

Qualia Surfing

by Richard Loosemore

What will happen to human civilization when we can go *qualia surfing*—collecting new experiences by transferring our consciousness back and forth between different substrates on a whim? As you wake up on a typical morning, your choice of activities for the day might include: becoming a tiger and going off to the jungle for some animal sex; changing into a body that can swim in the atmosphere of Jupiter; floating naked in interplanetary space in a body adapted for hard vacuum, while you stargaze through telescopic eyes that can see beyond the visible spectrum; swimming in a liquid-metal pool at noon on Mercury; porting yourself into a six-inch-tall body so you can go on a *Land of the Giants* expedition; or perhaps visiting a specially designed city where you check your memories at the door and live in a replica of, say, Restoration London or Classical Athens, all the while oblivious to the fact that you are in a simulation.

These activities are exotic enough in themselves, but the aspect I want to focus on here is their long term implications: how they might seep into every nook and cranny of our culture, redefining what it means to be human. My goal is to produce a quick survey of the different types and degrees of qualia surfing, so we can begin to see the larger picture and appreciate the role that this type of activity will play in our future.

Today, even the most utopian visions of the future contain a worm at their heart: the inevitable decline of humanity into a state of boredom and stagnation. But that depressing picture is partly caused by our failure to understand that we see the world from inside the bodies and minds that we currently possess, without realizing that when we become free to play with the design of our own minds and bodies, we open up almost infinite numbers of new worlds to explore. With luck, this short chapter will begin to reveal a more optimistic and expansive vision of the future.

The basic version of the qualia surfing idea is that at some point in the future people will freely transfer their consciousness into different substrates (where a *substrate* is either a biological brain or a computer), or modify their existing substrate in various ways, purely because they want to experience the sensations, feelings, points of view or knowledge that come with being another kind of living creature. I have labeled this “qualia surfing” because the term *qualia* refers to the philosophically inexplicable core properties of our sensations that can be known only to the sentient creature that experiences them—like the redness of the color red that is impossible to convey to someone who is blind. In today’s world we seek novelty by finding new patterns or combinations of our existing qualia—we go on vacation to new places and cultures so we can feel a new mix of colors, tastes, sounds and smells—but in the future we could modify the basic qualia that we experience, so we can know what it is like to perceive the

hue of ultraviolet light, or hear what a whale sounds like at infrasound frequencies, or be inside a centaur's body.

Fundamentals of qualia surfing

In previous papers (Loosemore 2009, 2012) I have described a theory of consciousness, one implication of which is that new qualia can be built by the right kind of modification to the brain's wiring (another implication of the theory, often taken as a *sine qua non* of mind uploading, is that if the functional machinery of a human brain is copied to another substrate, the consciousness of the original brain truly does exist in the new substrate). If that theory of consciousness is correct (as I will assume from now on), rewiring your brain to experience a new type of color would be relatively straightforward. It would involve building a new set of input pathways in parallel with the existing ones, and ensuring that those pathways end at new primitive concepts at a particular level of the concept system. This would not be a trivial piece of neuroengineering because the new wiring would have to be carefully meshed with the old, but there is no reason in principle why it could not be done by a sufficiently well choreographed nanobot ballet. To get new color qualia, the crucial aspects would be the parallelism of the wiring and the exact location within the brain of the new color concepts.

Similar new wiring could be set up for other senses, giving new primitive sensations in each modality. A simple extension of the auditory system would allow new octaves to be added above and below the normal range—although this would cause only a modest change to the way we experience sounds (we would sense a little more lowness and a little more highness). More dramatic auditory experiences could be generated by installing an entirely new auditory pathway, parallel to the existing one, connected to various kinds of sensor (radar, radio, microwave) in such a way that the signals ended in new sonic primitive concepts. Perceived from the inside, these new “sounds” would have the feel of something that was almost but not quite completely unlike ordinary sound. You could play a microwave version of Stravinsky's *Dances Concertantes* through this new modality, but it might sound both beautiful and... strange at the same time.

Taste, touch and smell could all be augmented in analogous ways: with new primary qualia, new degrees of resolution, or extensions to the existing ranges.

One thing that makes hearing different from seeing is that the processing machinery is quite different: a pair of serial channels in the case of sound, compared with a pair of massively parallel channels for the visual input. There are also big differences in the type of analysis done on the signals (among other things, visual input yields a detailed model of a three-dimensional world, whereas sound delivers only limited directionality). This observation leads to ideas about how entirely new channels could be designed, giving significant changes to the subjective experience. What about a five-dimensional sensory system, allowing us to experience spatial relationships directly in a 5D world? The input would have to come from a computer simulation, and the resolution might be limited unless

the brain's processing capacity could be augmented considerably—it seems reasonable to suppose that the processing needed would go up rather steeply as the dimensionality increased—but in principle it could be done. Someone could enter a simulated 5D world and feel its layout and processes in ways that only have a vague analogy to the feelings we get in a 3D world.

Abstract qualia

Something new happened in that last transition from the idea of new sensory primitives to new perceptions of the dimensionality of space. The person who wakes up with 5D vision will report new color qualia, because the input is coming from new receptors (albeit generated by the computer that is creating the 5D world)—but what's interesting is that she would also report some indescribable qualities in her perception of the new type of space. Would these new feelings in spatial cognition count as new, abstract qualia?

In the traditional philosophical literature the concept of qualia is usually illustrated by simple phenomena like color sensations, the feeling of pain, and so on. But the real meaning of the term is that certain feelings have a quality to them that resists objective description or comparison with other things, and those qualities can just as easily be found in some of the more abstract concepts we experience. If we compare sight and hearing, for example, we can observe that even something as simple as perceiving the *location* of an object is subtly different between the two senses. (To experience this, try walking around an empty room in pitch-black conditions, using the reflected sound of your voice to sense how far away the wall is: the subjective quality of that feeling of closeness is quite distinct from the feeling of distance you get when you can see the wall with your eyes.)

I have raised the subject of abstract qualia because there are some changes that could be made to an uploaded brain that would give sensations of newness over and above the simple sensory qualia. The insertion of a 5D vision system is one example, but if we look at other ways to augment the brain we quickly discover a large array of new possibilities, some of which involve clusters of sensations that do not include new primitive qualia.

New bodies

First on the list of these more complex qualia surfing activities would be any that involve a change of body. If nanobots were to create for me a tiger's body with an empty brain inside, the mapping between my original mind structure and the tiger brain would involve a lot of commonality—we both have tongue, and eyes, and four limbs—but the details of how these are shaped and wired into the brain would be very different.

There are two main ways to port my mind into a tiger brain. In the case of olfaction, new brain wiring could be added in parallel to my old human wiring (a parallel set of smell detectors processed by a new set of smell processing

pathways) or, alternatively, my existing brain wiring could be augmented to handle the tiger's more sophisticated smell faculty.

In the first case—new parallel wiring—I would feel a distinct set of new primitive qualia. As a tiger, I could take a big sniff and do some tiger-smelling of the world, but that experience would not resemble the ordinary human smells that I remember from yesterday. However, if my existing pathways were slightly augmented to match the design of the tiger olfactory pathways, I would still recognize the smell of, say, a damp forest, but in some measure it would have a greater detail or richness to it. Which of these modifications is chosen would be a matter of personal preference.

There is a third possibility: throw out all of my old sensory wiring (and not just the sensory wiring but also much of the cognitive wiring) and replace it with the tiger system, leaving the old human concepts intact but not connected to anything. This ought to give rise to a much stronger feeling of being a tiger, but make it harder to experience the side-by-side comparison of human-like and tiger-like existence—I might have vague ghost-like concepts of what it was like to be a human, but they would vanish at the touch.

After hanging out as a tiger for a while, I would return to my human body. What to do then? Presumably I would want to keep the memories of the tiger episode, but if I revert the wiring to its old state, I could end up with the same class of ghostly, intangible concepts, this time about tigerness. That might not be desirable, so instead I'd probably choose to keep a lot of the old tiger brain wiring, and, just as a human amputee is sometimes convinced that he can reach out with a non-existent limb, I might forever be able to remember those feelings of padding around on big soft feet, or swishing my tail, or extending my claws.

One implication is that collecting experiences as other kinds of creature would mean the gradual accumulation of new brain wiring. Our brains would have to grow to accommodate the machinery needed to recreate the sensations of being tigers, bats, whales and mermaids. This could get a little expensive on brain real-estate, so presumably we would have the choice to archive some of these experiences: we could keep some faint, ghost-like impressions all the time, but detailed memories of being other creatures could be put away like old photographs in the attic. Memorabilia to be visited occasionally.

Most of the creatures mentioned so far are natural biological types (assuming that mermaids, unicorns and centaurs are close to normal biology). But qualia surfing could just as easily involve constructed bodies robust enough to be comfortable in the atmosphere of Jupiter, a liquid-metal swimming pool on the surface of Mercury, the surface maelstrom on Venus, or deep interplanetary space.

Simple size modification of an otherwise regular human body would also be possible. There are limits to how small a human-replica body could be made while keeping a human-functional brain inside, but given what we suspect about the redundancy of processing in the brain, it seems likely that we could transfer into, say, a six-inch-tall person who could explore a *Land of the Giants* re-enactment. Physics would be subtly different at that level because many physical characteristics of the world are not scale-neutral; miniature people would discover

that their limbs moved differently, surfaces would be stickier, and falls from apparently fatal heights would be unproblematic (Dusenbery, 2011).

Finally, we could also transfer into bodiless computers, to recreate the *Minds* imagined by Iain M. Banks: entities found as the core sentience of houses, spacecraft, planets and orbitals, but capable of manifesting as androids or drones (Banks, 1987).

Classic uploading

Everything I have discussed so far involves staying in the real world but changing the form that we take, or the way that we experience it. I hope I have made it clear that even with this real-world restriction the number of different experiences available is vast, almost beyond imagination.

That said, though, we can step things up a level to what might be called the *classic uploading* scenario, where a person's mind is transferred into an entirely virtual body inside a virtual world. Provided that enough computing power is available to support the simulation, this kind of qualia surfing would be just as straightforward as the other kinds. The types of existence, and the types of sensory experience available, would be even more wonderfully diverse—in particular it would be possible to invent worlds in which the laws of physics were different.

There is little that needs to be said about classic uploading, except to suggest that it may be less common in the future than is widely supposed. Disappearing into your own universe will have its attractions, but part of human nature is that we want to have experiences we can share with others of our kind. I suspect that the range of activities chosen by the billions of people who make up our planetary population will eventually look like a thermodynamic distribution, with very many in the cooler (less adventurous, closer to human-normal) parts of the distribution, and then fewer and fewer numbers in the hotter, more exotic types of activity. My guess is that the human race will not suddenly disappear into virtual computer worlds, leaving a silent planet populated only by the robots who manage the host computers.

What is our purpose?

Let's pause for a moment to reflect on what it would mean to live in a future in which we were free to spend all of our time engaged in these kinds of pastime. We do the things we do because it gets us as close as possible to being happy--and making ourselves happy means, when all is said and done, the right kind of brain activity. We especially seek the pleasure of new experiences and discoveries, so qualia surfing would give us a supply of novelty that was almost inexhaustible.

Would there be limits to this novelty? Could the endless flood of new qualia become passé after a while? New qualia would be unique in ways we can't yet easily comprehend, so there really would be an infinite range of unique experiences to be had, but we must at least consider the possibility that an infinity of new qualia might not be satisfying enough.

Questions about limits to pleasure and boredom in utopia lead straight to questions about what motivation and pleasure really are. We sometimes talk about pleasure as if it were a monolithic thing—your brain either rewards you with a pleasure signal or not (Olds, 1956)—but in fact the real mechanism is likely to be more interesting. The exact nature of pleasure and motivation is still an open topic of research, but it seems clear that there is low level brain wiring that determines what kinds of activities give you pleasure, and what kinds of pleasure you experience when you find it.

Your own personal pleasure landscape is not fixed—it develops through your lifetime, and it is possible to shape it deliberately—but it does remain relatively stable in the short term. The set of things that give you pleasure today is a relatively stable part of who you are; if you happen to be a thrill-seeker, your first thousand years of qualia surfing might be dominated by a search for the most dangerous sports available, but if you are inclined toward the zen patience of a bonsai gardener, you might spend that time growing a giant redwood tree from seed.

How is this relevant to the question of whether qualia surfing might become boring? It is relevant because it opens up one more set of options for self modification: your motivation system itself can be changed. So a person who is addicted to adrenaline sports might decide that even though an activity like gardening has always seemed tedious, she would consider modifying her own motivation system to make gardening feel pleasurable. That kind of internal change would open up entirely new pleasures, so this meta-level qualia surfing (not changing the qualia directly, but changing your scope for appreciating new qualia) would further expand the landscape of possibilities.

One implication of this motivation-twiddling is that we could select any activity that humans have ever conceived, and even though we dislike that activity right now, we know that by reaching out and pressing a switch or slider we could alter our temperament and feel a delicious thrill at the idea of doing it. Do you find it horrible to imagine changing your sex? In a blink of an eye, you could overcome that feeling and thoroughly enjoy the transformation. Do you want to change yourself to become so passionate about mathematics that you could sit down in a library with a pile of books and slowly work your way through everything mathematical that has ever been written? Any conceivable activity could become pleasurable.

Domains

Perhaps the most complete form of qualia surfing would involve a visit to a *domain*—a place where a group of people go to create some way of life or historical epoch, but with the constraint that outside memories are checked at the door. A replica of London at the birth of the Enlightenment, when the very idea of modern science was being created. Or a fantasy domain inhabited by magical creatures. Once again, the range of possibilities is almost endless. The point of going into such a domain would be to experience the genuine thrill of a participant—you would not be a knowing, twenty-first-century interloper

watching others experience a fake seventeenth century. The domain could be set up in such a way that you were not, at any time, aware that you were in a simulation, and the key (apart from all the holodeck-like physical technology needed to support the place where this was happening) would be the temporary archiving of your previous memories so that you were not aware of the outside world. After a fixed amount of time you would emerge from the domain and resume your previous existence, but with the memories of your domain experience intact.

Archiving or modification of memories and knowledge would not be limited to domains, of course. Suppose you wanted to live for a while connected to the land, growing your own food, building and maintaining your property with your own hands, and patiently watching the trees grow over a period of thousands of years. Part of that way of life, for you, might involve a decision to temporarily forget the complexities of the world—not just getting rid of outside technology, but reducing the clutter in your mind so you did not even know that such things existed. There would be no need to erase the mental clutter, because it could be archived and then restored from time to time. And then one day, after enjoying the simple life for as long as you wish, you can decide to speed things up again, reload all the memories and jump into the next interstellar starship (or have yourself transmitted there as exabytes of data, to be reassembled at the far end)...

The greatest contribution made by these domains would be their capacity to eliminate the idea of stagnation. Yes, a mind that had spent eons looking for every conceivable pleasure might become tired. But a mind that could choose to temporarily forget that it had ever heard Bach's music would have the opportunity to hear it again for the very first time, in different circumstances.

In fact, a person could go into a domain in the form of a baby, and have another chance to rediscover all the things—the music, the art, the first kiss—that had already given them pleasure. But now consider: if substantial numbers of people felt inclined to become babies again (and surely they would), and if substantial numbers of people wanted to marry and have the experience of raising children (and surely they would), why not offer to put the two together in domains where their mutual interests converged? The children would be normal in every respect except that when they finally become adults they would be able to readopt their previous memories, adding one more lifetime to the stock of lifetimes they had already been through. Would this be a steady-state solution to the problem of a human civilization that wants to be as much like it was when it was in the past, but also not continue on the path of exponential population growth? It is possible.

The Utopia that came in from the cold

Utopia, for all its theoretical perfection, has almost become a laughing stock. It has acquired a reputation for being a place that would send the human race into a state of bored stagnation and decrepitude. But the picture of qualia surfing sketched here seems to indicate that the range of diversions and pleasures available to the human race would be almost limitless in scope. That picture depends on the ability to choose what experiences to have, what motivations to

feel, what memories to keep available and what environments to explore. And all of that emerges from the simple notion of getting inside the network of neurons, signals and other machinery of the human brain and either transforming it or moving it to another place.

References

- Banks, Iain M. (1987). *Consider Phlebas*. Macmillan, London.
- Dusenbery, David B. (2011). *Living at Micro Scale*. Harvard University Press.
- Loosemore, R.P.W. (2009). Consciousness in Human and Machine: A Theory and Some Falsifiable Predictions. In B. Goertzel, P. Hitzler and M. Hutter (Eds.) *Proceedings of the Second Conference on Artificial General Intelligence*. Atlantis Press, Amsterdam.
- Loosemore, R.P.W. (2012). Human and Machine Consciousness as a Boundary Effect in the Concept Analysis Mechanism. In: P. Wang & B. Goertzel (Eds), *Theoretical Foundations of AGI*. Atlantis Press, Amsterdam.
- Olds, J. (1956) "Pleasure Centers in the Brain." *Scientific American*. 195: 105-16.

Biographical note

Richard Loosemore is a lecturer in the Department of Mathematical and Physical Sciences at Wells College. He graduated from University College London as a physicist and from Warwick University as a cognitive scientist. His background includes work in artificial intelligence, cognitive science, physics, software development, philosophy, parapsychology and archaeology.

Keywords

future, qualia, consciousness, pastimes, entertainment, utopia, pleasure, motivation, domains.