

Industrial Priorities for Cognitive Robotics

David Vernon[†]
Carnegie Mellon University Africa
Rwanda
Email: vernon@cmu.edu

Markus Vincze
Technische Universität Wien
Austria
Email: vincze@acin.tuwien.ac.at

Abstract—We present the results of a survey of industrial developers to determine what they and their customers require from a cognitive robot. These are cast as a series of eleven functional abilities:

- 1) Safe, reliable, transparent operation.
- 2) High-level instruction and context-aware task execution.
- 3) Knowledge acquisition and generalization.
- 4) Adaptive planning.
- 5) Personalized interaction.
- 6) Self-assessment.
- 7) Learning from demonstration.
- 8) Evaluating the safety of actions.
- 9) Development and self-optimization.
- 10) Knowledge transfer.
- 11) Communicating intentions and collaborative action.

I. INDUSTRIAL REQUIREMENTS

While cognitive robotics is still an evolving discipline and much research remains to be done, we nevertheless need to have a clear idea of what cognitive robots will be able to do if they are to be useful to industrial developers and end users. The RockEU2 project canvassed the views of thirteen developers to find out what they and their customers want. The results of this survey follow, cast as a series of eleven functional abilities.

A. Safe, reliable, transparent operation

Cognitive robots will be able to operate reliably and safely around humans and they will be able to explain the decisions they make, the actions they have taken, and the actions they are about to take. A cognitive robot will help people and prioritize their safety. Only reliable behaviour will build trust. It will explain decisions, i.e. why it acted the way it did. This is essential if the human is to develop a sense of trust in the robot.

A cognitive robot will have limited autonomy to set intermediate goals to when carrying out tasks set by users. However, in all cases it defers to the users preferences, apart from some exceptional circumstances, e.g. people with dementia can interact in unpredictable ways and the robot will be able to recognize these situations and adapt in some appropriate manner.

The freedom to act autonomously will have formal boundaries and the rules of engagement will be set on the basis of

three parameters: safety for people, safety for equipment, and safety of the robot system. The rules may change depending on the environment and a cognitive robot will not exceed the limits of safe operation. The limits may be application specific, e.g., the robot should not deviate further than a given specification/distance/etc. A cognitive robot will use this type of knowledge to act responsibly and will ask for assistance when necessary (e.g. before it encounters difficulties). In particular, in emergency situations, the robot will stop all tasks to follow some emergency procedure. Ideally, if the user is deliberately trying to misuse the robot, e.g. programming it to assist with some unethical task, a cognitive robot will cease operation.

B. High-level instruction and context-aware task execution

Cognitive robots will be given tasks using high-level instructions and they will factor in contextual constraints that are specific to the application scenario when carrying out these tasks, determining for themselves the priority of possible actions in case of competing or conflicting requirements.

Goals and tasks will be expressed using high-level instructions that will exploit the robots contextual knowledge of the task. This will allow the robot to pre-select the information that is important to effectively carry out the task. The goals will reflect the users perspective. This means that all skills which implicitly define the goals are tightly linked to real-world needs and to the solution of specific problems, e.g., “get me a hammer”. The following guidelines will apply.

- Instructions will use natural language and gestures to specify the goals.
- Natural language will be relatively abstract but will be grounded in the codified organisational rules, regulations, and behavioural guidelines that apply to a given application environment. This grounding means that each abstract instruction is heavily loaded with constraints which should make it easier for the robot to understand and perform the task effectively.
- The goals should be specified in a formalised and structured way, where the designer defines them well and can verify them. For example, teach the robot the environment it is working in, follow a described route to reach each of the target locations and reach these positions to carry out the task. These clearly-specified tasks are tightly coupled with risks and costs, e.g. of incorrect execution.

[†]Much of the work described in this paper was conducted while the author was at the University of Skövde, Sweden. This research was funded by the European Commission under grant agreement No: 688441, RockEU2.

- It should be possible for the robot to be given goals in non-specific terms (e.g. assist in alleviating the symptoms of dementia), guidelines on acceptable behaviour (or action policies), and relevant constraints, leaving it to the robot to identify the sub-goals that are needed to achieve these ultimate goals.
- A cognitive robot will learn ways of measuring the success of outcomes for the objectives that have been set, e.g., creating a metric such as the owners satisfaction related not only to the directly specified objective but also the manner in which the job was done). It should be learn from these metrics.

A cognitive robot will consider the contextual constraints that are specific to the application scenario. It will determine the priority of potential actions, e.g., in case of competing or conflicting needs.

For example, the robot might know the procedure to be followed but the locations to be visited or the objects to be manipulated need to be specified (or vice versa). For example, when an automated harvester encounters a bale of straw, it can deal with it as an obstacle or something to be harvested, depending on the current task. For example, the robot might engage in spoken interaction with older adults until the goal is communicated unambiguously, using context to disambiguate the message and allow for the difficulties in dealing with different accents, imprecise speech, and poor articulation.

A cognitive robot will know what is normal, i.e. expected, behaviour (possibly based on documented rules or practices) and it will be able to detect anomalous behaviour and then take appropriate action.

The following guidelines will apply.

- It will be possible to pre-load knowledge about the robots purpose and its operating environment, including any rules or constraints that apply to behaviour in that environment.
- It will be possible to utilize domain-specific skill pools (e.g. from shared databases) so that the robot is pre-configured to accomplish basic tasks without having to resort to learning or development.
- The robot will continually improve its skills (within limits of the goals and safety, see above) and share these with other robots.
- The robot might assist the user by proposing goals from what it understood and the user makes the final selection.

The level of detail in the description required by a cognitive robot will decrease over time as the robot gains experience, in the same way as someone new on the job is given very explicit instructions at first and less explicit instructions later on. One should need to demonstrate only the novel parts of the task, e.g., pouring liquid in a container, but not the entire process.

It will be possible to instruct the robot off-line if there is no access to the physical site; e.g., using a simulation tool, with the robot then being deployed in the real scenario.

C. Knowledge acquisition and generalization

Cognitive robots will continuously acquire new knowledge and generalize that knowledge so that they can undertake new tasks by generating novel action policies based on their history of decisions. This will allow the rigor and level of detail with which a human expresses the task specification to be relaxed on future occasions.

A cognitive robot will build and exploit experience so that its decisions incorporate current and long term data. For example, route planning in a factory, hospital, or hotel should take into account the history of rooms and previous paths taken, or it might take another look to overcome high uncertainty. In general, the robot will overcome uncertainty in a principled manner.

A cognitive robot will generalize knowledge to new task by understanding the context of a novel task and extrapolating from previous experience. For example, a care-giving robot will reuse knowledge of a rehabilitation exercise, customizing it to another person. A welding robot will weld a new instance of a family of parts. In general, a cognitive robot will extract useful meaning from an interaction for a future and more general use, with the same or another user. This may extend to learn cultural preferences and social norms.

For example, in a domestic environment, a cognitive robot will learn how to do simple household tasks, e.g. how to grasp different objects and them bring to a person that wants them. This will be continuously extended, allowing the robot to do more complex things, including cooking.

D. Adaptive planning

Cognitive robots will be able to anticipate events and prepare for them in advance. They will be able to cope with unforeseen situations, recognizing and handling errors, gracefully and effectively. This will also allow them to handle flexible objects or living creatures.

A cognitive robot will be able to recognize that circumstances have changed to avoid situations where progress is impossible. It will also be able to recognize errors and recover. This may include retrying with a slightly different strategy. The learning process will be fast, ideally learning from each error.

A cognitive robot will be able to learn how to handle errors, how to react to situations where, e.g., a human is doing something unexpected or parts are located in an unexpected place.

A cognitive robot will be able to anticipate events and compensate for future conditions. For example, an automated combine harvester will be able to apply a pre-emptive increase of power to compensate for the demands caused when an area of high yield is encountered.

A cognitive robot will be able to learn about the environment it is in and modify the its current information accordingly. That is, it will adapt to changes in the environment, verifying that the environment matches with what is known, or there is a change and updates. This may require an update of the task but only after asking the user.

A cognitive robot will be able to manipulate flexible or live objects, e.g. living creatures such as laboratory mice. To do so means that the robot must be able to construct a model of their behaviour and adapt its actions as required, continually refining the model.

E. Personalized interaction

Cognitive robots will personalize their interactions with humans, adapting their behaviour and interaction policy to the users preferences, needs, and emotional or psychological state. This personalization will include an understanding of the person's preferences for the degree of force used when interacting with the robot. A cognitive robot will be able to adapt its behaviour and interaction policy to accommodate the user's preferences, needs, and emotional state. It will learn the personal preferences of the person with whom it is interacting. For example, an autonomous car will learn the preferred driving style of the owner and adopt that style to engender trust.

A cognitive robot will understand nuances in tone to learn a person's voice, detecting signs of stress so that it can react to it and review what it is doing. In the particular case of interaction with older adults, the robot will be able to understand gestures to help disambiguate words.

A cognitive robot will be able to extrapolate what has been taught to other situations. For example, it might remember that the user has certain preferences (e.g. to be served tea in the morning) and the robot will remember that preference. However, the robot will not allow these learned preferences to over-ride critical actions policies.

In cases where showing the robot what to do involves physical contact between the user and the robot, the robot will be able to learn the dynamics of the user, i.e. his or her personal preferred use of forces when interacting with objects in the environment.

A cognitive robot will be able to the psychological state of a user, e.g. based on the facial expressions, gestures, actions, movements. Based on this, it will be able to determine what they need by cross-referencing that with knowledge of the persons history.

A cognitive robot will be able to make decisions from a large body of observed data, thereby assisting people who typically make decisions based on learned heuristic knowledge but without a quantitative basis for this decision-making. For example, there is a need to provide farmers with a fact-based quantitative decision-making framework. A cognitive robot or machine would observe the physical environment and the farmer and provide a sound bases for making improved decisions.

F. Self-assessment

Cognitive robots will be able to reason about their own capabilities, being able to determine whether they can accomplish a given task. If they detect something is not working, they will be able to ask for help. They will be able to assess the quality of their decisions.

If a cognitive robot is asked to perform a certain task, it will be able to say whether it can do it or not. It will detect when something is not working and will be able to ask for help.

A cognitive robot will assess the quality of its decisions and apply some level of discrimination in the task at hand, e.g. being selective in its choice of fruit to harvest.

G. Learning from demonstration

Cognitive robots will be able to learn new actions from demonstration by humans and they will be able to link this learned knowledge to previously acquired knowledge of related tasks and entities.

Instructions will be communicated by demonstration, through examples, including showing the robot the final results, with the robot being able to merge prior know-how and knowledge with learning by demonstration. Some of this prior knowledge should be extracted from codified organisational rules, regulations, and behavioural guidelines.

The situation is analogous to training an intern or an apprentice: a trainer might ask "Has someone shown you how to do this? No? Okay, Ill show you how to do three, then you do 100 to practice (and to throw away afterwards). If you get stuck on one, call me, and Ill show you how to solve that problem".

A cognitive robot will learn and adapt the parameters to achieve the task. Today in the assembly of components, often robot assembly is not robotized because it requires too much engineering and it is too difficult for robots because it is based on traditional programming, tuning and frequent re-tuning of parameters.

Teaching will exploit natural language, gaze and pointing gestures, and by showing the robot what to do and helping it when necessary.

Actions will be expressed in high-level abstract terms, like a recipe, ideally by talking to it. For example, "go to hall 5 from hall 2 and pick up the hammer" or "open the valve".

When being taught, the robot should be anticipating what you are trying to teach it so that it predicts what you want it to do and then tries to do it effectively.

It will be possible to provide direct support for the robot, switching fluidly between full autonomy, partial autonomy, or manual control.

H. Evaluating the safety of actions

When they learn a new action, cognitive robots will take steps to verify the safety of carrying out this action. If a robot learns new action, it will be difficult to certify the new action. The process of generating a new action will involve interaction with the world and that may already be harmful. So, when learning a new action, there needs to be a step to verify the safety of carrying out this action. For example, showing a new action plus defining safety and success such that the robot can check if it achieved success.

I. Development and self-optimization

Cognitive robots will develop and self-optimize, learning in an open-ended manner from their own actions and those of others (humans or other robots), continually improving their abilities.

A cognitive robot will be able to use what it has learned to determine possible ways to improve its performance, e.g. through internal simulation at times when the robot is not working on a given task. It will also be able to learn from its mistakes, e.g., breaking china but learning from the effect of the action. A cognitive robot will learn to optimize the actions it performs (e.g. doing something faster) within the certified limits of safety and without increasing the risk of failure and associated costs.

J. Knowledge transfer

Cognitive robots will be able to transfer knowledge to other robots, even those having a different physical, kinematic, and dynamic configurations and they will be able to operate seamlessly in an environment that is configured as an internet of things (IoT).

A cognitive robot will be a crucial component of cyber-physical systems where the robot can be used, for example, as a way of collecting data from large experiments.

K. Communicating intentions and collaborative action

Cognitive robots will be able to communicate their intentions to people around them and, vice versa, they will be able to infer the intention of others, i.e. understanding what someone is doing and anticipating what they are about to do. Ultimately, Cognitive robots will be able to collaborate with people on some joint task with a minimal amount of instruction.

The need for people around a cognitive robot to be able to anticipate the robots actions is important because, if cognitive robots are to be deployed successfully, people need to believe the robot is trustworthy. A cognitive robot will be able to interact with people, collaborating with them on some joint task. This implies that the robot has an ability to understand what the person is doing and infer their intentions.

intelligence, autonomous systems, and cybernetics, among others — that can be deployed to satisfy these requirements in practical robots. It remains to complete this exercise.

II. CONCLUSION

Establishing functional requirements is an essential prerequisite to developing useful systems. This is as true of cognitive robotics as it is for any other domain of information and communication technology. However, the effort to give robots a capacity for cognition is made more difficult by the fact that cognitive science, as a discipline in its own right, does not yet have many established normative models that lend themselves to realization in well-engineered systems. The goal of the work described in this short paper is to re-assert the priority of user requirements in the specification of cognitive robot systems. The motivation underpinning this goal is that, having identified these requirements, we can then proceed to determine the scientific and technological tools and techniques — drawn from the disciplines of artificial