Unthought

THE POWER OF THE
COGNITIVE NONCONSCIOUS

N. Katherine Hayles

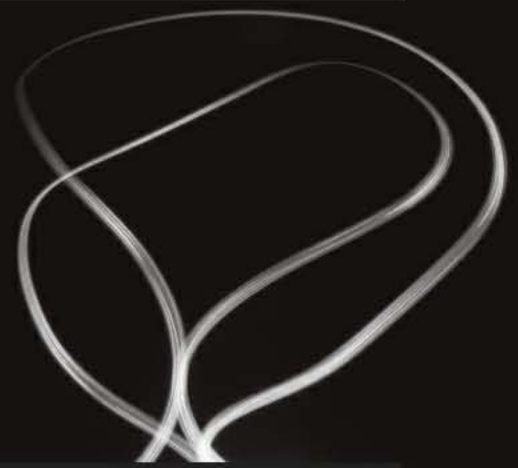

Transforming How We See the World

> When he looked at me with his clear, kind, candid eyes, he looked at me out of a tradition thirteen thousand years old: a way of thought so old, so well established, so integral and coherent as to give a human being the unself-consciousness of a wild animal, a great strange creature who looks straight at you out of his eternal present.

The epigraph, from Ursula Le Guin's science fiction novel *The Left Hand of Darkness*, describes the encounter of protagonist Genly Ai with Faxe, acolyte of the Zen-like cult of the Handdarata and their tradition of "unlearning" (57). "Given to negatives" (57), the Handdarata would immediately recognize "unthought" as indicating a kind of thinking without thinking. There is thought, but before it is unthought: a mode of interacting with the world enmeshed in the "eternal present" that forever eludes the belated grasp of consciousness.

"Unthought" may also be taken to refer to recent discoveries in neuroscience confirming the existence of nonconscious cognitive processes inaccessible to conscious introspection but nevertheless essential for consciousness to function. Understanding the full extent of their power requires a radical rethinking of cognition from the ground up. In addition, because the very existence of nonconscious cognitive processes is largely unknown in the humanities, "unthought" indicates the terra incognita that beckons beyond our received notions of how consciousness operates. Gesturing toward the rich possibilities that open when nonconscious cognition is taken into account, "unthought" also names the potent force of conceptualizing interactions between human and technical systems that enables us to understand more clearly the political, cultural, and ethical stakes of living in contemporary developed societies.

The first step toward actualizing this potential is terminological ground clearing about conscious, unconscious, and nonconscious mental processes.

"Thinking," as used in this book, refers to the thoughts and capabilities associated with higher consciousness such as rationality, the ability to formulate and manipulate abstract concepts, linguistic competencies, and so on. Higher consciousness is not, of course, the whole or indeed even the main part of this story: enhancing and supporting it are the ways in which the embodied subject is embedded and immersed in environments that function as distributed cognitive systems. From a cluttered desktop whose complicated topography acts as an external memory device for its messiness-inclined owner, to the computer on which I am typing this, to the increasingly dense networks of "smart" technologies that are reconfiguring human lives in developed societies, human subjects are no longer contained—or even defined—by the boundaries of their skins.

Part of the book's project is to analyze and explore the nonconscious cognitive assemblages through which these distributed cognitive systems work. In choosing the definite article (*the* cognitive nonconscious), I intend not to reify these systems but rather to indicate their systemic effects. When my focus is on individual subjects, I will use the more processually marked term “nonconscious cognitive processes.” The power of these assemblages, however, is maximized when they function as *systems,* with well-defined interfaces and communication circuits between sensors, actuators, processors, storage media, and distribution networks, and which include human, biological, technical, and material components. In these instances, I will refer to the cognitive nonconscious, a term that crucially includes technical as well as human cognizers. As noted in chapter 5, I prefer “assemblage” over “network” because the configurations in which systems operate are always in transition, constantly adding and dropping components and rearranging connections. For example, when a person turns on her cell phone, she becomes part of a nonconscious cognitive assemblage that includes relay towers and network infrastructures, including switches, fiber optic cables, and/or wireless routers, as well as other components. With the cell phone off, the infrastructure is still in place, but the human subject is no longer a part of that particular cognitive assemblage.

Although nonconscious cognition is not a new concept in cognitive science, neuroscience, and related fields, it has not yet received the attention that I think it deserves. For the humanities, its transformative potential has not yet begun to be grasped, much less explored and discussed. Moreover, even in the sciences, the gap between biological nonconscious cognition and technical nonconscious cognition still yawns as wide as the Grand Canyon on a sunlit morning. One contribution of this study is to propose a definition for cognition that applies to technical systems as well as biological life-forms. At the same time, the definition also excludes material processes such as tsunamis, glaciers, sandstorms, etc. The distinguishing characteristics, as explained in chapter 1, center on interpretation and choice—cognitive activities that both biological life-forms and technical systems enact, but material processes do not. A tsunami, for example, cannot choose to crash against a cliff rather than a crowded beach. The framework I propose, although it recognizes that material processes have awe-inspiring agency, comports neither with vitalism nor panpsychism. Although some respected scholars such as Jane Bennett and Steve Shaviro have given reasons why they find these positions attractive for their purposes, in my view they are not helpful in understanding the specificities of human-technical cognitive assemblages and their power to transform life on the planet.

I see this ongoing transformation as one of the most urgent issues facing us today, with implications that extend into questions about the development of technical autonomous systems and the role that human decision making can and should play in their operation, the environmental devastation resulting from deeply held beliefs that humans are the dominant species on the earth because of their cognitive abilities, and the consequent need for reenvisioning the cognitive capabilities of other life-forms. A correlated development is the spread of computational media into virtually all complex technical systems, along with the pressing need to understand more clearly how their cognitive abilities interact with and interpenetrate human complex systems.

As this framework suggests, another contribution of this study is to formulate the idea of a *planetary cognitive ecology* that includes both human and technical actors and that can appropriately become the focus for ethical inquiry. While traditional ethical inquiries focus on the individual human considered as a subject possessing free will, such perspectives are inadequate to deal with technical devices that operate autonomously, as well as with complex human-technical assemblages in which cognition and decision-making powers are distributed throughout the system. I call the latter cognitive assemblages, and part 2 of this study illustrates how they operate and assesses their implications for our present and future circumstances.

Here is a brief introduction to the book's plan and structure. Part 1 focuses on the concept of nonconscious cognition, with chapter 1 developing a framework for understanding its relation both to consciousness/unconsciousness and material processes. Chapter 2 summarizes the scientific research confirming the existence of nonconscious cognition and locates it in relation to contemporary debates about cognition. Chapter 3 discusses the "new materialisms" and analyzes how these projects can benefit from including nonconscious cognition in their frameworks. As nonconscious cognition is increasingly recognized as a crucial component of human cognitive activity, consciousness has consequently been scrutinized as incurring costs as well as benefits. We can visualize this dynamic as a kind of conceptual seesaw: the higher nonconscious cognition rises in importance and visibility, the lower consciousness declines as the arbiter of human decision making and the dominant human cognitive capability. Chapter 4 illustrates the costs of consciousness through an analysis of two contemporary novels, Tom McCarthy's *Remainder* (2007) and Peter Watts's *Blindsight* (2006).

Part 2 turns to the systemic effects of human-technical cognitive assemblages. Chapter 5 illustrates their dynamics through typical sites ranging from traffic control centers to piloted and autonomous drones. Chapter 6 focuses on autonomous trading algorithms, showing how they require and instantiate technical autonomy because the speeds at which they operate far transcend the temporal regimes of human decision making. This chapter also discusses the implications of these kinds of cognitive assemblages, particularly their systemic effects on destabilizing the global economy. Chapter 7 explores the ethical implications of cognitive assemblages through a close reading of Colson Whitehead's novel *The Intuitionist*. Chapter 8 expounds on the utopian potential of cognitive assemblages and extends the argument to the digital humanities, proposing that they too may be considered as cognitive assemblages and showing how the proposed framework of nonconscious cognition affects how the digital humanities are understood and evaluated.

In conclusion, I want to present a few takeaway ideas that I hope every reader of this book will grasp: most human cognition happens outside of consciousness/unconsciousness; cognition extends through the entire biological spectrum, including animals and plants; technical devices cognize, and in doing so profoundly influence human complex systems; we live in an era when the planetary cognitive ecology is undergoing rapid transformation, urgently requiring us to rethink cognition and reenvision its consequences on a global scale. My hope is that these ideas, which some readers may regard as controversial in part or whole, will nevertheless help to initiate conversations about cognition and its importance for understanding our contemporary situations and moving us toward more sustainable, enduring, and flourishing environments for all living beings and nonhuman others.

Part 1

The Cognitive Nonconscious
and the Costs of Consciousness

Chapter 1

Nonconscious Cognitions: Humans and Others

Rooted in anthropocentric projection, the perception that consciousness and advanced thinking necessarily go together has centuries, if not millennia, of tradition behind it. Recently, however, a broad-based reassessment of the limitations of consciousness has led to a correspondingly broad revision of the functions performed by other cognitive capacities and the critical roles they play in human neurological processes. Consciousness occupies a central position in our thinking not because it is the whole of cognition but because it creates the (sometimes fictitious) narratives that make sense of our lives and support basic assumptions about worldly coherence. Cognition, by contrast, is a much broader capacity that extends far beyond consciousness into other neurological brain processes; it is also pervasive in other life forms and complex technical systems. Although the cognitive capacity that exists beyond consciousness goes by various names, I call it nonconscious cognition.

Perhaps no areas are more rife with terminological disparities than those dealing with consciousness; rather than sort through centuries of confusions, I will try to make clear how I am using the terms and attempt to do so consistently throughout. “Consciousness,” as I use the term, comprises core or primary consciousness (Damasio 2000; Dehaene 2014; Edelman and Tononi 2000), an awareness of self and others shared by humans, many mammals, and some aquatic species such as octopi. In addition, humans and (perhaps) a few primates manifest extended (Damasio 2000) or secondary (Edelman and Tononi 2000) consciousness, associated with symbolic reasoning, abstract thought, verbal language, mathematics, and so forth (Eagleman 2012; Dehaene 2014). Higher consciousness is associated with the autobiographical self (Damasio 2012, 203–07), reinforced through the verbal monologue that plays in our heads as we go about our daily business; that monologue, in turn, is associated with the emergence of a self aware of itself as a self (Nelson, in Fireman, McVay, and Flanagan 2003, 17–36). Recognizing that the cognitive nonconscious (in his terms, the protoself) can create a kind of sensory or nonverbal narrative, Damasio explains how the narratives become more specific when melded with verbal content in higher consciousness. “In brains endowed with abundant memory, language, and reasoning, narratives . . . are enriched and allowed to display even more knowledge, thus producing a well-defined protagonist, the autobiographical self” (Damasio 2012, 204). Whenever verbal narratives are evoked or represented, this is the mental faculty that makes sense of them.[1]

Core consciousness is not sharply distinguished from the so-called “new” unconscious (in my view, not an especially felicitous phrase), a broad environmental scanning that operates below conscious attention (Hassin, Uleman, and Bargh 2005). Suppose, for example, you are driving while thinking about a problem. Suddenly the car in front brakes, and your attention snaps back to the road. The easy and continuous communication between consciousness and the “new” unconscious suggests that they can be grouped together as modes of awareness.[2]

In contrast, nonconscious cognition operates at a level of neuronal processing inaccessible to the modes of awareness but nevertheless performing functions essential to consciousness. The last couple of decades in neuroscientific research show that these include integrating somatic markers into coherent body representations (Damasio 2000), synthesizing sensory inputs so they appear consistent across time and space (Eagleman 2012), processing information much faster than can consciousness (Dehaene 2014), recognizing patterns too complex and subtle for consciousness to discern (Kouider and Dehaene 2007), and drawing inferences that influence behavior and help to determine priorities (Lewicki, Hill, and Czyzewska 1992). Perhaps its most important function is to keep consciousness, with its slow uptake and limited processing ability, from being overwhelmed with the floods of interior and exterior information streaming into the brain every millisecond.

The point of emphasizing nonconscious cognition is not to ignore the achievements of conscious thought, often seen as the defining characteristic of humans, but rather to arrive at a more balanced and accurate view of human cognitive ecology that opens it to comparisons with other biological cognizers on the one hand and on the other to the cognitive capabilities of technical systems. Once we overcome the (mis)perception that humans are the only important or relevant cognizers on the planet, a wealth of new questions, issues, and ethical considerations come into view. To address these, this chapter offers a theoretical framework that integrates consciousness, nonconscious cognition, and material processes into a perspective that enables us to think about the relationships that enmesh biological and technical cognition together.

Although technical cognition is often compared with the operations of consciousness (a view I do not share, as discussed below), the processes performed by human nonconscious cognition form a much closer analogue. Like human nonconscious cognition, technical cognition processes information faster than consciousness, discerns patterns and draws inferences and, for state-aware systems, processes inputs from subsystems that give information on the system's condition and functioning. Moreover, technical cognitions are designed specifically to keep human consciousness from being overwhelmed by massive informational streams so large, complex, and multifaceted that they could never be processed by human brains. These parallels are not accidental. Their emergence represents the exteriorization of cognitive abilities, once resident only in biological organisms, into the world, where they are rapidly transforming the ways in which human cultures interact with broader planetary ecologies. Indeed, biological and technical cognitions are now so deeply entwined that it is more accurate to say they interpenetrate one another.

The title of part 1, the cognitive nonconscious, is meant to gesture toward the systematicity of human-technical interactions. In part 2, I will refer to these as cognitive assemblages. *Assemblage* here should not be understood as merely an amorphous blob. Although open to chance events in some respects, interactions within cognitive assemblages are precisely structured by the sensors, perceptors, actuators, and cognitive processes of the interactors. Because these processes can, on both individual and collective levels, have emergent effects, I will use *nonconscious cognition(s)* to refer to them when the emphasis is on their abilities for fluid mutations and transformations. The more reified formulation indicated by the definite article (*the* cognitive nonconscious) is used when the systematicity of the assemblage is important. I adopt this form for my overall project because the larger implications of cognitive assemblages occur at the systemic rather than individual levels. As a whole, my project aims to chart the transformative perspectives that emerge when nonconscious cognitions are taken fully into account as essential to human experience, biological life, and technical systems.

Although my focus is on biological and technical cognitions that function without conscious awareness, it may be helpful to clarify my position relative to the cognitivist paradigm that sees consciousness operating through formal symbol manipulations, a framework equating the operations of human minds with computers. Clearly humans can abstract from specific situations into formal representations; virtually all of mathematics depends on these operations. I doubt, however, that formal symbol manipulations are generally characteristic of conscious thought. Jean-Pierre Dupuy (2009), in his study arguing that cognitive science developed from cybernetics but crucially transformed its assumptions, characterizes the cognitivist paradigm not as the humanization of the machine (as Norbert Weiner at times wanted to position cybernetics) but as the mechanization of mind: "The computation of the cognitivists . . . is symbolic computation. The semantic objects with which it deals are therefore all at hand: they are the mental representations that are supposed to correspond to those beliefs, desires, and so forth, by means of which we interpret the acts of ourselves and others. Thinking amounts, then, to performing computations on these representations" (Dupuy 2009, 13).

As Dupuy shows, this construction is open to multiple objections. Although cognitivism has been the dominant paradigm within cognitive science throughout the 1990s and into the twenty-first century, it is increasingly coming under pressure to marshal experimental evidence showing that brains actually do perform such computational processes in everyday thought. So far, the results remain scanty, whereas experimental confirmation continues to grow for what Lawrence Barsalou (2008) calls "grounded cognition," cognition supported by and entwined with mental simulations of modal perceptions, including muscle movements, visual stimuli, and acoustic perceptions. In part this is because of the discovery of mirror neuron circuits in human and primate brains (Ramachandran 2012), which, as Miguel Nicolelis (2012) has shown in his work on Brain-Machine-Interfaces (BMI), play crucial roles in enabling humans, primates, and other animals to extrapolate beyond bodily functions such as limb movements into prosthetic extensions.

One aspect of these controversies is whether neuronal processes can in themselves be understood as fundamentally computational. Dissenting from the computationalist view, Walter J. Freeman and Rafael Núñez argue that "action potentials are not binary digits, and neurons do not perform Boolean algebra" (1999, xvi). Eleanor Rosch, in "Reclaiming Concepts" (Núñez and Freeman 1999, 61–78) carefully contrasts the cognitivist paradigm with the embodied/embedded view, arguing that empirical evidence is strongly in favor of the latter. Amodal symbolic manipulation, as Barsalou (2008) characterizes the cognitivist paradigm, depends solely on logical formulations unsupported by the body's rich repertoire of physical actions in the world. As numerous researchers and theorists have shown (Lakoff and Johnson 2003; Dreyfus 1972, 1992; Clark 2008), embodied and embedded actions are crucial in the formation of verbal schema and intellectual comprehension that express themselves through metaphors and abstractions, extending out from the body to sophisticated thoughts about how the world works.

My comparison between nonconscious cognition in biological life-forms and computational media is not meant to suggest, then, that the processes they enact are identical or even largely similar, because those processes take place in very different material and physical contexts. Rather, they perform similar *functions* within complex human and technical systems. Although functionalism has sometimes been used to imply that the actual physical processes do not matter, as long as the results are the same (for example, in behaviorism and some versions of cybernetics), the framework advanced here makes context crucial to nonconscious cognition, including the biological and technical milieu within which cognitions take place. Notwithstanding the profound differences in contexts, nonconscious cognitions in biological organisms and technical systems share certain *structural* and *functional* similarities, specifically in building up layers of interactions from low-level choices, and consequently very simple cognitions, to higher cognitions and interpretations.

Exploring these structural parallels requires a good deal of ground clearing to dispense with lingering questions such as whether machines can think, what distinguishes cognition from consciousness and thought, and how cognition interacts with and differs from material processes. Following from these fundamental questions are further issues regarding the nature of agencies that computational and biological media possess, especially compared with material processes, and the ethical implications when technical cognitive systems act as autonomous actors in cognitive assemblages. What criteria for ethical responsibility are appropriate, for example, when lethal force is executed by a drone or robot warrior acting autonomously? Should it focus on the technical device, the human(s) who set it in motion, or the manufacturer? What perspectives offer frameworks robust enough to accommodate the exponentially expanding systems of technical cognitions and yet nuanced enough to capture their complex interactions with human cultural and social systems?

Asking such questions is like pulling a thread dangling from the bottom of a sweater; the more one pulls, the more the whole fabric of thinking about the significance of biological and computational media begins to unravel. Parts 1 and 2 pull as hard as they can on that thread and try to reweave it into different patterns that reassess the nature of human and technical agencies, realign human and technical cognitions, and investigate how these patterns present new opportunities and challenges for the humanities.

Thinking and Cognition

The first twist in knitting these new patterns is to distinguish between thinking and cognition. Thinking, as I use the term, refers to high-level mental operations such as reasoning abstractly, creating and using verbal languages, constructing mathematical theorems, composing music, and the like, operations associated with higher consciousness. Although Homo sapiens may not be unique in these abilities, humans possess them in greater degree and with more extensive development than other species. Cognition, by contrast, is a much broader faculty present to some degree in all biological life-forms and many technical systems. This vision overlaps with the position that Humberto Maturana and Francisco Varela articulated in their classic work on cognition and autopoiesis (1980). It also aligns with the emerging science of cognitive biology, which views all organisms as engaging in systematic acts of cognition as they interact with their environments. The field, named by Brian C. Goodwin (1977), has subsequently been developed by the Slovakian scientist Ladislav Kováč (2000, hereafter referred to as "FP"; 2007), who has been instrumental in codifying its principles and exploring its implications.

Cognition as formulated in cognitive biology employs some of the same terms as mainstream views but radically alters their import. Traditionally, cognition is associated with human thought; William James, for example, noted that "cognition is a function of consciousness" ([1909] 1975, 13). Moreover, it is often defined as an "act of knowing" that includes "perception and judgment" ("Cognition," in *Encyclopedia Britannica*, www.britannica.com/topic/cognition-thought-process). A very different perspective informs the principles of cognitive biology. Consider, for example, Kováč's observation that even a unicellular organism "must have a certain minimal knowledge of the relevant features of the environment," resulting in a correspondence, "however coarse-grained and abstract," between these features and the molecules of which it is comprised. He concludes, "In general, at all levels of life, not just at the level of nucleic acid molecules, a complexity, which serves a specific function . . . corresponds to an *embodied knowledge*, translated into the constructions of a system. The environment is a rich set of potential niches: each niche is a problem to be solved, to survive in the niche means to solve the problem, and the solution is the embodied knowledge, an algorithm of how to act in order to survive" ("FP," 59). In this view cognition is not limited to humans or organisms with consciousness; it extends to all life-forms, including those lacking central nervous systems, such as plants and microorganisms.

The advantages of this perspective include breaking out of an anthropocentric view of cognition and building bridges across different phyla to construct a comparative view of cognition. As formulated by Pamela Lyon and Jonathan Opie (2007), cognitive biology offers a framework consistent with empirical results: "Mounting evidence suggests that even bacteria grapple with problems long familiar to cognitive scientists, including: integrating information from multiple sensory channels to marshal an effective response to fluctuating conditions; making decisions under conditions of uncertainty; communicating with conspecifics and others (honestly and deceptively); and coordinating collective behavior to increase the chances of survival."[3] Kováč calls the engagement of a life-form with its environment its *onticity,* its ability to survive and endure in changing circumstances. He observes that "life incessantly, at all levels, by millions of species, is 'testing' all the possibilities of how to advance ahead" ("FP," 58). In a playful extension of this reasoning, he imagines a bacterial philosopher confronting the same issues concerning its onticity as a human, asking whether the world exists, and if so, why there is something rather than nothing. Like the human, the bacterium can find no absolute answers within its purview; it nevertheless pursues "its onticity in the world" and accordingly "is already a *subject,* facing the world as an object. At all levels, from the simplest to the most complex, the overall construction of the subject, the embodiment of the achieved knowledge, represents its *epistemic complexity*" ("FP," 59). The sum total of the world's epistemic complexity is continually increasing, according to Kováč, advanced by the testing of what he calls the beliefs of organisms: "only some of the constructions of organisms are embodied knowledge, the others are but *embodied beliefs*. . . . If we take a mutation in a bacterium as a new belief about the environment, we can say that the mutant would sacrifice its life to prove its fidelity to that belief" ("FP," 63). If it continues to survive, that belief becomes converted into embodied knowledge and, as such, is passed along to the next generation.

Comparing traditional and cognitive biology perspectives shows that the same words attain very different meanings. *Knowledge,* in the traditional view, remains almost entirely within the purview of awareness and certainly within the brain. In cognitive biology, on the contrary, it is acquired through interactions with the environment and embodied in the organism's structures and repertoire of behaviors. *Belief* in the traditional view is a position held by a conscious being as a result of experience, ideology, social conditioning, and other factors. In the cognitive biology view, it is a predisposition toward the environment that has not yet been confirmed through ongoing interactions testing its robustness as an evolutionary response to fluctuating conditions. Finally, *subject* in the traditional view is taken to refer to humans or at least conscious beings, while in the cognitive biology view it encompasses all life forms, even humble unicellular organisms.

Plant Signaling and Claims for Plant Intelligence

A convenient site to explore the complex interactions that arise when these perspectives on cognition confront traditional views of intelligence is the world of plants. In a recent *New Yorker* article, Michael Pollan summarizes research that explores homologies between "neurobiology and phytobiology," specifically that plants are "capable of cognition, communication, information processing, computation, learning and memory" (Pollan 2013, 1). The claims are made explicit in a 2006 article in *Trends in Plant Science* (Brenner et al.). Positioned as a review article, the piece is also a polemical manifesto aiming to establish the field of plant neurobiology, arguing that many of the complexities of plant signaling strongly parallel animal neurobiology. As the authors recognize, plant "intelligence" had become a lightning rod for controversy since the 1973 pop science book *The Secret Life of Plants* by Peter Tompkins and Christopher Bird, which made extraordinary claims with little evidence. As a result, many plant scientists wanted to distance themselves as much as possible from claims about plant "intelligence," including the assertion that plants are somehow attuned to human emotional states. Brenner et al. suggest that as a result, many plant biologists refused even to consider parallels between plant responses and animal neurology, practicing "a form of self-censorship in thought, discussion and research that inhibited asking relevant questions" (415).

However justified this comment, the Brenner article itself manifests rhetorical and argumentative strategies that exhibit a deep ambivalence. On the one hand, the authors want to document research showing how complex and nuanced are the mechanisms that underlie individual and communal plant behaviors; on the other, they inadvertently reinstall the privilege of animal intelligence by implying that the more plant signaling resembles animal neurobiology, the stronger the case that it is *really* intelligence. The ambivalence is apparent in the sidebar tracing the etymology of the term "neuron" back to Plato and the Greeks, where "'neuron' means "anything of a fibrous nature" (414). By this definition, plants clearly do have neurons, but in the usual sense of the term (cells with nuclei and axons that communicate using neurotransmitters), they do not. A similar ambivalence is apparent in how they define intelligence; by insisting on the word, they create a rhetorical tension between what they seem to be claiming and what they are actually saying. Offering first a definition of plant intelligence (from Trewavas 2005) as "'adaptively variable growth over the lifetime of a plant'" (414), they expand on it, adding an emphasis on processing information and making decisions: "an intrinsic ability to process information from both abiotic and biotic stimuli that allows optimal decisions about future activities in a given environment" (414).

In my view, this definition offers important clues for reenvisioning cognition (a trajectory I was already following before reading the Brenner article), as well as providing a case study in why it is better to avoid using "intelligence" for nonhuman (and technical) cognitions. As Pollan documents, "many plant scientists have pushed back hard" against what they (mis)understood to be the argument. He notes that thirty-six plant biologists issued a rebuttal to the Brenner piece, also published in *Trends in Plant Science*. The refutation opens with this salvo: "We begin by stating simply that there is no evidence for structures such as neurons, synapses or a brain in plants" (qtd in Pollan, 3). Pollan points out that "no such claim had actually been made—the manifesto had spoken only of 'homologous' structures—but the use of the word 'neurobiology' in the absence of actual neurons was apparently more than the scientists could bear" (3). This rather snide comment (revealing Pollan's own sympathies) does not, in my view, do justice to the complexities of the situation. The issue is not what plant scientists can bear, but how traditional views of intelligence interact with and complicate research that challenges (and perhaps also inadvertently reinstalls) the anthropocentric perspective of what intelligence is. Daniel Chamovitz, for example, while insisting on the remarkable abilities of plants to sense and respond to their environments, argues that "the question . . . should not be whether or not plants are *intelligent*—it will be ages between we all agree on what that term means; the question should be, 'Are plants aware?' and, in fact, they are" (2013, 170). Indeed, Pollan himself points out that "the controversy is less about the remarkable discoveries of recent plant science than about how to interpret and name them; whether behaviors observed in plants which look very much like learning, memory, decision-making and intelligence deserve to be called by those names or whether those words should be reserved exclusively for creatures with brains" (4).

For an analogy, I think of Gillian Beer's brilliant study in *Darwin's Plots: Evolutionary Narrative in Darwin, George Eliot, and Nineteenth-Century Fiction* (1983) tracing the struggle in Darwin's *The Origin of Species* between his view of evolution as a process with no foreordained end and the teleological worldview embedded in the Christian-oriented language he inherited and upon which he instinctively drew. Through a series of close readings, Beer traces in Darwin's metaphors, sentence structures, and rhetorical strategies his desire to articulate a new vision through language saturated with the old. A similar struggle informs the Brenner article; although it is true that the scientists who objected to the article's claims did misread it in a literal sense, they were reacting to the kind of ambivalence noted above between actual evidence and insinuations carried through such tactics as redefining "neuron." In this sense, they accurately discerned the article's double intent to draw upon the cachet of "intelligence" as an anthropocentric value while simultaneously revising the criteria for what constitutes intelligence.

Since plants make up 99 per cent of the planet's biomass, the issue is not trivial across a range of sites, including the question Christopher D. Stone ([1972] 2010) posed decades ago of whether trees should have legal standing. My own clear preference is to create a framework that is both robust and inclusive, and I see no way to exclude plants without sacrificing conceptual coherence (not to mention ignoring the wealth of evidence documenting their remarkable abilities to respond to changing environments).[4]

Nevertheless, assuming that one wanted to draw the line separating cognitive organisms from the noncognitive differently, most aspects crucial to my argument could still be included: the reevaluation of cognition as distinct from consciousness; the recognition that cognitive technologies are now a potent force in our planetary cognitive ecology; and the rapidly escalating complexities created by the interpenetration of cognitive technologies with human systems. These, in my view, are not debatable, while the arguments about plants occupy a less central (although still important) role in my own priorities. I recognize, then, that locating the boundary between the cognitive and noncognitive may be contested, and that different perspectives will lead to conclusions other than those that I endorse. The crucial point for me is less where the line is drawn than that the core issues mentioned above are recognized as critical to our contemporary situation. For me, another important point is the role that humanistic inquiry can play in this arena. Because reenvisioning cognition occurs along a broad interdisciplinary front fraught with linguistic as well as conceptual complexities, the humanities, with their nuanced understanding of rhetoric, argument, and interpretation, are well positioned to contribute to the debate.

I conclude this section with a brief acknowledgement of how complex plant cognition is, where "cognition" here refers to the ways plants sense information from their surroundings, communicate within themselves and to other biota, and respond flexibly and adaptively to their changing environments. Their "'sessile life style'" (Pollan, 4–5—"sessile" refers to organisms attached directly to a substrate, for example, corals and almost all plants) includes more than a dozen senses, among them kin recognition, detection of chemical signals from other plants, and analogues to the five human senses. Pollan explains how kin recognition has been observed to work: "Roots can tell whether nearby roots are self or other, and if other, kin or stranger. Normally, plants compete for root space with strangers, but, when researchers put closely related Great Lakes sea-rocket plants (*cakile edentual*) in the same pot, the plants restrained their usual competitive behaviors and shared resources" (Pollan, 5). It has long been known that plants emit and sense a wide variety of chemical signals; they also manufacture chemicals that deter predators and release others that have psychotropic effects for pollinators, encouraging them to revisit that particular plant again. As researchers continue to investigate the interplays between electrical and chemical signaling, gene structures, and plant behaviors, it becomes increasingly clear that, whatever one's position on the anthropocentrically laden word "intelligence," plants interpret a wide range of information about their environments and respond to challenges in remarkably nuanced and complex ways.

Technical Cognition

Cognitive biology, along with related research in phytobiology discussed above, opens the concept of cognition to a broad compass, and to that extent, it is consistent with the path I want to pursue here. However, these research endeavors miss the opportunity to think beyond the biological to technical cognition, despite redefining terms in ways that partially enable that extension. To illustrate, I turn to the view of cognition proposed by Humberto R. Maturana and Francisco J. Varela in their seminal work *Autopoiesis and Cognition: the Realization of the Living* (1980). Maturana and Varela are distinct from the science of cognitive biology, associated instead with the Chilean School of Biology of Cognition; nevertheless, their views are close enough to cognitive biology to show the modifications necessary to extend cognition to technical systems.

Although they agreed about the cognitive capabilities of living organisms, they disagreed about whether these capabilities could be extended to technical systems—Maturana dissenting, Varela embracing. The disagreement is understandable, for their vision of what constituted cognition made the extension to technical systems far from obvious. In their view, cognition is intimately bound up with the recursive processes whereby an organism's organization determines its structures, and its structures determine its organization, in cycles of what Andy Clark (2008) subsequently called continuous reciprocal causality (note, however, that Maturana and Varela would not have used the term *causality* because an essential part of their vision was the closed or autopoietic nature of the living). Cognition, for them, is nothing other than this informational closure and the recursive dynamics it generates. Their postulated informational closure of organisms makes the extension to technical systems problematic, as technical systems are self-evidently *not* informationally closed but accept information inputs of various kinds and generate information outputs as well. Exploring more fully the cognitive capacities of technical systems, then, requires another definition of cognition than the one they adopted.

In *The Embodied Mind: Cognitive Science and Human Experience* (1991), Varela and coauthors Evan Thompson and Eleanor Rosch extend these ideas into comparisons between the cellular automata (a kind of computer simulation) and the emergence of cognition within biological cells (1991, 150–52). Their definition of enaction is consistent with the approach that I follow, insofar as it recognizes that cognition emerges from context-specific (i.e., embodied) interactions. "We propose as a name the term *enactive* to emphasize the growing conviction that cognition is not the representation of a pregiven world by a pregiven mind but is rather the enactment of a world and a mind on the basis of a history of the variety of actions that a being in the world performs. The enactive approach takes seriously, then, the philosophical critique of the idea that the mind is a mirror of nature but goes further by addressing this issue from within the heartland of science" (1991, 9).

In his later work, Varela was also interested not only in computer simulations but in creating autonomous agents within simulations, an approach known as Artificial Life (Varela and Bourgine 1992). Several years ago pioneers in this field argued that *life* is a theoretical program that can be instantiated in many different kinds of platforms, technological as well as biological (von Neumann 1966; Langton 1995; Rosen 1991). For example, in an effort to show that technical systems could be designed to carry out biological functions, John von Neumann introduced the idea of "self-reproducing automata" (1966). More recently, John Conway's game of "life" (Gardner 1970) has often been interpreted as generating different kinds of species that can perpetuate themselves—as long as the computer does not malfunction or the electric current does not shut down. These caveats point to an insurmountable obstacle these researchers faced in arguing that life could exist in technical media, namely that such technical "life" can never be fully autonomous in its creation, maintenance, and reproduction. From the vantage of hindsight, I think this field of inquiry, although useful and productive in generating controversies and questions, was finally doomed to failure because technical systems can never be fully alive. But they *can* be fully cognitive. Their overlap with biological systems, in my view, should not be focused on "life itself" (as Rosen [1991] put it), but on cognition itself.

Following a path that has occupied me for several years, I offer a definition that will allow me to expand outward to include technical as well as biological cognition. *Cognition is a process that interprets information within contexts that connect it with meaning.* For me, the genesis of this formulation lay in Claude Shannon's theory of information (Shannon and Weaver 1948), in which he shifted the emphasis from a semantic basis for information to the selection of message elements from a set, for example, letters in an alphabet. This way of thinking about information has been enormously fruitful, as James Gleick has explained (2012), for it allowed the development of theorems and engineering practices that extended far beyond natural languages to information processes in general, including binary codes. From a humanities perspective, however, it had a major disadvantage. As Warren Weaver emphasized in his introduction to Shannon's classic work (Shannon and Weaver 1948), it appeared to sever information from meaning. Since the quest for meaning has always been central to the humanities, this meant that information theory would have limited usefulness for humanistic inquiries.

In retrospect, I think Weaver overstated the case in subtle but significant ways. As Shannon knew quite well, the process of selection, which he expressed as a function of probabilities, is not entirely divorced from a message's content and consequently from its meaning. In fact, the conditional probabilities of what message elements will follow their predecessors are already partially determined by the distribution of letters and their relative frequencies within a given language. In English and Romance languages, for example, there is a nearly 100 percent chance that a "q" will be followed by a "u," a higher than random chance that an "e" will be followed by a "d," and so forth. Shannon (1993) linked this idea to the redundancy of English (and other languages), and the theorems that followed were crucial for information compression techniques still in use for telephonic and other kinds of communication transmissions.

Nevertheless, to arrive at meaning, the constraints operating through selection processes are not enough. Something else is needed: context. Obviously, the same sentence, uttered in different circumstances, can change its meaning completely. The missing link between Shannon's view of information and context was supplied for me in a seminar given by the theoretical physicist Edward Fredkin, when he casually observed, "The meaning of information is given by the processes that interpret it" (Hayles 2012, 150). Although Fredkin gave no indication he thought this idea was particularly powerful, it hit me like a bolt of lightning. It blows the problem of meaning wide open, for processes occur within contexts, and *context* can be understood in radically diverse ways for different situations. It applies to utterances of natural language between humans, but it equally well describes the informational processes by which plants respond to information embedded in the chemicals they absorb, the behavior of octopi when they sense potential mates in their vicinity, and the communications between layers of code in computational media. In another context, the insight can also be related to how the brain processes sensory information, in which action potentials and patterns of neural activity may be experienced in different ways depending on which part of the brain engages them (see for example chapter 21, "Sensory Coding," in Kandel and Schwartz 2012, 449–74).[5]

Consistent with Fredkin's explosive insight is the processual and qualitative view of information (as distinct from the quantitative theory developed by Shannon) proposed by the French "mechanologist" Gilbert Simondon in the 1960s as part of his overarching philosophy focusing on processes rather than hylomorphic concepts (form and matter). For Simondon, reality itself is the tendency to engage in processes. A central metaphor for him is the concept of potential energy always tending to flow from a higher state to a lower one, but never coming to a stable equilibrium, only transitional metastable states. He called this flow "information" and thought it is inherently connected with meaning (Simondon 1989; see also Scott 2014; Iliadis 2013; and Terranova 2006). Similar to Fredkin's insight, information in this view is not a statistical distribution of message elements but the result of embodied processes emerging from an organism's embeddedness within an environment. In this sense, the processes that nonconscious cognition uses to discern patterns are constantly in motion, reaching metastable states as patterns are discerned and further reinforced when temporal matching with the reverberations between neural circuits cause them to be fed forward to consciousness. These processes of discerning patterns are always subject to new inputs and continuing transformations as the nonconscious and conscious contexts in which they are interpreted shift from moment to moment. In Simondon's terms, the transfer from one neural mode of organization to another can be conceived as a transfer from one kind of potential energy to another. The information coming to consciousness has already been laden with meaning (that is, interpreted in the relevant contexts) by the cognitive nonconscious; it achieves further meaning when it is re-represented within consciousness.

As we will see in chapter 5, interpretation within contexts also applies to the nonconscious cognitive processes of technical devices. Medical diagnostic systems, automated satellite imagery identification, ship navigation systems, weather prediction programs, and a host of other nonconscious cognitive devices interpret ambiguous or conflicting information to arrive at conclusions that rarely if ever are completely certain. Something of this kind also happens with the cognitive nonconscious in humans. Integrating multiple somatic markers, it too must synthesize conflicting and/or ambiguous information to arrive at interpretations that may feed forward into consciousness, emerging as emotions, feelings, and other kinds of awareness upon which further interpretive activities take place.

In automated technical systems, nonconscious cognitions are increasingly embedded in complex systems in which low-level interpretative processes are connected to a wide variety of sensors, and these processes in turn are integrated with higher-level systems that use recursive loops to perform more sophisticated cognitive activities such as drawing inferences, developing proclivities, and making decisions that feed forward into actuators, which perform actions in the world. In an important sense, *these multi-level systems represent externalizations of human cognitive processes.* Although the material bases for their operations differ significantly from the analogue chemical/electrical signaling in biological bodies, the *kinds* of processes have similar informational architectures. In addition, technical systems have the advantage of working nonstop 24/7, something no biological body can do, and of processing vast amounts of information much faster than humans can. It should not be surprising that human and technical nonconscious cognitions share attributes in common, because brains (deploying nonconscious cognition in their own operations) designed them.

Parsing Cognition

With this background, let us return to parse my definition more fully, since it is foundational for the arguments to follow. *Cognition is a process:* this implies that cognition is not an attribute, such as intelligence is sometimes considered to be, but rather a dynamic unfolding within an environment in which its activity makes a difference. For example, a computer algorithm, written as instructions on paper, is not itself cognitive, for it becomes a process only when instantiated in a platform capable of understanding the instruction set and carrying it out. *That interprets information:* interpretation implies a choice. There must be more than one option for interpretation to operate. In computational media, the choice may be as simple as the answer to a binary question: one or zero, yes or no. Other examples include, in the C++ programming language, commands such as "if" and "else" statements ("if" indicates that a procedure should be implemented only if certain conditions are true; "else" indicates that if these conditions are not met, other procedures should be followed). Moreover, these commands may be nested inside each other to create quite complex decision trees. Choice here, of course, does not imply "free will" but rather programmatic decisions among alternative courses of action, much as a tree moving its leaves to maximize sunlight does not imply free will but rather the implementation of behaviors programmed into the genetic code.

In *Cognitive Biology,* Gennaro Auletta (2011) writes that "biological systems represent the integration of the three basic systems that are involved in *any* physical process of information-acquiring: The processor, the regulator, and the decider" (200). In unicellular organisms, the "decider" may be as simple as the lipid membrane that "decides" which chemicals to admit and which to resist. In more complex multicellular organisms such as mammals and in networked and programmable media, the interpretive possibilities grow progressively more multileveled and open-ended. *In contexts that connect it with meaning:* the implication is that meaning is not an absolute but evolves in relation to specific contexts in which interpretations performed by the cognitive processes lead to outcomes relevant to the situation at that moment. Note that context *includes* embodiment. Lest I be misunderstood, let me emphasize that technical systems have completely different instantiations than biological life-forms, which are not only embodied but also embedded within milieus quite different from those of technical systems.[6] These differences notwithstanding, both technical and biological systems engage in meaning-making within their relevant instantiated/embodied/embedded contexts. For high-level cognitive processes such as human thought, the relevant contexts may be very broad and highly abstract, from deciding whether a mathematical proof is valid to questioning if life is worth living. For lower-level cognitive processes, the information may be the sun's angle for trees and plants, the location of a predator as a school of minnows darts to evade it, or the modulation of a radio beam by a radio-frequency identification (RFID) chip that encodes it with information and bounces it back. In this framework, all these activities, and millions more, count as cognitive.

A meta-implication is that humans do not have a lock on which contexts and levels are able to generate meanings. Many technical systems, for example, operate through communication signals such as radio waves, microwaves, and other portions of the electromagnetic spectrum inaccessible to direct human perception. To unaided human senses, the signals bouncing around the atmosphere are both imperceptible and meaningless, but to technical devices that operate in contexts relevant to them, they are filled with meaning. Traditionally, the humanities have been concerned with meanings relevant to humans in human-dominated contexts. The framework developed here challenges that orientation, insisting cognitive processes happen within a broad spectrum of possibilities that include nonhuman animals and plants as well as technical systems. Moreover, the meanings generated within these contexts, deeply worthy of consideration in their own right, are also consequential for human outcomes as well, from the flourishing of trees in rain forests to the communication signals emanating from a control tower to aircraft within its purview. This framework emphasizes that these different kinds of meanings are entangled together in ways that transcend any single human viewpoint and that cannot be bounded by human interests alone. As our view of what counts as cognition expands, so too do the realms in which interpretations and meanings emerge and evolve. All of these, this framework implies, count as meaning making and consequently should be of potential interest to the humanities, as well as to the social and natural sciences.

The Tripartite Framework of (Human) Cognition

Turning now specifically to human cognition, I develop this view with a tripartite framework that may be envisioned as a pyramid with three distinct layers (fig. 1, p. 40). At the top are consciousness and unconsciousness, grouped together as modes of awareness. As noted earlier, research on the "new" unconscious sees it as a kind of broad environmental scanning in which events are heeded and, when appropriate, fed forward to consciousness (Hassin, Uleman, and Bargh 2005). The new unconscious differs from the psychoanalytic unconscious of Freud and Lacan in that it is in continuous and easy communication with consciousness. In this view the psychoanalytic unconscious may be considered as a subset of the new unconscious, formed when some kind of trauma intervenes to disrupt communication and wall off that portion of the psyche from direct conscious access. Nevertheless, the psychoanalytic unconscious still expresses itself to consciousness through symptoms and dreams susceptible to psychoanalytic interpretation. The modes of awareness, designating the neurological functions of consciousness and the communicating unconscious, form the top layer of the pyramid.

The second part of the tripartite framework is nonconscious cognition, described in detail elsewhere (Hayles 2012). Unlike the unconscious, it is inherently inaccessible to consciousness, although its outputs may be forwarded to consciousness through reverberating circuits (Kouider and Dehaene 2007). Nonconscious cognition integrates somatic markers such as chemical and electrical signals into coherent body representations (Damasio 2000; Edelman 1987). It also integrates sensory inputs so that they are consistent with a coherent view of space and time (Eagleman 2012). In addition, it comes online much faster than consciousness and processes information too dense, subtle, and noisy for consciousness to comprehend. It discerns patterns that consciousness is unable to detect and draws inferences from them; it anticipates future events based on these inferences; and it influences behavior in ways consistent with its inferences (Lewicki, Hill, and Czyzewska 1992). No doubt nonconscious cognition in humans evolved first, and consciousness and the unconscious were subsequently built on top. Removed from the confabulations of conscious narration, nonconscious cognition is closer to what is actually happening in the body and the outside world; in this sense, it is more in touch with reality than is consciousness. It comprises the broad middle layer of the tripartite framework.

The even broader bottom layer comprises material processes. Although these processes are not in themselves cognitive, they are the dynamic actions through which all cognitive activities emerge. The crucial distinguishing characteristics of cognition that separate it from these underlying processes are choice and decision, and thus possibilities for interpretation and meaning. A glacier, for example, cannot choose whether to slide into a shady valley as opposed to a sunny plain. In contrast, as Auletta explains, “any biological system . . . produces variability as a response to environmental challenges and tries to integrate [these] aspects inside itself” (2011, 200). In general, material processes may be understood through the sum total of forces acting upon them. A special case is formed by criticality phenomena, structured so that even minute changes in initial conditions may change how the system evolves. Even here, the systems remain deterministic, although they are no longer predictable. There are many examples of material processes that can self-organize, such as the Belousov-Zhabotinsky (BZ) inorganic reaction. However, there remain crucial distinctions between such far-from-equilibrium systems and living organisms, for whom choices, decisions, and interpretations are possible. As Auletta points out, “biological systems are more than simply dissipative self-organizing systems, for the reason that they can negotiate a changing or nonstationary environment in a way that allows them to endure (to change in an adaptive sense) over substantial periods of time" (2011, 200). Material processes may however be harnessed to perform cognitive functions when natural or artificial constraints are applied in such a way as to introduce choice and agency into the system (Lem 2014), for example, through the interactions of multiple independent agents in complex environments.

Although the pyramidal shape of the tripartite framework may seem to privilege the modes of awareness over nonconscious cognitions and material processes, inasmuch as they occupy the top strata, a countervailing force is expressed through the pyramid volumes. The modes of awareness, precisely because they come at the top, reign over the smallest volume, a representation consistent with the roles they play in human psychic life. Nonconscious cognition occupies a much greater volume, consistent with the processes it performs as the neurological function mediating between the frontal cortex and the rest of the body. Material processes occupy a vast volume, consistent with their foundational role from which all cognition emerges.

Although the tripartite framework divides human processes into three distinct layers for analytical clarity, in reality complex recursive loops operate throughout the system to connect the layers to each other and connect different parts of each layer within itself. Each layer operates dynamically to influence the others all the time, so the system is perhaps better described as a dynamic heterarchy rather than a linear hierarchy, a view that animates and interconnects the system as it evolves in real time. Consequently, the structure sketched above is a first approximation. It is not so much meant to settle questions as to catalyze boundary issues and stimulate debates about how the layers interact with each other. That said, it nevertheless serves as a starting point to discuss issues of agency and to distinguish between actors and agents.

Because cognition in this framework is understood as inseparable from choice, meaning, and interpretation, it bestows special functionalities not present in material processes as such. These include flexibility, adaptability, and evolvability. Flexibility implies the ability of an organism or technical system to act in ways responsive to changing conditions in its environment. Whereas a ball thrown toward a window has no choice to alter its trajectory, a self-driving car can respond with a large repertoire of possibilities to avoid damage. As indicated above, flexibility is present in all living organisms to some extent, even those lacking central nervous systems.[7] Adaptability denotes developing capacities in response to environmental conditions. Examples include changed neurological functioning in plants, animals, and humans in response to environmental stresses or opportunities, such as the neurological changes human brains undergo through extensive interactions with digital media (Hayles 2012). Evolvability is the possibility to change the programming, genetic or technical, that determines the repertoire of responses. Genetic and evolutionary algorithms are examples of technical systems with these capabilities (Koza 1992), as are computers that can reconfigure their own firmware, rearranging logic gates to solve problems with maximum efficiency (Ling 2010). Biological examples are of course everywhere, as biologists from Darwin and Wallace on have confirmed. The important point is that material processes do not possess these capabilities in themselves, although they may serve to enhance and enlarge cognitive capabilities when enrolled as supports in an extended cognitive system.

Actors and Agents

It is fashionable nowadays to talk about a human/nonhuman binary, often in discourses that want to emphasize the agency and importance of nonhuman species and material forces (Bennett 2010; Grosz 2011; Braidotti 2013). To my mind, there is something weird about this binary. On one side are some seven billion individuals, members of the Homo sapiens species; on the other side sits everything else on the planet, including all the other species in the world, and all the objects ranging from rocks to clouds. This binary, despite the intentions of those who use it, inadvertently reinstalls human privilege in the vastly disproportionate weight it gives to humans. Some theorists in the ecological movement are developing a vocabulary that partially corrects this distortion by referring to the "more-than-human" (Smith 2011), but the implicit equivalence of the human world to everything else still lingers.[8]

Recognizing that binaries can facilitate analysis (their limitations notwithstanding), I propose another distinction to replace human/nonhuman: *cognizers versus noncognizers.* On one side are humans and all other biological life forms, as well as many technical systems; on the other, material processes and inanimate objects. At the very least, this distinction is more balanced in the relative weights it gives to the two sides than the very unbalanced human/nonhuman formulation. This binary (like all binaries) is not innocent of embedded implications. In particular, it foregrounds cognition as a primary analytical category. Skeptics may object that it too reinstalls human privilege, since humans have higher and more extensive cognitions than other species. However, this binary is part of a larger cognitive ecology emphasizing that *all* life forms have cognitive capabilities, including some that exceed human cognitions (smell in dogs, for example).

Moreover, because only cognizers can exercise choice and make decisions, they have special roles to play in our current environmental crises and the sixth mass extinction already underway. The one motivation that all life-forms share is the struggle to survive. As environmental stresses increase differentially, cognizers at all levels, from worms to humans, will make choices that tend to maximize their chances for survival. Admittedly, species with higher cognitive capabilities can supervene this motivation as it interacts with other priorities—as many humans are doing at present. Having an analytical category that emphasizes choice may help to foreground our common causes with other cognizers and draw our attention more vividly to the fact that we all make choices, and that these choices matter, individually and collectively. Moreover, the capabilities that cognition bestows—flexibility, adaptability, evolvability—imply that cognizers have special roles to play in our evolving planetary ecologies. Finally, this framework sets up the possibility that cognitive technologies may perform as ethical actors in the assemblages they form with biological life-forms, including humans.

For their part, noncognizers may possess agential powers that dwarf anything humans can produce: think of the awesome powers of an avalanche, tsunami, tornado, blizzard, sandstorm, hurricane. Faced with these events, humans utterly lack the ability to control them;

the best they can do is get out of the way. Moreover, since material processes are the underlying forces that nourish and give rise to life, they deserve recognition and respect in their own right, as foundational to everything else that exists (Strang 2014). What they cannot do, acting by themselves, is make choices and perform interpretations. A tornado cannot choose to plow through a field rather than devastate a town. Material processes, of course, respond to contexts and, in responding, change them. But because they lack the capacity for choice, they perform as agents, not as actors embedded in cognitive assemblages with moral and ethical implications.

I propose a further shift in terminology that clarifies the different roles performed by material processes and nonconscious cognizers. I suggest reserving the term *actors* for cognizers, and *agents* for material forces and objects. This latter category includes objects that may act as cognitive supports; it also includes material forces that may be harnessed to perform cognitive tasks when suitable constraints are introduced, for example, when electrical voltages are transformed into a bit stream within a computational medium.

Fueled by global capitalism, technical cognitive systems are being created with ever more autonomy, even as they become increasingly pervasive within developed societies. As David Berry (2015) among others points out, there is no technical agency without humans, who design and build the systems, supply them with power and maintain them, and dispose of them when they become obsolete. Nevertheless, the pockets within which technical systems operate autonomously are growing larger and more numerous. Examples include environmental monitoring systems, surveillance and communication satellites, digital search engines, and language learning systems, among many others. Perhaps an appropriate way to think about the growing autonomy of these systems is as punctuated agency, analogous to "punctuated equilibrium" (Gould 2007). Like punctuated equilibrium, punctuated agency operates within regimes of uneven activity, longer periods when human agency is crucial, and shorter intervals when the systems are set in motion and proceed on their own without direct human intervention.

Even within the autonomous regions, the effects of technical cognitions are not contained wholly within the technical systems. They interact with human complex systems to affect myriad aspects of human and biological life. In this respect, even the cognizer/noncognizer binary falls short because it fails to capture the powerful and subtle ways in which human and technical cognizers interact with each other as well as with noncognizing objects and material forces. Water is a good example (Strang 2014): on its own it exercises agency through such phenomena as waterfalls, rain, snow, and ice; incorporated into biological bodies, it provides fluids essential for life; run through a turbine, it contributes to the cognitions and effectiveness of a computerized hydroelectric power system. To express more adequately the complexities and pervasiveness of these interactions, we should resist formulations that reify borders and create airtight categories. The better formulation, in my view, is not a binary at all but interpenetration, continual and pervasive interactions that flow through, within, and beyond the humans, nonhumans, cognizers, noncognizers, and material processes that make up our world.

Why Computational Media Are Not Just

Another Technology

In *What Technology Wants*, Kevin Kelly (2010) argues that technologies develop along trajectories that he anthropocentrically identifies with "desire," including ubiquity, diversity, and intensity. As the provocation of his title indicates, his discussion fails to give a robust account of how human agency enters this picture. Nevertheless, there is a kernel of insight here, which I rephrase as this: technologies develop within complex ecologies, and their trajectories follow paths that optimize their advantages within their ecological niches. The advent of photography in the mid-to-late nineteenth century, for example, preempted the category of landscape description, and consequently literary novels readjusted their techniques, moving away from the pages of landscape description notable in late-eighteenth- and early-nineteenth-century novels and into stream of consciousness strategies, an area that photography could not exploit as effectively. As Cynthia Sundberg Wall has shown (2014, esp. chapters 1–3, 2–95), literary descriptive techniques are enmeshed within a cultural matrix of techniques of vision, including microscopes, telescopes, maps, and architectural diagrams. The dynamics of competition, cooperation, and simulation between media forms are powerful analytics for understanding technological change (Fuller 2007; Hansen 2015; Gitelman 2014).

In these terms, computational media have a distinct advantage over every other technology ever invented. They are not necessarily the most important for human life; one could argue that water treatment plants and sanitation facilities are more important. They are not necessarily the most transformative; that honor might go instead to transportation technologies, from dirt roads to jet aircraft. Computational media are distinct, however, because they have a *stronger evolutionary potential* than any other technology, and they have this potential because of their cognitive capabilities, which among other functionalities, enable them to simulate any other system.

We may draw an analogy with the human species. Humans are not the largest life-form; they are not the strongest or the fastest. The advantages that have enabled them to achieve planetary dominance within their ecological niche are their superior cognitive capabilities. Of course, we are long past the era when the Baconian imperative for humans to dominate the earth can be embraced as an unambiguous good. In an era of ecological crises, global warming, species extinction, and similar phenomena, the advent of the Anthropocene, in which human influences are changing geological and planetary records, is properly cause for deep concern and concerted political activism around climate change, preservation of habitats, and related issues.

The analogy with the cognitive capacities of computational media suggests that a similar trajectory of worldwide influence is now taking place within technical milieus. Fueled by the relentless innovations of global capital, computational media are spreading into every other technology because of the strong evolutionary advantages bestowed by their cognitive capabilities, including water treatment plants and transportation technologies but also home appliances, watches, eyeglasses, and everything else, investing them with "smart" capabilities that are rapidly transforming technological infrastructures throughout the world. Consequently, technologies that do not include computational components are becoming increasingly rare. Computational media, then, are not just another technology. They are the quintessentially *cognitive* technology, and for this reason have special relationships with the quintessentially cognitive species, *Homo sapiens*.

Note that this position should not be conflated with technological determinism. As Raymond Williams has astutely observed, such evolutionary potentials operate within complex social milieus in which many factors operate and many outcomes are possible: "We have to think of determination not as a single force, or a single abstraction of forces, but as a process in which real determining factors—the distribution of power or of capital, social and physical inheritance, relations of scale and size between groups—set limits and exert pressures, but neither wholly control nor wholly predict the outcome of complex activity within or at these limits, and under or against these pressures" (Williams 2003, 13). In fact, one can argue that the larger the cognitive components of technological systems, the more unpredictable are their specific developments, precisely because of the qualities conferred by cognition, namely flexibility, adaptability, and evolvability. As global capital continues to innovate ways in which computational media may be infused into other technologies, the e-waste created by their exponential growth increasingly poisons environments where they end up, disproportionately, in poor, underprivileged, and underfunded countries. Given that the cognitive capabilities of technical media are achieved at considerable cultural, social, political, and environmental costs, we can no longer avoid the ethical and moral implications involved in their production and use.

Technological Cognition and Ethics

As we have seen, *choice* in my framework has a very different meaning than in ethical theories, where it is associated with free will. What ethical approaches are appropriate to the former, which I will call CHOICEII (interpretation of information), as distinct from CHOICEFW (free will)? Bruno Latour (1992) touches on this question when he suggests that the "missing masses" of ethical actors (by analogy with the missing mass/energy that physicists need to explain the universe's inflation) are technical artifacts: "here they are, the hidden and despised social masses who make up our morality" (1992, 227). Using simple examples of seat belts and hydraulic door closers, Latour shows that technical artifacts encourage moral behavior (annoying buzzers that remind drivers to fasten seat belts) and influence human habits (speed bumps influencing drivers not to speed in school zones) (2002). In these examples, the technical objects are either passive or minimally cognitive. Even at this modest level, however, artifacts act as "mediators" influencing human behaviors, notwithstanding that they often sink into the background and are perceived unconsciously (Latour 1999, 2002; Verbeek 2011).

When artifacts embody higher levels of cognition, they can intervene in more significant and visible ways. Peter-Paul Verbeek develops a philosophical basis for thinking about technical systems as moral actors and suggests how to design technologies for moral purposes (2011, 135). The Fitbit bracelet (my example, not his) encourages fitness by monitoring heart rate, keeping track of workouts, noting calories burned, and measuring distances covered and stairs climbed. None of these devices absolutely compel obedience, as Latour acknowledges, because there are always ways to defeat their behavioral intent. Nevertheless, they have cumulative (and expanding) effects that significantly affect human social behaviors and unconscious actions.

Following Latour's lead in thinking about technical systems as "mediators," Verbeek develops the argument further by showing how technologies such as obstetric ultrasound not only open new areas for ethical consideration (for example, whether to abort a malformed or, even more distressing, a female fetus) but also reconfigure human entities in new ways (the fetus becoming a medical patient viewable by the physician). In the entangled web of human and technical actors, Verbeek argues, both humans and technics share moral agency and, implicitly, moral responsibility: "moral agency is distributed among humans and nonhumans; moral actions and decisions are the products of human-technology associations" (Verbeek 2011, 53).[9]

Like Verbeek, Latour emphasizes the unexpected effects of technological innovations, arguing that technological systems almost always modify and transform the ends envisioned in their original designs, opening up new possibilities and, in the process, entangling means and ends together so that they can no longer reasonably be regarded as separate categories.[10] The thrust of this argument, of course, is to defuse the objection that technological artifacts are merely the means for ends established by humans. Examples of technologies invented for one purpose and reappropriated for another are legion, from the typewriter, initially invented for blind people, to the Internet, originally intended as a place where scientific researchers could exchange results.

While Latour and Verbeek offer valuable guidance, to my mind their arguments do not go far enough. With technologies capable of significant decision making—for example, autonomous drones—it does not seem sufficient to call them "mediators," for they perform as actors in situations with ethical and moral consequences. One might argue, as Verbeek does, that distributed agency implies distributed responsibility, but this raises the prospect of a technological artifact being called to account for performing the actions programmed into it, a misplaced ethical judgment reminiscent of medieval animal trials in which starlings were executed for chattering in church and a pig was hanged for eating a communion wafer.

Ethical theories, for their part, are often intensely anthropocentric, focusing on individual humans as the responsible agents to whom ethical standards should apply, as in Emmanuel Levinas's complex notion of the Other's face (1998). Although some theories extend this to animals (for example, Tom Regan's suggestion [2004] that mammals over a certain age should be considered subjects of a life and therefore have ethical rights), few discuss the role of technical cognizers as responsible technical actors. Latour is certainly right to point to human-technical assemblages as transformative entities that affect ends as well as means, but he offers little guidance on how to assess the ethical implications of such assemblages. If, to use Latour's example, neither guns nor people are the agents responsible for gun violence but rather the gun-person collective they form (Latour 1999, 193), surely drone-with-pilot is a much more potent assemblage than either by itself/himself.

To assess such assemblages, we should move from thinking about the individual and CHOICEFW as the focus for ethical or moral judgment, and shift instead to thinking about CHOICEII and the consequences of the actions the assemblage as a whole performs. Jeremy

Bentham suggested a similar move when he wrote, "The general tendency of an act is more or less pernicious according to the sum total of its consequences, i. e., according to the difference between the sum of its good consequences and the sum of its bad ones" ([1780] PDF, 43). We need not subscribe to all the tenets of utilitarianism to accept this as an adequate framework in which to evaluate the effects of cognitive assemblages that include technical actors. Drone pilots cannot be considered simply as evil for killing other humans; even less so can the drone itself. Rather, they act within structured situations that include tactical commanders, lawyers, and presidential staff, forming assemblages in which technological actors perform constitutive and transformative roles along with humans. The results should therefore be evaluated *systemically* in ways that recognize not all of the important actors are human, an argument developed further in part 2. Moreover, drone assemblages are part of larger conflicts that includes suicide bombers, IEDs, military incursions, insurgent resistance, and other factors. The cognitive assemblages in such conflicts are differentially empowered by the kinds of technologies they employ as well as by how the humans enmeshed within them act. The consequences of the assemblages further interact with existing discourses and ethical theories in dynamic, constantly shifting constellations of opposing interests, sovereign investments, personal decisions, and technological affordances. Attempting to evaluate moral and ethical effects from the actions of individual people alone by focusing on CHOICEFW is simply not adequate to assess the complexities involved. As part 2 argues more fully, we need frameworks that explore the ways in which the technologies interact with and transform the very terms in which ethical and moral decisions are formulated.

We can see the inadequacy of remaining within individual-focused frameworks by considering the justification for designing robot weapons offered by Ronald C. Arkin, Regents Professor of computer science at Georgia Tech, compared with the drone theory of Grégoire Chamayou. Arkin, who has Defense Advanced Research Projects Agency (DARPA) grants to develop autonomous robot warriors for the battlefield, argues that robots may be morally superior to human warriors because they would be forbidden by their programming to commit atrocities, immune to emotional stress and the bad decisions that can accompany it, and able to direct their lethal encounters more precisely, minimizing collateral damage (Arkin 2010, 332–41). His critics attack these claims on a number of fronts; perhaps the most compelling is the objection that once robot warriors are available, they would likely be used more widely and indiscriminately than human warriors, where the prospect of putting one's troops "in harm's way" acts as a significant restraint on military and political leaders.

Evaluating the claims for robot morality requires a larger interpretive frame than the one Arkin uses. Leaving aside the question of whether robots would in fact be programmed to follow the rules of war established by international treaties (and whether these rules could ever make war "moral," an issue explored in part 2), I note that he treats the robots in the same terms as human individuals (but equipped with better sensors and decision-making capabilities) rather than as technical systems embedded in complex human-technical assemblages.

Grégoire Chamayou (2015) is subtler in interrogating how the specific rules of engagement for drone pilots cause conventional standards of appropriate behavior in warfare to be transformed and reinterpreted to accommodate the pilots' actions. For example, he points out that traditional accounts of war distinguish sharply between soldiers and assassins. Whereas the former are considered honorable because, by entering a field of combat, they establish who is an enemy combatant and also put their own bodies at risk, assassins are cowardly because they may strike targets who are not combatants and do not necessarily put themselves at risk in doing so. Applied to drone pilots, these views could force them to be counted as assassins rather than soldiers. To mitigate the situation, the US military has emphasized that drone pilots may be suffering from post-traumatic stress disorder and in this sense are putting themselves at risk as well. Although Chamayou has his own agenda and often is one-sided in his appraisals (as argued in part 2), his analyses nevertheless show that the consequences of human-technical assemblages include not only the immediate results of actions but also far-reaching transformations in discourses, justifications, and ethical standards that attempt to integrate those actions into existing evaluative frameworks.

The more powerful the cognitive capabilities of technical systems, the more far-reaching are the results and transformations associated with them. Drones are especially controversial examples, but technical cognitive systems employing CHOICEII are all around us and operating largely under the radar of the general public, including expert medical systems, automated trading algorithms, sensing and actuating traffic networks, and surveillance technologies of all kinds, to mention only a few. To analyze and evaluate their effects, we need robust frameworks that recognize technical cognition as a fact, allowing us to break out of the centuries-old traditions that identify cognition solely with (human) consciousness. We also need a more accurate picture of how human cognitive ecology works, including its differences from and similarities to technical cognition. Finally, we need a clear understanding of how cognizers differ from material processes, which includes a definition of cognition that sets a low threshold for participation but includes ways to scale up to much more sophisticated cognitions in humans, nonhuman life forms, and technical systems. Added together, these innovations amount to nothing less than a paradigm shift in how we think about human cognition in relation to planetary cognitive ecologies, how we analyze the operations and ethical implications of human-technical assemblages, and how we imagine the role that the humanities can and should play in assessing these effects.

In conclusion, let me address the role of humanistic critique. If thought in general is associated with consciousness, critique is even more so. Some may object that challenging the centrality of reason in cognitive processes undermines the nature of critique itself. Yet consciousness alone cannot explain why scholars choose certain objects for their critique and not others, nor can it fully address the embodied and embedded resources that humanities scholars bring to bear in their rhetorical, analytical, political, and cultural analyses of contemporary issues. Without necessarily realizing it, humanities scholars have always drawn upon the full resources of human cognitive ecologies (fig. 1), both within themselves and within their interlocutors. Recognizing the complexities of these interactions does not disable critique; on the contrary, it opens critique to a more inclusive and powerful set of resources with which to analyze the contemporary situations that confront us, including but not limited to the entanglements and interpenetrations of human and technical cognitive systems. That is the importance, and the challenge, of the cognitive nonconscious to the humanities today.

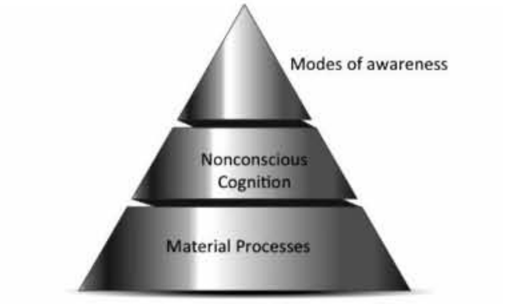

Figure 1. The tripartite framework of (human) cognition as a pyramid