[3]The Twisted Matrix: Dream, Simulation or Hybrid?[2] to appear in C. Grau (ed) Philosophical Essays on the Matrix (Oxford University Press,2004). Also to appear on the official Warner Bros Matrix website at:

http://whatisthematrix.warnerbros.com/rl_cmp/phi.html

The Twisted Matrix: Dream, Simulation or Hybrid[1]?

1. Ambivalence

"The Matrix is a computer-generated dreamworld built to keep us under control" Morpheus, early in The Matrix.

" In dreaming, you are not only out of control, you don't even know it…I was completely duped again and again the minute my pons, my amygdala, my perihippocampal cortex, my anterior cingulate, my visual association and parietal opercular cortices were revved up and my dorsolateral prefrontal cortex was muffled" " J. Allan Hobson, The Dream Drugstore, p.64

The Matrix is an exercise in ambivalence, and at the very heart of that ambivalence lies the Dream.

In our dreams, we are not in control. Real dreaming, unlike many popular philosophers' fictions, is an altered state, closely related to the states induced by chemical manipulations such as the use of (certain) medical or recreational drugs. The dreaming brain is not like the wakeful brain. Normal sensory input is blocked, attentional capacities are impaired or lost, memory is distorted, reasoning and logic are weakened, narratives run wild, self-reflection is dampened or destroyed, emotion and instinct are hyperstimulated, and forms of 'top-down' willed control and decision-making diluted and easily overwhelmed[2].

It may seem as if all of this is simply a direct effect of the blockage of sensory input, but this is not so. Instead, profound changes in neurochemical activity also occur, and these in turn compromise what J. Allan Hobson calls our 'critical self-awareness'. The result is that even though the goings-on in most dreams would cause us (were we awake) to suspect trickery or to question our sanity, in our dreams we simply accept them as normal, as real life! One way to keep people 'under control' as Morpheus put it, is to keep them (in this specific sense) *out of (self) control*. One of the issues I want to explore is: to what extent are the experiences of normal matrix-bound humans (matrixers) genuinely dream-like, where to be genuinely dream-like is (in part) to display this dampening of critical acuity? I shall call this normal, critically compromised dream-state 'uncritical dreaming'.

But there is, of course, another image of dreaming, and this is the image that has so far received most attention in these essays. This is the view of dreaming that links up with Descartes' famous 'malicious demon' thought experiments, and with standard philosophical discussions of what we can (and cannot) know. These explore the question, what can we really know on the basis of our actual experience? Dreaming, thus construed, is not really like dreaming at all: it is a state in which all the sensory experiences might be *just as they are in waking,* and our critical faculties as bright and active as ever. I shall call this kind of 'dreaming' '*industrial strength 'deception'*, so as to distinguish it from real dreaming.

The (apparent) deception practiced by the machines, we can now see, comes in two potential varieties. First, there is the industrial-strength version. Here all the sensory inputs that assail your brain are just as they would be were you living and moving in a world of persisting, external, independent people, cities, cars and objects, and you yourself are as alert and critical as ever. Second, there is the 'uncritical dreaming' version, where the flow of sensory images and data is actually not all as it should be[3]. Things may morph and change, scenes shift, identities alter. Here, weirdness may be rampant

yet generally unremarked, since your abilities to judge that all is not as it should be are fatally impaired, due to the critical-dampening effects of the neuro-chemical alterations distinctive of sleeping or drugging. The genius of the Matrix (I shall argue) is its ability to balance, both thematically and cinematically, on a knife-edge between these two versions of events.

The industrial strength version invites an important response, ably advanced by David Chalmers in his paper on this site. According to this response, industrial strength 'deception' is not really deception at all! Matrixers subject to this kind of manipulation really do inhabit (so the argument goes) pretty much the world they believe themselves to inhabit. 'Industrial strength matrixers', as I shall call them, really do live in cities and roam a planet much like earth. Later in this essay I will further defend this view, arguing that (still assuming industrial strength 'deception') matrix-based human intelligences would count as being as fully and richly embodied as you and I. Despite those (other) bodies we see suspended in the machine-feeding womb, industrial strength matrixers really do use their head and eyes to scan the visual scene and their legs to move around. According to this view, the body in the Matrix is not a dreamt body, at least not in any ordinary sense of dreaming. It is a real body, realized in the non-standard medium of bits of information. With the point made for the visible body, the parallel result for the wider world (of cities, cars, sky and dust) may become a little easier to swallow.

We can then return to the nature of dreaming. For part of the ambivalence at the heart of the Matrix is, I think, an abiding ambivalence about the nature of dreaming itself. Real dreaming, to repeat, involves profound changes to the cognitive system deployed in normal wakefulness, changes that systematically deprive us of much of our normal critical acuity. In real dreaming, activity in the dorsolateral prefrontal cortex is compromised: 'muffled' in that opening quote from J. Allan Hobson, Director of the Neurophysiology Laboratory at Harvard medical school.

Dorsolateral prefrontal cortex is (among other things) the 'executive brain' that helps us organize our thinking, critically assess our own gut responses and maintain at least a modicum of top-down control. The kind of state that Morpheus calls ' a computer-generated dreamworld' hovers uneasily between such true (profoundly cognitive acuity diminishing) dream states and something much closer to unknowing-yet-fully-awake participation in a form of multi-agent immersive virtual reality: an 'interactive virtual environment' as Morpheus also puts it.

The moral ambiguity that permeates the Matrix is rooted in this same balancing act between real dreaming and a multi-agent immersive virtual reality in which the neural states of the average matrix dwellers are (courtesy of the machines) really identical to those of awake, active humans. Drift towards the former reading, according to which the matrix-dwellers' cognitive states are neurologically akin to those critically diminished states distinctive of real human dreaming, and everything-*ABSOLUTELY EVERYTHING*- changes. Not only does the previous argument for true embodiment in the matrix fail, but the moral status of the machines' experiment is immediately and radically transformed. Instead of seeing the machines as maintaining a kind of innocent immersive virtual reality, wherein human embodiment and human intelligence is, in every way that matters, everything it always was, the ploy of the machines becomes more like that of pimps who keep their call-girls hooked on heroin. Thus cognitively diminished, the girls do not question their state, and are not able to plot a rebellion or plan an escape.

Drug-forcing pimps or master immersive reality programmers? It is by maintaining a studied ambivalence between these two readings that (I claim) the Matrix film sequence gains much of its power, its beauty and its profound ability to puzzle.

2. Real Dreaming (Asleep in the Matrix)

Neo believes, under the influence of Morpheus and of his own experiments, that matrix-bound humanity is in the grip of a

delusional dream. In the dream, the apparently sensed world seems real. It seems like a place where the body moves, where the eyes roam, where flesh meets flesh and sometimes lead and steel. But this, Neo comes to believe, is simply a dream, a device to keep us quiet while the machines patiently suck energy from the preserved slumbering bio-mass. To prove that it is a dream, Neo learns to subvert its logic: he becomes able to turn back bullets, to fly, to defeat agents, and all this because he wills it to be so.

In one important sense, however, this is no ordinary dream. For in this dream (as Neo understands it) there is real contact between multiple intelligences. When Neo speaks to Cypher, he does not speak merely to a construct of his own sleep-bound imagination, but to another sentient being, with genuinely independent memories, hopes, fears and skills. Moreover, these multiple intelligences can communally build persisting structures in their world. They can build worlds to live and act in. If Neo places a cup on a table in a certain room, it will be seen by Cypher when Cypher enters that room (unless someone else moves it away first). A simple bio-mass of individually dreaming humans could never achieve and maintain this kind of interpersonal and structural continuity and integration. So the machines must be doing something (a whole lot in fact) to keep things in line.

Nonetheless, there is clearly something dream-like going on, for only in a world not fully bound by the laws of (even simulated) earthly physics could Neo fly, or turn back those bullets. In a normal Virtual Reality simulation, you cannot bend the rules just by willing it. By the same token, video-gaming would be a whole different sport were the underlying code directly susceptible to the will of the players! Moreover, the movie is chock-full of images that morph and shift in ways not seen in waking life, again suggesting that this is not a perfect simulation of earthly physics, but something less stable, lacking in firewalls, and prone to direct subversion by the minds of enlightened matrixers. The average matrixer, of course, does not subvert, remains unenlightened, and seems to be almost sleepwalking through a mundane, yet not unpleasant, life. This

contrast is perhaps most striking in the scene at the end of the first movie, where Neo, increasingly enlightened, steps out of the 'phone box to see hordes of ignorant matrixers moving in trance-like, clockwork fashion, their images somewhat out of focus in a classic depth of field manipulation, while that of Neo is crystal-clear, alert, and bemused by their unquestioning, anesthetized progress through the world.

To try to clarify just what we are dealing with, it will help to first take a hard look at normal human dreaming. Then we can begin to plot some differences and to explore the space of options. A word of warning though. Familiar as they are, sleep and dreaming are complex, ill-understood phenomena. The sketch that follows is widely accepted and heavily rooted in the best contemporary neuroscience and psychopharmacology. But it is not written in stone, and much remains unclear.

The three dominant states for the human brain are waking, REM (Rapid Eye Movement) Sleep, and non-REM (NREM) sleep. Each state has clear physiological, pharmacological and experiential correlates.

In waking, we can occupy many states, from eyes-closed imagistic musing to eyes-open, alert engagement with a potentially threatening environment. The *option* of alertness and full critical engagement is, however, typically present, even if we are engaged momentarily in detached daydreaming.

In REM sleep our dreams (at least as evidenced by subsequent report) are vivid, but their logic is weak. Here is a typical enough report:

> I was at a conference and trying to get breakfast but the food and the people in line kept changing. My legs didn't work properly and I found it a great effort to hold my tray up. Then I realized why. My body was rotting away and liquid was oozing from it. I thought I might be completely rotted before

> the end of the day, but I thought I should still get some coffee if I still had the strength.
>
> Excerpt quoted in Blackmore (2004) p.340

Here is another description, this time from Helena Bonham-Carter, while she was expecting a baby with movie director Tim Burton:

> "I dreamed I gave birth to a frozen chicken. In my dream, I was very pleased with a frozen chicken"
>
> Quoted by Lynn Hirschberg (2003)

In NREM sleep, if we dream at all, the dreams (again, as evidenced by waking report) are more like faint and mundane thoughts or fuzzy rememberings.

All these states (waking, REM-sleep, NREM sleep) are correlated with specific patterns of neuro-chemical activity. A useful tool for displaying the pattern is Hobson's AIM model. Hobson is a leading sleep researcher interested in the relation between waking, sleeping and the kinds of altered state experienced during psychosis and drug-use. The AIM model characterizes different states as points in a three dimensional space, whose axes are:

1. Activation Energy
2. Input Source
3. Mode

Normal wakefulness is characterized by high activation (as measured by EEG for example) corresponding to fairly intense experience, external input sources (the brain is receiving and processing a rich stream of sensory signals from the world, rather than being shut down and largely re-cycling its own activity), and a distinctive mode. Mode here names a balance between brain chemicals, especially amines and cholines. Amines are neurotransmitters such as noradrenaline and serotonin, whose action is known to be essential

for normal waking consciousness (they are essential to the processes that enable us direct attention, reason things through, and decide to act). When these are shut off, and other neurotransmitters (cholines, such as acetylcholine) dominate, we experience delusions and hallucinations (if we are awake) and vivid, uncritical dreaming (if we are asleep). In this way it is the amine/choline balance that determines how signals and information (whether externally or internally generated) will be dealt with and processed. When the balance (as in waking) favors the amine-based (aminergic) system, we are rational, alert to our surroundings, easily able to direct our own actions and to rapidly and critically appraise our situation. When the balance favors the cholinergic system, our focus shifts inwards, emotion and analogical reasoning begin to dominate, and critical control and judgment wane. In REM sleep, the aminergic systems are totally deactivated and the cholinergic hyperactive. This is an extremely altered cognitive state. Only extreme forms of psychosis or serious medical or recreational drug use can induce this kind of state in non-sleeping humans. In normal waking states, the ratio of aminergic to cholinergic activity varies across a large continuum. In non-REM sleep, all the systems (aminergic and cholinergic alike) are dampened and (mostly) inactive[4].

This is not to suggest (far from it) that the best state for a human mind would be one of almost-complete aminergic dominance. Instead, the power, subtlety and beauty of wakeful human intelligence seems to have much to do with the precise details of the ever-shifting balance between the two systems. But in normal waking the mode (defined as the ratio between the activity of the two systems) leans towards the aminergic. Whereas in REM sleep, with acetycholine dominating, experience becomes increasingly dissociative, displaying "amnesia, hallucinations,bizarre mentation, anxiety,and loss of volition control" (Hobson, p. 91). All this, we now know, is matched by a shift in regional blood flow from (in waking) dorsolateral prefrontal cortex to (in REM sleep) subcortical limbic structures (some of which are mentioned in the opening quote). Here too, the psychological and the physiological march (unsurprisingly, surely) in step, with dorsolateral prefrontal cortex implicated in

analytic reason, inhibition and executive control, and the limbic structures dominating for emotion, instinct and association[5].

Bottom line: The kind of sleep in which we experience vivid dreaming is, typically at least, a state in which aminergic systems essential to critical reason are deactivated. This state is a far cry from normal wakefulness. The reason we (often) don't know we are dreaming is not because the dream simulates waking reality (the immersive Virtual Reality option) but because we are cognitively diminished in ways that block voluntary attention and critical engagement and that promote a kind of face-value acceptance. In REM sleep we are, in a real sense, drugged witless by our own brains. And the cure, as Hobson and Neo would probably both agree, is simple: it is called waking up!

Is it possible that the machines are electrically or chemically altering the states of the brains of their human power-cells, so as to partially or totally deactivate the aminergic system and the dorsolateral prefrontal cortex? To compromise these would be to compromise the matrixers capacities for critical engagement and analysis in a very profound way. Certainly, we sometimes see images of the humans suspended within the machine's grid in what seem to be advanced stages of REM sleep. If this is what is going on, those human brains are, in a fairly precise sense, permanently drugged. What Neo has achieved is then well compared (as other essayists have noted) to the state known as lucid dreaming. In lucid dreaming, a very few subjects are able to become aware that they are dreaming, without actually awakening. They may even be able to take control of the dream itself, forcing in directions previously requested by an experimenter or simply for their own enjoyment.

The full AIM (Activation/Input/Mode) profile of the lucid dreamer is still unclear, but Hobson speculatively suggests that one key may be a kind of prior-to-sleep priming in which the lucid dreamer prepares to recognize the delusional REM dream state as it develops. REM sleep actually enhances priming effects, in which (for example) prior exposure to one word makes recall of another quicker or more

likely. Pre-sleep preparations may, Hobson suggests, prime[6] a more complex association between the signature (fuzzy, delusional) character of REM sleep and the realization that you are indeed sleeping. A kind of positive feedback cycle can then take root, so that the primed realization of dreaming is fed by each encounter with new unlikely or delusional elements. At this point the mode balance may shift a little, so that some dorsolateral prefrontal involvement becomes possible. Not too much, or the vivid dream state would be lost. Not too little, or control would be lost and dreamt delusions would (as usual) be taken as real. This is a testable hypothesis, using PET neuroimaging to reveal fluctuations in neural regional blood flow during lucid dreaming. But it has not yet been tested.

Might Neo have been somehow primed, before the machine-induced sleep, in this kind of way? Or perhaps, by some kind of neuro-chemical accident, Neo, Morpheus and a few others are simply immune to the suppression of their aminergic systems? The Matrix as uncritical dream, and Neo as lucid dreamer. Such views at first seem to make good sense. But only if we assume that the state of the typical matrixer is indeed a real dream state. Such a reading, it should now be clear, is problematic for a number of reasons. For one thing, there is strong narrative continuity (major plot items do not just come and go without explanation). There is also good interpersonal agreement (what Neo does, Trinity sees, and so on). But most obviously of all, the typical matrixer simply does not seem to be unusually uncritical. For sure, they are not constantly reflecting on their life and analyzing its worth. But, and this is crucial, within the movie *truly unusual happenings are indeed usually spotted as such*. The security guards are as amazed as we are when Neo and Trinity, armed to the teeth, burst through the gates. When Neo stops bullets, or flies, the typical matrixer is astounded. This is not the reaction of an uncritical dreamer.

Colin McGinn, in his paper on this site, is at pains to highlight the many layers of dream-like quality that permeate the movie, and to depict Neo as a lucid dreamer. But in the chemically and cognitively precise sense just outlined, life in the Matrix is simply *not like life in a*

dream, and Neo is nothing like a lucid dreamer. The typical matrixer does not display the full cognitive signature of uncritical dreaming, and it therefore seems unlikely that the machines are actively maintaining the brains of their human power cells in the standard (aminergic-off, hyper-cholinergic) REM-sleep mode. Whatever else the machines may be doing, they do not seem to be acting merely like drug-pushing pimps. It is not that Neo is special because he has his wits *available*, but rather, because he is beginning to really *use them* and to question (guided by Morpheus) so much that he previously took for granted.

In section 4, I'll consider a midway option in which elements of real dreaming combine with elements of immersive simulation. For the moment, however, we are again face to face with the productive tension at the very heart of the Matrix. For McGinn is surely right, despite all this, to remark the genuinely dream-like, shifting and sometimes disconnected visual and dramatic qualities that repeatedly surface. The point I now want to make, and will return to at the end, is that this tension is distinctive of our normal waking life as well! For our waking experience is itself the product of that constantly shifting balance between the aminergic and cholinergic systems, a balance that alters and evolves minute-by-minute during any normal day. As Hobson puts it "it is as if we are designed to be rational (but cool) and irrational (but hot) by turns" (p.97). The Matrix film sequence, by repeatedly shifting between narrative and visual modes proper to critically engaged waking reality and then to delusional sleep, is able to explore the whole spectrum of positions in Hobson's AIM three-space. In so doing it gives us insight into the inherently unstable nature of human awareness itself.

3. Industrial Strength 'Deception' (Awake in the Matrix).

In real dreaming, we often believe ourselves to be in places we are not, doing things we are not. On the face of it, life in the Matrix is an endless dream. According to this very natural view, most humans in the matrix are doubly deceived. They are deceived about their physical surroundings, believing themselves to live in cities and to

roam the Earth, when "in fact" they are suspended in an energy-sucking machine-made womb. They are also (still on the face of it) deceived about their own bodies, believing themselves to be moving their limbs, flexing their muscles, and scanning the scene with their eyes and heads, when "in fact" their bodies are still, their heads are fixed, their eyes closed. Such, at least, is the dominant interpretation of the true state of Matrix-based humanity.

But there is an alternative. To bring it into view, we need to imagine that the machines are not simply guiding (and somehow, rather puzzlingly, making intersubjectively coherent) the real dreamings of the slumbering bio-mass. Instead, suppose that they have created a detailed simulation of the physics and structure of the normal human world, and that they are closely and continuously monitoring the neural activity in the brains of their human power cells. These brains are fed signals that correspond exactly to the ones they would receive were they awake and acting in the world, and the virtual world updated in ways that conform to those actions. Each day, these brains would go through just the same chemical cycles as normal human brains, moving systematically through Hobson's AIM space: from awake (aminergic systems highly active, so-called Sleep-On neurons inactive, critically alert) to N-REM sleep (forebrain Sleep-On neurons fire[7], aminergic and cholinergic systems inactive, more or less dreamless) to REM sleep (brain-stem REM-Sleep-On neurons fire, aminergic systems shut off, cholinergic systems highly active, vivid dreaming can occur, and critical and executive faculties extremely dulled).

Do matrixers really dream and awaken, or do they only dream they are dreaming and awakening? In the industrial strength version, it seems more accurate to say they really dream and awaken. The states of their brains[8] in matrix-sleep, on this version, differ from the states of their brains during matrix-waking in just the same ways as ours do between wake and sleep. By contrast if, in a real dream, we dream we fall asleep, there is no such neuro-chemical shift. So too, if (in a real dream) we dream we are awakening, that does not itself activate the dormant aminergic systems that would actually awaken us and

restore our critical acuity. Once again, the differences between the two versions are striking.

The machines, on the industrial strength version I am now pursuing, create a detailed immersive virtual reality that is sensitive to the actions of all the users, and they ensure that my actions encounter obstacles and generate systematic sensory feedback exactly as they do in the normal world. That would include, for instance, generating the whole panoply of signals distinctive of muscle tiredness after an arduous rock climb, and those distinctive of hunger, and of the satiation of hunger by food (recall Cypher's infamous steak), and so on. (Notice that on the real-dream version none of this is necessary: instead, Cypher only needs to believe he has enjoyed a steak. These are different states).

On this industrial strength version, I claim, matrixers are genuinely embodied and are able to eat, act, wake, sleep and dream in a world that is every bit as real as they imagine (though its deep physics is, as David Chalmers in his essay on this site points out, not quite as they think).

To make this stick (or even to make it begin to seem plausible) we can start by looking a bit harder at what it means to have a body, and at what having a body does for a mind like ours. With this understanding in place, it should be possible to see how the human intelligences in the industrial strength Matrix could be embodied intelligences through and through. They would be embodied not in virtue of those organic shells feeding the machines, but in virtue of the crucial role of eye-movements, head-movements, and limb-movements in altering the inputs to their brains and nervous system, and in virtue of the way the world presents itself both as a resource for action and as a source of limits on action.

Consider a fairly typical example of embodied action: solving a jigsaw puzzle. First, I arrange the pieces on the table in front of me, perhaps placing the predominantly red pieces (bits of an image of a rocket) in one pile, the green (bits of an image of the jungle) in

another, the pieces with one straight edge in a third, and so on. To solve the puzzle, I then combine a variety of tactics. One tactic involves repeatedly looking from the pieces to the half-completed puzzle. During these periods, my eyes make repeated movements (known as saccades) that bring different aspects of the scene into foveal view. (Human vision depends heavily on moving a small high resolution are, known as the fovea, around the scene, so as to retrieve information as and when needed). Another tactic involves picking pieces up and trying them out, to see if they really fit in certain locations. Yet another tactic involves reasoning about the shape of the missing pieces: there must be one there with a wiggly edge and half red and half green. So I again scan the scene, with this image in mind, hoping to find such a piece.

In this kind of problem solving, the body and world play important roles[9]. Instead of creating a full image of the half completed puzzle in my mind's eye, and then looking over the pieces, I repeatedly shift gaze from the real puzzle to the pieces. This saves my brain from encoding (no doubt badly) all that complex detail. And when I have isolated a candidate piece, I make the final decision by actually trying it out for fit. At this point, the world may fight back, refusing to allow a piece to fit, however much I want it to. This is an example of what Dreyfus and Dreyfus[10] call running up against a boundary condition in our attempts to cope with the world (I'll return to Neo's, and the agents', abilities to bend such rules in a moment). Moreover, notice that I started out by organizing the workplace in a way that then helps reduce my problem-solving load (making the various piles). All this is what Cognitive Scientist David Kirsh calls 'the intelligent use of space'.

Now all of these ploys and strategies are available, quite straightforwardly, to the average industrial strength matrixer. She can use the external world (as constituted by the machines detailed and action-responsive simulation) to reduce the problem-solving load for her own cognitive processing. She can use body and eye movements so as to leave lots of important detail 'in the world', retrieving information as and when needed for a specific action. She

can intervene so as to organize the workspace in ways that then persist, independently of further cognitive efforts on her part, and that (for example) save her searching for red pieces among the green by keeping all the red ones in a pile, and so on.

Consider, finally, the sense of presence, of where we are. There is a wonderful thought-experiment due to the philosopher Daniel Dennett. Dennett tells the story of a U.S. citizen who agrees to participate in a secret experiment. The citizen is Dennett himself, and in the experiment Dennett's brain is removed, kept alive in a tank of nutrients, and equipped with a multitude of radio links by means of which to execute all its normal bodily control functions. Dennett's body (which is to be used to explore a dangerous area) is equipped with receivers and transmitters, so that it can use its in-built sensors (eyes, ears, etc.) to relay information back to Dennett's brain. As the technicians in the story put it[11];

> "Think of it … as a mere *stretching* of the nerves. If your brain were just moved over an *inch* in your skull, that would not alter or impair your mind. We're simply going to make the nerves indefinitely elastic by splicing radio links into them."

There is a way of thinking of what the machines have done (in the industrial strength version) that is a lot like this. But instead of using the brain to control a standard body exploring distant and dangerous parts of the standard world, they have 'stretched the nerves' all the way into a fully immersive virtual reality, allowing the brain to control a kind of body-double avatar[12].

With his brain safely excised and re-located, and the radio links established, Dennett awakes. He sees the nurse, who leads him to the room where his brain is being kept. The experience that ensues is puzzling. There is Dennett, standing up, staring at his own brain. Or is he? Perhaps, he muses, the proper thought is "here I am, suspended in a bubbly fluid, being stared at by my own eyes …" Try as he may, Dennett cannot seem to place himself *in the tank*. It

continues to seem as if he is outside the tank, looking in. Dennett's point-of-view, as he moves, seems securely fixed outside the tank[13].

Where is Dennett? Is he *really* in the tank of nutrient, *really* outside the tank looking i*n,* or really no-place at all (or both places at once)? Such questions need have no clear-cut answers. But what does seem clear is that human location should not be taken to be a function of facts about the location of the brain. After all, wherever 'you' are, it surely isn't inside the top of your own head! Human presence, instead, is better understood as dependent upon our capacities for *dense, closed loop control.* By that I mean control (of some kind of body) such that as the body moves, the brain receives rich and detailed feedback. It is this kind of feedback cycle and closed loop control that supports skilful action. Skillful action then enables us, as the computer scientist Paul Dourish puts it, to engage in 'inhabited interactions'.[14] The difference between an inhabited and a non-inhabited interaction is just the difference between, for example, having to carefully plan, monitor and execute a reach for a coffee cup, and 'just reaching', as we expert coffee-cup grabbers do. Inhabiting the body, we are able to fluently use movement and action as parts of our own problem-solving routines.

Putting all this together, I can now offer a proposal for how to think about the body and the world:

1. The body is a controllable and inhabitable resource.

2. It is located (or its parts are located) in one or more co-ordinate spaces, and its actions (or the actions of its parts) evolve in time.

3. Experiences of dense, closed loop control involving this resource yield a robust sense of presence, and of 'inhabiting' the body.

4. It is a resource capable (via these 'inhabited interactions') of being skillfully used to transform a problem space, and to exploit properties and features of the world.

5. The world is the place where such embodied actions encounter boundary conditions and are forced to conform.

And finally:

6. Real space is wherever perception and embodied action occur.

Space, body and world are in this way all inter-defined. According to this formula, industrial strength deception is a contradiction in terms[15]. The world of the industrial strength matrixer is a real world. It acts as a boundary condition for skilled action, and it is populated by real bodies whose inhabited interactions play the very same problem-transforming roles as our own.

4. The Hybrid Matrix

There is a clear problem for any full Industrial Strength reading of the Matrix. Such readings make it hard to understand how Neo can (as Morpheus puts it) 'bend the rules'. If your brain was getting its inputs from, and feeding its outputs to, this kind of immersive virtual reality set-up, there should be no room to break the laws of physics just by willing it so. Worse still, rampant rule-bending seems to deprive the world of its ability to function as a boundary condition, and this would undermine my attempts to argue for genuine embodiment and presence inside the matrix. It is as if you really could make the jigsaw puzzle piece fit just by wishing it so, in which case the world is surely not playing the cognition-enhancing role I described.

By contrast, the real dreaming model makes rule-bending easy to understand, perhaps along the lines of lucid dreaming mentioned in section 2. But this model fails to account for the kinds of preserved critical acuity that we *do* see in the matrix: the fact that Neo's flying is seen by everyone as something remarkable, as proof of superhuman prowess, and is not simply accommodated courtesy of dampened critical and executive processing.

Certainly, with a bit of ingenuity, we could probably come up with patches for each of these models. For example, someone who favors the industrial strength model could depict Neo as a kind of 'psychological hacker' whose willpower somehow alters the underlying code, bypassing rich and detailed restraints that really do apply to the average person. Dreyfus and Dreyfus (in their paper on this site) offer a version of this, in which Neo's belief that a spoon is bending forces the system to conform. Similarly, someone who favors the 'real dreaming' option might argue that the machines somehow link all the sleepers into a single web, maintaining the standard (critical acuity diminishing) REM sleep chemistry but thus forcing the sleepers to dream a single dream.

Between these two extremes, however, lies some of the most interesting ground of all, the ground of what I am calling the 'Hybrid Matrix'. On the hybrid model, the matrixers' world is indeed a kind of immersive virtual reality, but *one that has been rather lazily programmed*. Instead of recreating a deep and fully constraining physics, we can imagine that the machines' simulation is patchy, and depends on a lot of quick and dirty tricks[16]. For example, instead of running a complete, continuous full simulation of all locations and objects, they may only have programmed detail to unfold and update where one or more matrixers happen to look (a standard move, in fact, in ordinary Virtual Reality simulations). They may also not have bothered much about fine-grained continuity. Perhaps minor objects can come and go quite freely. On this model, the machines (like cinematographers!) just make sure that nothing major, and in anyone's focal attention, behaves strangely. We humans are surprisingly oblivious to unexpected scene-changes and non-central

continuity errors anyway, as a large recent literature on "change blindness" clearly demonstrates[17]. It is almost as if we are built to live in a lazily programmed world! In addition, lazy programming, as we all know, is also a royal invitation to hacking. The lazily programmed Matrix is at once an eminently hackable Matrix, as Neo, Trinity and Morpheus know so well.

To further support lazy programming, we might even imagine that the machines have tilted the chemical balances in the hosts' brains just enough to make them even more unlikely to attend to much fine detail, or to pursue rigorous and sustained environmental examinations. Such tilting would not yield genuine chemically-sleeping brains, but they would not be fully alert brains either. All this would just underline the guiding politics of the Matrix, which is a politics of awakening from dull, unthinking conformity and thus of escape from invisible, corrosive, but surprisingly violable constraints.

Would a lazily programmed Matrix still count as a real world, according to the argument of the previous section? I think it would, just so long as the lazy program was stable enough, and powerful enough, to impose some boundary conditions and cognitively exploitable order in most normal circumstances. It would be, as David Chalmers has suggested, just like a real world with a rather lazy God and a surprisingly patchy ultimate physics!

Which model is correct? Is the Matrix a consensual dream-world, a multi-agent immersive virtual reality simulation or a lazily programmed hybrid? I do not think we should seek an answer to this question. The power, beauty, and philosophical depth of the Matrix all derive from its ability to show us our world under many guises. At times, it shows us our world as a genuine dream-world, dominated by (strangely consensual) delusions. Such a world is ultimately unconstraining but hard to fathom, and maximally resistant to critical attention. At other times, it shows us our world as boundary condition, as a hard-edged arena for rational thought and embodied action. At still other times, it shows us a hybrid world, poised unsteadily between the two extremes. Just as normal

wakefulness comes in many grades, characterized by the shifting balance between the aminergic and cholinergic systems, so the movie constantly shifts from one state to another, morphing between delusional dream and immersive virtual reality. By flipping between, and mixing among, these two perspectives it finally reveals our own world as a potent cocktail of genuine boundary conditions, delusions, and mutual constructions.

This is where we need to end. Ours is a world in which much of what we ordinarily think constrains us is not truly binding. But this freedom does not reveal our world as a simple dream world, but rather as a real world, rich with the possibility of renewal and reconfiguration. By refusing to conform to any single interpretation, the narrative, structure and filmic texture of the Matrix sequence all encode the same message: take nothing for granted, don't write yourself in stone, just wake up.

References

Blackmore, Susan (2004) *Consciousness: An Introduction* (Oxford University Press, NY)

Chalmers, David (2003) "The Matrix as Metaphysics" (on this site).

Clark, Andy. (1997). *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.

Clark, Andy (2003) *Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence* (Oxford University Press, NY)

Dennettt, D. (1981) "Where Am I?" in Dennett, D. . *Brainstorms*. Sussex: Harvester Press, 1981

Dourish, P (2001) *Where The Action Is: The Foundations of Embodied Interaction* (MIT Press, Cambridge, MA)

Dreyfus, H and Dreyfus, S (2003) "The Brave New World of the Matrix" (on this site)

Hirschberg, L (2003) "Drawn to Narrative" *New York Times Magazine*, Nov 9, 2003.

Hobson, J. Allan (2001) *The Dream Drugstore: Chemically Altered States of Consciousness* (MIT Press, Cambridge, MA)

Kirsh, D. (1995). *The Intelligent Use of Space. Artificial Intelligence.* 73: 31-68

Roberts, A, Robbins, T and Weiskrantz, L (1998) *The Prefrontal Cortex: Executive and Cognitive Functions* (Oxford University Press, 1998)

Siegel, J (2003) "Why we sleep" *Scientific American* November 2003 p. 92-97

Simons, DJ, & Levin, DT (1997). "Change blindness". *Trends in Cognitive Science*, 1, 261-267

---

[1] Thanks to David Chalmers, Christopher Grau, and Tyler Waite for helpful comments and suggestions.

[2] This list is based on the Table 3.1 "Physiological Basis of Differences between Waking and Dreaming' in J. Allan Hobson (2001) p.57

[3] I do not mean to suggest that the distinction between real dreaming and industrial strength 'deception' will help us evade skeptical

uncertainty. Rather, my concern will be to explore how making this distinction affects our view of what the machines are doing, and of the moral status of their manipulations.

[4] The above paragraph condenses, and slightly oversimplifies, the views found in Hobson (2001), Blackmore (2004). See also Roberts, Robbins and Weiskrantz (1998) and Siegel (2003).

[5] Once more, this is in no way to privilege one system over another for effective reason and intelligence, which demonstrably depends on the proper temporally evolving balance of the two, but just to note the different contributions made by each.

[6] One thing may be said to prime another when exposure to the first makes the occurrence of the second more likely. The term is mostly used in psychological studies in which unconsciously perceived stimuli make subsequent conscious choices faster or incline them in one direction rather than another.

[7] For more on this, see Siegel (2003).

[8] There is an interesting question here concerning what we should take to be the brain of an industrial strength matrixer. Is it the brain of the human in the machines power grid? Or is it the brain that is

specified in the immersive reality simulation itself (e.g. the one we would see were we to observe matrix-bound neurosurgery!)? One way to proceed is to think of the brain in the grid as the ultimate (and unexpected) physical realization of the brain in the simulation (for a defense of this line, see Chalmers paper on this site). This is ok if we really are dealing with the full strength version, since any practice of neurosurgery inside the matrix would need to be synched either to real changes in the brains in the grid, or (at the very least) to changes in the input and output signals that correspond to those that such surgery would have induced.

[9] For more on this, and many examples from development and robotics, see Clark (1997)

[10] See their paper on this site. Their view, like my own and that of David Chalmers, is that there is an important sense in which industrial strength matrixers really are embodied and coping with a wider world. Dreyfus and Dreyfus, however, pursue an interesting final twist concerning our ability to (as they put it) 'open up new worlds'. I highly recommend reading both these essays these essays in full.

[11] From D. Dennettt (1981)

[12] Technologies of this stripe are by no means inconceivable. Just a year or two ago, the neuroscientist Miguel Nicolelis conducted a study to understand the way signals from cerebral cortex control the motions of a monkey's limbs. An Owl Monkey had 96 wires implanted into its frontal cortex, feeding signals into a computer. As the monkey's brain sent signals to move the monkey's limbs, this "neural wire-tap" was used to gather data about the correlations between patterns of neural signal and specific motions. The correlations were not simple. But the patterns, though buried, were there in the signals. Once these mappings were known, the computer could then predict the intended movements directly from the neural activity. The computer could then use the neural signal to specify the movements of a distant robot arm (an electro-mechanical prosthesis in an MIT laboratory six hundred miles distant). The system used a haptic interface, part of a multi-sensory Virtual Reality system used to touch, feel and manipulate computer generated objects. The machines, we can imagine, have simply taken this technology to the natural limit, developing advanced neural wiretaps allow the matrixers to explore and act upon a common virtual world. For a large-scale exploration of these new technologies, and of what it ultimately means to be human, see Clark (2003)

[13] The feeling shifts, however, when Dennett's body is subsequently trapped by a rockslide, entombed far beneath the earth's surface. At first, Dennett feels tapped beneath the surface. But then the radio links themselves begin to give way, rendering him blind, deaf and incapable of feeling. The shift in point-of-view was immediate;

> Whereas an instant before I had been buried alive in Oklahoma, now I was disembodied in Houston … as the last radio signal between Tulsa and Houston died away, had I not changed location from Tulsa to Houston at the speed of light?
> Dennett (1981) p.317

[14] See Chapter 4 of Dourish (2001). This idea is drawn from phenomenology, and has roots in the work of Heidegger, Merleau-Ponty and others.

[15] Christopher Grau (personal communication) asks whether such a view is too strong, amounting in effect to a simple redefinition of the real, rather than a substantial account. But my claim is not that any consistently imagined world counts as real (a claim that David Chalmers may actually be closer to making). That's why I stress the importance of genuine (not merely imagined) boundary conditions, and of the agent actually being able to offload computational work

onto the environment. These are not mere matters of what the agent thinks they are doing, but matters of fact. My line, roughly, is that to perceive a real world is to perceive a genuinely useable cognitive resource. Even a lazily programmed matrix (see section 4) might provide for that, for example by allowing people to really find out that one jigsaw puzzle piece (simulated) doesn't fit into one space (simulated), and by allowing the use of intelligent saccades directed at a stable scene (kept stable by real-world physics or good simulation) as a problem solving tool. As the amount of lazy programming (and thus instability and unreliability) increases, this 'signature of the real' gets eroded. All this is the case whether or not the agents actually notice anything. It is not, on my account, a matter of seeming to use a stable external world as a cognitive resource, but of actually doing so.

[16] Special thanks to David Chalmers for encouraging me to expand on this possibility.

[17] For a review, see Simons and Levin (1997)