The Consciousness Machine: Bridging the Hard and Easy Problems of Consciousness

Eric Saund, Ph.D.

The nature of our own conscious experience endures as one of the profound mysteries of human existence. Scientific understanding is advancing steadily toward explaining how brain activity underpins mental phenomena. Yet there persists a gulf between objective explanations of mind/brain functions, and the subjective experience that belongs to each of us internally. Can this gulf ever be bridged? Perhaps it can, through a technology we might call a Consciousness Machine. A Consciousness Machine is roughly a cross between a video game and a real-time tracking simulation of the mind of its user. Its purpose is to create an intimate link between intellectual understanding of how one's own mind/brain is functioning, and the actual experience one is feeling. Could such a device actually work? Let us perform a thought experiment and try to discern an answer.

The problem of apprehending consciousness has famously been separated by philosophers of mind into an "Easy Problem" and a "Hard Problem" (Chalmers, Scientific American, 1995, 2002). The so-called Easy Problem is to understand the phenomenon in objective terms using the standard processes of science. Instruments such as Functional Magnetic Resonance Imaging and studies of normal subjects and brain injury patients produce observational data. Data is compared with predictions from mathematical models and computer simulations to validate and invalidate hypotheses. Accumulated knowledge supports technologies that put theory to work for practical applications in medicine and elsewhere. The Easy Problem represents a major undertaking in modern science that integrates dozens of disciplines. Clearly, the "Easy Problem" is actually quite difficult.

Let us assume that progress on the Easy Problem will eventually reveal the neural correlates of consciousness. Various parts of the brain are known to specialize in speech, vision, emotions, and all other aspects of mental life involving the higher nervous system. Extrapolating forward, specific brain regions, activity patterns, and encoded data structures will presumably be shown to not only correlate with conscious experience, but even to be their causal basis. What the neural correlates of consciousness will look like cannot be known at this time simply because the nature and mechanisms of consciousness itself are so poorly understood. Current theories of consciousness place different emphasis on various of its aspects, including physical embodiment; homeostasis; situational awareness; attention; world models; knowledge and awareness of self; information complexity; and cognitive architecture.

The other problem of consciousness, the "Hard Problem," takes this name because it is less accessible to scientific methods. The Hard Problem is to provide a meaningful account of subjective experience. Subjective experience, or "qualia," is owned exclusively by the being having that experience. For this reason, it cannot even be proven directly that qualia even exists for anyone but oneself.

The Hard Problem extends past the outer reaches of scientific explanation into the realm of Philosophy. This territory sustains important debates about the human condition. For example, given explanations of choice behavior in terms of brain mechanisms, does free will exist? If so, under what epistemic framing? Is our own subjective experience, say, of contemplating among alternatives and then making a choice, actually an illusion? Are we no more than very complex automatons?

The Easy Problem has not been solved, and it is not certain that the Hard Problem can ever be conquered through science. Many people are repelled by the notion that their thoughts and feelings might be reducible to mere neural activity in their brains. No one takes well to being called an automaton. Today, it cannot even be definitively proven that consciousness is fully explainable through the mechanisms of our physical universe. Metaphysical accounts of the human mind can be the source of hope and optimism, but also ignorance and superstition. Human progress demands that we try to understand ourselves, including the nature of our conscious minds, as fully and rigorously as possible.

Thus there exists a quandary. As the mysteries of Easy Problem are unraveled to explain conscious experience in objective terms, can this knowledge ever penetrate the barrier of subjective experience to impact the Hard Problem?

The answer may come through technology that does not exist but whose basis is being developed today. The challenge is to build a tight loop connecting subjective experience, with accessible representations of explanations of how that experience is being produced. The challenge will be met when experiencing, believing, and understanding, become a unified state of mental operation—a conscious experience in its own right, as it were. In fact, experiencing, believing, and understanding do become unified routinely in our day-to-day lives whenever we slip into a zone of competent performance in a job, game, hobby, book, or other mentally immersive act. More precisely then, the challenge of penetrating the Hard Problem of consciousness is for the topic of the immersive engagement to be apprehension of the functional mechanisms underpinning the subject's (user's) own subjective experience, moment to moment.

We may label the device that provides this experience a *Consciousness Machine*. A Consciousness Machine will be a hybrid of several existing and emerging technologies. In some respects, it will be like a video game, with the topic of the game being the user's own mental processes. The analogy is thus:

- **Immersive:** Video games are immersive experiences that cognitively and emotionally transport the user to other realms, some realistic, some fantastical. A Consciousness Machine will immerse the user in a visualization of an objective reality, namely, what their own brain/mind is doing and more importantly, how it is doing it.

- **Simulation:** Video games frequently involve simulations of action by physical objects and characters. In order to transport its user to a perceptual environment presenting an accurate and authentic account for what his or her mind is doing, a Consciousness Machine will embody a tracking simulation of that user's own mind functioning in real time.

- **Direct Manipulation:** Video games are engaging because the user exercises practiced skills at taking actions to achieve results. A Consciousness Machine should permit its user to perform deliberate, direct manipulation of their own brain functions in order to test and observe the subjective experiential effects.

The hypothesis of the Consciousness Machine proposal is that this hybrid device will create a way of knowing of an entirely new kind, one where intuitively subjective feeling and rigorously objective thinking meet on an objectively grounded basis to mutually inform one another about the self.

Could such a such a Consciousness Machine be built? And if so, could it effectively bridge the gap between its user's objectively-based, "Easy Problem" level understanding of their own mind/brain, and their "Hard Problem" qualia experience?

Let us consider technical challenges, constraints, and requirements for Consciousness Machine technology. These considerations will further inform the likelihood of its effectiveness.

Challenge I: The "Easy Problem," Models, and Simulation

The most fundamental challenge to building a Consciousness Machine is to understand consciousness itself, in other words, to solve the "Easy Problem." Merely to identify the neural correlates of consciousness will probably not be sufficient. A live-action display of neural correlates might be informative, but it falls short of a causal model that can make predictions and support interventions. Beyond identification of correlations, solutions to the Easy Problem will lead to computational theories of consciousness which in turn provide the bases for *functional* models that, if correct, should predict input/output relations to some interesting and reasonable levels of accuracy. To enable a Consciousness Machine, a functional model must operate to track and even anticipate the user's mental states in real time. Thus it would most likely be realized as a function*ing* computer simulation.

Computer simulation has become an integral element of modern scientific and technological methodology spanning basic science through product design and manufacturing. Simulations allow scientists to test hypotheses and interpret observational data. Simulations allow designers to rapidly explore design parameters, detect flaws and failure modes, and predict performance. Simulations enable machine and system controllers to formulate appropriate actions to achieve desired operation for what is known in control parlance as “the plant.” In neuroscience, simulation of neural circuits is becoming a cornerstone for understanding the brain.

Fortunately, a Consciousness Machine will probably not require detailed simulation at the neural circuitry level. Just as simulations for weather forecasting operate over aggregate chunks of atmosphere in a spatial grid, the main contours of conscious awareness may well be modeled as a complex system comprised of merely hundreds or thousands of functional blocks and connections, comparable to the number of distinct brain regions and nuclii found in a brain atlas.

The technology underlying computer simulation is advancing rapidly across multiple threads of computer science including parallel hardware architectures, programming languages, compilers, and algorithms. The internet’s demand for server farms is driving the cost for amassing processors ever lower. As a matter of progress in simulation technology, then, the raw computational infrastructure for building real-time simulations of mental function appears to be falling into place.

The steps toward mind/brain simulation will be difficult and incremental. As functioning models are developed, their scopes of competency cannot be expected to immediately span all domains of mental experience. Each of the various aspects of mental life may be included, neglected, or else only crudely modeled in early, restricted Consciousness Machines, which will necessarily be designed to illuminate only some limited but interesting aspect of the mind’s operation and awareness.

For example, a Consciounsess Machine focused on visual perceptual awareness might emphasize color sensation; perceptual organization of texture, shape, and movement; cues to scene and object classification; and the perceptual-motor routines underlying familiar activities. It might then extend to encompass mechanisms for recognizing faces and associating them with the user’s social mileau. Or it might include pathways for visual detection of threats and deployment of reactions and strategies to avoid them, along with the emotions they invoke. A different limited Consciousness Machine might concentrate on the gustatory aspects of life, and would thereby model, at the lowest levels, thirst and hunger, thence proceeding through perception of taste and smell, to appreciation of cuisine, and progress to the social aspects of food preparation and the dining experience. In many instanations of a Consciousness Machine, modeling of the user’s knowledge and feelings about themself and their relation to the external world will be of paramount interest.

The forms that simulation models of consciousness will take depend ultimately on the facts of how consciousness functions. It may be that consciousness can be largely understood in terms of which brain regions and nuclei are active from moment to moment. Or, it may turn out that effective theories will center on complex temporal synchronization, phase locking, and entrainment across highly distributed, hologram-like representations. Conscious experience might turn out to be highly predictable given strong enough models and measurements. Or, its complexity may prove to be chaotic like the weather. It is possible that solutions to the Easy Problem will reveal that a Consciousness Machine as envisioned here cannot be built, either because of practical intractabilities, or because of some yet unknown fundamental principle. Such pessimistic outcomes are, however, by no means established at this time.

Challenge II: Abstraction and Visualization

One objective of a Consciousness Machine is to present the user with an objective account of the mechanism of their own mind/brain in action that is readily comprehensible and interpretable. The machine is successful only when intellectual understanding can be intuitively and effortlessly linked to the subjective awareness it corresponds to. The human mind is often described as the most complex and sophisticated apparatus in the known universe. Billions of neurons connect through billions of synapses through billions of firings per second. It is fair to ask, how could a person ever comprehend that?

Hope lies in the principle of abstraction. Abstraction simply refers to separating principles and patterns from specific instances and details. Formulation of appropriate abstractions is a key practice in science. For example, the field of thermodynamics abstracts the microscopic collisions of many many air molecules in terms of the gross properties of temperature, pressure, and entropy. Solutions to the Easy Problem will necessarily involve abstractions for mental states and processes. Some of these we can already glimpse, such as top-down and bottom-up directives to attention, associative recall from memory, and pre-rehearsal of actions. Others have yet to be discovered.

To exploit such abstractions, a Consciousness Machine must employ effective means for communicating in terms of them. Here, the burgeoning field of Data Visualization will play a critical role. The internet is driving the technology. Every day, petabytes of data are generated and stored about business transactions, browser clicks, search terms, phone calls, texts, television viewing, voting history, the list is endless. Concurrently, an industry is growing to comprehend and make productive use of this data. The tools mix statistical data analysis with data visualization methods including synthesis of animated graphics, multi-dimensional colored charts, and sensorially-rich narrative explanations. The popular buzzword is "Big Data Analytics."

Data visualization is grounded in both technical analysis and the art of creative design. In order to be readily comprehensible, data presentations must be designed to match the human mind's perceptual skills. To portray a trend over time, a list of numbers is ineffective because it requires inspecting, decoding, storing, and comparing numerals sequentially. By contrast, the slope of a graph is perceived immediately because it more directly engages geometric features our visual systems are wired to detect automatically and effortlessly.

A key ingredient to a Consciousness Machine will therefore be the design of evocative and readily comprehensible visualizations of abstractions of mental function. (By "visualization," I do not intend to restrict sensory modality to the visual displays; audition and haptics certainly play important roles in immersive environments.) It is possible that effective visualizations will adhere closely to technical expressions such as UML diagrams (Unified Modeling Language used in design of computer programs and other complex system), or to electronic circuitry. Or, it is possible that more intuitively graspable portrayals of mental activity will take a more metaphorical character. Consider sitting in an airplane cockpit and zooming through thought space, or standing as a conductor before an orchestra of players representing elements of one's own mental function.

Whatever data visualization designs are employed for a Consciousness Machine, technology is being rapidly developed that can deliver evocative images, sounds, and touch sensations to users. Full-field visual displays have become fast enough to respond to head motions without inducing motion sickness due to vestibular/visual delays. Recent commercial investments and acquisitions reflect bets that these devices will become as commonplace in the computer gaming industry as handheld controllers are today. It is known that full-field visual displays have overwhelming power to drive the mind's awareness of one's place and surroundings. Placed in a computer-generated scene in which the user is standing on the precipice of a cliff, it becomes extremely terrifying to take a step, even if one possesses full knowledge that they are standing on the solid floor of a laboratory. The power of sensory immersion will play an important role in a Consciousness Machine's ability to engage the user's full perceptual faculties to readily process and understand well-designed data visualizations of their own mental functioning.

Challenge III: Synchronization

In order to achieve intimate connection between externalized representations of mind/brain function and subjective experience as it happens, the real-time simulation undergirding a Consciousness Machine must accurately track what the user is actually thinking and feeling. Just as weather simulations can only extrapolate a few days ahead and must be constantly updated with current measurements, a Consciousness Machine will rely strongly on a feed of signals and measurements from its user. This could require heavy instrumentation.

Technology for monitoring brain activity has shown dramatic progress over the past several decades. It falls into four major categories.

First, direct measurements of brain activity can be taken by placing electrodes directly on or within brain tissue. This method is obviously highly invasive and dangerous, and occurs only during brain surgery or in the treatment of serious medical conditions. This approach is quite unlikely for a healthy person's Consciousness

Machine.

Second, electrodes placed on the scalp and face can sense electrical activity of assemblies of neurons through the skull. In conjunction with skin galvanic measurements and other surface sensors, EEG techniques are commonly used in polygraph and biofeedback devices. Because their measurements capture gross activity produced by only a fraction of cortical brain tissue, EEG is quite limited in the comprehensiveness of the brain activity data it can collect. Moreover, interpreting the signals has proven extremely difficult. The most meaningful discoveries pertain to cyclic variations in collective neural firing. Coupled with electrodes attached to the face, EEG is effective for example at detecting stages of sleep.

Third, non-invasive observation of the face, eyes, and body has reached a turning point through computer vision technology. Popular psychology abounds with guides to reading "body language." This practice is more art than science. But at least anecdotally, under certain situations, it appears likely that skilled human observers are indeed able to discern covert intent and knowledge on the part of others. Popular examples are the skills of reading gestures and facial expressions revealing "tells" in card games, and "tipping pitches" in baseball. Under laboratory conditions, eye tracking apparatus can be effective at interpreting subjects alertness levels, focus of attention, and topical interests in visual displays. To the extent that body language is a true expression of underlying thoughts and motivations, detailed and systematic analysis might offer an accessible window into the mind. Current advances in computer vision are achieving real-time tracking of body pose and facial expression. It is possible that subtle signals of this nature might some day be "read" to reveal limited conscious and unconscious dimensions of a user's mental activity.

Fourth, brain activity measurement has made remarkable progress in the form of FMRI (Functional Magnetic Resonance Imaging) (Scientific American, 2008, 2012). This approach measures side-effects of neural activity, namely changes in blood flow resulting from metabolic processes. FMRI machines are capable of sensing deep brain as well as cortical activity. Their spatial and temporal resolution depends on the strength and proximity of magnets situated around a subject's head in a large, tightly enclosing chamber. These factors conflict with the goals of relaxation, comfort, and ease of movement that might be important to users of a Consciousness Machine. Yet, FMRI measurements of brain activity patterns have reportedly discovered correlations with various states of consciousness, including sleep, meditation, anesthesia, brain damage, and controlled doses of psycho-active drugs.

Because of the number, size, and speed of neurons that comprise brain function, the Synchronization challenge for building a Consciousness Machine remains extremely daunting. Only future technological developments will tell if even the most sophisticated of brain/mind simulations can be accurately tracked to a user's actual thoughts and feelings in real time.

Challenge IV: Direct Manipulation

A fourth challenge to building a Consciousness Machine is to close the loop between hypothesis and proof. Any person shown exotic visualizations claiming to reflect the inner workings of their mind might or might not find them convincing. An easily persuadable person might readily assent that indeed, certain lighted displays and engaging sounds are remarkably coincident with their thoughts and feelings. But the ideal user of a Consciousness Machine will be a skeptic. They will demand proof that the asserted claims about how their qualia is produced stand up to critical testing. The Consciousness Machine hypothesis asserts that convincing proof is possible through the method of Direct Manipulation. Direct Manipulation is fundamentally about closing the perception/action loop.

The significance of the linkage between perception and action was exposed in a classic 1963 experiment by Alan Hein and Richard Held. Pairs of kittens were raised in a highly controlled environment. After birth, they lived in the dark, except for three hours per day when they were placed in a painted chamber. One kitten, known as Active Kitten, was free to roam around the chamber. Active Kitten was connected through a special collar and linkage apparatus to its companion, Passive Kitten. Passive Kitten rode in a cart that mirrored Active Kitten's wanderings. Thus, Active Kitten and Passive Kitten obtained equivalent visual experience. The difference is that Active Kitten controlled the experience through its own movements, while Passive Kitten had no control. Passive Kitten was allowed to walk during the periods of darkness and suffered no purely motor deficits as a result of the experimental conditions. But when tested at the conclusion of the training period, only Active Kitten demonstrated competent visual-motor coordination. Passive Kitten was for example unable to walk toward a visual cliff without stepping off (with glass protecting the dropoff). The widely accepted conclusion is that in order to obtain full and meaningful perceptual-motor mastery of sensory signals, it is important that those signals be connected through experience to feedback from actions taken by the motor systems. What Active Kitten had that Passive Kitten did not, was Direct Manipulation.

In Human-Computer User Interfaces, Direct Manipulation refers to displays and controls that permit a user to modify content without recourse to indirect or secondary underlying representations. The classic example is WYSIWYG (What You See Is What You Get) text editors. Microsoft Word is a WYSIWYG document editor because the document that results from printing or rendering in a presentation app appears identical to the representation of the document manipulated by the author in the editing environment. By contrast, the document formatting language, LaTeX, popular in scientific circles, employs a "source document" that includes not only the text of the document proper, but also meta-instructions that control appearance and formatting. An experienced user of LaTeX learns to mentally bridge the gap, to anticipate how meta-commands will affect appearance on a page. Similarly a skilled musician plays an intended note without thinking. The finger, mouth, or body motion required to generate a desired sound becomes automatic, so that mentally they simply produce the note they are thinking of.

In a Consciousness Machine, direct experience of qualia can be only indirectly represented through visualizations of mind/brain mechanisms. Whatever visualization language is employed, it can be only an intermediary representation, like the keys on a piano in relation to the sounds produced by the strings. To achieve Direct Manipulation, the user should be enabled to deliberately manipulate the representation through motor commands, and thereby produce the subjective experience immediately.

This effect implies two conditions. First, the user must achieve the skill of effortless association between the indirect representation and the corresponding subjective experience. This is like learning to play the piano, or to ride a bicycle, or steer a car, where in the latter cases the indirect representation is a physical instrument and the goal experience is producing a musical passage or advancing on a roadway. Second, direct manipulation of the mind/brain mechanism visualization should reliably activate the corresponding sensory, perceptual, emotive, or cognitive experience.

This second condition raises a very significant potential challenge, for it suggests that the Consciousness Machine's representation for the user's mental state should be required to elicit direct influence on the user's actual brain activity. The technical situation in this regard is similar to that of brain monitoring and synchronization. We possess a few blunt instruments, but nothing close to precision control.

Direct electrical stimulation of brain tissue has long been known to produce striking behavioral and psychological effects. Indeed, therapeutic Deep Brain Stimulation is now used to treat Parkinson's disease and other neurological disorders. But no one would suggest implanting electrodes for the purposes of illuminating the nature of consciousness in normal individuals. Less invasively, the method of Transcranial Magnetic Stimulation (TMS) is used to directly affect brain activity for purposes of medical diagnosis, therapy, and research studies into functional operation of cognitive mechanisms. TMS affects only relatively large areas of brain located just beneath the skull; unforeseen technical breakthroughs would be needed to achieve controlled perturbations of the intricate circuitry that in all likelihood gives rise to consciousness.

The Direct Manipulation challenge can also be attacked through more conventional psychological avenues. The brain's wiring is complex and plastic. Complexity implies that the mental phenomena we experience—sensations, percepts, thoughts, and emotions—are subject to myriad external factors, all of which can be manipulated if sufficiently well enough understood. Consider the surprise induced by visual, auditory, and tactile illusions; consider the bedazzlement achieved by the perceptual and cognitive deceptions of professional stage magicians. Plasticity implies that associations and habits can be trained and overwritten. Although simple behaviorism was overtaken by computational theories of mind that ground modern cognitive science, straightforward stimulus-response conditioning remains a true factor and an important tool in psychology.

For a Consciousness Machine, it is possible in principle that an interesting degree of Direct Manipulation could be achieved by clever exploitation of complexity and plasticity. To take a crude example, suppose that a certain emotional state could be associatively trained to follow a particular sound—like Pavlov's dog. Then, a Consciousness Machine might be able to offer the user a cause-and-effect pairing. A visualization displays the brain mechanisms that equate to the emotional state, and an interface action induces that state. The responsible brain region would be indirectly and involuntarily exited by the trained associative sound. In a sense, the modern gaming industry is already quite sophisticated at exploiting human psychological foibles, as witnessed by the Pavlovian conditioning and related effects that drive the design of sights and sounds pervading gambling casinos. Obviously, the feasibility of this "psychological" approach to Direct Manipulation is subject to a much deeper understanding of the mind/brain's complexity, plasticity, and manipulability in the sense of solving the Easy Problem.

Challenge V: Learning

A Consciousness Machine cannot be, and is not intended to be, a magical device in which the user simply sits down and immediately secrets of the self are revealed. Instead, competency must be built in stages, through learning. And learning must take place in two directions.

Because the purpose of a Consciousness Machine is to link intellectual understanding of the user's mind with subjective experience produced in, by, and of that mind, the user of a Consciousness Machine must build a substantial body of knowledge and skills. They must build conceptual knowledge of how a mind achieves consciousness in principle. They must build knowledge of the specific parameters, patterns, and peculiarities of their own mind within the space of human variation. They must acquire perceptual skills to readily interpret data visualizations portraying their mental function in action. They must acquire motor skills to operate the controls.

At this point it is difficult to estimate just how much effort, talent, and capacity will be required to master the requisite knowledge and skills to operate a Consciousness Machine to the point of authentically unifying intellectual understanding with subjective experience. For reference, a rule of thumb is that 10,000 hours is required to become an expert across a broad array of subject matter, and this generally takes ten years. But, some exceptionally talented individuals are able to master sports, video games, music, languages, and meditation within just a few years of practice and study, especially if undertaken during critical periods of brain plasticity and readiness. Optimistically, popular classes in meditation and self-hypnosis are able to train many people to achieve some level of enhanced conscious state of self-awareness with a handful of intensive sessions.

In the counter direction, a Consciousness Machine will have to be tuned to the mind/brain of its individual user. Each of us has our own particular personality, thinking style, knowledge, habits, skills, and mental deficiencies. For a machine to accurately model and simulate mental function of any particular user, it will have to acquire and implement accurate detailed knowledge of these peculiarities.

The technology of Brain/Computer Interfaces presents an early model for mutual learning between human users and machines. In a typical application such as controlling a robotic prosthesis, a rich stream of signals is collected from surface and implanted electrodes situated on and in the body and nervous system. In order for these signals to be interpreted to generate intended and coordinated action, a period of training is required. In early devices, the burden was placed on users to figure out which combinations of twitches and efforts would get the machine to move as intended. But the clear trend in the field is toward incorporating modern machine learning methods to map the complex input signal array to higher level user goals for machine posture and actions.

Another technical tributary for mutual learning may lie in the field of Intelligent Tutoring Systems. By design, these systems implement models of the learner, and in particular, estimation of where the learner lies in relation to mastery of some domain of study such as Geometry or a skilled trade. A good teacher maintains broad and deep understanding of the ways in which students assimilate instruction and sometimes struggle with common misunderstandings. And through experience, they learn how to interpret individual students' responses to diagnose each one's placement in this space. Intelligent tutoring systems currently do not emulate human teachers in this way, but they are designed under theories of cognitive function incorporating long and short term memory, skills, the formulation and execution of plans, and interpretation of input from visual and linguistic modalities. These design principles and methodology would inform the design of, and perhaps even contribute technology to, a Consciousness Machine.

The five main challenges to building a Consciousness Machine, Modeling and Simulation, Abstraction and Visualization, Synchronization, Direct Manipulation, and Learning, are all quite formidable, but across all fronts, the interplay of scientific investigation, technological invention, and economic and social motivators has driven remarkable progress. The most formidable technical barriers come from access to physical brain tissue for purposes of both measurement, and externally-driven manipulation. What is yet unknown is whether these barriers can be sufficiently circumvented by virtue of increasingly sophisticated understanding of their function, thereby permitting achievement of the end goal of linking understanding with subjective experience by leveraging largely indirect physical means.

Prognosis

A complete and comprehensive theory of consciousness should penetrate not only the so-called Easy Problem of understanding the brain or possibly other computational mechanisms that can produce a conscious mind, but ideally, also the Hard Problem of engaging such a conscious mind in comprehension of its own subjective experience. The Consciousness Machine is at this point a speculative proposal, but, frequently, directed investigation is preceeded by thought experiments. In case the Consciousness Machine thought experiment achieves validity, it is possible that the device itself could eventually become key apparatus for testing theories of consciousness. As long as the Hard Problem remains ensconced exclusively in the form impenetrable qualia, our understanding of this most mysterious aspect of nature will remain unresolved.

The Consciousness Machine proposal raises new questions for discussion for scientists and philosophers. If a Consciousness Machine modeled the mind of its user with sufficient fidelity, would the machine itself be a conscious entity? Would it be a mental clone of its user? Could a Consciousness Machine finally definitively answer the perennial question, Do you see the color green in the same way I do? Would the device be capable of super-learning injection of knowledge and skills to its user? Would this be ethical? Informed answers to these questions may lie only in the actual building of Consciousness Machines, if and when the challenges can be met.

As the science and technology for a Consciousness Machine progresses, a more direct step toward its realization and appreciation would be to move beyond the realm of thought experiment, to an envisionment. Envisionments are commonly used by academic and commercial teams attempting to project their efforts into the future. The instruments of envisionment include mockups, storyboards, collection and study of contributing technologies, and implementations of pieces of the greater endeavor. The attractive thing about envisioning a Consciousness Machine is, the most important tool is not any advanced or exclusive technology, but an informed imagination operated by a conscious being.

More to Explore

"The Feeling of What Happens: Body and Emotion in the Making of Consciousness," Antonio Damasio, Harcourt Brace & Company, 1999.