RockEU2
Robotics Coordination Action for Europe Two

Grant Agreement Number: 688441

01.02.2016 – 31.01.2018

Instrument: Coordination and Support Action

Industrial Priorities for Cognitive Robotics

Markus Vincze, Technische Universität Wien, Austria

David Vernon, University of Skövde, Sweden

Deliverable D3.1

Lead contractor for this deliverable:	University of Skövde
Due date of deliverable:	February 01, 2017
Actual submission date:	February 20, 2017
Dissemination level:	Public
Revision:	1.0

History of Changes

Name	Status	Version	Date	Summary of actions made

Executive summary

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. ROCKeu2 canvassed the views of thirteen developers[1] to find out what they and their customers want. The results of this survey follow, cast as a series of eleven functional abilities.

1. 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.
2. 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.
3. 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.
4. 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.
5. Personalized interaction. Cognitive robots will personalize their interactions with humans, adapting their behavior and interaction policy to the user's 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.
6. 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.
7. 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.
8. 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.
9. 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.
10. 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).
11. 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.

[1] The developers and the questions put to them are listed in Appendix 1.

Content

1. Introduction

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. ROCKeu2 canvassed the views of thirteen developers[2] to find out what they and their customers want. The results of this survey follow, cast as a series of eleven functional abilities.

2. Industrial priorities for cognitive robotics

1. 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 user's 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.

2. 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 robot's 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 user's 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.

[2] The developers and the questions put to them are listed in Appendix 1.

- 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.
- 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 owner's 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 robot's 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.

3. 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.

4. 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.

5. Personalized interaction.

Cognitive robots will personalize their interactions with humans, adapting their behaviour and interaction policy to the user's 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 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 person's 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.

6. 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.

7. 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 (see No. 2).

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, I'll 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 I'll 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.

8. 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.

9. 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.

10. 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.

11. 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 robot’s 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.

Appendix I

The following is a list of questions to be put to industrial developers and users to determine their what they want in a cognitive robot. The people that where interviewed are also listed.

Questions

1. Can you give any examples of some robot actions you would call cognitive?
2. What would you like a cognitive robot to be able to do (that it can't do today)?
3. Traditionally, robots just followed a script of actions for a given task/goal. In what situations would you want the robot to deviate from the script?
4. How would you like to specify the actions of a cognitive robot, i.e. teach it to do something?
5. How much freedom should a cognitive robot have to select its actions
6. How would you like to specify the goals of a cognitive robot?
7. How much freedom should a cognitive robot have to set its own goals?
8. What should a cognitive robot learn?
9. What is the minimum requirement for a robot to be called cognitive?

Interviewees

Tim Guhl, KUKA Systems GmbH (*) 2/8/2016

Patrick Courtney, Tec-Communication (*) 2/8/2016

Rich Walker, Shadow Robot Company (*) 23/8/2016

Maja Rudinac, Robot Care Systems (*) 30/8/2016

Slawomir Sander, KUKA Systems GmbH () 30/8/2016

David Ball, Bosch () 30/8/2016

Andrew Graham, OC Robotics () 7/9/2016

Mauricio Calva, Chevron (*) 12/9/2016

Amit Kumar Pandey, Softbank Robotics (*) 12/9/2016

Ugo Cupcic, Shodaw Robot (*) 12/9/2016

Daniel Wäppling, ABB (*) 19/9/2016

Ekkehard Zwicker, GE Inspection Robotics (*) 19/9/2016

Thilo Steckel, CLAAS E-Systems KGaA mbH & Co KG (*) 28/9/2016

(*) By Skype

(**) By email

Appendix II

The following are the answers that were received in response to a list of questions that were put to industrial developers and users to determine their what they want in a cognitive robot. The interviewees and the nine interview questions are listed in Appendix I. The responses are organized under nine headings corresponding to the interview questions.

Characteristic cognitive robot actions

1. Autonomously formulating intermediate goals that serve to achieve a goal that has been provided by a user, and doing so in a manner that automatically factors in contextual constraints that are specific to the application scenario.
2. Doing something that was not explicitly pre-programmed, including actions that are generalised versions of pre-programmed actions (even modest versions).
3. Sensor guided actions, industrial use but limited autonomy, e.g., fastening (nut and bolt) applications to cope with uncertainties for adjusting the robot motion.
4. Recognising that something unforeseen has happened or is about to happen and automatically adapting to compensate and recover, if necessary.
5. Learning from its own actions and those of others, either people or other robots.
6. Reasoning about the situation in pursuit of a goal including reasoning over uncertainty and acquired knowledge. For example, a robot might decide to look around the corner to see if it can go there ("cognitive navigation") or it tries to pick up something. In doing this, it reasons about how to achieve the goal and stores the results of this reasoning and generalises from these actions ("cognitive grasping"). In essence, it overcomes uncertainty in a principled manner. For example, a exploring household robot might be searching for a box of cereal box. Having failed to see it in the kitchen, it then looks in other rooms reducing the probability in each room until it goes back to the kitchen and then finds it.
7. Making decisions and saying how good they are. This would raise the user interaction to the level of reasoning about the situation rather than exchanging details about the action or task. For example, in tele-operation the robot would take care of getting the object interaction correct.
8. Applying some level of discrimination in the task at hand, e.g. selectively harvesting fruit.
9. Learning to converse with humans about what it is doing and any problems it is having in carrying out its tasks.
10. Any action that involves the robot selectively identifying the information that matters in performing a task, ignoring all other data. This applies both to actions currently being executed and future anticipated actions. For example, when looking after an elderly person, a care-giving robot should attend only to the most important and relevant issues and it should be able to infer what these are. This selective attention should operate at two levels: (a) ensuring the action is carried out effectively and (b) ensuring that the actions reflect the priorities inherent in the robot's responsibilities.
11. Any action that demonstrates an ability to reuse previous knowledge, adapting it appropriately for the new surroundings or circumstances. For example, a care-giving robot might reuse knowledge of a certain rehabilitation exercise, customizing it to suit the person it is now taking care of.
12. Learning new objects and recognising them.
13. Learning new objects and then reasoning or discovering how to use it. This means learning affordances and object function and their potential use.
14. Learning the use of an object from the human example. Linking knowledge to related objects. Using this knowledge in a problem-solving scenario over multiple steps.
15. Remember to recall facts at the correct time and related to the correct situation. This could be reminders or things to initiate activity in humans.

16. Sensor driven actions in toy robots, robot dog programme using stimuli to be active and get interest, though it ended up throwing itself on the floor because that created greatest response to sensors.
17. Robot driving to a plant and recognising an incidence by gas sensor and tacking a decision on how to react.
18. Optimise the quality of whatever process the robot is performing.
19. Collaborative space: avoid harming a human, anticipate what the human is doing. Robot helps the human worker with assisting in a task, do not need to instruct it in detail, when the human needs help.
20. Any action that exploits semantics and derives from context-driven decisions, e.g. dealing with a bale of straw either as an obstacle to be avoided or an object to be handled.

Desirable abilities for cognitive robots

1. The ability 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.
2. The ability to understand the context of a novel task and extrapolate from existing knowledge to deal with that task; e.g., understand a scene, being able to know what is in the scene, what can the robot do with the objects and operating in the given context.
3. The ability to achieve a given goal even when there are unforeseen changes in the environment; this means the robot needs to be able to recognize that circumstances have changed and adapt accordingly.
4. The ability to understand error and to recover from errors. This may include retrying and optimizing. The learning process has to be fast, ideally from each error.
5. The ability to operate seamlessly in an environment that is configured as an internet of things (IoT), i.e. as an element in a cyber-physical system where the robot can be used, or example, as a way of collecting data from large experiments.
6. The ability to adapt, in the sense of being able to automatically extend a current capability or reuse its knowledge to deal with a new situation. For example, if a robot can weld a car door of a specific type, it would be able to adapt automatically to a new type of door, either by exploiting available CAD data or by visual inspection.
7. The ability to adapt, in the sense of being able to detect subtle changes in the work environments and automatically adjust its operation to cater for these changes.
8. The ability to communicate with people around it so that they can anticipate the robot's actions and intentions. This is essential if the robot can act autonomously because people need to believe the robot is trustworthy.
9. The ability to perceive the world through several modalities, especially vision, so that it can make choices. These choices should be sensible, i.e. within the robot's capabilities.
10. The ability to learn about the environment and 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.
11. The ability to learn from example demonstrations.
12. The ability to deliberately learn and improve around a problem. E.g., recognising a task has been done purely and practice/optimise to improve, for example grasping reliability).
13. The ability to develop, i.e. to learn in an open-ended manner throughout its lifetime.
14. The ability to transfer knowledge from one robot to another robot with a different physical (kinematic and dynamic) configuration.
15. Grasping: being able to grasp, understand how well an object is grasped, weight distribution may change and react to it, how to put it down or insert, how to merge/assembly with other objects.

Learning form grasping and action: re-inforce to do any actions better and better, improve its capabilities on its own.

16. The ability to manipulate flexible objects, e.g. living creatures such as mice, where it is not possible to determine in advance how the object will behave. This means that the robot has to be able to construct a model of this behaviour and adapt its actions as required, continually refining the model.
17. The ability to reason about its own capabilities. The ability to explain what it can and cannot do. If the robot is asked to perform a certain task, it can tell if it can do it or not. Extends to the ability that the robot notices if it is doing something not well and stops to ask for help. E.g., error bars are bigger than they should be, better not do it but ask for assistance. This is a sort of introspection at the robot's interface system.
18. The ability to move things around to make it possible to complete some action, e.g. moving obstructing leaves and small branches out of the way when harvesting fruit or removing weeds.
19. The ability to determine the priority of potential actions, possibly by inferring their consequences, in situations where there are several completing or conflicting needs.
20. The ability to adapt its interaction policy to allow for the individual needs, preferences, or emotional state of users. Start to interact with person before they are too old, start with little, learning from doing the job and from observing the user and its own actions.
21. The ability to engage in spoken interaction with users, going beyond simply responding to commands spoken by a user so that it can engage in natural discourse by understanding the thought pattern of humans. This is important because, for example, elderly people speak to care-giving robots and they expect the robot to understand. The problem is not confined to understanding the semantics of the spoken interaction; there is also a need for more robust speech processing that can deal with different accents, imprecise speech, and poor articulation, possibly caused by some disability.
22. Perform the visit of a flat/premises and guiding a customer, highlight features of objects/flat, learn from the questions of the customer to update the speech.
23. The ability to navigate in complex environments: indoor as well as outdoor, alternative methods for localisation besides GPS, e.g., in tanks of ships or repetitive structure.
24. Understand a scene, being able to know what is in the scene, what can the robot do with the objects and operating in the given context. Compare what is necessary to do to solve a task with what is actually seen and what needs to be done to solve the task given this input.
25. Learning from doing, e.g., breaking china but learning from the effect of the action.
26. Cognitive abilities linked to decision taking: pre-programmed replaced with a sequence of decisions that result in new decisions. Use it to overcome obstacles in perception, navigation, taking the decision to adapt the procedure, a kind of decisional autonomy at procedure level.
27. Some form of implicit programming with high level instructions. The robot has the context to program the task. Similar to a new employee coming to the production line. E.g., robot knows the task is arc welding, context = angle for welding and other process parameters are given, programming task is where to weld.
28. The ability to anticipate up-coming events and compensate to deal with them, e.g. pre-emptively increase power in a combine-harvester to compensate for the demands caused when an area of high yield is encountered.
29. A cognitive robot should be very safe and very reliable. Specifically, it should be able to achieve Performance Level d (the probability of a dangerous failure per hour is 0.00001% to 0.0001%) or achieve Safety Integrity Level 3 (SIL 3), i.e. exhibit Safety Availability of 99.90% - 99.99% and Probability of Failure on Demand Average of 0.001 to 0.0001.

Situations in which a robot should deviate from a scripted program of actions

1. When it detects a possible safety risk, in which case it needs either to adapt – go off-script and take unforeseen actions that still need to be within the safety specifications – or degrade gracefully.
2. When the robot detects there is something wrong with respect to achieving the task and it recognizes that it needs to execute an alternative action to complete the task. For example, when harvesting fruit, the robot detects an obstacle and it adapts by deciding to pick the fruit from a different angle. One possible way to achieve this is by embedding intelligence at the bottom, i.e. in the system components, rather than in the script or the top controlling layer of the system. This implies the need for a strategy to managing errors rather than handling errors on an individual basis.
3. When there is an opportunity to learn; e.g. when there is an opportunity to exploit the detection of some novel or anomalous event and then infer an explanation resulting in a course of action that improves performance in the future.
4. If a robot needs to deviate from a script, it should be enabled to do so but the extent of the deviation should be constrained by the operating conditions of the application domain. These constraints need to be codified and already are in many modern workplaces.
5. The robot needs to be able to deviate from the script whenever the user interacts in an unexpected manner or does not interact at all. While it is possible to define an action protocol for almost all situations, some users, e.g. people with dementia, can interact in unpredictable ways and often they can be fearful, emotional, and even irrational in these circumstances. The robot needs to be able to recognize these situations and adapt in some appropriate manner.
6. The robot needs to know what is normal, i.e. expected, behaviour (perhaps based on documented rules or practices) and it needs to be able to detect anomalous behaviour or incidents and then autonomously formulate a policy to take appropriate action. This might be done by availing of a shared database of scenarios (possibly cloud-based) derived from many examples contributed by other robots
7. In emergency situations: the robot has to stop all tasks to follow emergency guideline. Similarly in any situation where the user/human demands a "stop" or other measure to interrupt the robot's behaviour.
8. Non-critical actions such as delivering meassgaes in different ways to be regarded as less stereotyped or in entertainment scenarios.
9. In a situation where something unforeseen is encountered, e.g., a human or object not expected there: first avoid harming humans, second avoid damaging any equipment, third avoid damage to the robot system.
10. Scripts to robot should be high level, so it should not deviate at that high level, but rather at trajectory implementation level.
11. Depending on the safety aspect specific actions may allow to deviate within certain limits, safety for environment and robot need to be assured.
12. Deliberately trying to misuse the robot, programming it to execute suicide, robot not to do this if it detects it.
13. Permit to adaption is given by the user; or the user defines if adaption of a decision is allowed. Whenever the user allows the robot to deviate within given constraints. The open issue: How to set the constraints? There is a need for visualisation methods to cope with many different situations.
14. Conflict between selling points: always same result = predictable. Moving towards cognitive abilities, this is gone. The cognitive ability has to become an improvement. The challenge is how to validate the “cognitive” automated system. It is fine if the quality becomes better, but needs to be assured. Cycle time may be fixed but depends on customer.

Ways to specify a robot's actions or teach it to do something

1. In same way as you might teach a child to do something, by showing the robot what to do and helping it out when necessary. Showing could be by demonstration as well as guiding the robot. Using high level language: like a recipe for cooking, in the same way for an industrial task, then the high-level logic takes care of carrying out the actions.
2. Instructions should 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 (see previous section).
3. Actions should be expressed in high-level abstract terms, ideally by talking to it. For example, "go to hall 5 from hall 2 and pick up the hammer".
4. Understand nuance in tone to learn a person's voice, does it show signs of stress to react to it and robot reviews what it is doing. In particularly in older adults, understand gestures to aid word findings issues.
5. The level of detail in the description should 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. Only demonstrating the novel parts, e.g., pouring liquid in a container, but not the entire process. Or a kind of autnomous tele-operation: high-level commands like open the valve, but then keeping the human in the loop to organise the main action while the robot does the actual action.
6. It should be possible to assume that the robot is able to exploit previous instructions and can automatically generalise from these.
7. From a drop-down list of capabilities that have been certified. Whatever is selected it must be certified that the robot will not do harm (paper clip thought experiment).
8. In cases where showing the robot what to do involves physical contact between the user and the robot, the robot should 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.
9. The robot should 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 robots needs to remember that preference and update its action policy accordingly to accommodate that preference. However, the robot must not allow these learned preferences to over-ride critical actions policies, e.g. those that are derived from instructions given by a doctor caring for the user.
10. By allowing it to learn from the experience of other robots in similar situations.
11. Natural language and demonstration, teach at symbolic level and robot should ground given the perceived scene, use gaze and pointing to specify different things, more like teaching a child.
12. Feed the robot with the environment that it is expected to operate in, based on that information the robot should go and fulfil the task. User only provides the system with the available information of the environment it will operate in, defines the points to obtain data, or follow a specific line of motions. The robot then adapts to the specific environment or stops to ask the user on how to proceed if deviations too large.
13. Teach the robot off-line (if there is no access to the site), best with a simulation tool to train a system and the system is then deployed in real scenario all by itself. Support to directly the robot and switching fluently between autonomy or partial or manual control. Ideal
14. 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.

Necessary limitations on a robot's freedom to select its actions

1. A cognitive robot should have reasonable but limited freedom, provided its actions are safe and they satisfy the parameters of the application environment. These parameters are often available

in the form of codified rules, regulations, and behavioural guidelines; as already mentioned, these should be exploited by the robot.

2. A cognitive robot should have freedom to select actions or an alternative strategy to a previously prescribed one if it works equally well or better but this freedom is constrained by the application domain and codified acceptable behaviours for that domain. For example, as an extreme example, ordinarily a robot should never drive over a human but a firefighting robot might choose to drive over a dead body when evacuating someone who is alive.
3. Freedom should be limited to executing actions that are within the range of certified and safe actions to achieve a goal. It seems plausible to negotiate between different actions.
4. A cognitive robot should have as much freedom as it needs for to function in a normal environment, adapting to its surroundings to execute actions, but it should not have the freedom to plan long term strategy. While freedom of action is necessary for cognitive development, it is essential that there is an element of human control to ensure that the user remains safe. This is particularly important where the outcome of the selected action is uncertain.
5. When physically interacting with a human, the robot should not select an action that should not be done, i.e., an action that puts human in danger or harming the human.
6. When interacting with humans, robots should be free to select actions but need to take into account privacy, ethical, safety aspects. Make sure the person is not harmed.
7. Rule of engagement as basic rule: e.g., 5 meters limit to deviate from task trajectory, needs handshake to do more, or observe people who are not expected, then stop and ask back to verify validity of task, or if wind speed exceeds a certain limit, etc. If it cannot follow these rules, it needs a safe return, or wait safely, until the user clarified the situation.
8. Multiple robots may self-organise about a common goal, goal is set but individual actions could be shared and vary within the set goal.
9. The robot should have sufficient freedom to adapt its behaviour to comply with ethical guidelines. For example, in farming we interfere in the ecology of the environment – plants, ground water, animals, air – and robots should have some freedom to, for example, protect wild animals when harvesting by adapting their actions.

Ways to specify a robot's goals

1. The goals should be defined by clearly-specified tasks. Tasks are tightly coupled with risks and costs, e.g. of incorrect execution.
2. The goals should be specified in a user's terms, using relatively abstract language that is 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 perform effectively. Besides natural language gestures could be used to specify the goals.
3. The goals should be specified as loosely as possible and in terms that are high-level, e.g. "Get me a hammer". Again, this implies that the robot has to know what rules apply, what actions and behaviours are acceptable based on codified practices.
4. The goals should be specified in a formalised and structured way, where the designer defines them well and can verify them. E.g.,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.
5. The ultimate goals of a cognitive robot must always reflect the robot's purpose, e.g. in caring for the elderly or in supporting rehabilitation, but the robot should be able to adapt these goals as the user's circumstances change or when the robot is introduced into a new environment. As a hypothetical example, a mobile washing machine with end-effectors might adopt its action policy depending on whether it was operating in a house full of teenagers or the home of an elderly couple but it remains a device for washing clothes and nothing else.

6. It should be possible to pre-load knowledge about the robot's purpose and its operating environment, including any rules or constraints that apply to behaviour in that environment. It should 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. This pre-loaded set of skills essentially defines the goals of the robot. However, the robot should continually improve its skills and share these with other robots.
7. Goals should be specified in terms that reflect the user's 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. The level of difficulty of these problems might be quite low at the beginning, and the goal would then be quite limited goals. The level of difficultly can be subsequently increased and the limits relaxed as the related skills are demonstrated to be robust.
8. The robot may assist the user by proposing goals from what it understood and the user makes the final selection. Natural language, gestures, could be assisted by touch screens or similar. modalities.
9. Set task at structure level, tank by tank, equipment needs to be prepared for the inspection work, motion in between tanks.

Necessary limitations on a robot's freedom to set its own goals

1. A robot should never redefine the main high-level goal given by the user. Even if idle, it should ask to do cleaning or whatever.
2. A robot should not have the ability to set its own goals, except where these are intermediate goals that support an overall goal set by a user in situations where the robot is adapting to unforeseen events. Thus, the robot should have freedom to determine how the overall goal is achieved but not what the goal is.
3. This freedom is always subject to constraints that are specific to the application domain. For example, robot birds used to keep pest bird away from certain locations, e.g. airports, harbours, rubbish tips, and farms, should be able to choose how far away to chase the pest birds based on the robot's knowledge of environment and the application requirements.
4. The user should be asked if a deviation is required, e.g., when energy level is low, to interrupt and go recharging.
5. The more a robot can be depended on using this type of knowledge to act responsibly and to have the ability to ask for assistance when necessary (e.g. before it encounters difficulties), the more freedom it can be given.
6. In the short to medium term, there should be no freedom to change the ultimate goal of a cognitive robot and any freedom should only be to exercise sufficient flexibility to achieve the required adaptation to new environments or unexpected events. In other words, any freedom to set goals must not alter the purpose for which the robot was designed. In the long term, 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 subgoals that are needed to achieve these ultimate goals.
7. Freedom should have boundaries: the rule of engagement set on three parameters: safety for people, safety for equipment/plant, and safety of the robot system. The rules may change depending on the environment, e.g. inside no wind, or with/without tether, this would limit autonomy; robot should not get outside the limits of safe operation.
8. Goals towards the operators, users needs, not necessarily time constrained, e.g., once a day walk the dog or vacuum, i.e., time scheduling is free.
9. Interest in achieving full autonomy has declined in the farming sector in recent years. There are many reasons for this, e.g. a decline in the market with a drop in yield and an uncharacteristic simultaneous drop in price, resulting in a need to re-assert priorities and this has resulted in a move away from fully autonomous machines. What is needed are partially autonomous systems

that have a capacity for artificial intelligence, e.g. rule engines for setting parameters and adapting to changing environments, e.g. assessing likelihood of rain and resultant moisture content, balancing this against the extra energy required to dry the grain, and then providing advice on whether or proceed or not to proceed with harvesting.

What a cognitive robot should learn

1. A cognitive robot should learn how to learn how to carry out a task and how to generalise so that the rigour of the task specification can be relaxed on future occasions.
2. To begin with, the robot needs to acquire a model of the world that's simpler that the world actually is and it should proceed to learn by building a better model by adapting it based on experience and third-party knowledge. The situation is analogous to training an intern or an apprentice: a trainer might ask "Has someone shown how to do this? No? Okay, I'll 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 I'll show you how to solve that problem".
3. A cognitive robot should learn personal preferences of the person with whom it is interacting and adapt its behaviour accordingly. For example, an autonomous car need to learn the preferred driving style of the owner and adopt that style to engender trust.
4. It should also learn the situations in which the preferred style should not be adopted, e.g. it running the red light, based on codified general rules and regulation. In other words it should not learn anomalous styles that contravene the norms of acceptable behaviour.
5. A cognitive robot needs to learn the trade-offs that are acceptable in certain circumstances, i.e. what degree of performance is actually required in the current situation.
6. The robot should learn to optimise 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.
7. If it learns new action, it will be difficult to certify the new action. The process of generating a new action will interact with the world and 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. E.g., showing a new action plus defining safety and success such that the robot can check if it achieved success.
8. It learns to do simple household things, grasp, how to grasp different objects, and then bring to person. This could be continuously be extended to start doing more complex things up to cooking, etc.
9. It should transfer learning from other robots,e.g., all robots in a home environment learn all the object present there typically.
10. It should learn the particular needs of individual people, based on a knowledge of human psychology. Often, people do not know what is the best course of action for them to take and the cognitive robot should be able to determine this, in the context of the purpose for which it has been design (i.e. the skill set and knowledge with which it has been provided) and its experience of interacting with that individual. In doing this, it should avail of the huge amount of knowledge available on the internet so that the robot's strategy for solving a problem or carrying out a task can exploit the experience of other agents – human and robot – to identify solution strategies that are even better than humans on their own might be capable of creating.
11. In general, a cognitive robot should learn how to infer the psychological state of a user (e.g. based on the facial expressions, gestures, actions, movements) and from that what they need by cross-referencing that with knowledge of the person's history. However, the strategy in developing cognitive robots should be to build competencies incrementally, only proceeding to a more advanced level once a more straightforward ability has been mastered.
12. 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.
13. Learn additional information about the environment it is in and modify the original information that has been provided. I.e., it adapts to changes in the environment, verifies that the environment

matches with what is known, or there is a change and updates. This may require an update of the task (only with asking the user, see above).

14. Learn to share a task with the human. E.g., executing a procedure for an inspection or other task, such that the human is only carrying out the more cognitive task of following and interpreting measurements and does nt need to execute the motion or mobility or robot task.
15. Robots need to learn the right things. People could teach it bad things. This would need a framework to define what is right.
16. Means of measuring outcomes for the objectives that have been set: context driven response, e.g., clearing the dishes, learn outcome from process to learn what is success, create a metric such as the owner's satisfaction (but also how happy the user is to learn subtleties, e.g., how noisy or fast it made this, not the directly specified objective but implicit part of it or how well it did the job), compare metric to a new situation and learn from this.
17. Feedback mechanisms that help to extrapolate from one task to a different task, one plate to put away to another activity, e.g., helping to arrange pieces on the mantelpiece.
18. It would be a tremendous marketing impact to let the system learn and adapt the parameters to achieve the task. Today in the assembly of components: often robot assembly is not roboticed because it is too much engineering and too difficult for robots. All based on traditional programming, tuning of parameters, requires frequent re-tuning.
19. Classical challenges with pin-picking: doing it fast and with other parts than simple geometries.
20. Semi-unstructured environment, error handling cannot be written out, the robot needs to learn over time how to react to situations, e.g., human is doing something unexpected, or parts are somewhere unexpected.
21. In some sectors, e.g. farming, people typically made decisions based on learned heuristic knowledge but they lack a quantitative basis for this decision-making. There is a need to provide farmers with a fact-based quantitative decision-making framework to complement the farmers valuable heuristic knowledge. Ideally, a cognitive robot or machine would observe the physical environment and the farmer and provide a sound bases for making improved decisions.

Minimum requirements to qualify as a cognitive robot

1. The robot should help you to use it.
2. The robot should learn how to solve new tasks based on previous experience.
3. The robot should have the ability to detect problems and it should be continually looking for alternative ways of doing things to avoid situations where progress is impossible.
4. The robot should be able to surprise its designer (in a good way) by selecting an innovative action or set of actions that avoid potential problems when carrying out a task.
5. Take a decision based on a given situation. Today all is pre-programmed.
6. It should also be able to explain its 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.
7. It should be able to combine perception and action in a non-linear way, that is, a change in perception produces a substantial change in action. It uses not only instantaneous data but using long term data, e.g., route planning takes into account the history of rooms and paths travelled.
8. It should be able to pre-select the information that is most important to the task at hand and then act on it.
9. It should be able to adapt in a basic manner to the environment and the individual user's preferences and needs;
10. It should be able to accumulate new knowledge and reuse previous knowledge.
11. Learn and re-plan, even if it is from pre-set start and goal, using the experience to improve its behaviour. Over all learning it needs to be within the necessary safety limits.

12. Avoiding any harm to humans. Second, avoid any harm to itself or the environment.
13. Do something that has not been explicitly scripted, any kind of improvement algorithm that does not depend on simple if then else rules, some form adaption to the environment and perceptions.
14. The robot can improve over time given the action/goal constraints. The learning part is done towards optimising performance.
15. One element of learning and improvement, transferring experience from one instance to another, more than just sensory feedback on one instance.
16. It should be aware of the context in all its actions and it should have the ability to continually improve its performance, measured in terms of time, quality, and cost for the task at hand.