systems can be in terms of resource management, learning, and adaptation. Questions such as what new control methods are needed, how can generic platforms be developed, how can they then be specialized for critical applications, and how can we have some assurance that flexible, adaptive, learning-endowed cognitive control systems will operate reliably and consistently over extended time periods . . . these remain to be addressed by researchers in controls in collaboration with other disciplines.

Control Response to Rare and Sudden Events

Currently, almost all control systems are designed around structured nominal conditions. At the lowest level, a PID controller will regulate to a setpoint, using an error signal to determine how to move a valve or a motor. Although mathematically much more sophisticated, a multivariable predictive controller is conceptually similar—it processes sensor data with a fixed algorithm (in this case, model-based) and provides an output to a lower-level controller or an actuator. Little else is required for the operation of the control loop under nominal conditions, but what about sudden, and unmodeled, events: sensor or actuator failure, a drastic change in the plant, or a major disturbance?

Automation systems have strategies in place, from redundant devices to fault detection systems to safety shutdown systems, to deal with many such eventualities, but there is a qualitative difference between how expert human operators will respond to an unforeseen event and how today's automation systems respond. Partly as a result of training (often heavily reliant on simulators), pilots and process plant operators can continue the operation of an affected complex system in situations that would be beyond the scope of a fully automated system, based on the best of off-the-shelf technology.

One recourse, of course, is to explicitly model emergency conditions and to "program" appropriate responses to each. To undertake such a project for all conceivable situations would be impossible, but this strategy does not need to be an all-or-nothing one. So questions arise: Can one develop a systematic control design methodology weighing the human resource effort required for the design of fail-safe algorithms with performance when sudden events occur, given likelihoods of events as best they can be estimated? Is there a continuous progression of controls capability with increased human design effort? Can the design effort be automated or adapted online to such sudden events? Can control systems learn online when faced with rare events? Is this knowledge interchangeable through a rare event database with all local control systems feeding knowledge into this database? Therein lies more grist for the cognitive control research mill.

Acknowledgments

This section is based on the discussion on cognitive control at the Berchtesgaden Workshop. Participants in the group included Michael Branicky, Munther Dahleh, Klaus Diepold, Ulrich Konigorski, Henk Nijmeijer, Christoph Stiller, Dawn Tilbury, Georg von Wichert, and Janan Zaytoon. Their constructive comments and contributions to this section are highly appreciated. We also thank Anuradha Annaswamy for comments on an earlier version.

Selected recommendations for research in cognitive control:

- Control strategies for the adaptive management of cognitive resources in real-time systems need to be developed. Cognitive control systems will need to aggregate and consolidate information, balance long-term and immediate priorities, and shift attention dynamically as circumstances dictate.
- Human operators are still the preferred recourse for responding to rare and sudden adverse events. Research is needed to develop automation systems that can exhibit humanlike capabilities in such situations.
- Modeling and estimation take on added dimensions in cognitive control, with representations of self, the environment, objectives, and other elements required. Such representations must often be developed from partial and uncertain information.

References

[1] euCognition, the European Network for the Advancement of Artificial Cognitive Systems [Online], 2010. Available at http://www.eucognition.org/euCognition_2006-2008/definitions.htm.

[2] CoTeSys. Cluster of Excellence: Cognition for Technical Systems [Online], 2010. Available at http://www.cotesys.org.