livered from granule cell bodies into dendritic process, where mRNA levels remain elevated for several hours. New evidence based on local infusion of Arc antisense oligodeoxynucleotides shows that *stabilization* of LTP requires a window of Arc translation (Messaoudi et al. 2004). In addition, a handful of mRNA species are stored in dendrites in a translationally dormant state. One of these mRNA encodes the α-subunit of CaMKII, a key enzyme in synaptic plasticity. LTP is associated with translocation of preexisting CaMKII mRNA into dendritic spines (Havik et al. 2003), and mutant mice lacking dendritic CaMKII mRNA have impaired late LTP and memory function (Miller et al. 2002). Furthermore, in isolated synapses, BDNF application drives the translation of both CaMKII and Arc mRNA (Kanhema et al., submitted; Yin et al. 2002). This raises the possibility that BDNF stimulates synaptic stabilization through translation of transient (Arc) and resident (CaMKII) dendritic mRNA. As components of the post-synaptic density, these gene products are positioned to modulate glutamate receptor signaling and synaptic structure.

When in sleep would this take place? In slow-wave sleep (SWS), synchronized population events such as hippocampal sharp-waves and dentate spikes could be involved (Bramham 1998; Buzsaki 1989). The frequency of these events is increased during SWS, and this state is associated with global increases in protein synthesis (Ramm & Smith 1990) Depolarization of dendrites during these population events may facilitate BDNF release, which then drives dendritic mRNA translation. In REM sleep, synaptic inputs synchronized to the positive theta peak could have a similar role. Furthermore, cholinergic transmission during REM sleep could contribute to modulation of dendritic protein synthesis (Feig & Lipton 1993).

In addition to stabilization, sleep-related modification of synaptic strength may involve activity-dependent decay of glutamatergic transmission and depotentiation. Decay of LTP appears to be an active, NMDA receptor-dependent processes extending over several days (Villarreal et al. 2002). Synapses transiently potentiated during waking that fail to undergo protein synthesis-dependent stabilization may be depotentiated. Protection from depotentiation is protein synthesis-dependent (Woo & Nguyen 2003). These mechanisms of selective stabilization and destabilization may contribute to *enhancement* of memory consolidation during sleep.

Modulation of LTP induction during sleep. The target review also raises the question of new acquisition during sleep. Studies of LTP have shed some light on this issue. For example, LTP is powerfully modulated by sleep-wakefulness state. While readily induced during REM sleep, LTP induction is suppressed in an all-or-none manner during SWS (Bramham & Srebro 1989; Leonard et al. 1987). There must exist powerful endogenous mechanisms to inhibit new LTP induction during particular stages or events in SWS. The mechanism and physiological significance of this suppression is unclear. As detailed in the review, post-trial learning is associated with periods of enhanced REM sleep important for retention. Perhaps LTP induction is generally facilitated during these prolonged REM epochs. However, in the only work examining this issue so far, no changes were found in the magnitude or duration of LTP induced during enhanced REM sleep (Bramham et al. 1994). Thus, if *new potentiation* is facilitated during enhanced REM sleep, this potentiation is likely to be a function of REM-epoch duration and reactivation of network nodes.

Sleep is optimizing

Thomas L. Clarke

Institute for Simulation and Training, University of Central Florida, Orlando, FL 32826. **tclarke@ist.ucf.edu**

Abstract: It is suggested that Walker's consolidation-based enhancement of memory during REM sleep corresponds to the simulated annealing technique used for function optimization, and that robotic and AI design could benefit from inclusion of a deliberate REM-like memory optimization phase.

Scene 1: A robotics laboratory. A frustrated John Doe sits at a computer connected to a humanoid robocup soccer player. Across the room is his friend, Matthew Walker, who is visiting from the medical school.

JOHN: I don't get it. I've been able to train my robot to have all the basic skills, but it can't play worth a damn. Sometimes it makes the wrong decision, sometimes it just takes too long to choose. Either way, the other player just drives the ball by. . . .

MATTHEW: What sort of architecture do you use for memory and learning?

JOHN: It uses a neural network-based system. Recognition functions use Hebbian learning; actions are learned using the back-propagation algorithm.

MATTHEW: So it's basically patterned after the brain?

JOHN: Pretty much.

MATTHEW: So when does it sleep?

JOHN: Sleep!?

MATTHEW: Yes, sleep. My research shows that the initial acquisition and stabilization of memory do not depend on sleep, but a further process, consolidation-based enhancement, depends crucially on sleep. Since this is true of the brain, it should be true of your neural-based soccer player.

JOHN: Nonsense! Why would a humanoid robot using silicon circuits to simulate neural activity need sleep? Ridiculous! Sleep is a defect of meatware! I sleep as little as I can get away with. Avoiding those defects is one reason for our robotics research!

MATTHEW: Still, your neural networks are patterned after the functioning of the brain, so I don't think it unreasonable that they also share some of the defects as you say.

[Tom Clarke enters.]

TOM: What's the shouting about?

JOHN: Matthew here thinks that if I make my robots sleep, they will play soccer better.

TOM: Well, wouldn't they? Don't you play better after a sleep?

JOHN: But I'm – a human. My robots are silicon and algorithms . . .

TOM: I've read Matthew's latest paper in *BBS* and there's a lot to it that parallels techniques of mathematical optimization.

JOHN: [*Sputtering*] Matthew's a psychiatrist. How could psychiatric research contribute to humanoid robot design?

TOM: I think if you pay attention, quite a bit. You've heard of simulated annealing? Well, I happen to think that the REM portion of sleep is a kind of simulated annealing for the brain.

JOHN: I use simulated annealing as an optimization technique all the time. You add noise to a system being optimized, then slowly reduce the noise and the system settles into an optimum state. If the noise is reduced slowly enough, the system is guaranteed to fall into the global optimum and not get stuck in a less than optimal local state. Hence the term "annealing."

MATTHEW: Noise? Global optimum? What does this have to do with sleep?

TOM: Well I've always thought that dreaming – the random dis-

https://doi.org/10.1017/S0140525X05240025 Published online by Cambridge University Press

connected, illogical content of dreams – is a kind of mental noise that subjects your mind to several sessions of simulated annealing each night.

MATTHEW: Ah. I see what you are getting at. At a cognitive level the process of REM sleep and dreaming serves to stabilize and enhance memories through a process of optimization. Now that you mention it, it seems like rather an obvious idea. There must be an extensive literature on this.

TOM: Oddly, no. I've searched the web for the combination of "simulated annealing" and "REM sleep" and only get 34 hits. There doesn't seem to be much discussion of this except in a few theses. There is literature on optimization at the neural level via simulated annealing, but not at the level of memory content. I'll send you some references (e.g., Beckerman 1998; Cameron 1988; Hinton & Sejnowski 1986; Malhotra 2003; Robocup Soccer League 2004).

JOHN: So how is this going to help my robot win?

TOM: You will have to put your robot into the equivalent of sleep after training sessions. Don't just turn it off, but let it run with random inputs. I'm not sure if random inputs alone will be sufficient, I think there must be feedback from memory to sensory input. But I've never worked out a good architecture to simulate this.

MATTHEW: Don't forget to disable the outputs. Animals in REM sleep are effectively paralyzed.

JOHN: Very interesting. I think I'm getting some ideas of how I can rewire my robot to have REM-like states.

Scene 2: Several months later.

MATTHEW: Congratulations, John. Your robot won the cup. Don't you find that it performs so much better after a good night's sleep?

TOM: I could say something about how the REM sleep – er, simulated annealing – provides precomputed feasible solutions to the NP Hard problems of real world soccer play, but I think I'll quote Douglas Adams instead. "The endless dancing shapes and patterns would reach far deeper into our minds than we could manage by reason and logic. . . . Logic comes afterwards. It's how we retrace our steps. It's being wise after the event. Before the event you have to be very silly" (Adams 2002).

ACKNOWLEDGMENT

Thanks to Randall Shumaker for his support and challenging words.

Where is the classic interference theory for sleep and memory?

Anton Coenen

Department of Biological Psychology, Nijmegen Institute for Cognition and Information, Radboud University Nijmegen, 6500 HE Nijmegen, The Netherlands. **a.coenen@nici.ru.nl**

Abstract: Walker's target article proposes a refinement of the well known two-stage model of memory formation to explain the positive effects of sleep on consolidation. After a first stage in which a labile memory representation is formed, a further stabilisation of the memory trace takes place in the second stage, which is dependent on (REM) sleep. Walker has refined the latter stage into a stage in which a consolidation-based enhancement occurs. It is not completely clear what consolidation-based enhancement implies and how it can be dissociated from a stage for memory-stabilisation. A more serious consideration, however, is whether a second stage in memory consolidation that is solely dependent on sleep, is really necessary. The classical, passive, interference theory is able to explain adequately the findings related to the effects of sleep and memory, and can lead perhaps better to an understanding of the highly variable data in this field.

The idea that sleep is a favourable brain condition for memory consolidation has been going around since the days of Sir John Hughlings Jackson. In 1881 he predicted that sleep had a function in stabilising the vulnerable memory trace by making "permanent rearrangements during so-called dreamless sleep" (see Taylor 1958). Although not all researchers are able to show this phenomenon, generally, subjects remember more of a learning task when tested after a period of sleep than they do when tested after an equal period of wakefulness (Benson & Feinberg 1977; Grosvenor & Lack 1984). Historically, this positive effect of sleep on memory is interpreted in terms of the "interference theory," implying that during sleep there is less new learning interfering with memory storage than during wakefulness (Underwood 1966). In the period that the labile memory trace is stabilised, it is still vulnerable for interference. Hence, in the absence of interpolated learning, in a situation with less retroactive inhibition, the remembrance is better when sleep intervenes between original learning and recall than when this interval is filled with wakefulness during an ongoing uptake of new information. This is the classical, passive, way to explain the positive effect of sleep on memory.

An alternative interpretation, however, is that sleeping might play an active, rather than a passive, role in the stabilisation of the labile memory trace (Ekstrand 1967). An experimental upswing came when animal research suggested that REM sleep in particular, rather than sleep per se, could be important for the positive effect of sleep on memory. Two lines of research dominated the field: the study towards the effects on learning on sleep with the question whether REM sleep should be increased after a period of intensive learning, and the question of whether deprivation of REM sleep after learning should impair memory consolidation. Especially popular in animal research were the deprivation experiments, and a huge number of rats were put on inverted flowerpots to deprive them from REM sleep. Inconsistent and controversial results were found, which could be ascribed mainly to inappropriate deprivation techniques and the lack of adequate controls. Studies to establish that intensive learning was followed by an increase of REM sleep were also inconclusive. Even less positive evidence could be found in human research. These disappointing outcomes resulted in a decline of these kinds of studies in the late eighties and the beginning of the nineties.

However, the findings in the paper by Karni et al. (1994) heralded an experimental revival. These authors showed, in humans, an improvement in performance of a perceptual skill dependent upon the occurrence of REM sleep. Although the focus was again on an active role of REM sleep and not of sleep per se, Giuditta et al. (1995) gathered evidence for the "sequential hypothesis" on sleep function. According to this hypothesis, REM sleep, occurring later in the night, should serve the ultimate completion of consolidation, but the earlier occurring slow-wave sleep should function in the first step towards stabilisation of the memory trace. Since then, the role of non-REM sleep has also been explored, but up to now sound evidence for an active role of non-REM sleep, as well as REM sleep, in memory organisation is weak and contradictory (Vertes & Eastman 2003).

It is against this background that Walker's model must be regarded. He has adjusted the existing two-stage model of sleep and memory, which was formulated in 1970 by Bloch. Bloch suggested a discrete two-stage model in which the acquisition of new information is followed by a short period of processing, and true consolidation occurs later, during (REM) sleep. Walker's two-stage model is quite reminiscent of Bloch's model, but it is refined in the sense that he proposes a first stage in which a time-dependent consolidation-based stabilisation takes place, followed by a second stage in which a sleep-dependent consolidation-based enhancement takes place. Walker based this refinement on data obtained with an extended version of the Jenkins and Dallenbach (1924) paradigm, among others. In this paradigm, subjects were trained to perform a motor task skill, either in the morning or in the evening. Subjects trained in the morning were tested in the evening and retested the following morning. Whereas the first (evening) test showed no significant change in performance from the training session, the second (morning) test, which was admin-

https://doi.org/10.1017/S0140525X05240025 Published online by Cambridge University Press