BOOKS & ARTS

Choose your own reward

Does human creativity stem from a process that turns arbitrary ideas into goals like food and sex?

Why Choose This Book? How We Make Decisions
by Read Montague
Dutton: 2006. 331 pp. $24.95, £15.99

Andy Clark

Why don't our brains get as hot as the processors in our personal computers? And what does that tell us about biological computation, the nature of choice, the value of value signals, and the power of ideas? These questions may seem rather disparate, but Read Montague's provocative and accessible treatment of them in *Why Choose This Book?* displays a deep and unexpected unity. The story goes something like this. Biological computation is constrained to be efficient: to confer close to maximal computational power for close to minimal energetic expenditure. The outward sign of this efficiency is said to be the mere warmth of the human brain compared with the searing (wasted) heat of those computer processors. To be efficient, biological computations need to be equipped, so the argument goes, with some kind of measure of their own value, in relation to the in-built goals of maintaining life and reproductive success. Such measures allow the system to expend energy only on those computations that matter most. How such a measure might work remains problematic, but once it is in place, general principles of thrifty processing, such as the slow use of power, compression of data, conservation of wiring, and frugal use of bandwidth and communication, are all recruited to the mix.

But it is the goals and value signals that play the lead role in Montague's story. He introduces us to the guiding principle that will link efficiency to choice and to the power (and pathologies) of ideas. That principle, familiar enough in cognitive scientific circles but here tweaked and nuanced in novel and potentially transformative ways, goes by the unpromising name of 'reinforcement learning'. In reinforcement learning, goal states are approached by sensitivity to signals that predict rewards (the attainment of goals). But the system is not simply hardwired to regard only some fixed set of signals as reward signals; it obtains flexibility by learning associations between experienced signals and temporally removed (but consistently associated) rewards. Past experience of what signal leads to what reward is thus combined with present feedback (what's here now, and what is it worth?) to generate choices that (ideally) maximize total future reward.

R. BAMBER/ALAMY

Neat trick: activities such as gambling and card games can supplant the basic rewards of food or sex.

Moving all this along is the 'reward-prediction error signal', which carries information about how well the actual rewards tally with the predicted rewards. When the actual reward exceeds the predicted one, it makes sense to upgrade the stored value of the states that predicted the unexpectedly greater reward. In the brain, dopamine neurons provide at least one means of mechanistically encoding just such a reward-prediction error signal. Bursts of dopaminergic activity result when the reward exceeds the predicted reward; pauses in activity mean that the reward falls short of the predicted reward; and unchanged activity means the reward was as expected.

But what counts as a reward anyway? The most obvious rewards are the basic biological achievements of life maintenance (such as the ingestion of a tasty and nourishing morsel) and reproduction (or rather its precursor, sexual intercourse). Montague is motivated, however, by a strong desire to unravel the mechanistic underpinnings of what he describes as a uniquely human 'superpower': the capacity to make choices that seem to value biologically arbitrary objects, achievements and actions. Examples of such biologically arbitrary goal states mentioned in the text include solving Fermat's last theorem and committing group suicide in the belief that a spaceship hidden in a comet's tail will then take you to 'the next level'. What makes all this possible, in Montague's model, is the capacity of ideas themselves to act as reward signals, hijacking the prediction-error systems implemented by dopamine neurons in the brain. When this happens, the dopamine outputs start to act as error signals that encourage the rest of the brain to learn and to make decisions in ways that increase the chances of acquiring some biologically arbitrary reward.

Given the potentially biologically catastrophic consequences of such re-tooling of mere thoughts as rewards, Montague suggests that powerful filtering processes control what gets into the reward slot. But such processes can be fooled — in ways that the book describes in compelling and often sinister detail — by damage, by drug abuse, and perhaps even by some forms of advertising and branding (brands are just cues that predict rewards). Montague's proposal is that biologically arbitrary goals can somehow plug into a kind of 'special status reward socket', and thus become a basic, primary reward, like food or sex. He does not claim that these ideas become associated, either directly or indirectly, with food or sex; rather,

they plug directly into the 'socket' normally occupied only by the most basic high-status rewards. If we humans have indeed learnt such a powerful trick, it is no surprise that it fuels so much that is both good (creative and expansive) and ill (pathological and restrictive) in our species. Montague begins by laying out this possibility, then follows it deep into the fascinating territories of creative thought, addiction, obsessive–compulsive disorder, Parkinson's disease, and then on to the psychosocial realms of trust and regret.

Despite its attractions, there are some important mechanistic gaps in the story, as Montague acknowledges. For example, it isn't clear why or how one idea might win out over another in the bid to occupy a high-status reward socket, or how the occupation itself is accomplished. Nor is it really clear when such occupation should be deemed pathological rather than creative. I was also left wondering whether the basic idea of each symbol and each computation carrying its own value 'tag' — the difference, Montague argues, between standard computational models using 'meaningless symbols' and the hyper-efficient, value-rich computations said to be characteristic of biological nervous systems — is sufficiently clear and workable. Exactly how do these computation-value or symbol-value pairs work, and how do they transform mere symbol processing into meaning? Do they compose? Two computations whose individuals values are low might together constitute a complex computation whose value to the organism is high, but Montague suggests no way of systematically predicting such combined values from the values assigned to the parts.

Perhaps I am missing something, but it repeatedly struck me that Montague's overall vision is both rather more radical, and rather less mechanistically clear, than his book suggests. The prospective reader should be aware that the story on offer actually departs quite a long way from the basic computational theory of the mind. It builds in value and computation right down to the cellular level, and (more generally) systematically blurs the usual distinctions between life, mind and information processing. This blurring is evident, for example, in the puzzling idea that each individual neuron, in the quest for efficient interneural communication, might need to contain up to 100 million 'dynamic models' of other neurons and neuronal subsystems.

These are not really complaints, however. The book spans several seldom-bridged worlds, from neuroscience to psychiatry, economics and social psychology, and does so with wit, precision and elegance. It succeeds in many of its goals. Above all, it left me feeling I had actually learnt something about myself: a thinking, feeling, choosing, yet painfully vulnerable chemically modulated learning machine. ■

Andy Clark is at the School of Philosophy, Psychology and Language Sciences, University of Edinburgh, Edinburgh EH8 9JX, UK.

NASA KSC

James Webb (second left, with President John F. Kennedy) led NASA's Apollo space programme.

Shooting for the Moon

The Man Who Ran the Moon: James Webb, JFK, and the Secret History of Project Apollo
by Piers Bizony
Thunder's Mouth/Icon Books: 2006.
256 pp. $24.95/£16.99

Alex Roland

James Webb was the antithesis of the enterprise that made him famous. Administrator of NASA during the Apollo space programme, Webb commanded a sprawling, anonymous team of 400,000 workers. Yet he himself was colourful, singular, eccentric even — a politician and an individualist leading an army of technicians. Apollo was 'big science', what Webb and others called 'large-scale technology'. It dwarfed the Manhattan Project, for example, in cost, size and complexity. In contrast to the Wright brothers, who achieved atmospheric flight using equipment from their bicycle shop, the flight to the Moon integrated civil servants and industry contractors in a nationwide web of activity. They built unprecedented facilities such as the launch centre at Cape Canaveral and the manned-spacecraft centre at Houston Texas, and dispensed some $25 billion (worth several times as much today) over a decade or more. Such undertakings demand the subordination of the individual to the cause. Webb was certainly a team player, but he was also the pilot who steered his own course and imposed his personality on Apollo.

Webb's background and the purportedly "secret history" revealed in *The Man Who Ran the Moon* by Piers Bizony have long been well known. An accomplished Democratic operative and former budget director for President Harry Truman, Webb accepted the NASA position after as many as 17 other prospects had turned it down. He ran NASA (not the Moon) like a chief executive, handling the politics of external relations and leaving the technical management of the programme to trusted subordinates. He survived power struggles, budget battles and conflicts with two presidents, John F. Kennedy and Lyndon B. Johnson, over maintaining an overall space programme balanced between manned and unmanned activities. And he won the right to distribute some of NASA's money to support socially constructive programmes such as university development. Smart, energetic, gregarious and iconoclastic — his official limousine was a Checker cab — he cut a colourful swath through the bureaucratic maze of Washington.

Then came the Apollo 204 tragedy. On 27 January 1967, three astronauts died in a gruesome fire during a routine ground test of the Moon capsule at Cape Canaveral. In the ensuing investigation, an internal NASA report came to light that criticized the capsule manufacturer, North American Aviation (NAA), for shoddy work. Caught out on the witness stand without foreknowledge of the report, Webb began to lose control of events. Worse still, he appeared to misrepresent the fact that he and his leading deputies had overturned the recommendation of their own review panel and awarded NAA the contract in the first place. This revelation tarred him with the brush of the Bobby Baker scandal, a rat's nest of lobbying, bribes and even organized crime.

Webb survived the hearings but his reputation was damaged and his power diminished. New executives were brought in to restore