

RH: Cognition Everywhere

Cognition Everywhere:

The Rise of the Cognitive Nonconscious and the Costs of Consciousness

N. Katherine Hayles

A massive shift is underway in our intellectual and cultural formations. Many different streams of thought are contributing, coming from diverse intellectual traditions, holding various kinds of commitments, and employing divergent methodologies. The differences notwithstanding, they agree on a central tenet: the importance of nonconscious cognition, its pervasiveness and computational potential, and its ability to pose new kinds of challenges not just to rationality but to consciousness in general, including the experience of selfhood, the power of reason, and the evolutionary costs and systemic blindnesses of consciousness.

The implications for interpretation are profound. Interpretation is deeply linked with questions of meaning; indeed, many dictionaries define interpretation in terms of meaning (and meaning in terms of interpretation). For the cognitive nonconscious, however, meaning has no meaning. As the cognitive nonconscious reaches unprecedented importance in communication technologies, ambient systems, embedded devices, and other technological affordances, interpretation has become deeply entwined with the cognitive nonconscious, opening new avenues for exploring, assessing, debating and resisting possible configurations between interpretive strategies and the cognitive nonconscious. In a later section I will identify sites within the humanities where these

contestations and reconfigurations are most active and comment upon the strategies emerging there. Whatever one makes of these changes, one conclusion seems inescapable: the humanities cannot continue to take the quest for meaning as an unquestioned premise for their ways of doing business. Before we arrive at this point, some ground clearing of terminology is necessary, as well as consideration of how the cognitive nonconscious differs from and interacts with consciousness.

Rethinking the Cognitive Nonconscious

One way into understanding the cognitive nonconscious is through Stanisław Lem's Summa Technologiae,¹ a work that, as far as the Anglophone world is concerned, has been caught in a time-warp. Published in Polish in 1964, Summa was never completely translated into English prior to its present appearance in 2013. Lem presciently understood that our society was facing what he called an "information barrier," a deluge of information that would overwhelm scientific and technological enterprises unless a way was found to automate cognition. He observed that formal languages, such as mathematics and explicit equations, do not deal well with complexity. For example, equations for gravitational interactions do not have explicit solutions when as few as three bodies are involved. Nevertheless, there are many instances when complex problems are solved effortlessly by nonconscious means. For example, when a rabbit chased by a coyote leaps over a chasm, the feat would require many equations and considerable time to solve explicitly, but the animal does it instantly without a single calculation.

Reasoning that translating tasks into formal languages may be unnecessary for solving complex problems, Lem proposes a form of evolutionary computation programmed in natural media, an “information farm” in which systems could successfully perform cognitive modeling functions without consciousness. He suggests modeling a dynamic system by first creating a “diversity generator” such as a fast-running stream carrying along rocks of various sizes. Then, to match the target system’s momentum, one places some kind of barrier or “sieve” that selects only rocks of a specific size and velocity. Other “sieves” select for different variables, and the process continues until the desired match is achieved.² One could imagine expanding this kind of modeling by using living cells, which in carrying out division, excretion, and osmosis employ many different kinds of “sieves” as selection devices. In fact, contemporary experiments by Leonard Adleman show it is possible to use DNA sequences to solve complex topological challenges similar to the traveling salesman problem, another example of how the cognitive nonconscious can be harnessed to arrive at solutions difficult or impossible to achieve by explicit means.

What kind of processes do such systems entail, and what is implied by calling them cognitive? First, these systems operate within evolutionary dynamics, that is, they are subjected to fitness criteria that select certain states out of the diverse range available. Second, they are adaptive; they change their behaviors as a result of fitness challenges such as homeostasis for the cell. Third, they are complex, composed of parts interacting with each other in multiple recursive feedback loops, or what Andy Clark calls “continuous reciprocal causation.”³ Consequently, they exhibit emergence, results that cannot be predicted and that exceed the sum of their parts. Fourth, they are “constraint

driven,” which implies that the individual agents’ behaviors are guided by simple instincts or rules that constrain them to certain productive paths, such as the sieves mentioned above. Generally, they enact the artificial life mantra: “From simple rules to complex patterns or behaviors.” Together, these properties enable such systems to perform modeling and other functions that, if they were performed by a conscious entity, would unquestionably be called cognitive.

To avoid confusion, I will reserve “thinking” for what conscious entities such as humans (and some animals) do, and “cognition” as a broader term that does not necessarily require consciousness but has the effect of performing complex modeling and other informational tasks. On this view, we can say that while all thinking is cognition, not all cognition is thinking. In this respect the cognitive nonconscious is qualitatively different from the unconscious, which communicates with consciousness in a number of ways.⁴ Accordingly, I will call consciousness/unconsciousness “modes of awareness.” By contrast, the cognitive nonconscious operates at a lower level of neuronal organization not accessible to introspection.

Nonconscious cognitive systems are distinct from the processes that generate them because they show an “intention toward” not present in the underlying material processes as such. For example, a termite mound is a complex architectural structure that emerges as a result of pheromone trails laid down by individual termites enacting simple behavioral rules.⁵ It has an “intention toward,” namely the protection and preservation of the colony. Another example is a beehive, an emergent result created when individual bees position themselves at a certain distance from their fellows and, moving in a circle, spit wax. Since the adjacent bees are doing the same, the wax lines press against each

other and form a hexagon, the polygon with the most efficient packing.⁶ The “intention toward” is instantiated in the beehive, the emergent result of individual bees acting as autonomous individual agents, each of which need only perform a few simple behaviors to achieve an effect greater than the sum of the beehive’s parts.

In contrast to the cognitive nonconscious, material processes operating on their own rather than as part of a complex adaptive system do not demonstrate emergence, adaptation, or complexity. For example, a glacier sliding downhill generally lacks adaptive behavior (it cannot choose a shady versus a sunny valley), has negligible emergent capacity, and its path can be calculated precisely if the relevant forces are known. The distinction between material processes and complex systems may not always be so clear cut. Indeed, this framework, positing a tripartite structure of conscious thinking, nonconscious cognition, and material processes, catalyzes boundary questions about the delineations between categories as active sites for interpretation and debate.

Enlarged beyond its traditional identification with thought, cognition in some instances may be located in the system rather than an individual participant, an important change from a model of cognition centered in the self. As a general concept, the term “cognitive nonconscious” does not specify whether the cognition occurs inside the mental world of the participant, between participants, or within the system as a whole. It may operate wholly independently from consciousness, as in the cases of bees and termites, or it may be part of the larger system such as a human, where it mediates between material processes and the emergence of consciousness/unconsciousness. Alternatively, it may be instantiated in a technological device such as a computer. Nonconscious cognition, then,

operates across and within the full spectrum of cognitive agents: humans, animals, and technical devices.

The Costs of Consciousness

Along with an expanded sense of cognition come reassessments of consciousness, the purposes it serves, and the costs it entails. Most researchers recognize (at least) two levels of consciousness, a lower level called core or primary consciousness, and a higher level called extended or secondary consciousness. Humans share core consciousness with other primates and (arguably) a wide range of mammals and other animals as well. According to Thomas Metzinger⁷, a contemporary German philosopher, core consciousness creates a mental model of itself that he calls a “phenomenal self-model” (PSM) (107); it also creates a model of its relations to others, the “phenomenal model of the intentionality relation” (PMIR) (301-5). Neither of these models could exist without consciousness, since they require the memory of past events and the anticipation of future ones. From these models, the experience of a self arises, the feeling of an “I” that persists through time and has a more or less continuous identity. The PMIR allows the self to operate contextually with others with whom it constructs an intentionality relation.

The sense of self, Metzinger argues, is an illusion, facilitated by the fact that the construction of the PSM and the PMIR models are transparent to the self (that is, the self does not recognize them as models but takes them as actually existing entities). This leads Metzinger to conclude, “nobody ever was or had a self” (1). In effect, by positioning the self as epiphenomenal, he reduces the phenomenal experience of self back to the underlying material processes. Philosopher of consciousness Owen Flanagan,

following William James, tracks a similar line of reasoning: “the self is a construct, a model, a product of conceiving of our organically connected mental life in a certain way.”⁸ Who thinks the thoughts that we associate with the self? According to Flanagan (and James), the thoughts think themselves, each carrying along with it the memories, feelings, and conclusions of its predecessor while bearing them toward its successor.

Antonio Damasio holds a somewhat similar view, in the sense that he considers the self to be a construct created through experiences, emotions, and feelings a child has as she grows rather than an essential attribute or possession. Damasio, however, also thinks that the self (illusion though it may be) evolved because it has a functional purpose, namely to create a concern for preservation and well-being that propels the organism into action and thus guarantees “that proper attention is paid to the matters of individual life.”⁹ Owen Flanagan agrees: consciousness and the sense of self have functions, including serving as a clearinghouse of sorts where past experiences are recalled as memories and future anticipations are generated and compared with memories in order to arrive at projections and outcomes. In Daniel Dennett’s metaphor, consciousness and the working memory it enables constitute the “workspace” where past, present, and future are put together to form meaningful sequences.¹⁰

Meaning, then, can be understood at the level of core consciousness as an emergent result of the relation between the PSM and the PMIR—that is, between the self-model and the models the self constructs of objects which it has an “intention toward.” Damasio puts it more strongly; there is no self without awareness of and engagement with others.¹¹ The self thus requires core consciousness, which constructs the PSM and the PMIR; without consciousness, a self could not exist. In humans (and some animals),

the core self is overlaid with a higher-level consciousness capable of metalevel reasoning, including interrogations of meanings that call for interpretations.

In addition to concern for the self, a crucial role of consciousness, which occurs at both the core and metalevel, is creating and maintaining a coherent picture of the world. As Gerald Edelman and Giulio Torino put it, “Many neuropsychological disorders demonstrate that consciousness can bend or shrink, and at times even split but it does not tolerate breaks of coherence.”¹² We can easily see how this quality would have adaptive advantages. Creating coherence enables the self to model causal interactions reliably, make reasonable anticipations, and smooth over the gaps and breaks that phenomenal experiences present. If a car is momentarily hidden by a truck and then reappears, consciousness recognizes this as the same car, often at a level below focused attention. This very quality, however, also frequently causes consciousness to misrepresent anomalous or strange situations.

A number of experiments in cognitive psychology confirm this fact. In one now-famous situation,¹³ subjects are shown a video of players passing a basketball and are asked to keep track of the passes. In the middle of the scene, someone dressed in a gorilla suit walked across the playing area, but a majority of subjects report that they saw nothing unusual.¹⁴ In another staged situation, a man stops a passerby and asks for directions.¹⁵ While the subject is speaking, two workmen carrying a vertical sheet of wood pass between them, momentarily blocking the view. When they pass, the interlocutor has been replaced by another person, but the majority of subjects do not notice the discrepancy. Useful as is the tendency of consciousness to insist on coherence, these experiments show that one cost is the screening out of highly unusual events.

Without our being aware of it, consciousness edits to make them conform to customary expectations, a function that makes eyewitness testimony notoriously unreliable. Even in the most ordinary circumstances, consciousness confabulates more or less continuously, smoothing out the world to fit our expectations and screening from us the world's capacity for infinite surprise.

A second cost is the fact that consciousness is slow relative to perception. Experiments by Benjamin Libet and colleagues show that before subjects indicate that they have decided to raise their arms, the muscle action has already started.¹⁶ Although Daniel Dennett is critical of Libet's experimental design, he agrees that consciousness is belated, behind perception by several hundred milliseconds, the so-called missing half-second.¹⁷ This cost, although negligible in many contexts, assumes new importance when cognitive nonconscious technical devices can operate at temporal regimes inaccessible to humans and exploit the missing half second to their advantage.

Finally there are the costs, difficult to calculate, of possessing a self aware of itself and tending to make that self the primary actor in every scene. Damasio comments that "consciousness, as currently designed, constrains the world of imagination to be first and foremost about the individual, about an individual organism, about the self in the broad sense of the term."¹⁸ The anthropocentric bias for which humans are notorious would not be possible, at least in the same sense, without consciousness and the impression of a reified self that consciousness creates. The same faculty that makes us aware of ourselves as selves also partially blinds us to the complexity of the biological, social, and technological systems in which we are embedded, tending to make us think we are the most important actors and that we can control the consequences of our actions

and those of other agents. As we are discovering, from climate change to ocean acidification to greenhouse effects, this is far from the case.

Neural Correlates to Consciousness and the Cognitive Nonconscious

Damasio and Edelman, two eminent neurobiologists, have complementary research projects, Damasio working from brain macrostructures on down, Edelman working from brain neurons on up. Together, their research presents a compelling picture of how core consciousness connects with the cognitive nonconscious. Damasio's work has been especially influential in deciphering how body states are represented in human and primate brains through "somatic markers," indicators emerging from chemical concentrations in the blood and electrical signals in neuronal formations.¹⁹ In a sense, this is an easier problem to solve than how the brain interacts with the outside world, because body states normally fluctuate within a narrow range of parameters consistent with life; if these are exceeded, the organism risks illness or death. The markers, sending information to centers in the brain, help initiate events such as emotions—bodily states corresponding to what the markers indicate—and feelings, mental experiences that signal such sensations as feeling hungry, tired, thirsty and frightened.

From the parts of the brain registering these markers emerge what Damasio calls the protoself, "an interconnected and temporarily coherent collection of neural patterns which represent the state of the organism, moment by moment, at multiple levels of the brain" (174). The protoself, Damasio emphasizes, instantiates being but not consciousness or knowledge; it corresponds to what I have been calling the cognitive nonconscious. Its actions may properly be called cognitive in my sense because it has an

“intention toward,” namely the representation of body states. Moreover, it is embedded in highly complex systems that are both adaptive and recursive. When the organism encounters an object, which Damasio refers to as “something-to-be-known,” the object “is also mapped within the brain, in the sensory and motor structures activated by the interaction of the organism with the object” (169). This in turn causes modifications in the maps pertaining to the organism and generates core consciousness, a recursive cycle that can also map the maps in second-order interactions and thereby give rise to extended consciousness. Consciousness in any form only arises, he maintains, “when the object, the organism, and their relation, can be re-represented” (160). Obviously, to be re-represented, they must first have been represented, and this mapping gives rise to and occurs within the protoself. The protoself, then, is the level at which somatic markers are assembled into body maps, thus mediating between consciousness and the underlying material processes of neuronal and chemical signals.

This picture of how consciousness arises finds support in the work of Nobel Prize winner neurologist Gerald M. Edelman and his colleague Giulio Tononi.²⁰ Their analysis suggests that a group of neurons can contribute to the contents of consciousness if and only if it forms a distributed functional cluster of neurons interconnected within themselves and with the thalamocortical system, achieving a high degree of interaction within hundreds of milliseconds. Moreover, the neurons within the cluster must be highly differentiated, leading to high values of complexity (146).

To provide a context for these conclusions, we may briefly review Edelman’s theory of neuronal group selection (TNGS), which he calls “neural Darwinism.”²¹ The basic idea is that functional clusters of neurons flourish and grow if they deal effectively

with relevant sensory inputs; those less efficient tend to dwindle and die out. In addition to the neural clusters, Edelman (like Damasio) proposes that the brain develops maps, for example, clusters of neurons that map input from the retina. Neural groups are connected between themselves through recursive “reentrant connections” (45-50, esp. 45), flows of information from one cluster to another and back through massively parallel connections. The maps are interconnected by similar flows, and maps and clusters are also connected to each other.

To assess the degree of complexity that a functional neuronal cluster possesses, Edelman and Tononi have developed a tool they call the functional cluster index (CI).²² This concept allows a precise measure of the relative strength of causal interactions within elements of the cluster compared to their interactions with other neurons active in the brain. A value of $CI = 1$ means that the neurons in the cluster are as active with other neurons outside the cluster as they are among themselves. Functional clusters contributing to consciousness have values much greater than one, indicating that they are strongly interacting among themselves and only weakly interacting with other neurons active at that time.

From the chaotic storm of firing neurons, the coherence of the clusters mobilize neurons from different parts of the brain to create coherent maps of body states, and these maps coalesce into what Edelman calls “scenes,” which in turn coalesce to create what he calls primary consciousness (in Damasio’s terms, core consciousness). Edelman’s account adds to Damasio’s the neuronal mechanisms and dynamics that constitute a protoself from the underlying neurons and neuronal clusters, as well as the processes by

which scenes are built from maps through recursive interactions between an organism's representations of body states and representations of its relations with objects.

It is worth emphasizing the central role that continuous reciprocal causation plays in both Damasio's and Edelman's accounts. Thirty years ago, Humberto Maturana and Francisco Varela intuited that recursion was central to cognition,²³ a hypothesis now tested and extended through much-improved imaging technologies, microelectrode studies, and other contemporary research practices.

Let us now turn to the processes by which re-representation occurs. Recalling Damasio's strong claim that there is no consciousness without re-representation, re-representation is clearly a major function of the protoself, site of the cognitive nonconscious and the processes that give rise to core and higher consciousness. In his theory of "grounded cognition," Lawrence W. Barsalou in an influential article gives a compelling account of how re-representation occurs in what he calls "simulation," "the re-enactment of perceptual, motor, and introspective states acquired during experience with the world, body, and mind."²⁴

In particular, sensory experiences are simulated when concepts relevant to those experiences are processed and understood. He marshals a host of experimental evidence indicating that such mental re-enactments are integral parts of cognitive processing, including even thoughts pertaining to highly abstract concepts. The theory of grounded cognition "reflects the assumption that cognition is typically grounded in multiple ways, including simulations, situated action, and, on occasion, bodily states" (619). For example, perceiving a cup handle "triggers simulations of both grasping and functional actions," as indicated by fMRI scans (functional magnetic resonance images). The

simulation mechanism is also activated when the subject sees someone else perform an action; “accurately judging the weight of an object lifted by another agent requires simulating the lifting action in one’s own motor and somatosensory systems” (624). In order for a pianist to identify auditory recordings of his own playing, he must “simulate the motor actions underlying it” (624). Perhaps most surprising, such simulations are also necessary to grasp abstract concepts, indicating that thinking is deeply entwined with the recall and reenactment of bodily states and actions. The importance of simulations in higher-level thinking shows that biological systems have evolved mechanisms to re-represent perceptual and bodily states, not only to make them accessible to consciousness but also to support and ground thoughts related to them. These thoughts in turn feed back to affect somatic states. We can now appreciate the emphasis that Damasio places on re-representation, for it serves as an essential part of the communication processes between the protoself and consciousness, and it also invests abstract thought with grounding in somatic states.

Although the reenactment mechanisms differ with present-day computational methods, the idea of recursion is central in artificial media as well. Whereas with biological organisms bodily states provide the basis for higher-level thinking, with artificial media recursion operates along a hierarchy that moves from simple to complex, local individual agents operating according to a few simple rules to global systemic patterns of complexity. For technical devices, an “intention toward” is necessary but not sufficient: a hammer and a finance trading algorithm are both designed with an intention in mind, but only the trading algorithm demonstrates nonconscious cognition. What makes the difference? Nonconscious cognition operates through many of the same

strategies employed by biological organisms, including emergence (call it the “termite strategy”), re-representation (as in grounded cognition), evolutionary dynamics to bootstrap cognition, and a variety of other mechanisms. Some nonconscious cognitive devices have sensors and actuators, so they can interact with their environments and perform actions in the world. Others live in artificial environments structured to judge performances according to predetermined fitness criteria, allowing only the most successful agents to propagate into the next generation. Although the range of technical devices demonstrating nonconscious cognition is too broad to cover here, the next section will give a sense of their range and diversity.

Technical Devices and the Cognitive Nonconscious

Evolutionary computations, so called because they instantiate evolutionary algorithms in a variety of artificial media, have by now been extensively studied. John R. Koza and coauthors have created genetic algorithms to carry out a variety of tasks, for example designing electric circuits.²⁵ Their work demonstrates typical strategies to achieve nonconscious cognition. The seed program generates an array of very simple circuits. The performance of each circuit is tested according to how well it carries out certain tasks. The most successful are selected and “married” to each other (that is, their circuits are combined to create hybrids) which are used to create the next generation, circuits somewhat similar to the parents but with minor variations among the “children.” The most successful of these are again selected and again propagate with minor variations, and so on through hundreds or thousands of generations. Eventually circuits evolve that can achieve what Koza and his colleagues call “human-competitive” results

(1), which they define as designs publishable in a peer-reviewed professional journal or circuits judged worthy of a patent.

Similar techniques have been used with algorithms designed to compose music. Such programs are typically given predefined grammars, but they also can modify these or even create their own grammar. As noted by John A. Maurer in “A Brief History of Algorithmic Composition,”²⁶ a genetic algorithm system created by David Cope, called “Experiments in Musical Intelligence” (EMI), works with a large database of different composition strategies, which it can draw on and/or modify. Creating from scores fed into it, it can also create its own database and compose based on that. Compositions in the style of a number of composers have been created in this fashion, including Bach, Mozart, Brahms, and others. There are also genetic programs that start with a small number of functions such as transposition, note generation, and creating or modifying time values. The program then randomly combines these functions, which are judged according to some fitness value. The most successful are “married,” as with Koza’s genetic algorithms, and produce “children” which are evaluated against fitness values in turn to identify the next pair of parents, and so forth.

Recently programs have been developed that act as critics judging how commercially successful a given composition (or movie) is likely to be. Christopher Steiner discusses the case of Polyphonic HMI,²⁷ a company that developed algorithms to evaluate the likely commercial success of a song. The algorithm works by using Fourier transforms and other mathematical functions to isolate and analyze tempos, melodies, beats, rhythms and so forth, creating a three-dimensional visualization showing how similar the song is to songs that have made it big in the past. Mike McCready, creator of

the Polyphonic HMI, used the program to assess an album by Norah Jones, then an unknown artist, and discovered it had extraordinarily large fitness values. Subsequently, the album went on to sell twenty million copies and won eight Grammy Awards (83).

Other kinds of evolving agents employ learning processes similar to embodied biological organisms, using their experiences in the world as physical beings to learn, draw inferences, achieve simple linguistic skills, and interact with humans. Rodney Brooks' "Cog," a head and torso robot, exemplifies this kind of approach (begun in 1994, Cog was retired in 2003). Brooks advocates what he calls cheap tricks, emergent results caused by the interactions of different systems within the robot, often giving the appearance of human-level intelligence without, however, possessing any conscious awareness.²⁸ Another version of a language-learning device is Tom Mitchell's "NELL" (Never-Ending Language Learning), a program that scans "wild" (i.e., unstructured) text on the internet 24/7 and draws inferences from it with a minimum of human supervision (<http://www.cmu.edu/homepage/computing/2010/fall/nell-computer-that-learns.shtml>).²⁹ Less exotic everyday software demonstrating some of the same properties are programs that draw inferences from databanks about a user's preferences, for example programs used by Amazon to make suggestions for future purchases ("We think you might like . . .").

In the financial markets, automated trading algorithms, which now account for about 70 percent of all trades, also operate in highly competitive ecologies.³⁰ The faster algorithms can detect pending orders from their slower competitors and "front run" their orders, for example by purchasing a desired stock at a lower price and then, within milliseconds, turning around and offering it at a slightly higher price, which the slower

algorithm now has no choice but to buy at the new price. Such algorithms typically have several trading strategies from which to choose, and they will opt for the one that yields the best final result. In addition, capitalizing on market regulations, they also use their temporal advantages to force other algorithms to be charged fees while raking in rebates for themselves (the so-called maker and taker fees and rebates).

The example of trading algorithms demonstrates that, when nonconscious cognitive devices penetrate far enough into human systems, they can potentially change the dynamics of human behaviors. As Neil Johnson and his collaborators at Nanex argue (a firm specializing in studying the behaviors of automated trading algorithms), the effect of automated trading algorithms has transformed the stock market from a mixed human-machine ecology to a machine-machine ecology.³¹ As algorithms account for more and more trades, the major exchanges (now for-profit corporations themselves) shape their practices accordingly, for example by offering to sell at a premium rack space next to their servers, thereby shaving milliseconds off the transmission time, a temporal interval in which further financial advantages can be gained. The exchanges have also multiplied the kinds of bids that can be submitted, giving algorithms more ways to turn milliseconds into megadollars.

The ways in which built environments affect human cognition have of course been extensively studied in architecture, geography, economics, political science, group psychology, and a host of other fields. It is scarcely news that humans are affected not only by social exchanges with each other but also by their interactions with their environments. What is (relatively) new is the extent to which the built environment instantiates nonconscious cognition; as the number of such devices grows, so do their

effects on human systems. Moreover, the effects are not merely cumulative but exponential, for increasingly devices operate not just singly but ecologically in niches and groups.

This tectonic shift greatly magnifies the effect of the technical cognitive nonconscious on human systems with which it interacts. The general trend is for more and more communication to flow among intelligent devices, and relatively less among devices and humans. In part this is because of the slow speed at which humans can process information relative to devices, and in part because the population of devices is growing much faster than the population of humans. The internet company Cisco estimates that by 2015, there will 24 billion intelligent devices connected to the internet; by contrast, the present human population of the planet is estimated at 7.1 billion. Compared to the rate at which the human use of the internet is growing, the rate at which intelligent devices are joining the internet is orders of magnitude higher.

As an example, consider the smart house, where the lighting system connects with the heating system which connects with the entry/exit system and so on.³² Because these systems are aware of what the others are doing, they achieve a degree of coordination that has qualitatively different effects on the human occupants than if each was separate. Another example is the self-driving car, now in development, that has sensors and actuators capable of monitoring the environment and reacting according. Moreover, this capability catalyzes the development of smart roads that can communicate directly with the car systems. Just as human cognition is massively affected by sociality, so the nonconscious cognition of intelligent devices operates in different ways when devices connect and communicate with one another.

Interactions between Humans and the Cognitive Nonconscious of Intelligent Devices

Because computational media operate in microtemporal regimes inaccessible to humans, some cultural critics are concerned that the “missing half-second” between perception and conscious awareness may be exploited for capitalistic purposes. A grand chess master takes about 650 milliseconds to recognize he is in checkmate; most people’s responses, less finely tuned, require about a second or more for perceptions to register in consciousness. By contrast, computer algorithms (in automated stock trading, for example) can operate in the one to five millisecond range, about three orders of magnitude faster than humans. One of the ways in which the cognitive nonconscious is affecting human systems, then, is opening up temporal regimes in which the costs of consciousness become more apparent and more systemically exploitable.

Luciana Parisi and Steve Goodman sketch these consequences in their discussion of “affective capitalism.”³³ “Affective capitalism is a parasite on the feelings, movements, and becomings of bodies, tapping into their virtuality by investing preemptively in futurity. Possessed by seductive brand entities you flip into autopilot, are abducted from the present, are carried off by an array of prehensions outside chronological time into a past not lived, a future not sensed. We term this mode of affective programming “mnemonic control,” a deployment of power that exceeds current formulations of biopower” (164). In terms I have been using, computational media can address the protoself at time scales below those at which conscious/unconscious modes of awareness operate, so that by the time they process the protoself’s input, they are already preconditioned to pay more attention to one consumer brand than to another. The effects

are similar to “subliminal advertising” in the 1950s and 1960s, but now, through the rapid development of computational media, are operating at temporalities, sensory modalities, and diverse environmental inputs that would have been unimaginable half a century ago. Mark B. N. Hansen’s forthcoming book *Feed Forward* addresses in depth the implications of these temporal effects of twenty-first century media.³⁴

Of course, not all uses of the cognitive nonconscious are exploitive or capitalistic in their orientations. Often nonconscious cognitive devices are designed to enhance productivity, open new avenues for research, and increase safety and well-being for humans immersed in or affected by them, for example in the computational media essential to the operations of major airports, where they increase safety as well as throughput. In the case of computational media involved in the digital humanities, faster processing speeds allow questions to be posed that simply could not have been asked or answered using human cognition alone. As the digital humanities increasingly penetrate the traditional humanities, misunderstandings of what computational media can and cannot do abound, especially among scholars who have made little or no use of computational media in their own research other than email and internet searches.

The spectrum of humanistic practices altered by the engagement with computational media is too vast to be adequately discussed here, so I will focus on one aspect of special interest in this journal issue: the interplay between description and interpretation. Sharon Marcus, answering critics who contest her and coauthor Stephen Best’s call for “surface reading,” takes on the charge that “pure” description is impossible, because every description already implicitly assumes an interpretive viewpoint determining what details are noticed, how they are arranged and narrated, and

what frameworks account for them. Rather than arguing this is not the case, Marcus turns the tables by pointing out that every interpretation necessitates description, at least to the extent that descriptive details support, extend, and help to position the interpretation.³⁵ Although not the conclusion she draws, her argument can be taken to imply that description and interpretation are recursively embedded in one another, description leading to interpretation, interpretation highlighting certain details over others. Rather than being rivals of one another, then, on this view interpretation and description are mutually supportive and entwined processes.

This helps to clarify the relation of the digital humanities to traditional modes of understanding such as close reading and symptomatic interpretation. Many print-based scholars see algorithmic analyses as rivals to how literary analysis has traditionally been performed, arguing that digital-humanities algorithms are nothing more than glorified calculating machines. But this implication misunderstands how algorithms function. Broadly speaking, an algorithmic analysis can be either confirmatory or exploratory. For confirmatory projects the goal is not to determine, for example, what literary drama falls into what generic category, but rather to make explicit the factors characterizing one kind of dramatic structure rather than another. Often new kinds of correlations appear that raise questions about traditional criteria for genres, stimulating the search for explanations about why these correlations pertain. When an algorithmic analysis is exploratory, it seeks to identify patterns not previously detected by human reading, either because the corpora is too vast to be read in its entirety, or because long-held presuppositions constrain too narrowly the range of possibilities considered.

One might suppose that algorithmic analyses are primarily descriptive rather than interpretative, because they typically produce data about what the subject texts contain rather than what the data mean. However, just as interpretation and description are entwined for human readers (as Marcus's argument implies), so interpretation enters into algorithmic analyses at several points. First, one must make some initial assumptions in order to program the algorithms appropriately. In the case of Tom Mitchell's Never-Ending Language Learning project at Carnegie Mellon, mentioned above, the research team first constructs ontologies to categorize words into grammatical categories. In Timothy Lenoir and Eric Gianella's algorithms designed to detect the emergence of new technology platforms by analyzing patent applications, they reject ontologies in favor of determining which patent applications cite the same references.³⁶ The assumption here is that co-citations will form a network of similar endeavors, and will lead to the identification of emerging platforms. Whatever the project, the algorithms reflect initial interpretive assumptions about what kind of data is likely to reveal interesting patterns. Stanley Fish to the contrary, there are no "all-purpose" algorithms that will work in every case.³⁷

Second, interpretation strongly comes into play when data are collected from the algorithmic analysis. When Matthew Jockers found that Gothic literary texts have an unusually high percentage of definite articles in their titles, for example, their interpretation suggested this was so because of the prevalence of place names in the titles (The Castle of Otranto, for example).³⁸ Such conclusions often lead to the choice of algorithms for the next stage, which are interpreted in turn, and so forth in recursive cycles.

Employing algorithmic analyses thus follows a similar pattern to human description/interpretation, with the advantage that the nonconscious cognition operates without the biases inherent in consciousness, where presuppositions can cause some evidence to be ignored or underemphasized in favor of other evidence more in accord with the reader's own presuppositions. To take advantage of this difference, part of the art of constructing algorithmic analyses is to keep the number of starting assumptions small, or at least to keep them as independent as possible of the kinds of results that might emerge. The important distinction with digital humanities projects, then, is not so much between description versus interpretation but rather the capabilities and costs of human reading versus the advantages and limitations of nonconscious cognition. Working together in recursive cycles, conscious analysis and nonconscious cognition can expand the range and significance of insights beyond what either alone can accomplish.

Staging the Cognitive Nonconscious in the Theater of Consciousness

If my hypothesis is correct about the growing importance of the cognitive nonconscious, we should be able to detect its influence in contemporary literature and other creative works. Of course, since these products emerge from conscious/unconscious modes of awareness, what will be reflected is not the cognitive nonconscious in itself, but rather its restaging within the theater of consciousness. One of the sites where this staging is readily apparent is in contemporary conceptual poetics. Consider, for example, Kenneth Goldsmith's "uncreative writing." In Day, Goldsmith re-typed an entire day (September 1, 2000) of the New York Times; in Fidget, he recorded every bodily movement for a day; in Soliloquy, every word he spoke for a week

(but not those spoken to him); and in Traffic, traffic reports, recorded every ten minutes over an unnamed holiday, from a New York radio station. His work, and his accompanying manifestos, have initiated a vigorous debate about the work's value. Who, for example, would want to read Day? Apparently not even Goldsmith himself, who professed to type it mechanically, scarcely even looking at the page he was copying. He often speaks of himself as mechanistic,³⁹ and as the "most boring writer who ever lived."⁴⁰ In his list of favored methodologies, the parallel with database technologies is unmistakable, as he mentions "information Management, word processing, databasing, and extreme process . . . Obsessive archiving & cataloging, the debased language of media & advertising; language more concerned with quantity than quality."⁴¹ Of course we might, as Marjorie Perloff does, insist there is more at work here than mere copying.⁴² Still, the author's own design seems to commit him to enacting something as close to Stanley Fish's idea of algorithmic processing as humanly possible—rote calculation, mindless copying, mechanical repetition. It seems, in other words, that Goldsmith is determined to stage nonconscious cognition as taking over and usurping consciousness, perhaps simultaneously with a sly intrusion of conscious design that a reader can notice only with some effort. That he calls the result "poetry" is all the more provocative, as if the genre most associated with crafted language and the pure overflow of emotion has suddenly turned the neural hierarchy upside down. The irony, of course, is that the cognitive nonconscious is itself becoming more diverse, sophisticated, and cognitively capable. Ultimately what is mimed here is not the actual cognitive nonconscious but a parodic version that pulls two double-crosses at once, at both ends of the neuronal spectrum: consciousness performed as if it was nonconscious, and the nonconscious

performed according to criteria selected by consciousness. As Perloff notes, quoting John Cage, “If something is boring after two minutes, try it for four. If still boring, try it for eight, sixteen, thirty-two, and so on. Eventually one discovers that it’s not boring at all but very interesting” (157). Consciousness wearing a (distorted) mask of the cognitive nonconscious while slyly peeping through to watch the reaction—that’s interesting!

Another example of how the cognitive nonconscious is surfacing in contemporary creative works is Kate Marshall’s project on contemporary novels, which she calls “Novels by Aliens” (focusing on “the nonhuman as a figure, technique and desire,” Marshall shows that narrative viewpoints in a range of contemporary novels exhibit what Fredric Jameson calls the “ever-newer realisms [that] constantly have to be invented to trace new social dynamics.”⁴³ In Colson Whitehead’s *Zone One*, for example, the viewpoint for the Quiet Storm’s highway clearing project involves an overhead, far-away perspective more proper to a high-flying drone than to any human observer. The protagonist, Mark Spitz, collaborates with Quiet Storm in part because he feels “lust to be a viewpoint.”⁴⁴ Although Marshall herself links these literary effects to such philosophical movements as speculative realism, it is likely that both speculative realism and literary experiments in nonhuman viewpoints are catalyzed by the expansive pervasiveness of the cognitive nonconscious in the built environments of developed countries. In this view, part of the contemporary turn toward the nonhuman is the realization that an object need not be alive or conscious in order to function as a cognitive agent.

Katherine Hayles 2.6.14 20.22

Deleted: new

Madigan Haley 2.6.14 15.24

Comment [1]: Please provide info for this quote, which doesn't show up in Amazon edition

Reframing Interpretation

Today the humanities stand at a crossroad. On one side the path continues with traditional understandings of interpretation, closely linked with assumptions about humans and their relations to the world as represented in cultural artifacts. Indeed, the majority of interpretive activities within the humanities arguably have to do specifically with the relation of human selves to the world. This construction assumes that humans have selves, that selves are necessary for thinking, and that selves originate in consciousness/unconsciousness. The other path diverges from these assumptions by enlarging the idea of cognition to include nonconscious activities. In this line of reasoning, the cognitive nonconscious also carries on complex acts of interpretation, which syncopate with conscious interpretations in a rich spectrum of possibilities.

What would it mean to say that the cognitive nonconscious interprets? A clue is given by physicist Edward Fredkin, when in a seminar he casually announced, “The meaning of information is given by the processes that interpret it.”⁴⁵ When Claude Shannon first formulated information theory, Warren Weaver declared that it had nothing to do with semantic meaning, for Shannon defined information as a function of probability.⁴⁶ Although very significant results have followed Shannon’s version of information,⁴⁷ for the humanities, a theory totally divorced from meaning has little to contribute. Fredkin’s approach, however, suggests that flows of information occur within contexts, and those contexts frequently offer multiple opportunities for interpretation. One reason that digital technologies have become so pervasive and important is that they are constructed to make interpretive choices as clear-cut as possible (because digital technologies use discrete digital encoding, rather the continuous signals that analogue

technologies use). In many instances, however, ambiguities remain, and substantive choices have to be made. Medical diagnostic systems, automated satellite-imagery identification, ship navigation systems, weather-prediction programs, and a host of other nonconscious cognitive devices interpret ambiguous information to arrive at conclusions that rarely if ever are completely certain. Something of this kind also happens with the protoself in humans. Integrating multiple somatic markers, the protoself too must synthesize conflicting and/or ambiguous information to arrive at interpretations that feed forward into the relevant brain centers, emerging as emotions, feelings, and other kinds of awareness in core and higher consciousness, where further interpretive activities take place.

What advantages and limitations do these two paths offer? The traditional path carries the assumption that interpretation, requiring as it does consciousness and a self, is confined largely if not exclusively to humans (perhaps occasionally extended to some animals). This path reinforces the idea that humans are special, that they are the source of almost all cognition on the planet, and that human viewpoints therefore count the most in determining what the world means. The other path recognizes that cognition is much broader than human thinking and that other animals as well as technical devices cognize and interpret all the time. Moreover, it also implies that these interpretations intersect with and very significantly influence the conscious/unconscious interpretations of humans, which themselves depend on prior integrations and interpretations by the protoself. The search for meaning then becomes a pervasive activity among humans, animals, and technical devices, with many different kinds of agents contributing to a rich ecology of collaborating, reinforcing, contesting and conflicting interpretations.

One of the costs of the traditional path is the isolation of the humanities from the sciences and engineering. If interpretation is an exclusively human activity and if the humanities are mostly about interpretation, then there are few resources within the humanities to understand the complex embeddedness of humans in intelligent environments and in relationships with other species. If, on the contrary, interpretation is understood as pervasive in natural and built environments, the humanities can make important contributions to such fields as architecture, electrical and mechanical engineering, computer science, industrial design, and many other fields. The sophisticated methods that the humanities have developed for analyzing different kinds of interpretations and their ecological relationships with each other then pay rich dividends for other fields and open onto any number of exciting collaborative projects.

Proceeding down the nontraditional path, in my view much the better choice, requires a shift in conceptual frameworks so extensive that it might as well be called an epistemic break. One of the first moves is to break the equivalence between thought and cognition; another crucial move is to reconceptualize interpretation so that it applies to information flows as well as to questions about the relations of human selves to the world. With the resulting shifts of perspective, many of the misunderstandings about the kinds of interventions the digital humanities are now making in the humanities simply fade away. In closing, I want to emphasize that the issues involved here are much larger than the digital humanities in themselves. Important as they are, focusing only on them distorts what is at stake in talking about “Interpretation and Its Rivals.” The point, as far as I am concerned, is less about methods that seem to be rivals to interpretation—a formulation that assumes “interpretation” and “meaning” are stable categories that can be

adequately discussed as exclusively human activities—than it is about the scope and essence of interpretation in itself.

Duke University

Notes

¹ Stanisław Lem, Summa Technologiae, trans. Joanna Zylińska (Minneapolis: Univ. of Minnesota Press, 2014).

² Manuel De Landa, “The Geology of Morals, a Neomaterialist Interpretation,” 1995, <http://www.t0.or.at/delanda/geology.htm>.

³ Andy Clark, Supersizing the Mind: Embodiment, Action, and Cognitive Extension (London: Oxford Univ. Press, 2008), 28.

⁴ The so-called “new unconscious” is understood as a mental faculty constantly surveying the environment and communicating contextual clues that influence behavior, goals, and priorities; see Ran Hassin, James S. Uleman, and John A. Bargh, eds., The New Unconscious (Oxford: Oxford Univ. Press, 2005).

⁵ Tristram Wyatt, Pheromones and Animal Behavior: Communication by Smell and Taste (Cambridge: Cambridge Univ. Press, 2003).

⁶ J. P. Gullan and P. S. Cranston, The Insects: An Outline of Entomology (New York: Wiley-Blackwell, 2010), 321.

⁷ Thomas Metzinger, Being No One: The Self-Model Theory of Subjectivity (Cambridge, MA: A Bradford Book, 2004), 107-305.

⁸ Owen Flanagan, Consciousness Reconsidered (Cambridge, MA: A Bradford Book, 1993), 177.

⁹ Antonio Damasio, The Feeling of What Happens: Body and Emotion in the

Making of Consciousness (New York: Mariner Books, 2000), 303.

¹⁰ Daniel C. Dennett, Consciousness Explained (New York: Back Bay Books, 1992), 139-70; 256-66.

¹¹ Damasio, The Feeling of What Happens, 194.

¹² Gerald M. Edelman and Giulio Tononi, A Universe of Consciousness: How Matter Becomes Imagination (New York: Basic Books, 2000), 194.

¹³ You may see the video: <http://www.youtube.com/watch?v=vJG698U2Mvo>

¹⁴ Daniel J. Simons and Christopher S. Chabris, The Invisible Gorilla: How Our Intuitions Deceive Us (New York: Harmony Books, 2011), 8; Simons and Chabris, "Gorillas in Our Midst," Perception 28 (1999): 1059-74.

¹⁵ Simons and Chabris, Invisible Gorilla, 59.

¹⁶ Benjamin Libet and Stephen M. Kosslyn, Mind Time: The Temporal Factor in Consciousness (Cambridge MA: Harvard Univ. Press, 2005), 50-55.

¹⁷ Dennett, Consciousness Explained.

¹⁸ Damasio, The Feeling of What Happens, 300.

¹⁹ Damasio, The Feeling of What Happens.

²⁰ Edelman and Tononi, A Universe of Consciousness.

²¹ Edelman, Neural Darwinism: The Theory of Neuronal Group Selection (New York: Basic Books, 1987).

²² Edelman and Tononi, A Universe of Consciousness, 122-23.

²³ Humberto R. Maturana and Francisco J. Varela, Autopoiesis and Cognition: The Realization of the Living (Dordrecht: D. Reidel Publishing, 1980).

-
- ²⁴ Lawrence W. Barsalou, "Grounded Cognition," Annual Review of Psychology 59 (2008): 618.
- ²⁵ John R. Koza, Martin A. Keane, Matthew J. Streeter, William Mydlowec, Jessen Yu, and Guido Lanza, Genetic Programming IV: Routine Human-Competitive Machine Intelligence (New York and Berlin: Springer 2005).
- ²⁶ John A. Maurer, "A Brief History of Algorithmic Composition," ccrma.stanford.edu/~blackrse/algorithm.html.
- ²⁷ Christopher Steiner, Automate This: How Algorithms Came to Rule Our World (New York: Penguin, 2012), 77-83.
- ²⁸ Rodney A. Brooks and Anita M. Flynn, "Fast, Cheap and Out of Control: A Robot Invasion of the Solar Systems," Journal of the British Interplanetary Society 42 (1989):478-85.
- ²⁹ Matt Gardner, Partha Pratim Talukdar, Bryan Kisiel, and Tom Mitchell, "Improving Learning and Inference in a Large Knowledge-base using Latent Syntactic Cues," Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (2013). <http://rtw.ml.cmu.edu/rtw/publications>. For NELL, see: <http://www.cmu.edu/homepage/computing/2010/fall/nell-computer-that-learns.shtml>
- ³⁰ Scott Patterson, Dark Pools: High-Speed Traders, A. I. Bandits, and the Threat to the Global Financial System (New York: Crown Business, 2012).
- ³¹ Neil Johnson, Guannan Zhao, Eric Hunsader, Hong Qi, Nicholas Joshson, Jing Meng and Brian Tivnan, "Abrupt Rise in New Machine Ecology Beyond Human Response Time," Scientific Reports: Nature Publishing Group 2627 (2013): 1-11.

³² Molly Edmonds and Nathan Chandler, “How Smart Homes Work,” 2008.

<http://home.howstuffworks.com/smart-home.htm>.

³³ Luciana Parisi and Steve Goodman, “Memonic Control,” in Patricia Ticineto Clough and Craig Willse, eds., Beyond Biopolitics: Essays on the Governance of Life and Death (Durham NC: Duke Univ. Press, 2011), 163-76.

³⁴ Mark B. N. Hansen, *Feed Forward* (Chicago: University of Chicago Press, forthcoming fall 2014).

³⁵ These remarks refer to Marcus’s talk at the “Interpretation and Its Rivals Conference” at the University of Virginia in September 2013.

³⁶ Timothy Lenoir and Eric Giannell, “Technology Platforms and Layers of Patent Data,” in Mario Biagioli, Peter Jaszi, and Martha Woodmansee, eds. Unmaking Intellectual Property: Creative Production in Legal and Cultural Perspective (Chicago: Univ. of Chicago Press, 2011), 359-84.

³⁷ Stanley Fish, 2012. “Mind Your P’s and B’s: The Digital Humanities and Interpretation,” New York Times Opinionator 2012.

http://opinionator.blogs.nytimes.com/2012/01/23/mind-your-ps-and-bs-the-digital-humanities-and-interpretation/?_php=true&_type=blogs&_r=0.

³⁸ Scott McLemee, “Crunching Literature,” Inside Higher Ed 2013.

<http://www.insidehighered.com/views/2013/05/01/review-matthew-l-jockers-macroanalysis-digital-methods-literary-history>.

³⁹ Kenneth Goldsmith, “Conceptual Poetics,” 2008,

<http://www.poetryfoundation.org/harriet/2008/06/conceptual-poetics-kenneth-goldsmith/?woo>.

⁴⁰ Quoted in Marjorie Perloff, Unoriginal Genius: Poetry by Other Means in the New Century (Chicago: Univ. of Chicago Press, 2012), 149.

⁴¹ Goldsmith, “Conceptual Poetics.”

⁴² Perloff, Unoriginal Genius, 146-65.

⁴³ [Fredric Jameson, “Realism and Utopia in *The Wire*,” *Criticism* 52.3-4 \(Summer/Fall 2010\): 359-372.](#) Please provide page number for this quotation from Antimonies of Realism.

⁴⁴ Kate Marshall, “The View from Above,” Modern Language Convention, Jan. 14, 2014.

⁴⁵ Quoted in N. Katherine Hayles, “Cybernetics,” in W. J. T. Mitchell and Mark B. N. Hansen, eds., Critical Terms for Media Studies (Chicago: Univ. of Chicago Press, 2010), 145-156.

⁴⁶ Claude E. Shannon and Warren Weaver, The Mathematical Theory of Communication (New Jersey: Bell Labs, 1948).

⁴⁷ James Gleick, The Information: A History, a Theory, a Flood (New York: Vintage, 2012).

Katherine Hayles 2.6.14 20.20

Formatted: Font:Italic

Katherine Hayles 2.6.14 20.21

Formatted: Font:Italic