

Research



**Cite this article:** Winfield AFT, Blackmore S. 2021 Experiments in artificial culture: from noisy imitation to storytelling robots. *Phil. Trans. R. Soc. B* **377**: 20200323. <https://doi.org/10.1098/rstb.2020.0323>

Received: 14 June 2021  
Accepted: 21 September 2021

One contribution of 17 to a discussion meeting issue ‘The emergence of collective knowledge and cumulative culture in animals, humans and machines’.

**Subject Areas:**  
behaviour, cognition

**Keywords:**  
artificial culture, cultural evolution, memetics, social learning, embodiment, collective robotics

**Author for correspondence:**  
Alan F. T. Winfield  
e-mail: alan.winfield@brl.ac.uk

# Experiments in artificial culture: from noisy imitation to storytelling robots

Alan F. T. Winfield<sup>1</sup> and Susan Blackmore<sup>2</sup>

<sup>1</sup>Bristol Robotics Laboratory, University of the West of England, Bristol BS16 1QY, UK  
<sup>2</sup>Department of Psychology, University of Plymouth, Plymouth PL4 8AA, UK

AFTW, 0000-0002-1476-3127; SB, 0000-0003-0845-2223

This paper presents a series of experiments in collective social robotics, spanning more than 10 years, with the long-term aim of building embodied models of (aspects of) cultural evolution. Initial experiments demonstrated the emergence of behavioural traditions in a group of social robots programmed to imitate each other’s behaviours (we call these Copybots). These experiments show that the noisy (i.e. less than perfect fidelity) imitation that comes for free with real physical robots gives rise naturally to variation in social learning. More recent experimental work extends the robots’ cognitive capabilities with simulation-based internal models, equipping them with a simple artificial theory of mind. With this extended capability we explore, in our current work, social learning not via imitation but robot–robot storytelling, in an effort to model this very human mode of cultural transmission. In this paper, we give an account of the methods and inspiration for these experiments, the experiments and their results, and an outline of possible directions for this programme of research. It is our hope that this paper stimulates not only discussion but suggestions for hypotheses to test with the Storybots.

This article is part of a discussion meeting issue ‘The emergence of collective knowledge and cumulative culture in animals, humans and machines’.

## 1. Introduction

In this paper, we describe two sets of experiments with small groups of real robots, conducted over the course of more than 10 years, in the Bristol Robotics Lab. The long-term aim of these ongoing experiments is to explore aspects of the question ‘how do we have culture?’ in a new way, by modelling the low-level processes and mechanisms of cultural evolution with robots. In this paper we adopt Mesoudi’s definition of culture: ‘information that is acquired from other individuals via social transmission mechanisms such as imitation, teaching or language’ [1]. We outline two sets of experiments—the first already completed and the second in preparation—with a focus on two of these transmission mechanisms: imitation and language.

The first set of experiments we describe were directly inspired by the thought experiment in [2, p. 106], which imagines a group of robots capable of imitating each other. Referred to as Copybots, their ability to imitate actions with variation makes them very simple meme machines. Another source of inspiration was Gabriel Tarde who proposed ‘a remarkable sociological research project’ [3] when he wrote

If we wish to make sociology a truly experimental science and stamp it with the seal of exactness, we must, I believe ... write out with the greatest care and in the greatest possible detail the succession of minute transformations in the political or industrial world, or some other sphere of life, ... in (our) native town or village, beginning in (our) own immediate surroundings (quoted in [3, p. 511]).

Barry & Thrift [3] write that ‘Tarde’s project was not, as far as we are aware, ever carried out’. When we started to work on what became known as the

Artificial Culture project, we realized that we could set up a free-running group of robots (an artificial society) and literally observe, record and analyse every minute detail of the robots' interactions with each other.

A second and more recent set of experiments extends our robots' cognitive capabilities with simulation-based internal models. A simulation-based internal model (literally a robot with a simulation of itself, inside itself), allows a robot to be able to ask itself 'what if' questions. This capability has been described as a functional imagination [4], as it enables a robot to 'imagine' the consequences of its actions (and—in our implementation—the reaction of others to those actions). Our experimental implementation of a simulation-based internal model, which we refer to as a consequence engine (CE), has proven to be remarkably powerful. Our experiments with the CE were inspired by both the simulation theory of cognition [5,6] and Dennett's 'Tower of Generate-and-Test' [7]. Both simulation and the loop of generate-and-test are present in the architecture of the CE.

In our current work, also using the CE, we aim to explore social learning not via imitation but robot–robot storytelling in an effort to model this very human mode of cultural transmission. Although we are not aiming at evolving language, we have nevertheless been influenced by both the seminal Talking Heads experiments of Steels [8] and work to evolve mechanisms of communication in a swarm of robots [9]. Instead, and in addition to the CE, our 'Storybots' are being equipped with the means to communicate via speech, and what Penn *et al.* [10] call the 'spectacular scaffolding provided by language'.

Our method for both sets of experiments is to build a working model or, as we prefer to describe it, an *embodied simulation* consisting of a group of autonomous robots, in which the robots are programmed with simple behaviours and interact with each other in an artificial arena. The arena is equipped with a system that allows each robot's movements to be tracked and recorded, alongside a time-stamped record of each robot's internal decisions sent to the logging system via a local WiFi network. In this way, we are able to capture Tarde's 'minute transformations' for analysis.

Physical embodiment is important to us for several reasons: first, because experiments with real robots are noisy and unpredictable. Even though our robots are seemingly identical, small differences between the motors and sensors mean that each robot will move and sense in slightly different ways. These unintended heterogeneities serve to model differences between conspecific animals and humans. And the noise will prove to be of critical importance. Unlike computer simulations, the noise, stochasticity and physics come for free, just as they do for animals and us. Second, and perhaps most importantly, robots—like animals—have physical bodies that constrain how they behave and 'think' [11]. Robots also see each other only from their own first-person perspectives. Yes, our robots have distinctly non-human minds [10]—albeit of a kind so simple that the term 'mind' is hardly appropriate—but, we contend, they have enough in common with animals and humans to allow us to plausibly model interesting aspects of social learning and behavioural evolution. The work of this paper fits, we believe, within the microevolution strand of the science of cultural evolution [12], and although simulation models of cultural evolution are not new [13], we believe that our approach using robots in an embodied individual-based simulation is novel. Physical

embodiment and noise together provide the 'natural phenomena' [14] that can be exploited in cumulative cultural evolution.

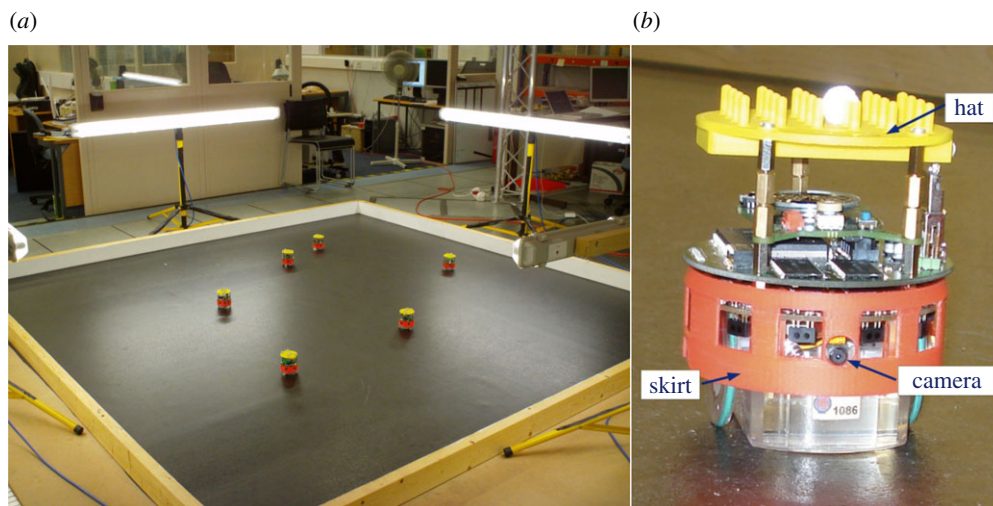
This paper proceeds as follows. In §2, we outline the Copybots and the key findings from the first set of experiments. Then, in §3 we describe the consequence engine—the key innovation of our second generation of experimental work. We illustrate the CE and the kind of emergent behaviour that is typical of real-robot embodied simulations with the pedestrian experiment, before then introducing the Storybots. We conclude the paper in two parts: in §4a we outline Dennett's tower of generate-and-test before then showing how it provides a unifying framework in which we can classify all of the robots of this paper. Finally in §4b, we discuss experimental possibilities for the Storybots hoping this will stimulate research questions in cultural microevolution that might feasibly be explored.

## 2. Copybots

In a series of experiments, we implemented social learning in a group of robots [15]. Simple wheeled robots (figure 1) were programmed to learn socially, from each other, by imitation. These miniature robots—called e-pucks—are extremely simple compared with animals; they have just two wheels and so cannot interact with objects in their environment except by colliding with them. Their sensorium is equally limited—they 'see' only with a single 640 × 480 resolution camera, and something approaching a sense of touch is provided by eight short-range infrared proximity sensors mounted around their body radius [16]. From a behavioural perspective the robots are a blank canvas. They have no built-in or innate behaviours, all must be programmed [17].

In these experiments, imitation was strictly embodied. Robots have no access to each other's internal states, instead robots observed each other using their onboard sensors and, on the basis of only visual sense data from a robot's own camera and perspective, the learner robot inferred another robot's pattern of movements. (This contrasts with the social learning in Bredeche & Fontbonne [18], in which robots learn by sharing internal parameters when they encounter each other.) We 'seed' each Copybot with initial behaviours, which are self-contained movement sequences (or 'dances') which we refer to as memes (Dawkins [19] defines a meme as 'that which is imitated'). We then free run the Copybots with each robot alternating between enacting memes and watching (and learning) those memes.

Not surprisingly embodied robot–robot imitation is imperfect. A combination of factors including the e-puck robots' low-resolution camera, variations in ambient lighting, heterogeneities among the robots, multiple robots sometimes appearing within a learner robot's field of view, and of course having to infer another robot's movements by tracking the relative size and position of that robot in the learner's field of view, lead to imitation errors. Furthermore, some memes are easier to learn by imitation than others (think of how much easier it is to learn the steps of a slow waltz than the tango by watching your dance teacher). The fidelity of embodied imitation for robots, just as for animals, is a complex function of four factors: (i) the behaviours being imitated, (ii) the robots' sensorium and morphology, (iii) environmental noise and (iv) the inferential learning algorithm (for a brief outline of how the algorithm works see appendix A).



**Figure 1.** (a) Artificial culture laboratory showing six robots in the arena. (b) An e-puck robot, fitted with a red skirt which makes it easier for robots to see each other, and a yellow hat which provides a matrix of pins for the reflective spheres that allow the tracking system to identify and follow each robot.

But rather than being a problem, noisy imitation was our aim. We are interested in the dynamics of social learning, and in particular the way that memes evolve as they propagate across the collective, by social learning. Noisy social learning means that behaviours are subject to variation as they are copied from one robot to another. Multiple cycles of imitation (robot B socially learns behaviour  $m$  from A, then robot C learns the same behaviour  $m'$  ( $m$  mutated), from robot B, and so on), gives rise to behavioural heredity. And if robots are able to select which learned behaviours to enact we have the three Darwinian operators for evolution, except that this is behavioural, or memetic, evolution.

Our experiments demonstrate that embodied behavioural evolution does indeed take place. If selection is random, that is robots select which behaviour to enact from those already learned—with equal probability—then we see several interesting findings. First, if by chance one or more high fidelity copies follow a poor fidelity imitation, the large variation in the initial noisy learning can lead to a new behavioural species, as shown here in figure 2, thus demonstrating that noisy social learning can play a role in the emergence of new—and potentially useful—behaviours in behavioural (i.e. cultural) evolution [20]. Second, we observe that behaviours adapt to be easier to learn, i.e. better ‘fitted’ to the sensorium and morphology of the robots [21], a result which appears to mirror the findings of Kirby *et al.* [22], that artificially evolved languages evolve to be easier to learn.

A third finding from this series of experiments is perhaps the most unexpected. When we ran the same embodied behavioural evolution with three memory sizes—no memory, limited memory and unlimited memory—the limited memory case led to the most ‘stable’ population of behaviours across the robot collective, i.e. a smaller number of larger clusters of related memes; in other words, a small number of relatively persistent behavioural types. In figure 3, we see one cluster of 12 closely related memes. This result suggests the intriguing conclusion that forgetting may be a significant collective trait in behavioural evolution [21], and might also be related to what is referred to as conformist social learning, in which learners are more prone to act as others do [23].

A related series of experiments combine social and individual learning. We extended reinforcement learning with


imitation, so that robots could observe and learn, by imitation, from more ‘experienced’ individual learners (for an outline description see appendix B). Reinforcement learning is a well-known approach to machine learning based on trial-and-error interactions between an agent and its environment [24]. As above, the imitation is strictly embodied, and an imitating robot has no access to the internal state of an observed robot. In a series of experiments, we saw that robots with imitation-enhanced reinforcement learning learned faster than those with reinforcement learning alone. Not a surprising result; social learning is very much faster than individual learning, and robots, just like animals, can benefit from learning socially from more experienced others [25]. Such social learning in animals, for instance when juveniles typically prefer to copy older individuals who are more experienced, is one of the social learning strategies reviewed in Kendal *et al.* [26]. However, we were surprised to observe that errors in the imitation phase sometimes led to robots learning even faster. It appears that imitation errors that arise while copying another robot can lead to faster learning.

This work has, perhaps for the first time, studied embodied social learning, by imitation, in real-robot collectives. The work has value in extending techniques for robot–robot learning. But its primary purpose is to model and illuminate low-level processes and mechanisms of behavioural evolution. Embodied social learning provides minimal but sufficient biological plausibility and as outlined here, embodiment leads naturally to imperfect imitation, which appears to play an important role in the dynamics of behavioural evolution.

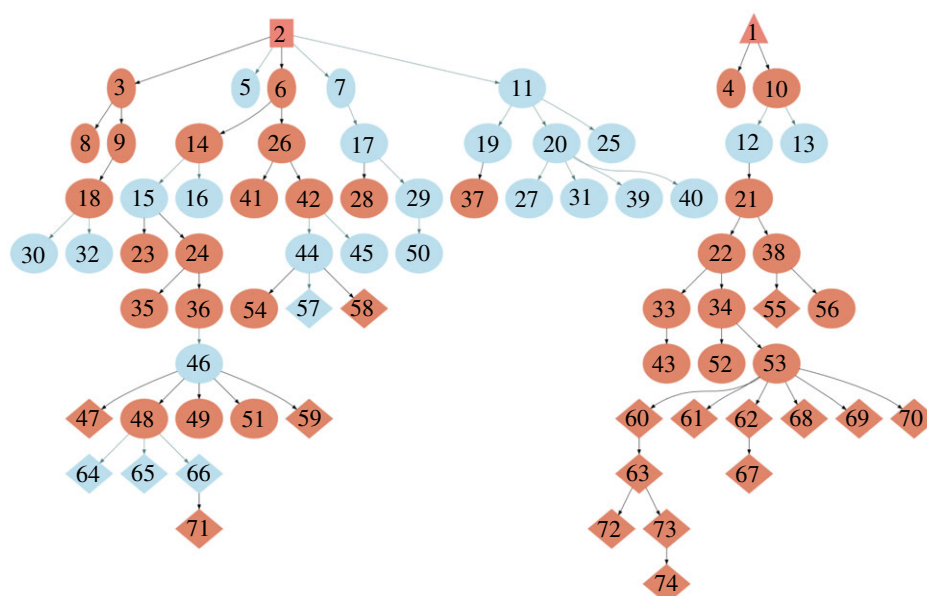
### 3. Storybots

In more recent work, we have extended the robots’ cognition with a simulation-based internal model. Robots equipped with a simulation-based internal model have the ability to simulate (or ‘imagine’) the future actions of both themselves and others, and the consequences of those actions.

Figure 4 shows a block diagram of a robot equipped with a consequence engine (CE), which consists of the blue boxes and dataflows on the left. The simulator at the heart of the CE contains three components: a model of the world, which must be initialized to mirror the robot’s immediate environment



**Figure 2.** Trajectory plot: two-robot movement-meme evolution in which all copied memes are stored, and meme selection is random with equal probability. The experiment starts with e-puck A (a) in teacher mode and e-puck B in learner mode (b), following a pre-programmed movement trajectory (seed) that describes a triangle (1) with sides of 15 cm. Each meme is shown in a different colour. By chance the triangle meme (1) is imitated imperfectly by e-puck B as an angular figure-of-eight meme (2). This is followed by a high-fidelity imitation (2)  $\rightarrow$  (3), and again (3)  $\rightarrow$  (4), so the figure-of-eight meme becomes dominant in the two robots' collective memory. Adapted from Winfield & Erbas [20].



**Figure 3.** Behavioural evolution map following a four-robot experiment with limited memory; each robot stores only the most recent five learned behaviours. The four robots are each seeded with two behaviours: a triangle dance (labelled 1) and a square dance (labelled 2), from which all memes are descended. Orange nodes are high fidelity copies, light blue nodes are low fidelity copies. The 20 behaviours in the collective memory of all four robots at the end of the experiment are highlighted as diamonds. Note one cluster of 12 closely related memes (55), (60–63), (67–70) and (72–74) all descended from (1). Adapted from Erbas *et al.* [21].

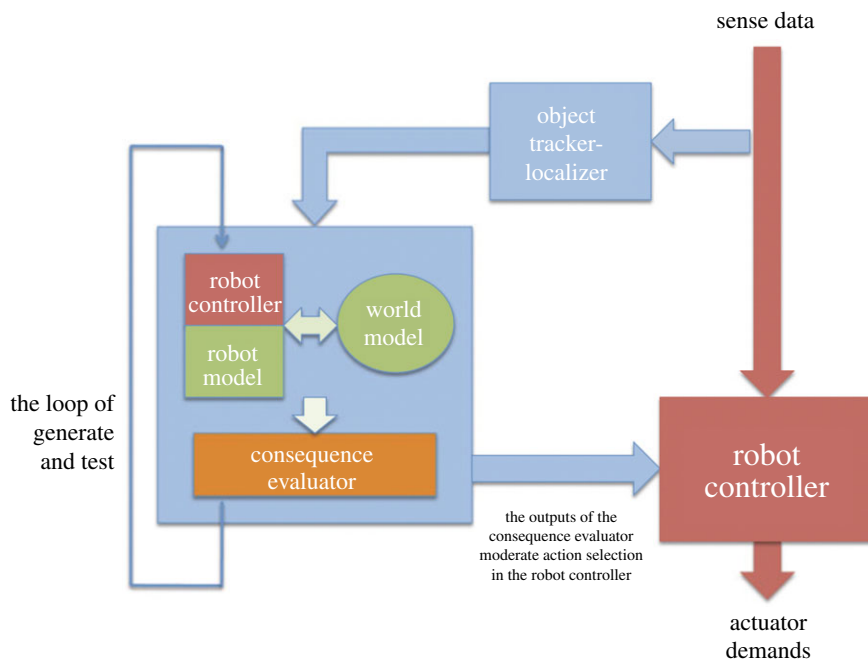
including the objects and actors in it, as it is now, via the 'object tracker-localizer'; a model of the robot itself; and an exact copy of the robot's controller. The loop of generate-and-test shown on the left, generates each of the robot's next possible actions, then 'runs' the simulator for each of those actions in turn. The consequence evaluator determines the anticipated outcome for each of those actions, so that the robot's action selection can be appropriately moderated (what counts as appropriate depends on whether the CE's primary purpose is keeping the robot safe, or behaving ethically, etc.). In our experiments, the CE will typically generate and test 30 next possible actions and, for each action, simulate 10 s into the future. The complete generate-and-test cycle will be repeated every 0.5 s.

The CE has proven to be a remarkably powerful piece of cognitive machinery. With it we have experimentally demonstrated (i) robots that can make simple ethical decisions in order to pro-actively prevent another robot (acting as a proxy human) from coming to harm [28,29]; (ii) robots with enhanced safety [30]; and (iii) robots capable of the imitation


of goals [31]. We have also argued that the CE provides a robot with a simple artificial theory-of-mind [32]. These experiments were conducted with real physical e-puck (figure 1) and NAO robots (figure 6).

The pedestrian experiment provides an elegant example of emergent behaviour in two robots, each equipped with a CE, and programmed with the goal of approaching and then passing each other safely. Figure 5 shows the trajectories of the two robots: blue, starting from the left, and green, starting from the right. In real-robot experiments, four times out of five blue and green pass each other as two pedestrians would, each stepping to her left (or right) (see figure 5a). But one time in five, both blue and green step toward each other, and—just like humans—engage in a short dance before resolving and proceeding on their way (figure 5b).

More recently we have theorized that the CE may also be co-opted as a mechanism for robot-robot storytelling [27]. The CE provides a robot with the cognitive machinery to be able to ask 'what if?' questions. These could be very



**Figure 4.** Block diagram of a robot equipped with a consequence engine. The robot control dataflows are shown in red (right); the consequence engine and its dataflows in blue (left). Adapted from Winfield [27].



**Figure 5.** The pedestrian experiment—two trials showing robot trajectories. Two robots, blue and green, are each equipped with a CE. Blue starts from the right, with a goal position on the left, while at the same time green starts from the left with a goal position on the right. (a) We see the typical behaviour in which the two robots pass each other without difficulty, normally because one robot—anticipating a collision—changes direction first, in this case green. (b) Here both robots make a decision to turn at the same time, green to its left and blue to its right; a ‘dance’ then ensues before the impasse is resolved. Adapted from Winfield [32].

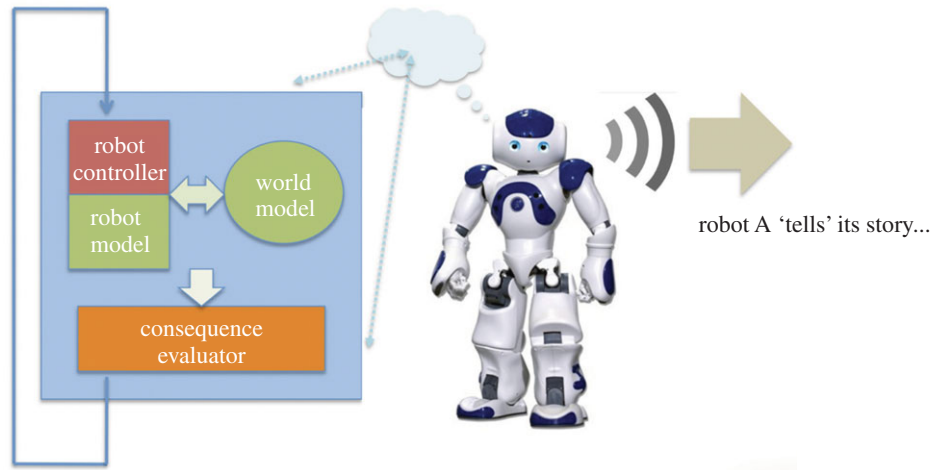
simple questions like ‘what if I turn left?’ But consider now that instead of the robot using the predictions of the CE to perform (or not perform) some action, it narrates that action together with its imagined outcome to another robot also equipped with a CE, as in ‘If I turned left, I would collide with the wall’. The robot would, in effect, be thinking out loud. A more accurate term might be ‘narrated cognition’, although we prefer to use the shorthand ‘story’.<sup>1</sup> Figure 6 illustrates this process for robot A.

If robot B, the listener, is equipped with a microphone and speech recognition process it is able to listen to robot A’s story, as shown in figure 7. Because robot B has the same internal modelling machinery as A—they are conspecifics—it is capable of ‘running’ the story it has just heard within its own internal model. In order that this can happen we need to modify the robot’s programming so that the what-if sequence it has heard and interpreted is substituted for an internally generated what-if sequence. Once that substitution is made, robot B is able to run A’s what-if sequence (its story) in exactly the same way it runs its own internally generated next possible

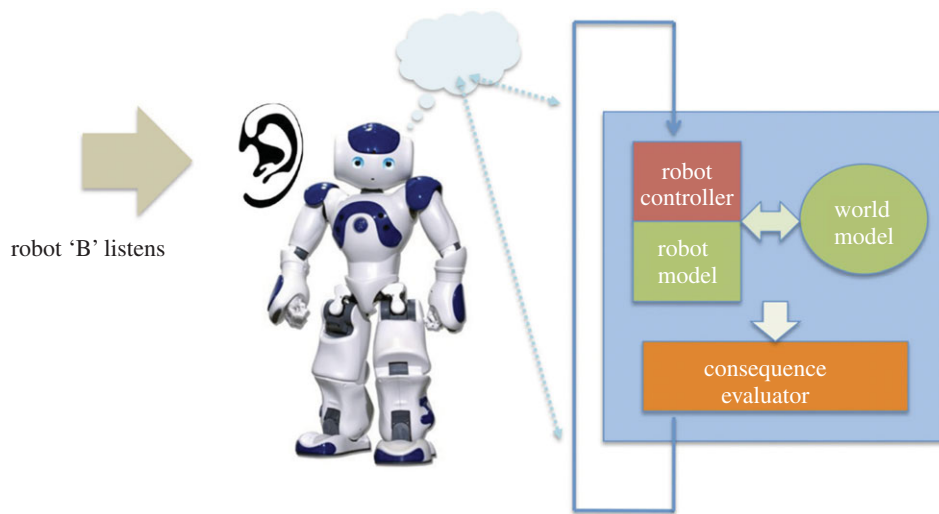
actions, simulating and evaluating the consequences. Robot B is therefore able to ‘imagine’ robot A’s story. Does this story mean anything to robot B? Arguably it does, as B is able to simulate and therefore ‘experience’ the sensory inputs, and consequences (if any) of listening to A’s story.

Note that the humanoid NAO robots, shown in figures 6 and 7, do not have human-like intelligence even though their appearance might suggest otherwise. Like the e-puck Copybots, all behaviours must be programmed from scratch. The NAO robots do, however, have the advantage of microphones and loudspeakers, alongside a library of functions for speech recognition and synthesis. This makes them much more suitable as Storybots.

If we provide not just two, but a group of robots with a rich physical environment they can explore then we are providing the robots with something they can tell each other stories about. And, for the same reasons that our Copybots’ imitation is noisy, so will our Storybots experience imperfect communication, so the stories will mutate as they are told and re-told. The architecture of the CE



**Figure 6.** Robot A, the storyteller, ‘narrativizes’ one of the ‘what-if’ sequences modelled by its generate-and-test machinery. First, an action is tested in the robot’s internal model (left); second, that action—which is not executed for real—is converted into speech and spoken by the robot. Adapted from Winfield [27].



**Figure 7.** Robot B, the listener, uses the same ‘what-if’ cognitive machinery to ‘imagine’ robot A’s story. Here the robot hears A’s spoken sequence, then converts it into an action (or series of actions) which is simulated in its own internal model. Adapted from Winfield [27].

and its simulation-based internal model opens the possibility that we can replay and visualize any episode in a robot’s ‘imagination’, thus adding further detail to Tarde’s ‘minute transformations’ and allowing us to inspect the robots’ mental representation of stories as they pass from robot to robot.

## 4. Discussion and conclusion

### (a) Dennett’s tower of generate-and-test: a unifying framework

As mentioned in the introduction, Dennett’s tower of generate-and-test directly inspired the second series of experiments outlined above, culminating in the Storybots. Dennett’s tower also provides us *post facto* with a single framework to unify all of the experimental work outlined in this paper.

Dennett [7] proposes a conceptual framework, the tower of generate-and-test, for thinking about design options for brains. Each floor of the tower uses the three Darwinian operators: copy,

generate variations, test outcomes—repeat. Each floor builds on the outcome of the previous ones. The framework provides a way of seeing how humans, as a cultural species, emerged from creatures with no cumulative culture, using the same ‘generate and test’ process all the way up. The ground floor is inhabited by Darwinian creatures. Variation is provided by the more or less random recombination and mutation of genes, and selection is brutal—design-by-death [33]. All living things are Darwinian creatures.

Some of these creatures emerged with conditionable plasticity; that is, not all their behaviour was genetically determined. Occupying the first floor of Dennett’s tower, these Skinnerian creatures try out a variety of responses to their environment selecting only actions that are reinforced for repeating. Dennett named them after Skinner’s comment that, ‘Where inherited behaviour leaves off, the inherited modifiability of the process of conditioning takes over’ [34, p. 83]. Most plants and animals are Skinnerian, as well as Darwinian, creatures.

On the third floor, we find Popperian creatures. These are able to think through possible actions and their likely consequences before carrying them out. This requires internal

**Table 1.** Classifying the robots of this paper within Dennett's conceptual tower of generate-and-test.

robot	Darwinian	Skinnerian	Popperian	Gregorian
basic Copybots [20]	✓			
Copybots with imitation-enhanced reinforcement learning [25]	✓	✓		
robots with consequence engine (CE) [30,31]	✓		✓	
Storybots [27]	✓		✓	✓

models of (relevant features of) the environment, as well as of their own body and behaviour. Variation comes from imagining different actions; selection is by imagined consequences. As Popper remarked, when imagining outcomes we 'let our conjectures, our theories, die in our stead' [35]. We are not alone in being Popperian creatures; most mammals, birds, fish and reptiles can learn through both classical and operant conditioning, and can contemplate the consequences of at least some of their actions.

On the fourth floor are Gregorian creatures. As far as we know, we are the only Gregorian creatures, at least on this planet. Gregory introduced the idea of 'tools of Mind' or 'mind-tools' by which he meant 'aids to measuring, calculating and thinking' [36, p. 48], including tools like scissors or levers as well as spoken and written words, and ways of counting. Language makes possible long trains of thought, the ability to look ahead, and the sharing of tools that enhance intelligence. These tools are built up over generations by creatures that can copy information from each other, building up culture. Gregorian creatures are therefore meme machines [2] as well as Darwinian, Skinnerian and Popperian creatures. This ability to imitate and learn from others makes possible what Dennett calls the 'deliberate, foresightful generate-and-test known as science' [7, p. 380].

Dennett's tower of generate-and-test contrasts with other modular-mind frameworks, for instance Mithen's cathedral model [37], in three important respects: (i) it is not a flat model of intelligence modules but instead a nested hierarchy, (ii) it defines a number of key transitions in the evolution of mind, and (iii) in the third transition, from Skinnerian to Popperian, it introduces the crucial innovation of an internal model. Dennett's tower is commented on in several papers (including [38,39], and more recently [40]). Godfrey-Smith's paper [40] both critiques and extends Dennett's framework, and notably points out an omission in the original framework, that an internal model requires a mechanism for consequence evaluation in order to be of value. In developing our consequence engine (CE), we too realized that we needed to implement just such a mechanism.

Let us now place our robots within the floors of Dennett's tower. The basic Copybots are Darwinian creatures alone, their design having been selected from many other possible designs we might have chosen—a form of design-by-death. The Copybots with imitation-enhanced learning are also Skinnerian creatures. Both types can imitate, which might suggest they are Gregorian, but this would be a misclassification as the Copybots have no internal model—the defining characteristic of Popperian, and hence also Gregorian creatures. By contrast, all of our robots with a CE are Popperian; the CE enables them to generate and test hypotheses about what to do next. They are, however, not strictly Skinnerian because we have not added reinforcement learning (although this is perfectly

feasible). The Storybots proposed in figures 6 and 7 finally take us from the Popperian to Gregorian level.

While the Copybots imitate by visually observing movement memes, lacking a CE they cannot predict and evaluate the consequences of those imitated behaviours. The Storybots, on the other hand do not imitate behaviour directly. Their method of learning from others is mediated by the mind-tool of language. Table 1 summarizes the classification outlined here.

Applying the framework of Dennett's tower does raise the interesting question of whether intelligence must necessarily be achieved by building each level of generate-and-test on top of the preceding one, or whether robot intelligence might skip one or more, to achieve cumulative culture more directly. The fact that we are able to skip the Skinnerian level is a consequence of hard coding the capabilities summarized in table 1. Strictly speaking, this is a serious break with both the evolutionary origins and cumulative nature of each level in Dennett's framework. Given, however, that artificially evolving each capability in succession is far beyond what can be achieved in evolutionary robotics at present [41], and that we have demonstrated the implementation of Skinnerian learning, the break is—we believe—justified.

### (b) What can we learn from the Storybots?

In order to address this question, consider first what we might learn from the Storybots as presented in §3. These robots are equipped with a CE plus the speech synthesis and recognition capabilities shown in figures 6 and 7. Like the other robots with a CE outlined at the start of §3, the basic Storybots have a short-term memory which is used to retain the evaluated consequences of each generated-and-tested action, in order that the most appropriate action can be selected. But at the end of each complete cycle of generate-and-test that short-term memory is cleared ready for the next cycle. These robots do not learn, which may seem surprising given the capabilities demonstrated by the CE. However, even these basic Storybots (like the first Copybots we tested) can usefully allow us to explore how stories vary as they are told and re-told by several robots. As part of this experiment, we would first need to find the optimal balance between zero and perfect fidelity robot-robot speech-transmission (as we had to do with the Copybots for vision-based imitation) such that we do see a reasonable level of variation; this would require for instance adjusting speech loudness, microphone sensitivities and attending to directionality so that the listener robot is facing the narrating robot.

Next consider memory. It would be straightforward to equip each Storybot with a long-term (episodic) memory. The memory would store discrete events (things that







**Figure 10.** The work areas of the two e-puck robots used in the imitation-enhanced reinforcement learning experiment. Note the small gaps in the bottom left of each area.

**Competing interests.** We declare we have no competing interests.

**Funding.** The work outlined in §2 was funded by EPSRC grant ref EP/E062083/1. The work outlined in §3 was in part funded by EPSRC grant ref EP/L024861/1.

**Acknowledgements.** The authors gratefully acknowledge project co-investigators James Bown, Robin Durie, Frances Griffiths, Jenny Tennant Jackson and Alistair Sutcliffe, and are especially grateful to Mehmet Erbas who conducted most of the experimental work. The CE was implemented by Christian Blum, Wenguo Liu and Dieter Vanderelst, and the pedestrian experiment was coded and run by Mathias Schmerling and Chen Yang. We are also grateful for conversations with Daniel Dennett and the late Richard Gregory who directly inspired the Storybots when he declared ‘when your robots start telling each other stories *then* you’ll really be onto something’. We are very grateful to the both the anonymous reviewers and editor, for their thoughtful comments on the first draft of this paper, from which many improvements followed.

## Endnotes

<sup>1</sup>We also wish to honour Richard Gregory who first suggested story telling robots—see Acknowledgements.

<sup>2</sup>Which would provide an interesting model of remembering by re-imagining [42].

<sup>3</sup>Noting that Conway [44] uses the term ‘working self’.

## Appendix A. Imitation in the copybots

In this appendix, we outline how the Copybots learn movement-memes by imitation. This algorithm solves the so-called ‘correspondence problem’, a term which refers to the learner’s problem of translating a set of perceptual inputs to motor actions that correspond with the perceived actions of the teacher [45].

To simplify the process, each movement-meme consists only of turns (in which the e-puck robot rotates on the spot) and straight line segments of a given length. Thus the ‘triangle’ meme seeded into e-puck A, in figure 2 (meme 1), is described by the list of three pairs (60, 15), (60, 15), (60, 15), where (60, 15) means ‘rotate 60° then move 15 cm’.

The e-pucks are fitted with red ‘skirts’, so that when the learner robot observes the teacher robot with its camera the

learner ‘sees’ the teacher’s skirt as a red rectangle in its field of vision. While the teacher is enacting the movement-meme, the red rectangle will both move within the learner’s field of vision and sometimes get larger or smaller—as the teacher moves either nearer to or further away from the learner. The learner robot uses vision processing to estimate the position of the teacher robot, relative to itself, as  $x,y$  coordinates. Such a list of coordinates is shown on the left of figure 8. This list is the input to the algorithm.

The first stage of the algorithm ‘detect turns’ is to identify when the teacher robot is turning, by finding pairs of similar  $x,y$  coordinates. These are circled in figure 8. The turns mark the beginning and end of each straight line move. The second stage is to then use a line-fitting (regression) algorithm to estimate each straight line move. An illustration is shown in figure 9. The final stage is to put each turn and estimated line together as a reconstructed trajectory—a list of pairs of turn angles and distances moved. In this way, the imitation algorithm enables the learner robot to infer the teacher robot’s sequence of moves. For a more detailed explanation, see Erbas *et al.* [21].

## Appendix B. Imitation-enhanced reinforcement learning

In this experiment, each robot has its own work area, as shown in figure 10, and must—using individual (reinforcement) learning—learn how to navigate from the top right-hand corner, to the bottom left-hand corner of its area. Learning this way is slow, taking several hours. But in this experiment the robots also have the ability to learn socially, by watching each other. Periodically one of the robots will stop its individual learning and drive itself out of its own area, to the small opening at the bottom left corner of the other robot’s work area. There it will stop and simply watch the other robot while it is learning for a few minutes. Using the same movement imitation algorithm outlined in appendix A, the watching robot will (socially) learn a fragment of what the other robot is doing, then combine this knowledge into what it is individually learning. The robot

then runs back to its own work area and resumes its individual learning. We call the combination of social and individual

learning ‘imitation-enhanced learning’. For a more detailed explanation, see Erbas *et al.* [25].

## References

- Mesoudi A. 2011 *Cultural evolution: how Darwinian theory can explain human culture and synthesize the social sciences*. Chicago, IL: University of Chicago Press.
- Blackmore S. 1999 *The meme machine*. Oxford, UK: Oxford University Press.
- Barry A, Thrift N. 2007 Gabriel Tarde: imitation, invention and economy. *Econ. Soc.* **36**, 509–525. (doi:10.1080/03085140701589497)
- Marques H, Holland OE. 2009 Architectures for functional imagination. *Neurocomputing* **72**, 743–759. (doi:10.1016/j.neucom.2008.06.016)
- Hesslow G. 2002 Conscious thought as simulation of behaviour and perception. *Trends Cogn. Sci.* **6**, 242–247. (doi:10.1016/S1364-6613(02)01913-7)
- Hesslow G. 2012 The current status of the simulation theory of cognition. *Brain Res.* **1428**, 71–79. (doi:10.1016/j.brainres.2011.06.026)
- Dennett D. 1995 *Darwin's dangerous idea*. Harmondsworth, UK: Penguin.
- Steels L. 1999 *The talking heads experiment*. Brussels, Belgium: Laboratorium Antwerpen.
- Marocco D, Nolfi S. 2007 Emergence of communication in embodied agents evolved for the ability to solve a collective navigation problem. *Connect. Sci.* **19**, 53–74. (doi:10.1080/09540090601015067)
- Penn DC, Holyoak KJ, Povinelli DJ. 2008 Darwin's mistake: explaining the discontinuity between human and nonhuman minds. *Behav. Brain Sci.* **31**, 109–130. Discussion 130–178. (doi:10.1017/S0140525X08003543)
- Pfeifer R, Bongard J. 2006 *How the body shapes the way we think: a new view of intelligence*. Cambridge, MA: MIT Press.
- Mesoudi A, Whiten A, Laland KN. 2006 Towards a unified science of cultural evolution. *Behav. Brain Sci.* **29**, 329–347. (doi:10.1017/S0140525X06009083)
- Acerbi A, Mesoudi A, Smolla M. 2020 Individual-based models of cultural evolution: a step-by-step guide using R. OSF Preprints. (doi:10.31219/osf.io/32v6a)
- Derex M. 2021 Human cumulative culture and the exploitation of natural phenomena. *Phil. Trans. R. Soc. B* **377**, 20200311. (doi:10.1098/rstb.2020.0311)
- Winfield AFT *et al.* 2011 First steps toward artificial culture in robot societies. *Procedia Comput. Sci.* **7**, 130–132. (doi:10.1016/j.procs.2011.09.063)
- Mondada F *et al.* 2009 The e-puck, a robot designed for education in engineering. In *Proc. 9th Conf. on Autonomous Robot Systems and Competitions, LIS-CONF-2009-004*, pp. 59–65. Castelo Branco, Portugal: Instituto Politécnico de Castelo Branco.
- Winfield AFT. 2012 *Robotics a very short introduction*. Oxford, UK: Oxford University Press.
- Bredeche N, Fontbonne N. 2021 Social learning in swarm robotics. *Phil. Trans. R. Soc. B* **377**, 20200309. (doi:10.1098/rstb.2020.0309)
- Dawkins R. 1976 *The selfish gene*. Oxford, UK: Oxford University Press.
- Winfield AFT, Erbas MD. 2011 On embodied memetic evolution and the emergence of behavioural traditions in robots. *Memetic Comput.* **3**, 261–270. (doi:10.1007/s12293-011-0063-x)
- Erbas MD, Bull L, Winfield AFT. 2015 On the evolution of behaviors through embodied imitation. *Artif. Life* **21**, 141–165. (doi:10.1162/ARTL\_a\_00164)
- Kirby S, Cornish H, Smith K. 2008 Cumulative cultural evolution in the laboratory: an experimental approach to the origins of structure in human language. *Proc. Natl Acad. Sci. USA* **105**, 10 681–10 686. (doi:10.1073/pnas.0707835105)
- Eriksson K, Coultas JC. 2012 *Theory of conformist social learning*. Boston, MA: Springer US.
- Kaelbling L, Littman M, Moore A. 1996 Reinforcement learning: a survey. *J. Artif. Intell. Res.* **4**, 237–285. (doi:10.1613/jair.301)
- Erbas MD, Winfield AFT, Bull L. 2014 Embodied imitation-enhanced reinforcement learning in multi-agent systems. *Adapt. Behav.* **22**, 31–50. (doi:10.1177/1059712313500503)
- Kendal RL, Boogert NJ, Rendell L, Laland KN, Webster M, Jones PL. 2018 Social learning strategies: bridge-building between fields. *Trends Cogn. Sci.* **22**, 651–665. (doi:10.1016/j.tics.2018.04.003)
- Winfield AF. 2018 When robots tell each other stories: the emergence of artificial fiction. In *Narrating complexity* (eds R Walsh, S Stepney), pp. 1–11. Berlin, Germany: Springer.
- Winfield AFT, Blum C, Liu W. 2014 Towards an ethical robot: internal models, consequences and ethical action selection. In *Advances in autonomous robotics systems* (eds M Mistry, A Leonardis, M Witkowski, C Melhuish), pp. 85–96. Berlin, Germany: Springer.
- Vanderelst D, Winfield A. 2018 An architecture for ethical robots inspired by the simulation theory of cognition. *Cogn. Syst. Res.* **48**, 56–66. (doi:10.1016/j.cogsys.2017.04.002)
- Blum C, Winfield AFT, Hafner VV. 2018 Simulation-based internal models for safer robots. *Front. Robot. AI* **4**, 74. (doi:10.3389/frobt.2017.00074)
- Vanderelst D, Winfield A. 2017 Rational imitation for robots: the cost difference model. *Adapt. Behav.* **25**, 60–71. (doi:10.1177/1059712317702950)
- Winfield AFT. 2018 Experiments in artificial theory of mind: from safety to story-telling. *Front. Robot. AI* **5**, 75. (doi:10.3389/frobt.2018.00075)
- Blackmore S. 2010 Dangerous memes; or what the Pandorans let loose. In *Cosmos and culture: cultural evolution in a cosmic context* (eds S Dick, M Lupisella), pp. 297–318. Washington, DC: NASA.
- Skinner BF. 1953 *Science and human behavior*. New York, NY: Macmillan.
- Popper K. 1977 Natural selection and the emergence of mind. Lecture delivered at Darwin College, Cambridge, 8 November. [https://www.informationphilosopher.com/solutions/philosophers/popper/natural\\_selection\\_and\\_the\\_emergence\\_of\\_mind.html](https://www.informationphilosopher.com/solutions/philosophers/popper/natural_selection_and_the_emergence_of_mind.html).
- Gregory RL. 1981 *Mind in science: a history of explanations in psychology and physics*. Cambridge, UK: Cambridge University Press.
- Mithen S. 1996 *The prehistory of the mind: a search for the origins of art, religion and science*. Phoenix, AZ: Orion Books.
- de Villiers T, Cilliers P. 2004 Narrating the self: Freud, Dennett and complexity theory. *S. Afr. J. Phil.* **23**, 34–53. (doi:10.4314/sajpem.v23i1.31383)
- Beisecker D. 2006 Dennett's overlooked originality. *Minds Mach.* **16**, 43–55. (doi:10.1007/s11023-006-9017-3)
- Godfrey-Smith P. 2018 Towers and trees in cognitive evolution. In *The philosophy of Daniel Dennett* (ed. B Heubner), pp. 225–249. Oxford, UK: Oxford University Press.
- Hart E, Le Goff L. 2021 Artificial evolution of robot bodies and control: on the interaction between evolution, learning and culture. *Phil. Trans. R. Soc. B* **377**, 20210117. (doi:10.1098/rstb.2021.0117)
- Rathbone CJ, Conway MA, Chris JA. 2011 Remembering and imagining: the role of the self. *Conscious Cogn.* **20**, 1175–1182. (doi:10.1016/j.concog.2011.02.013)
- Migliano AB, Vinicius L. 2021 The origins of human cumulative culture: from the foraging niche to collective intelligence. *Phil. Trans. R. Soc. B* **377**, 20200317. (doi:10.1098/rstb.2020.0317)
- Conway MA. 2005 Memory and the self. *J. Memory Lang.* **53**, 594–628. (doi:10.1016/j.jml.2005.08.005)
- Nehaniv CL, Dautenhahn K. 2002 The correspondence problem. In *Imitation in animals and artifacts* (eds CL Nehaniv, K Dautenhahn), pp. 41–61. Cambridge, MA: MIT Press.