

No Ghost in the Machine

Alan M. Kadin

Princeton Junction, New Jersey, USA

amkadin@alumni.princeton.edu

The prevalent pre-scientific paradigm for understanding nature focused on design or intention, even for inanimate objects. This approach was debunked by Newton for physics, and by Darwin for biology. But belief in the unique supernatural nature of human intelligence is still widespread. I argue that biological intelligence is due to simple evolved structures based on neural networks, without the need for any new physical mechanisms (quantum or classical) or a “ghost in the machine”. Humans see agency and intent everywhere, because we are programmed to do so. The conscious mind may turn out to be a virtual reality simulation that is largely illusory. Furthermore, these structures may be emulated in artificial neural networks, to create true artificial intelligence.

“Everything should be made as simple as possible, but not simpler.”

[Attributed to Albert Einstein](#)

I. Introduction

This year’s essay contest for the [Foundational Questions Institute](#) on [“Wandering Towards a Goal”](#) contrasts “mindless mathematical laws” with “aims and intentions”, and asks how this apparent paradox can be resolved. I argue below that if one properly understands the scientific basis for physics, biology, psychology, and computer science, there is no paradox. But the superficial philosophical implications of each of these may be misleading. For example, the mathematical equations of simple physical systems lead to causality and determinism, but real physical systems have noise, and complex physical systems tend to be chaotic and unpredictable. The paradigm of causal control in complex systems is an illusion, and reality is more subtle.

Similarly, biological evolution is predicated on random mutations, but leads to complex structures and organisms that appear to be “intelligently designed”. So biological design, too, is an illusion, which is explainable in terms of blind adaptation to complex environments. This biological tendency toward complexity may seem incompatible with the thermodynamic tendency toward increasing entropy in a closed physical system near equilibrium, but biological systems are neither closed nor in equilibrium.

Regarding human intelligence, people feel that they are independent rational agents, but in reality, most of human behavior is subconscious and irrational. In classical computer science, a computer is an arithmetic engine with external memory, but most modern computational problems related to “artificial intelligence” are more akin to matching patterns than to arithmetic calculations, and the same is true for natural intelligence.

Further, I argue that consciousness itself reflects an evolved brain structure that is not uniquely human, and provides an adaptive system capable of making rapid decisions based on simplified models and incomplete data. An analogy may be with a dynamic virtual reality environment that integrates and synthesizes data from diverse sensor inputs and generates motor outputs that are reflected within the same virtual environment. I further suggest that a similar consciousness engine may be emulated artificially.

The title of this essay comes from the 1967 book by Arthur Koestler, [“The Ghost in the Machine”](#), which in turn came from a phrase used by philosopher Gilbert Ryle to describe [mind-body dualism](#). My view is that the subjective experience of consciousness reflects the brain activity associated not with the entire brain, but only that small portion that is projected into this self-conscious virtual reality construct.

II. Physics, Mathematical Determinism, and the Illusion of Control

In pre-modern physics, as exemplified by the work of [Aristotle](#), the behavior of an object, animate or inanimate, was believed to be driven by its natural tendencies, in a qualitative rather than a mathematical sense. This general doctrine is known as [teleology](#), from the Greek for end, i.e., goal. Fire tends to rise, while massive objects fall because that is their nature. This approach was a partial attempt to get away from the anthropomorphic tendency of pre-scientific peoples that all objects in nature are driven by spirits.

A further general principle in pre-modern physics was that humans are the center of the universe, with the earth and the heavens having fundamentally different laws. For example, objects on earth tend to move in straight lines, while objects in the heavens tend to move in perfect circles. But in order to get this [geocentric model](#) to work, Ptolemy modeled the motion of planets using epicycles, i.e., circles within circles. When this was done properly, it worked quite well, but it was both complicated and arbitrary. Isaac Newton took simplification and unification much further, by developing a quantitative theory of universal gravitation on both the earth and the heavens, and by developing the mathematics (differential equations) needed to do the calculations. The paradigm for classical Newtonian physics is the “[Clockwork Universe](#)”, where everything in the universe follows mathematical trajectories that are fully deterministic and completely predictable. This was very influential in the sciences and beyond, including with the doctrine of [philosophical determinism](#), that the future is completely defined and can never be altered.

It is widely believed that this clockwork paradigm held until quantum mechanics introduced indeterminacy in the early 20th century, but this is incorrect. The entire fields of [statistical mechanics](#) and thermodynamics were developed in the 19th century to deal with uncontrolled thermal noise in classical many-body systems, which are microscopically deterministic but macroscopically random. This led to the concept of [entropy](#), which is a measure of the probability of a given physical many-body configuration. For a closed system in equilibrium, entropy must always increase, leading to more probable configurations. While the microscopic equations of classical physics are time-reversible, irreversibility associated with increasing entropy enables one to define the arrow of time.

More recently in the 20th century, classical nonlinear dynamical systems were shown to be highly sensitive to initial conditions in a way that is practically unpredictable (the “[butterfly effect](#)”). The mathematical theory of deterministic [chaos](#) was developed to describe the dynamics of some of these classical nonlinear systems. Furthermore, macroscopic systems often include complex [feedback](#) loops that largely decouple the macroscopic behavior from microscopic degrees of freedom. While textbooks emphasize the mathematical certainty of simple dynamical systems, the behavior of real classical systems is more subtle and much less predictable.

This is illustrated in the block diagrams shown in Fig. 1. Fig. 1L shows a classical causal sequence of A causing B causing C. In contrast, Fig. 1R shows a more complex set of interactions: each component affects the others, but also is affected by them via back-action. Furthermore, each component also acts on itself (recursion), and is affected by noise sources N. For example, consider Fig. 1R in the context of [global warming](#). The atmosphere might be system A, the biosphere might be system B, and the carbon reservoir in fossils fuels might be system C. The noise N represents those factors, local and global, which are not controlled. The interactions between these systems can provide the basis for a complex model, which is still oversimplified – there is also radiation input/output and the influence of the oceans. While the near-term trends are clear, the accumulation of uncertainties and model inadequacies make long-term predictions less reliable.

Philosophical determinism is frequently applied to human interactions, where there are not even good mathematical models. It makes little sense to derive human-level determinism from mathematical trajectories of planets or atoms.




Fig. 1. Comparison of simple causal sequence to more complex set of interacting systems. (L) Causal sequence of A causing B causing C. (R) Interacting systems A, B, C, and N, where N represents uncontrolled factors such as noise.

The philosophical doctrine of [free will](#) is often contrasted with determinism, on the level of individual humans. Where determinism suggests that the future is completely predictable, free will suggests that one's future is undetermined until one makes a decision. But both determinism and free will are really straw men; as individuals, we are under varying degrees of external and internal constraints. You are free to choose your breakfast, but the climate of the earth is beyond your control. Furthermore, we tend to see ourselves as free agents, but we actually have much less control than we think we do. Control is just another illusion.

Quantum mechanics incorporates indeterminacy even on the microscopic level, which is different from classical uncertainty. However, the major effect of quantum uncertainty on the macroscopic level is simply to introduce another source of noise ([quantum noise](#)), in addition to thermal noise. There are other more exotic quantum effects, but it is unclear whether quantum entanglement has any significant impact on macroscopic systems at ambient temperatures. Quantum measurement theory focuses on the role of the observer in changing the quantum state, and it has been [suggested](#) that some of these paradoxical aspects may provide a basis for aspects of the human mind and consciousness, but this would seem to be yet another example of anthropocentrism. In contrast, I have proposed (Kadin 2012, 2015) an alternative picture of quantum mechanics that avoids the paradoxes of entanglement and measurement. This is addressed briefly in the End Notes.

III. Biology, Evolution, and the Illusion of Design

Living organisms have long been regarded as different from the rest of nature. Indeed, it was widely believed that biology did not follow the same physical laws; this is the basis for the discredited doctrine of [Vitalism](#). Biological systems appear to be intricately designed machines, without a self-evident set of instructions. It was not until the middle of the 20th century that the DNA [genetic code](#) was discovered, providing a digital program of instructions for growing an animal, plant, or microorganism. But the greatest breakthrough in biology occurred in the 19th with Darwin's Theory of [Natural Selection](#), commonly known as evolution. Evolution is both the simplest and the most profound concept in all of science; all it requires are exponential reproduction and genetic noise, interacting with an environment. But prior to Darwin, this was hidden for centuries, because of the illusion of [intelligent design](#).

As shown in Fig. 2, a guided design (Fig. 2T) follows a specific plan towards a goal, whereas in natural selection (Fig. 2B), all possible changes to the current design are generated automatically, but only those that provide improved adaptation to the environment and reproductive success are preserved for future generations. These diagrams appear somewhat

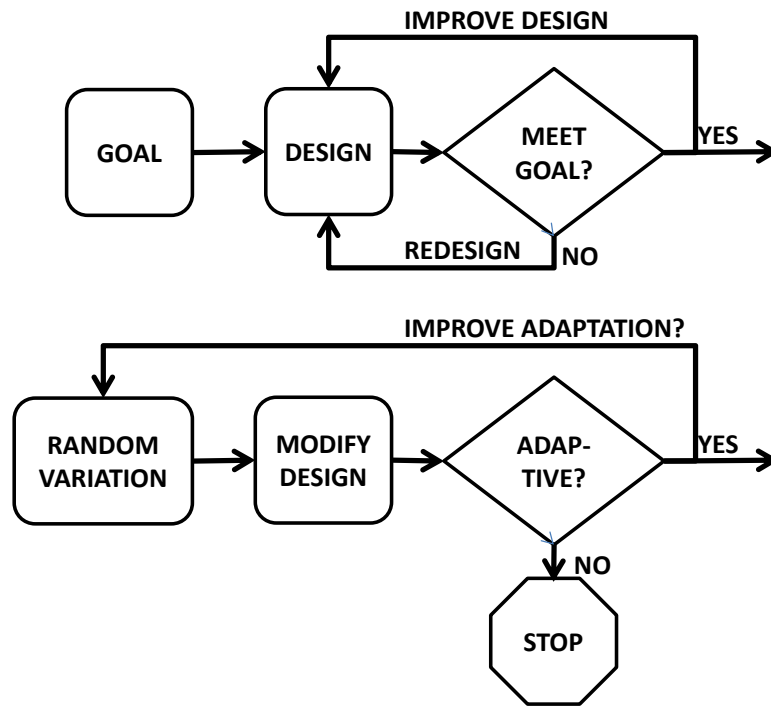


Fig. 2. Comparison of guided and unguided design pathways. (T) Guided design pathway, whereby a product is designed for a functionality or market. (B) Unguided design pathway, as in biological evolution of an organism.

similar, but have critical differences. In particular, only a guided design can produce a radical redesign in a single step. In contrast, the unguided design of natural selection can only make minor modifications per generation, each of which must be adaptive. But evolution can often be quite fast, since it is both massively parallel (all organisms) and serial (once per generation). Although biological systems may initially appear to be well designed, they are locally rather than globally optimized. Their design limitations follow from their earlier history. This point was emphasized by Stephen Jay Gould (1980) in [“The Panda’s Thumb”](#). The thumb of the panda is not a thumb at all, but rather a modified wrist bone. The thumb bones had been eliminated in the earlier evolution of the panda’s ancestors.

Random [mutations](#) produce the variation needed to generate environmental adaptation. Random [“genetic drift”](#) in small populations may sometimes play a role in evolution, although evolution is mostly deterministic and can be modeled mathematically. There is no underlying goal or intent, apart from survival. The primary feature that distinguishes biological systems from physical systems is exponential reproduction, based on the digital code of DNA.

Evolution is driven by random variation and guided only by the environment for a given organism, where the environment includes a multiplicity of other organisms. There has been a move toward increasing complexity in evolution, but this is not complexity for its own sake. Indeed, excessive complexity is not favored by natural selection; like theories of physics, following Einstein’s dictum, life should be as simple as possible, but no simpler. Complex organisms developed only because they are better adapted to compete successfully in complex environments.

If animal bodies have evolved complexity to compete more successfully, the same should be true for animal brains. We are accustomed to comparing the sizes of brains, but brain organization is arguably more critical. Brain structures evolved because they improved environmental fitness and survival of progeny. These include not only those structures associated with sensory input or muscular coordination, but also those associated with intelligence and consciousness. Brain structures involved with intelligence (adaptive learning in complex environments) may be quite widespread in the animal kingdom, particularly among vertebrates. Finally, while human consciousness involves language and abstraction, it may not be qualitatively different from that in “lower” organisms.

IV. Psychology, Cognition, and the Illusion of Agency

We all feel that we can understand the human mind based on our own thinking. We identify ourselves as individual agents who perceive the world around us and are in control of our own actions. Indeed, [Rene Descartes](#) took this as the basis for his philosophy: “[Cogito, ergo sum](#)”, or in English, “I think, therefore I am.” But this makes the mind a distinct object separate from the body and from the physical world, i.e., [mind-body dualism](#). This outlook is embedded in our sense of self and in our language for describing it, but there are good reasons to believe that this is mostly an illusion. Research in psychology and cognitive science has shown that our self-perceptions can be quite deceptive.

Historically, one response to this dualism was to go to the other extreme and deny the existence of an internal mind. This is the basis for [behaviorism](#), which deals with an animal or even a human as a “black box” with inputs and responses. More recent approaches in cognitive science have recognized that there are internal cognitive structures (both conscious and unconscious), and have identified fallacies of thinking that reflect aspects of these internal structures.

This was explored in the psychological research of [Daniel Kahneman](#), as summarized in “[Thinking, Fast and Slow](#)” (2011). This showed by careful experiments that people are not really the rational free agents they think they are. This has been particularly influential in providing the basis for [behavior economics](#). Kahneman identified two distinct systems at work in the human mind: System 1 (“fast thinking”) is the unconscious mind that does things automatically without us having to think about them. System 2 (“slow thinking”) is the conscious mind, which requires deliberate attention and thought. Most of our actions are actually done by System 1, even though we firmly believe that they were rational decisions made by our conscious mind. System 2 operates with a simplified model and a coherent narrative, and when the model appears inconsistent (“[cognitive dissonance](#)”), the perceptions may be altered to maintain a consistent picture.

Kahneman uses the analogy with a Chairman of the Board of a corporation, who thinks that he runs the entire operation. Most of the work is actually done by his underlings, generally without even asking the chairman. The underlings give the chairman credit for successful operation, building up his ego, and he believes it. In this context, the chairman represents the conscious mind, which has an inflated sense of its own importance, and is largely unaware of the existence of the unconscious mind. This is summarized in the block diagram in Fig. 3.




Fig. 3. Block diagram of control in corporation or human mind. In the corporate example, C may represent the Chairman of the Board, U the underlings, and E everyone else. In a model of the mind, C may represent the conscious mind, U the unconscious mind, and E the environment. (N represents random noise.) In both cases, top-down control and agency may be illusory.

It seems that the human mind is preprogrammed to identify agency, both in ourselves and in others (and even in inanimate objects!). These agents are central to a simplified model of the world, which filters all our perceptions. Further, events do not just happen in isolation, but represent a causal sequence, a coherent story in progress, with heroes, villains, and value judgements. The conscious mind operates on this simplified model, although we are conditioned to believe that this is the real world. Simple goals are identified which promote basic urges such as survival, pursuit of pleasure, and avoidance of pain. The brain organization that supports this evolved to enable rapid decisions in complex dynamic environments, with incomplete information. In a fight-or-flight situation, indecisiveness is not adaptive. I am not

suggesting that we are living in “[The Matrix](#)”, but perhaps we are each living in our own Matrix, with only a partial mapping to those of other people and to the real world.

V. Computation, Neural Nets, and the Illusion of Intelligence

Ever since digital computers were first developed in the 1950s, they were commonly thought of as “electronic brains”. But traditional computers are actually quite different from brains, both in structure and capabilities. However, recent research in alternative computer architectures, combined with research on the brain, has shown that “[neuromorphic](#)” computer architectures may finally be emulating brains more closely (see also Monroe 2014).

This difference can be seen in Fig. 4, which compares a traditional [von Neumann computer architecture](#) to a neural network. The von Neumann architecture in Fig. 4L extended a simple mathematical model of a computer proposed by [Alan Turing](#), and was brought to practical fruition under the direction of [John von Neumann](#) at the Institute for Advanced Study in Princeton (see Dyson 2012). This architecture consists of an arithmetic engine with a control program and memory. Virtually all practical computers since then have had the same basic structure.

Compare this to a basic design of an [artificial neural network](#) (or neural net) shown in Fig. 4R. This represents an array of artificial neurons (which may be electronic elements) organized into layers and connected with “synapses”. The strength of a given synapse may be changed according to a procedure of iterative training or learning, rather than an imposed program. The collection of synapse strengths correspond to memory, but are distributed throughout the system, rather than localized in a single module. Furthermore, there is no central processing unit; the processing is also distributed. This structure is similar in several respects to the interconnections between biological neurons, although it is much simpler.

A traditional von Neumann processor works very well in mathematical calculations, whereas brains do mathematics rather poorly. On the other hand, brains have evolved to match patterns, and do this very well. Pattern matching is not limited to image recognition, but occurs in all sensory processing, memory retrieval, language translation, and a host of correlation and optimization problems. Recent research has shown that neural nets with many “hidden layers” between the input and output can be trained to be particularly efficient in learning to match patterns; this is known as “[deep learning](#)”. More hidden layers enable more abstract correlations among inputs, which in turn enables more flexible recognition of a wider variety of complex patterns. The learning process itself is adaptive in a way that is similar to evolution. The environment selects those variations that are most effective in matching the training patterns.




Fig. 4. Computer architectures for (L) classic von Neumann computer and (R) artificial neural network (from Wikipedia).

The field of “[artificial intelligence](#)” is almost as old as computers themselves, but it has long fallen far short of its goals. The traditional method of artificial intelligence is to devise a list of rules about a particular topic, and program them into a conventional computer. But knowledge of a fixed set of rigid rules is not what we generally mean by intelligence. Indeed, biological organisms with very simple nervous systems, such as worms or insects, behave as if they are “hard-wired” with rule-driven behavior, and are not regarded as intelligent at all. Furthermore, sophisticated responses to dynamic environments cannot be fully programmed in advance; there must be a strong element of learning involved. In contrast, neural nets can learn how to [recognize a cat](#), for example, without being told explicitly what defines a cat. The newer “deep learning” approach to artificial intelligence is starting to have a major impact on technology (see, for example, a [recent article](#) on the use of deep learning for Google Translate). In traditional computers, the greatest difficulty is found in writing and debugging the software. In contrast, in both natural intelligence and the newer approaches to machine learning, the software is generated automatically via learning. No programmer is necessary; as with evolution, this is unguided.

VI. Minds, Dreams, and the Illusion of Consciousness

The most persistent illusion associated with human consciousness is that there must be an immaterial spirit. But this is clearly a remnant of pre-modern religious thinking, where everything is driven by immaterial spirits. Efforts to assert that somehow consciousness emerges from brains of a certain scale or complexity (see e.g., [Teilhard de Chardin](#)) are misdirected. Biological vision did not arise simply from scale or complexity; it required evolution of structures with specific functionalities. Why should we expect consciousness to be different? Several philosophers, such as [Daniel Dennett](#) (2015), have been arguing for some time that [consciousness is an illusion](#), but they have been less clear as to the actual structure behind the illusion. There have also been various proposals, mostly not very specific, for [machine consciousness](#) (Manzotti 2013).

Consider a preliminary outline of the structure of a consciousness organ, which might be based on neural networks. While many aspects responsible for consciousness are hidden from view (either external or internal), some key functional requirements should be clear. Consciousness involves a self-identified agent following a continuous, causal narrative in a dynamic environment. The environment must integrate various sensory modalities with relevant emotions and memories. This is shown in Fig. 5T, which shows consciousness as a “[virtual reality](#)” (VR) construct created from filtered input data, and representing a simplified dynamic model of the reality presented to an individual. In addition, this must also be linked to a short-term dynamic memory module, containing the recent past, and a predictive model of one or more near futures, as shown in Fig. 5B. A clock time-stamps the present frame and shifts it to the past, while selecting a possible future. This ensures that perceived time is a central element in consciousness.

For example, a conscious visual representation of a rose is not just a portion of a larger two-dimensional image. Rather, it is an object embedded in a three-dimensional space, which represents the concept of a rose and links to abstract attributes and memories of images and aromas. This may be analogous to a PDF file of a scanned document which is subjected to [optical character recognition](#) (OCR). Such a file contains an image of a given page, but the image of each character is also associated with the functional representation of the character. Now imagine that the document also contains embedded links to definitions of words and other relevant documents, and one starts to have the richness necessary for consciousness. This would require that the consciousness organ represent the central nexus of a network linking the filtered inputs and the control outputs. Indeed, [recent research](#) (Lahav 2016) has found direct evidence for a network of this type in the cerebral cortex, which may represent such a conscious nexus. In contrast to a static environment of a document, the conscious environment is highly dynamic, and is rapidly updated in time.

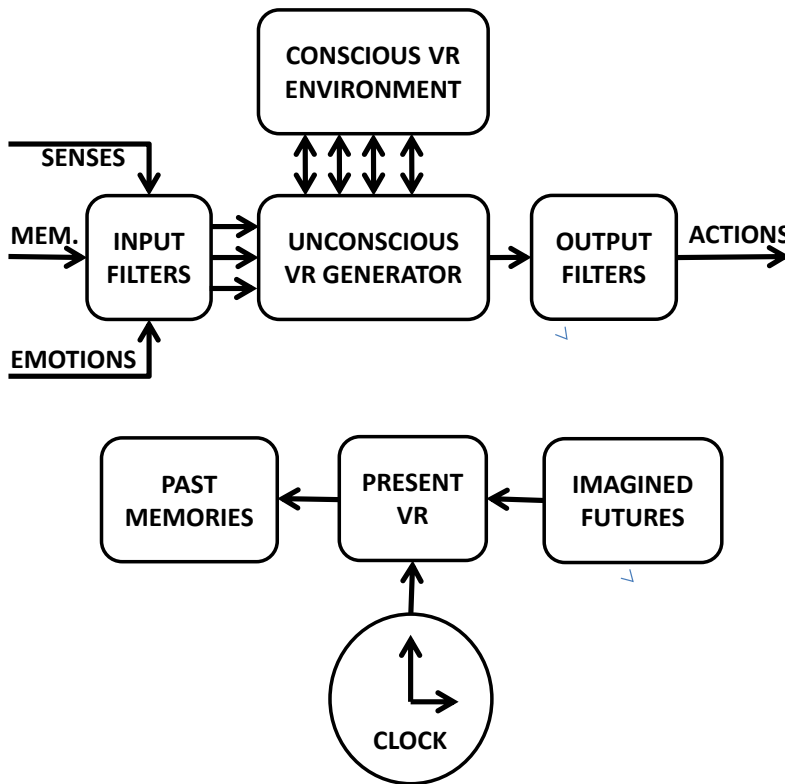


Fig. 5. Suggested organization of autonomous natural or artificial intelligence. (T) Unconscious mind integrates pre-filtered senses, emotions, and memories, and generates VR environment with active agent. (B) Present VR frame is time-stamped and shifted into the past, and a future VR is selected to overlay the present.

But the most critical aspect to explain is the subjective sense of agency. I would suggest that this is due to an adaptive neural net, primed from infancy to recognize self-agency and causality, and also to recognize external agents. The dynamic environmental model is built around such actors. The sense of self is merely the internal sensation of activation of brain circuits in the consciousness nexus. Furthermore, this is a case of dynamic and adaptive learning; the mind learns to generate and refine a simplified model which maintains effective interaction with the environment.

Perhaps the strongest evidence for a dynamic VR generator occurs during dreams, when input and output are isolated from the real world. Dreams include self-agency and coherent narratives in complex environments. [Allan Hobson](#) (2014) has recently suggested that dreams may represent adaptive training and optimization of the consciousness engine, whereby memories are consolidated or forgotten, as appropriate. This may be in accord with other [recent research](#) that has shown that brain interconnections grow dramatically during daytime activity but are selectively pruned back at night.

An important aspect of investigating consciousness is how we can identify it in either artificial or biological systems. Consciousness may be an evolved adaptive structure that is widespread among animals, even without language and other abstractions. Such a conscious structure may never have been present in any artificial intelligence system to date, but it could be constructed in the near future. How can we tell? Alan Turing proposed the “[Turing Test](#)” (1950) for artificial intelligence based on asking a series of questions, but a non-verbal method of exploring consciousness might be preferable. Science fiction writer Philip K. Dick wrote “[Do Androids Dream of Electric Sheep?](#)”, which later became the basis for the film “[Blade Runner](#).” Perhaps we will know that we have made conscious machines when we can observe them dream.

VII. Conclusions

The stated question in this essay contest, “*How can mindless mathematical laws give rise to aims and intention?*,” implicitly assumes that human behavior should be ultimately derivable from particle physics. This is entirely wrong, on

several levels. The problem is that there are several illusions implicit in our thinking that impede our understanding, illusions of control, design, agency, intelligence, and consciousness. The conscious mind is an evolved, adaptive brain structure that hides more than it reveals.

A second point is that the paradigm of natural selection is central not only to biology, but to psychology as well. Neural networks are capable of learning and adaptation to complex environments, and the conscious mind represents a simplified dynamic model of the environment. Goals and intentions are abstract representations of adaptive programs that can promote individual well-being and success. Human behavior can be predictable and causal, subject to statistical variation. There is no ghost in the machine, and no paradox.

Third, true artificial intelligence and consciousness can be emulated in a properly designed electronic system, probably within about 20 years. Such an autonomous system can distinguish self, other agents, and objects, and can create a simple causal narrative of changes in its environment. Further, most of our food animals (except perhaps shellfish) may also be conscious. Together, these will profoundly impact our sense of what it means to be human. We evolved in small family bands of hunter-gatherers; our conscious minds may be too simple to deal with a complex modern world. Some sort of augmented artificial intelligence may be required for humans to survive to the next century. Let us hope that this can incorporate more of the positive human attributes of rationality and cooperation, rather than the negative human attributes of war and xenophobia. Notwithstanding micro-determinism, the future of humanity is still undetermined. Only time will tell.

References

- Daniel C. **Dennett**, 2015, “Why and How Does Consciousness Seem the Way It Seems?”, in *Open Mind*, ed. T. Metzinger, available [online](#). See also [TED Video 2003](#), “The Illusion of Consciousness”.
- Philip K. **Dick**, 1968, *Do Androids Dream of Electric Sheep?*, Doubleday, New York. See also [Wikipedia entry](#).
- George **Dyson**, 2012, *Turing’s Cathedral: The Origins of the Digital Universe*, Random House, NY. See also [book review](#).
- Albert **Einstein**, Quotation about Simplicity, see [online](#). This seems to be a 1950 paraphrase of a more complicated statement by Einstein from 1934 – see [here](#) for history.
- Stephen Jay **Gould**, 1980, *The Panda’s Thumb: More Reflections in Natural History*, Norton, New York. See also [Wikipedia entry](#).
- Allan **Hobson**, et al., 2014, “Virtual Reality & Consciousness Inference in Dreaming”, *Frontiers in Psych.*, available [online](#).
- Alan M. **Kadin** and Steven B. Kaplan, 2016a, “Proposed Experiments to Test the Foundations of Quantum Computing,” submitted to *International Conference on Rebooting Computing*, available [online](#).
- Alan M. **Kadin**, 2016b, “Non-Metric Microscopic Formulation of Relativity”, submitted to [Gravity Research Foundation](#), available [online](#).
- Alan M. **Kadin**, 2015, “Remove the Blinders: How Mathematics Distorted the Development of Quantum theory”, submitted to FQXi essay contest, available [online](#).
- Alan M. **Kadin**, 2013, “Watching the Clock: Quantum Rotation and Relative Time”, submitted to FQXi essay contest, available [online](#).
- Alan M. **Kadin**, 2012, “The Rise and Fall of Wave-Particle Duality”, submitted to FQXi essay contest, available [online](#).
- Daniel **Kahneman**, 2011, *Thinking, Fast and Slow*, Farrar Strauss & Giroux, New York. See also [Wikipedia entry](#).
- Arthur **Koestler**, 1967, *The Ghost in the Machine*, Macmillan, New York. See also [Wikipedia entry](#).
- Nir **Lahav**, et al., 2016, “K-Shell Decomposition Reveals Hierarchical Cortical Organization of the Human Brain”, *New Journal of Physics*, available [online](#), including embedded video.
- Gideon **Lewis-Krauss**, 2016, “The Great AI Awakening”, in *New York Times Magazine*, Dec. 14, 2016, available [online](#).
- Riccardo **Manzotti**, 2013, “Machine Consciousness: A Modern Approach”, in *Natural Intelligence: The INNS Magazine*, vol. 2, no. 1, p. 7, July 2013, available [online](#).
- John **Markham**, 2012, “How Many Computers to Identify a Cat? 16,000”, *New York Times*, available [online](#).
- Don **Monroe**, 2014, “Neuromorphic Computing Gets Ready for the (Really) Big Time,” *Proc. ACM*, available [online](#).
- Carl **Sagan**, 1990, “Cosmos: A Personal Voyage”, *PBS Television Series*, 1990 Update, source of quote on “extraordinary claims”. See also [Wikiquote entry](#).
- Alan **Turing**, 1950, “Computing Machinery and Intelligence,” *Mind*, vol. 256, p. 433, available [online](#).
- Christopher **Wanjek**, 2017, “Sleep Shrinks the Brain – and That’s a Good Thing”, *Scientific American*, available [online](#).
- Steven **Weinberg**, 2017, “The Trouble with Quantum Mechanics,” *New York Review of Books*, available [online](#).

End Notes:

I have argued here that direction and intelligence in nature follow from adaptation in a complex dynamic world based on the foundations of classical physics, without the need for anything either quantum or supernatural. Elsewhere, I have argued that the foundations of quantum mechanics have been profoundly misunderstood, and that quantum mechanics properly provides a unified foundation for all of physics. The comments below describe some implications of this for the nature of time and for the future of computing.

A. Quantum Waves and the Nature of Time

Time is central to physics, biology, psychology, and computation. But is it the relativistic spacetime of Einstein, or the subjective time of our internal chronometers? I have suggested (Kadin 2016b), and in an earlier FQXi essay (Kadin 2013), that time is defined on the microscopic level by quantum waves, and everything else follows from that.

In the Newtonian clockwork universe, time is abstract and universal. In Einstein's universe, time becomes part of abstract [spacetime](#), which is non-universal and inhomogeneous. Einstein focused on light waves, any one of which has a frequency f and a wavelength λ , and a universal speed of light $c = f\lambda$. But taken collectively, light waves (electromagnetic or EM waves) can have any value of f ; there is no characteristic frequency for EM waves. But consider a quantum wave (or [de Broglie wave](#)) for an electron of mass m_e . This has a minimum frequency $f = m_e c^2/h$, which represents a characteristic frequency f_c ($h = \text{Planck's constant}$). Similarly, an electron has a characteristic length, the Compton wavelength $\lambda_c = c/f_c = h/m_e c$. So an electron is a natural clock and ruler which defines the local calibration of time and space. These change in a gravitational field, giving rise to the trajectories of general relativity. These trajectories may be computed using classical formalisms, without reference to abstract spacetime metrics.

Any system of physical units has three independent units (plus the fundamental electric charge e); for the standard metric system ([SI](#)) these are MKS, for meter, kilogram, and second. In the universal system described above, the natural units are f_c , λ_c , and h . Quantized spin provides the basis for h , and is Lorentz-invariant. Other parameters are defined in terms of these: $c = f_c \lambda_c$ and $m_e = hf_c/c^2$.

Within this picture, an electron is a real extended wave in real space, with quantized total spin. There is no intrinsic uncertainty, and no entanglement. Composites such as a proton or an atom are not quantum waves at all; they are bags of confined quantum waves, which inherit quantization from their constituents. Such a picture provides a simple, unified system with local reality all the way from atoms to galaxies and beyond.

Time is not a physical dimension; rather, it is simply a parameter that governs motion and change. Our most accurate clocks are atomic clocks, which inherit their properties from the fundamental electron clock. This defines the time of classical physics, which in turn defines our macroscopic clocks, as well as our biological and psychological clocks. Not all of these are quite so accurate, but they derive from the same physics. The direction of time follows simply from increasing microscopic entropy and macroscopic adaptation.

B. Quantum Computing is *not* the Future of Computing

There has recently been much attention to the potential of quantum computing, which promises the ability to perform computational tasks that are virtually impossible for any classical computer, regardless of the architecture. It has even been suggested that modern classical supercomputers will soon become obsolete due to ["quantum supremacy"](#). This essay proposes that classical computers with non-classical (neuromorphic) architectures may enable true artificial

intelligence and even consciousness -- why not quantum computers for this role? The reason is that I do not believe that the promised quantum computing is possible, on fundamental grounds (see Kadin 2016a).

The key point is that the promised exponential enhancement in performance is based on the presence of quantum superposition and entanglement among N coupled quantum bits (or qubits), as implied by the orthodox Hilbert-space mathematical formalism for quantum states. I have questioned whether this mathematical formalism is correct (Kadin 2012, 2015), and suggested some new experiments that can address this question. One can see the essential role of entanglement from the following simple argument. In any classical computing architecture, one enhances performance by the use of parallelism in hardware bits. If the hardware consists of N parallel bit slices, it can operate (ideally) N times as fast. In sharp contrast, for a quantum computing system, if there are N entangled qubits, this enables an effective computational parallelism by a factor of 2^N , which increases exponentially. For example, if $N=300$, 2^N is greater than the number of atoms in the known universe. Clearly, no classical supercomputer with merely billions of bit slices can compete. The ability to obtain exponentially scaled equivalent performance from linear growth in hardware provides the entire motivation for quantum computing.

Such a claim of exponential performance is fantastic, in both senses of the word. [Carl Sagan](#) once said (1990) that *“Extraordinary claims require extraordinary evidence”*. In fact, no such evidence exists, and people should be extremely skeptical. The primary reason why this has been accepted is that the orthodox theory of quantum mechanics, obscure and confusing though it may be, has been defended by most of the smartest minds in physics for decades. But early in the 20th century, both Albert Einstein and Erwin Schrödinger questioned the foundations of quantum mechanics. Recently, Steven Weinberg (2017) has also questioned these foundations.

Given the billions of dollars that are now being invested in quantum computing research by governments and corporations, I expect that this question will be settled within 20 years. I am not suggesting that quantum computing research is useless; on the contrary, it provides important insights into the isolation of nanoscale systems from environmental noise, which will be essential in the continued evolution of “classical” computer technology toward the atomic scale. But if I am correct, this will radically disrupt the orthodox understanding of quantum mechanics, leading to the adoption of a new quantum paradigm, with major long-term implications for the future of physics.