Tobin and collaborators' results represent the first direct observational evidence that small-scale disk fragmentation can lead to the formation of binary and multiple star systems. The authors observe two of the expected outcomes of gravitational instability of a protostellar disk: namely, a hierarchical configuration of protostars and a spiral-shaped disk.

The strength of this study is the attention to detail in Tobin and colleagues' analysis. For instance, the authors show that the disk of the protostar system is gravitationally unstable by calculating the 'Toomre Q parameter' (ref. 8). Specifically, they determine the size at which the disk would fragment because of a gravitational instability, and conclude that it would be unstable at radii of between 150 and 320 AU. This range agrees with the observed separation between the two protostars at the centre of the system and the third member in the spiral arm. Moreover, the authors find that this fragmentation probably occurred recently (within the past few thousand years), which is consistent with the young age of the protostar system.

It is important to emphasize that both of the fragmentation mechanisms are plausible and expected, rather than being mutually exclusive. Determining the frequency at which each mechanism occurs will require follow-up studies. Fragmenting disks like the one observed by Tobin and colleagues are probably not rare — rather, they are waiting to be studied in more detail using the powerful (sub-)millimetre-wavelength telescopes that are now available. ■

Adele Plunkett *is at the European Southern Observatory, Vitacura, Casilla 19001, Santiago, Chile.*
e-mail: aplunket@eso.org

1. Raghavan, D. *et al. Astrophys. J. Suppl.* **190,** 1–42 (2010).
2. Tobin, J. J. *et al. Nature* **538,** 483–486 (2016).
3. Evans, N. J. II *et al. Astrophys. J. Suppl.* **181,** 321–350 (2009).
4. Mathieu, R. D. *Annu. Rev. Astron. Astrophys.* **32,** 465–530 (1994).
5. Duchêne, G. & Kraus, A. *Annu. Rev. Astron. Astrophys.* **51,** 269–310 (2013).
6. Reipurth, B. *et al.* in *Protostars and Planets* Vol. VI (eds Beuther, H. *et al.*) 267–290 (Univ. Arizona Press, 2014).
7. Pineda, J. E. *Nature* **518,** 213–215 (2015).
8. Toomre, A. *Astrophys. J.* **139,** 1217–1238 (1964).

ARTIFICIAL INTELLIGENCE

Deep neural reasoning

The human brain can solve highly abstract reasoning problems using a neural network that is entirely physical. The underlying mechanisms are only partially understood, but an artificial network provides valuable insight. SEE ARTICLE P.471

HERBERT JAEGER

A classic example of logical reasoning is the syllogism, "All men are mortal. Socrates is a man. Therefore, Socrates is mortal." According to both ancient and modern views[1], reasoning amounts to a rule-based mental manipulation of symbols — in this example, the words 'All', 'men', and so on. But human brains are made of neurons that operate by exchanging jittery electrical pulses, rather than word-like symbols. This difference encapsulates a notorious scientific and philosophical enigma, sometimes referred to as the neural–symbolic integration problem[2], which remains unsolved. On page 471, Graves *et al.*[3] use the machine-learning methods of 'deep learning' to impart some crucial symbolic-reasoning mechanisms to an artificial neural system. Their system can solve complex tasks by learning symbolic-reasoning rules from examples, an achievement that has potential implications for the neural–symbolic integration problem.

A key requirement for reasoning is a working memory. In digital computers, this role is served by the random-access memory (RAM). When a computer reasons — when it executes a program — information is bundled together in working memory in ever-changing combinations. Comparing human reasoning to the running of computer programs is not a far-fetched metaphor. In fact, a venerable historical alley leads from Aristotle's definition of syllogisms to the modern model of a programmable computer (the Turing machine). Alan Turing himself used 'mind' language in his groundbreaking work[4]: "The behaviour of the computer at any moment is determined by the symbols which he is observing and his 'state of mind' at that moment."

The authors' neural system cannot and need not be programmed — instead, it is trained.

Although there are clear parallels between human reasoning and the running of computer programs, we lack an understanding of how either of them could be implemented in biological or artificial neural networks. Graves and colleagues take a substantial step forward in this quest by presenting a neuro-computational system that shows striking similarities to a digital computer.

The authors' system consists of several modules, all of which are entirely non-symbolic and operate by exchanging streams of purely analog activation patterns — just like those recorded from biological brains. There are two main modules: a 'memory' comprised of a large grid of memory cells, each of which can have a particular numerical value that is akin to a voltage; and a 'controller', which is an artificial neural network. The controller can access selected locations on the memory grid, read what it finds there, combine that with input data and write numerical values back to selected memory locations. The two modules interact in many respects like the RAM and central processing unit of a digital computer.

Graves and colleagues demonstrate the capabilities of their system by putting it through several tasks that require rational reasoning, such as planning a multi-stage journey using public transport. Such tasks are fairly easy to solve using the symbolic computer programs of artificial intelligence, but have so far been rather out of reach of artificial neural networks.

A digital computer solves a given task by executing a program that has been written for that purpose. By contrast, the authors' neural system cannot and need not be programmed — instead, it is trained. During training, the system is presented with a large number of solved examples of the task at hand. With each new presentation, the system slightly adapts its internal neural wiring so that its response moves gradually closer to the given task's solution.

The analog, smoothly adaptable nature of the authors' neural system is the key to its ability to be trained. Mathematically speaking, the system is a 'differentiable function', which has led to the authors calling it a differentiable neural computer (DNC). A digital computer is not differentiable and could not be trained in any similar fashion.

A DNC is a mathematical object that boasts tens of thousands of adjustable parameters. Training such a monster raises a plethora of mathematical, numerical and run-time issues. Only in the past few years has machine-learning research overcome these obstacles, through a compendium of techniques that have become branded as deep learning[5]. The authors' training of a DNC is a splendid demonstration of the power of deep learning.

Graves *et al.* steer clear of grand claims about their work's implications for the neural–symbolic integration problem and, with due caution, suggest possible mappings of DNC structures to those of biological brains. This is wise, because the debates fought out in this arena are fierce and without winners. Instead, the authors establish an undeniable technical anchor point that will help to ground the debates — they have shown that certain non-trivial, central aspects of symbolic reasoning can be learnt by artificial neural systems.

With regard to practical exploits, deep-learning methods have so far excelled in tasks that require limited or no working memory, such as image recognition[6] and sentence-wise language translation[7]. Whether or not DNCs will bring about practical advances in big-data technologies remains to be seen. The authors' demonstrations are not particularly complex as demands on rational reasoning go, and could be solved by the algorithms of symbolic artificial intelligence of the 1970s. However, those programs were handcrafted by humans and do not learn from examples.

For the time being, the DNC by itself cannot compete with state-of-the-art methods in digital computing when it comes to logical data mining[8]. But a flexible, extensible DNC-style working memory might allow deep learning to expand into big-data applications that have a rational reasoning component, such as generating video commentaries or semantic text analysis. A precursor to the DNC, the neural Turing machine[9], certainly sent thrills through the deep-learning community. ■

Herbert Jaeger *is at Jacobs University Bremen, 28759 Bremen, Germany.*
e-mail: h.jaeger@jacobs-university.de

1. Newell, A. *Cogn. Sci.* **4,** 135–183 (1980).
2. Hammer, B. & Hitzler, P. (eds) *Perspectives of Neural–Symbolic Integration* http://doi.org/fsrb8m (Springer, 2007).
3. Graves, A. *et al. Nature* **538,** 471–476 (2016).
4. Turing, A. M. *Proc. Lond. Math. Soc.* Ser. 2, **42,** 230–265 (1937).
5. LeCun, Y., Bengio, Y. & Hinton, G. *Nature* **521,** 436–444 (2015).
6. Szegedy, C. *et al. Proc. IEEE Conf. Computer Vision Pattern Recognition* http://dx.doi.org/10.1109/CVPR.2015.7298594 (2015).
7. Bahdanau, D., Cho, K. & Bengio, Y. *Int. Conf. Learning Representations* Preprint at http://arxiv.org/abs/1409.0473 (2014).
8. De Raedt, L. & Kimmig, *A. Machine Learn.* **100,** 5–47 (2015).
9. Graves, A., Wayne, G. & Danihelka, I. Preprint at http://arxiv.org/abs/1410.5401 (2014).

This article was published online on 12 October 2016.

PALAEONTOLOGY

Ancient avian aria from Antarctica

A discovery of the sound-producing vocal organ known as the syrinx in a bird fossil from the end of the 'age of dinosaurs' highlights the anatomical basis for myriad aspects of avian social and behavioural evolution. **See Letter p.502**

PATRICK M. O'CONNOR

The origin of modern birds and their tremendous diversification continue to inspire many studies aimed at elucidating the functional, ecological and evolutionary context of the group's success. The past three decades of palaeontological research have deepened our understanding of birds as close relatives of the predatory theropod dinosaurs, with a wealth of anatomical evidence supporting this idea. Remarkable non-avian dinosaur fossils exist with several hallmark 'avian' features, including a diversity of feather types[1] and a complex lung air-sac system[2], both of which have a broad distribution among non-avian theropod dinosaurs such as *Velociraptor*. A study on page 502 by Clarke *et al.*[3] reports the identification of an avian-specific feature — the sound-producing syrinx — in a Late Cretaceous (69 million to 66 million years ago) fossil of a bird recovered from the Antarctic Peninsula.

Modern birds are a tremendously diverse group, numbering about 10,000 living species[4]. Thanks to an impressive range of anatomical and behavioural specializations, birds have spread to almost every environment on the planet. Although considerable ink has been spilled regarding the roles of avian feathers, wings, feeding systems and body-size differences as drivers of this diversification, the fossil record has produced only limited information regarding the evolution of avian social and behavioural systems. The evolution of the syrinx, that one avian-specific feature responsible for the vast array of sounds produced by birds, is sorely understudied.

The work of Clarke and colleagues represents one of only a handful of examples from the entire avian fossil record in which the mineralized cartilage rings of the syrinx, which support the sound-producing soft-tissue membranes, have been identified. So far, the sound-producing system for non-avian dinosaurs has been thought to be the muscular tube in the throat known as the larynx. Such inferences rest on generalized comparisons with other (non-avian) ancestral groups known as tetrapods, and there is a relative dearth of studies that even consider the origin of the syrinx.

Clarke *et al.* assigned the fossil to the bird