

Notes

- Intelligence == ability to predict the future
 - Brains evolved to quickly determine best option in vast search space (evolution!!)
 - for robots, researchers implemented an **trial and error algorithm**
- Prediction of events requires:
 - Experience (also requires correlating current actions with historical context)
 - Understanding of "how the world works" (basic laws of physics + nature)
 - Capacity to evaluate one's own actions in the context of those of others
 - Necessary to pick "best" action from a search space
- Planning future actions required self to recreate an abstract model of the world (action space) inside the robot brain to run mental trials and errors to reveal the best action
 - Sounds like reinforcement learning!!!
- Robot actions constrained by physical appearance and restriction
 - Similar to an FSM (finite state machine) -> future subset of all possible activations determined by prior particular activation
 - Actions that are most likely to succeed are 'close' to other actions in behavior space
- Future research potential: "Could these intuitive trial-and-error strategies be used to discover more general problem-solving methodologies for ... uncertain environments?"
 - Yes, but finding optimal way to determine correct action in behavior space will take longer
- True AI will be modeled with some kind of evolutionary algorithm, since good strats are improved on by replication w/ variation and selection

Ideas

Optimize for some modified RL algorithm that can evolve to create split-second decisions
For training, assign fitness values to every action in behavior space

Todo

Understand RL more conceptually

Keep reading about instinctual learning (both in humans and in artificial agents)

ARTIFICIAL INTELLIGENCE

Robots with instincts

An evolutionary algorithm has been developed that allows robots to adapt to unforeseen change. The robots learn behaviours quickly and instinctively by mining the memory of their past achievements. [SEE LETTER P.503](#)

CHRISTOPH ADAMI

Intelligence, by some accounts, is synonymous with the ability to predict the future. Because doing so quickly can often mean the difference between life and death, our brains have evolved to be able to search the vast number of potential futures easily. How is such a feat accomplished? On page 503 of this issue, Cully *et al.*¹ attempt to answer this question by demonstrating that robots can learn to recover quickly and robustly from physical damage — a sudden event that requires them to adopt a new behavioural strategy to continue functioning. The robots (a six-legged mobile robot and a robotic arm; Fig. 1) use a trial-and-error algorithm that lets them tap into the experiences they have accumulated over a simulated lifetime, to quickly find optimal compensating behaviours as if by instinct (see Supplementary Video 1 in ref. 1).

Accurate prediction of events in complex environments requires experience, an understanding of 'how the world works', and the capacity to evaluate one's own actions in the context of those of others. It can be argued that the further out in time an organism or a machine can make accurate predictions of the future, the more intelligent it is. Using this definition, even simple organisms have some intelligence: microbes such as *Escherichia coli*, for example, make predictions about where they must move to find higher concentrations of sugars, and squirrels anticipate the winter by stashing away nuts.

Among animals, humans have arguably the highest level of intelligence, because we can anticipate events hundreds, thousands or even millions of years in the future — albeit largely in domains that do not involve the actions of people, such as planetary orbital dynamics. How can we begin to understand the cognitive underpinnings of a predictive capability that is, to a smaller or larger extent, inherent in all forms of life on Earth? One way, following Richard Feynman's dictum "What I cannot create I do not understand", is to recreate intelligence in a machine or robot.

Attempts to create robot intelligence have come and gone with limited success in the past half-century, and it seems as if the goal of creating a machine with human-like intelligence




Figure 1 | Adaptive machines. Cully *et al.*¹ have designed an algorithm that allows robots to develop strategies for overcoming the effects of damaged limbs. Two robots were used: a, a hexapod (width 50 centimetres); b, a robotic arm (length 62 cm).

remains elusive — even as great strides are being made. Notwithstanding the successes of chess-playing programs, IBM's artificially intelligent computer Watson, and the advent of algorithms for self-driving cars, true robot intelligence still eludes us.

Previous studies in the field of robot cognition^{2,3} have suggested that the ability to plan future actions hinges on the ability to recreate a model of the world inside the robot brain — an abstract version, but one that is accurate enough for mental trials and errors to quickly reveal the best strategy to adopt. But even supposing that these model representations⁴ can be generated, how can the vast 'space' of likely future actions be searched quickly and efficiently?

Cully *et al.* subjected their robots to several different unforeseen changes in the machines' morphology (akin to damage), and then asked them to find movement strategies that would compensate for the injury. Before being injured, the robots used an algorithm to establish a baseline of possible actions, which they used after injury to try out moves that were likely to be successful before deciding on any particular compensatory behaviour. Even though the range of possible behaviours (the behaviour space) for a robot might theoretically be infinite, this baseline can be established because, in reality, a robot's actions are constrained by its morphology.

A hexapod robot such as that studied by the authors is controlled by 36 parameters, but most of the strategies (sequences of motor

activations in a 36-dimensional space) make no sense. Within the robot's 'embodiment'⁵ — the way in which the robot's body is realized — only a small subset of activations can follow any particular prior activation. In other words, the robot's embodiment dramatically reduces the number of potential strategies, so that sensible actions occupy a severely reduced behaviour space (think of a line instead of a sphere). This reduced space is actually searchable in real time.

The authors created the set of all possible behaviours by having each robot perform many thousands of motions (sequences of motor activations) and recording the 'fitness values' of each sequence. The fitness could be as simple as the distance travelled by the robot. Collating this database is time-intensive, but it is analogous to what happens in the natural world, in which living organisms have a lifetime to acquire such information. The robots synthesized new behaviours from this data set using a set of special-purpose machine-learning algorithms that assume that — even in changed circumstances — the actions that are most likely to succeed are 'close' to other such actions in a suitably defined behaviour space.

Although these machine-learning algorithms are unlikely to be similar to those used, for example, by mammalian cognitive systems, they share a common premise: that a behaviour space that is dramatically reduced through embodiment, and that is learned from experience, can be searched quickly through

ANTOINETTE CULLY/UPMC

cannot create

search strategy

trial and error. If we return to the analogy of a one-dimensional line as opposed to a sphere of possible strategies, only two directions have to be attempted for the line before the preferred direction is clear, whereas in a sphere six directions must be sampled. Given that the robot's behaviour space is 36-dimensional, it is clear that the 'flattening' of the space of options can have dramatic effects.

Could these intuitive trial-and-error strategies be used to discover more-general problem-solving methodologies, of the kind that require planning in uncertain environments? It is difficult to imagine that the method could easily be scaled up to such a level; this particular algorithm was hand-designed by the authors, whereas the 'algorithm' our brains use is the result of millions

of years of Darwinian tinkering and pruning.

Given the failure of past efforts to design robots that display the quick, intuitive and situation-appropriate behaviour of even the smallest rodents, perhaps it is time to give up on the idea that we can design brains, and instead place our hopes in the power of adaptive and evolutionary algorithms. Indeed, the core algorithm that generates the map of possible high-performance behaviours in Cully and colleagues' study is inherently evolutionary, because good strategies are improved on by replication with variation, and selection.

We may never understand our brains in terms of information-processing concepts, but we do understand how to harness the power of evolution. We should therefore let evolution

create for us what we do not understand, one more time. ■

Christoph Adami is in the Department of Microbiology and Molecular Genetics, and the Department of Physics and Astronomy, Michigan State University, Michigan 48824, USA.

e-mail: adami@msu.edu

1. Cully, A., Clune, J., Tarapore, D. & Mouret, J.-B. *Nature* **521**, 503–507 (2015).
2. Bongard, J., Zykov, V. & Lipson, H. *Science* **314**, 1118–1121 (2006).
3. Adami, C. *Science* **314**, 1093–1094 (2006).
4. Marsteller, L., Hintze, A. & Adami, C. *Neur. Comput.* **25**, 2079–2107 (2013).
5. Pfeifer, R. & Bongard, J. *How the Body Shapes the Way We Think: A New View of Intelligence* (MIT Press, 2007).

need to simulate in areas where error cannot be afforded

CELL BIOLOGY

Polarized transport in the Golgi apparatus

Proteins can be transported in either direction across a cellular organelle called the Golgi apparatus. It emerges that CDC42, a molecule that confers cell polarity, acts to control the directionality of transport in the Golgi. [SEE LETTER P.529](#)

AKIHIKO NAKANO

Polarity is a part of life on all scales, from head-to-tail polarity in whole organisms to polar orientation in cells. Even intracellular organelles, such as the Golgi apparatus, can be polarized. Proteins synthesized by another organelle, the endoplasmic reticulum, mainly enter the Golgi on one side — dubbed the *cis* side — and, after appropriate modifications, exit from the opposite, *trans*, side. On page 529 of this issue, Park *et al.*¹ identify CDC42 as a protein that regulates the directionality of polarized protein transport within the Golgi apparatus.

The Golgi apparatus, which prepares proteins for shuttling to their final destination in the cell, is a mysterious organelle. More than 100 years have passed since its discovery, but there are still many controversies regarding the mechanism that underpins its function. One debate famous among cell biologists concerns how cargoes such as secretory proteins are transported across the polarized structure of the Golgi from *cis* to *trans*² — the anterograde direction. Various models have been put forward to explain this directionality, but several lines of evidence^{3,4} support 'cisternal maturation'. In this model, flattened, membrane-bound sacs called cisternae, which make up the Golgi, form on the *cis* side of the organelle and move towards the *trans* side as they mature,

carrying proteins produced by the endoplasmic reticulum with them as they go.

However, this mechanism alone cannot

explain complicated protein-sorting events in the Golgi; these mediate transport not only from *cis* to *trans*, but also in the opposite direction. Such retrograde transport involves vesicles coated in coat protein complex I (COPI)^{3,4} that shuttle between cisternae. But it is unclear whether these vesicles are involved in anterograde cargo transport. Tubular structures that connect cisternae have also been posited to be involved in intra-Golgi transport^{3,4} (Fig. 1).

Previous work⁵ by the group that undertook the current study revealed that the COPI coat regulates not only vesicle formation but also the formation of tubules in the Golgi — although how it does this is not known. Park *et al.* demonstrated that COPI binds to proteins being transported in both the




Figure 1 | Bidirectional transport. The Golgi apparatus, which contains a stack of membrane-bound sacs called cisternae, prepares proteins for shuttling around the cell. Proteins mainly move through the Golgi in the anterograde direction, from the *cis* to the *trans* side, travelling from their site of synthesis in the endoplasmic reticulum towards the cell membrane. However, proteins can also move in the opposite, retrograde, direction. Proteins can be transported across the Golgi in three ways: through 'cisternal maturation', in which protein-containing cisternae move from *cis* to *trans* as they mature; in vesicles coated in coat protein complex I (COPI), which mediate retrograde transport; and in tubules, in which COPI binds to proteins to promote transport in either direction. Park *et al.*¹ report that the CDC42 protein competes with COPI for binding of retrograde, but not anterograde, cargoes in tubules, and so enhances anterograde and inhibits retrograde transport.