

HORIZON

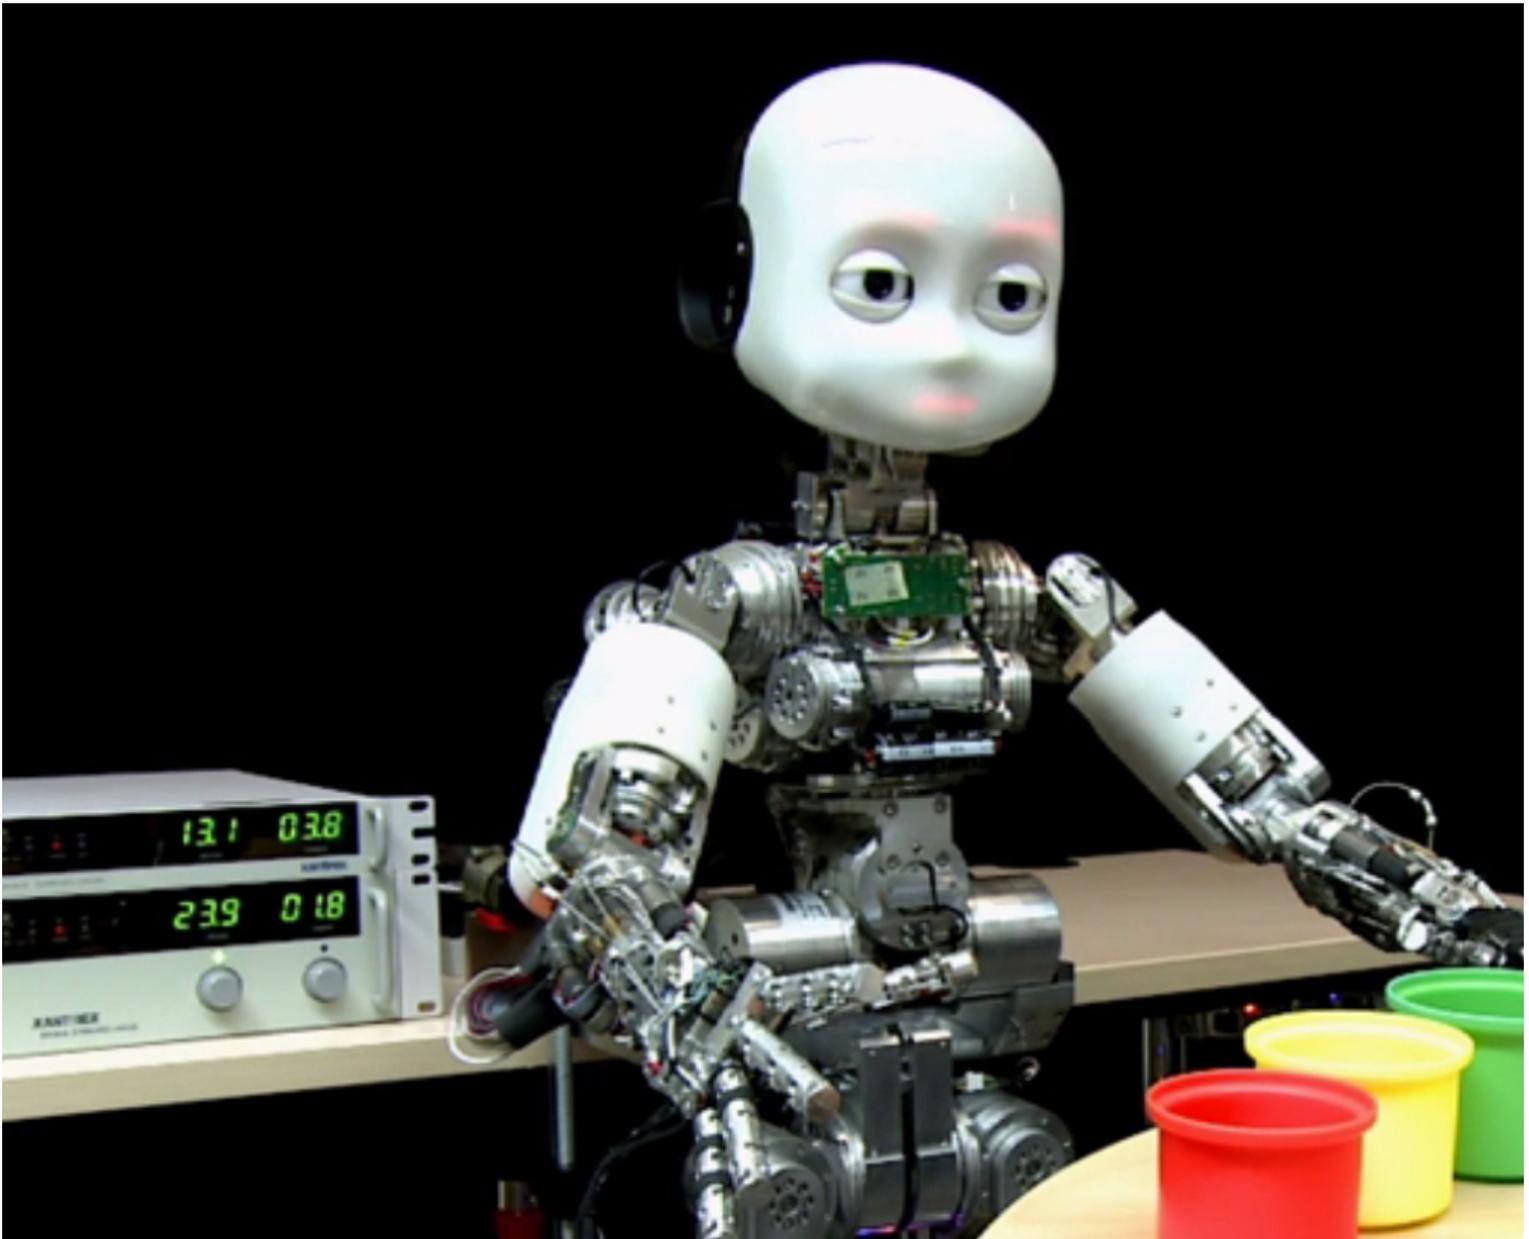
The EU Research & Innovation Magazine

ICT

Teaching robots to talk is child's play

17 August 2015

by Joanna Roberts



Researchers use artificial neural networks to enable the iCub robot to generalise the concept of a cup. Image courtesy ITALK

Repetition, finger counting and a period in kindergarten could be on the curriculum for tomorrow's robots, thanks to European researchers who are teaching robots to speak and count using the same methods they would use to teach children.

Today's talking robots, including systems such as Apple's Siri, have gained their speech capability by essentially swallowing a dictionary: words and grammar rules are uploaded into the system. They deduce what is being asked – and the answer – by some combination of these rules.

See also

- [Digital age 'desperately' needs ethical and legal guidelines](#)
- [We want to end the de-industrialisation of Europe – Prof. Jürgen Rüttgers](#)
- ['Earworm melodies with strange aspects' – what happens when AI makes music](#)
- [Computers learning to read, watch and understand](#)
- [Dreaming robots and creative computation – the future of AI takes shape](#)

The problem with this approach, according to Professor Angelo Cangelosi from the University of Plymouth, UK, is that such systems don't truly understand what you, or they, are saying.

'Siri can give us the impression that it knows what it is talking about, but clearly it has no understanding of the replies that it gives. With robots, we can't really rely on a chatterbox that replies back meaningful sentences without a real understanding.

'If I have a robot in my home and I need the robot to go and pick up something for me or make a cup of coffee for me, I really want the robot to understand what that means.'

Prof. Cangelosi led the EU-funded ITALK project, which instead used an approach known as developmental robotics, where robots begin as a 'black box' and are taught to speak just as you would teach a child.

'Like in a child, you start small, you gradually learn to name individual objects, then you name actions, then you name properties like adjectives and you start to combine two or three words together in a short sentence,' he said.

'We have done experiments in which you can tell a robot "red cup, yellow cup" and you can show them a purple cup and the robot can generalise the concept of cup and use the right adjective.'

This learning ability is made possible by using a type of computer system known as an artificial neural network, which is structured to replicate the human brain. Rather than being pre-programmed with rules, these networks develop and adjust the rules as they go along.

When the robot learns something, the system creates connections, for example between the sound of a word and the image of an object. As it repeats the experience, the connection strengthens.

Researchers have been able to build robots that can verbally describe objects they have never seen before. Video courtesy of ITALK

Teaching robots in this way is very time- and labour-intensive, but Prof. Cangelosi says it may be possible to fast-track the process by replicating another method by which children learn: interaction.

'There is an idea now of a kindergarten for robots. So if you could leave the robot for weeks in a rich environment, the robot can increasingly learn without you having to be a teacher. The robot kindergarten allows the robot to accumulate learning over time.'

Finger counting

Prof. Cangelosi and his team have also been teaching robots to count, and in the process discovered that, like children, the robots learn quicker if they use their fingers.

What's more, if the robots learn numbers by just counting out loud, the numbers ended up stored in their brains in a random order. However, when the robots learned by finger-counting, the numbers were stored in their brains in a sequential list from one onwards. The effect was that finger-counting robots were able to do addition without any extra tuition.



'The robot kindergarten allows the robot to accumulate learning over time.'

Professor Angelo
Cangelosi, University of
Plymouth, UK



'The robot can do additions, because if you ask a robot who can finger count to do two plus two, the robot will do one and two with the fingers,' explained Prof. Cangelosi.

'Then you say "add two" and the robot will simply continue to three, four and it will say "four", because its number representation is constrained by its finger representation. Therefore you get for free the fact that you have four fingers opened.'

Long-term, the idea is to create a commercial robot with a basic level of education which also has the ability to keep learning, meaning that it will be able to pick up new words and adapt to the particular vocabulary and colloquialisms of its owner.

This result is some way off, however, as current models of robot have only succeeded in learning around 20 words.

Dr Katerina Pastra from the Cognitive Systems Research Institute in Greece speculates that it will be around 10 to 15 years before we can hold a natural language conversation with a robot.

She coordinates the EU-funded POETICON++ project, which builds on the work of ITALK and a previous POETICON project. It aims to enable robots to use their language ability not only for communication, but also to generalise their knowledge to new situations.

This means that if a robot has experience of performing one action, such as picking up a cup, it should be able to generalise the concept of 'pick up' and transfer it to other situations, such as picking up a book, when asked to do so.

'Until now, robots have learned how to form a specific action or how to recognise a specific object in their environment,' said Dr Pastra. 'We have not seen any robots with generalisation abilities that can actually build on things they have learned and use this knowledge in new situations in a creative way as humans do.'

'We believe that language is a key parameter for endowing agents with such ability.'

The POETICON++ researchers have been developing two important elements to help a robot to generalise its knowledge. One of these is a kind of dictionary of actions, in which each goal - such as 'to cut' or 'to push' - is linked to a collection of specific actions, which the robot can be taught to perform.

This is part of a computational model of a semantic memory, similar to that in humans, which is designed specifically to store general information that can be applied to other situations.

'Up until now robots only have temporary storage of things they learned,' said Dr Pastra. 'What we developed is a semantic memory, in which robots can have generalised knowledge of object categories, of actions, of the relations between objects and actions, the effect of their actions, their goals and so on. It can be used whenever the robot has to deal with a new situation.'

The other is a reasoner that runs on this semantic memory to help the robot make sense of what it is being asked to do, or what it observes, and deal with uncertainty and unexpected situations in its environment.

By categorising actions as a series of smaller sub-actions, robots can apply their knowledge to novel situations. Video courtesy of POETICON++

Researchers are now reaching the stage where the robot's intellectual capability is more advanced than its physical capability.

'The paradox is that we have already reached a stage where the robot can understand to a significant extent what it is being asked to do and deal with unexpected situations creatively; however the behavioural reaction (motoric implementation) is still largely challenging,' said Dr Pastra.

'So, in a way, the brain skills advanced faster than the bodily skills of the robot.'

More info

[ITALK](#)

[Poeticon++](#)