CAPACIOUS

Unthought Meets the Assemblage Brain

~~~~~~~

N. Katherine Hayles and Tony D. Sampson

A Dialogue Between N. Katherine Hayles and Tony D. Sampson

It has become such a refrain for affect studies—Spinoza's "what can a body do?"—that it is sometimes repeated half-absentmindedly: as if all questions involving the realm of consciousness have been largely settled, as if they are somehow located elsewhere, as if they are entirely different sets of questions than those asked of a body. But of course they're not. Perhaps that's one of the reasons why the place of "nonconscious" is so particularly intriguing to take up: troubling any too-ready line drawn around, through, or between body-mind-and-world. "Just *where/ when* is the nonconscious?" is, at least initially, as pertinent as the "what-can-it-do?" question—as it comes out of the cognitive sciences to intersect with and complicate the affect studies' refrain.

And that's what makes this dialogue between N. Katherine Hayles and Tony Sampson so fascinating and theoretically rich. Following upon their most recent books—Hayle's *Unthought: The Power of the Cognitive Unconscious* (University of Chicago, 2017) and Sampson's *The Assemblage Brain: Sense Making in Neuroculture* (University of Minnesota, 2016), the convergences and divergences that emerge and weave throughout this conversation are quite revealing.

What transpires in the unmediated space-time excess that moves, at once, between and alongside cognition and recognition, between and alongside formation and information, between and alongside prehension and comprehension? Nietzsche

said that we sometimes need to recall consciousness to its necessary modesty. But non-consciousness (technological or otherwise) is another sort of beast, a materially more immodest and rangy one—not responsive to just any call (after all, who/what precisely would be on the line? Except everything and no single thing). And, like all things addressed by way of affect, it matters deeply what avenues of inquiry and what particular aims are brought to bear on the foundational questions—in this case, it is the biosemiotics of human-computer interaction for Hayles and the ethology of more-than-human assemblages for Sampson. In that lively wedge of distinction between their angles of approach, this dialogue offers a more widely conceptualized world of the "doings" for affect studies.

It is beyond delightful that Kate and Tony agreed to engage in this spirited conversation. We already have plans for further dialogues between key figures working within or in near-vicinity of affect studies for future issues of *Capacious*.

—Greg Seigworth, co-editor of this journal

Assemblage Brain (alternative cover), Francesco Tacchini, 2015

TS: I'm keen to begin by noting some convergences in our work. For example, let's discuss how both of our approaches begin with a desire to jettison the computational brain thesis. I realize that you do this from "within" the cognitive frame, and my starting place is broadly affect theory/new materialism, but even if there's a degree of palpable divergence here, this rejection of the computational brain points to some clear examples of common ground that runs through each book.

The problem for me is that although there's a suggested break from computational approaches in cognitive brain sciences, and people like Damasio and LeDoux have helped shift debate to a far more interesting materialist approach that takes into account the environment, soma, affect, emotions and feelings, many of the information metaphors from the old paradigm are still intact. This is a point made in theoretical neuroscience in Bennett and Hacker's critique of Damasio and LeDoux, and gets repeated in HCI theory, wherein there's a similar move away from cognitive approaches based on human-information coupling toward a focus on a situated and embodied phenomenology of user experience.[1]

I note that on p. 218 of your book there's a footnote on the kind of materialism you are committed to.[2] I see this as materialism +information. Indeed, the processing of information (interpreting, choosing, and deciding) seems absolutely key to the important categorization you go on to make between cognizers and noncognizers.

Can I begin by asking you about why you see the computational brain as problematic? Then ask why you find information (and information processing, flows etc.) so important to your work on the nonconscious?

KH: My problem with the hypothesis that the brain manipulates symbols to accomplish its tasks is a lack of evidence for this thesis, and the lack of any reference to what the brain actually has to work with, namely the body, extended nervous system outside the brain, organs such as skin and viscera, etc. Nevertheless, I can understand why theorists want to make the connection between computational cognition and biological cognition, but I think it cannot be done by saying both work with symbols. A much better approach, I think, is biosemiosis: the creation, exchange, processing, storage and transmittal of information both within an organism and between an organism and the environment. Clearly, this also requires a definition of information very different from that of the Shannon-Wiener theory. Jesper Hoffmeyer, in developing the idea of biosemiosis and the semio-

sphere, uses the Peircian model of sign/object/interpretant. The movement from a binary logic to a triadic mode crucially includes the interpretant, the "someone" for whom the sign-processes have meaning—and the ground level of meaning for lifeforms is survival.

To make the connection to computational media, it is necessary at the outset to emphasize the enormous differences in embodiment between computational media and biological organisms. In addition, computers do not evolve but are designed for purposes. Thus, they have no innate imperative to survive but rather operate in artificial environments to fulfill their purposes (or better, since this assumes the computer knows it has a purpose, its design mandates). Of special importance in this regard is the "if/else" (in Fortran, "if/then") command: if a certain set of criteria are present, do the following; if not, then do something else. This command is fundamental to computational semiosis because it establishes the temporality that sign-exchange implies.

It also opens a path to talk about how computers achieve meaning.

Here John Dewey's theory of meaning is useful, because it does not center on anthropomorphic criteria but instead emphasizes that the meaning of an action can be understood in terms of its consequences. For computational media, the "if/else" command is precisely aimed at the consequences of a computer operating in its environment to achieve something. That environment includes its dataset, the source code (and associated other code layers), the operating system, any sensors and actuators present in the system, and other affordances. When a computational system makes a decision about what actions to perform (execute commands if, else do something different), those decisions constitute an anticipation of what the consequences will be and thus constitute meaning-making for that system.

Meaning-making for biological organisms can be understood in similar terms. Even one-celled organisms are capable of creating meaning from their actions, because they process information in terms of their systemic dynamics and make decisions about what to do based on that information, their surrounding environments, and their sensory/biological capacities. They too anticipate consequences, and this is true even for minimally cognitive lifeforms such as plants.

The key components here, as I argue in *Unthought*, are cognition, interpreting information in contexts, and connecting information with meaning. In this age

when cognitive assemblages are crucial to everyday life, meaning cannot be restricted only to activities that humans undertake. It must be broadened to include nonhuman lifeforms and, equally important, networked and programmable machines.

TS: Absolutely, I agree with you that nonhumans importantly need our consideration and I think it's exciting that you extend sense making to algorithms and plants. To some extent, this hints at the kind of assemblage theory I explore in *The Assemblage Brain.* But can we stay with information for a moment? I have a few related points to make before we move on to look at cognition.

As I see it, your commitment to biosemiotics relies on an assumption that what occurs at the genetic level (where information is encoded onto physical matter) also emerges, albeit as a translation, at higher levels (the semiotic mind and perhaps even distributed consciousness). I see how this evolutionary emergence of information from the biosemiotic to the phenomenological semiotic corresponds to the bridge Damasio similarly develops between the protoself and the coreself (a major feature of your notion of the nonconscious). So we effectively move on from the hardware/software machines of the computational brain thesis to a series of interpretation machines taking part in an information dialogue between codes at different levels.

It's interesting that you say there's no scientific evidence to support the computational brain thesis. Can you please briefly outline what evidence shows us how encoding/decoding processes move up through these levels? Is there a specified location or network of neurons where the production of signs ("information") occurs in the biology of the brain? I can see how information processing works in computational media as datasets, code, operating systems etc., but where is the biosemiotic equivalent located in the brain? Is it something that has a simple location? Is there an fMRI scan, for example, that reveals this kind of information coding/decoding, representational storage, processing, and transmission at work? I say this, because although I welcome the departure from Shannon-Wiener, I wonder if this move from the symbol manipulations of computational cognition to the sign manipulation of biosemiotics is really radical enough. In short, is there an alternative to what still seems to be information theory analogies applied to biology?

My book is, as you might guess, critical of this emergent evolutionary leveling up process in which each level seems to transcend the next, like a staircase leading to consciousness and perhaps leading all the way up to a collective social consciousness. I initially follow thinking in HCI that considers information as an

inadequate way to conceive of experiences that are "felt" before they are thought. For example, Donald Norman describes visceral felt experiences that arrive before reflective thoughts, emotions, and behaviors (an idea that in many ways relates to affect theory). Similarly, I wonder how a model of information interpretation accounts for the emergence of feelings and moods in broader societal relations, like mass panic, for example.

Another way of looking at the problem of information interpretation is to consider how biosemiotics, for example, responds to the Whiteheadian notion of feeling or prehension. Whitehead also looks at plants and what he calls the "sense of conformation" a plant experiences when it responds to light or warmth (Debaise 2017, 46). This leads to the question of what is the "form" of feeling, which is conceived of as an "immediate" form of experience. In a Whiteheadian mode then, we find a vital theory of experience that takes into account a temporal sense of the *event* rather than a phenomenological representation.[3]

My point is that your materialism is crucially interwoven with information, which I think binds your notion of assemblages to a cognitive theoretical frame that is still essentially adhering to a kind of computational cognition (sign rather than symbol). That's why I'm interested in what you think of attempts to move beyond/away from this human-information coupling model. Bennett and Hacker's notion that, for example, a better metaphor for the brain might be the ocean. It's a playful metaphor, but one that I think offers compelling temporal alternatives to information, including waves, rhythms, and fluid flows, which can be displaced and distributed as intensities.

Anatomy of a Purkinje neuron inspired by Santiago Ramón y Cajal, Dorota Piekorz, 2017

KH: Tony, thanks for your comments. Let me address first your comment about where the production of signs occurs in a biosemiotic perspective. The major advantage of sign-processing over symbol-processing as formulated in the computationalist hypothesis is that signs are not arbitrarily interjected onto biological processes. Rather, as Hoffmeyer makes clear, the point is to interpret biological processes that act as signs signifying something to the cell. For example, the proteins on a cell membrane are folded in specific ways so they can detect alien bacteria; when these appear, the proteins engage in sign processes that signal a meaning to the cell, namely that the invading bacteria are non-self rather than self and need to be attacked.

The new elements here are 1) the notion that such processes (which occur at every level, from the cellular level on up) have the potential to act as signs that may be iconic and/or indexical rather than symbolic, and 2) that signs always require an interpretant, the "someone" for whom the sign signifies something of consequence. Such sign-processes do have a hierarchy of networks, but they also operate semi-autonomously at level-specific sites. Moreover, the consequences of those sign-processes are never completely accounted for (or exhausted) by the upward messages; that is, their content always exceeds the information sent forward to the next level up the hierarchy. In addition, these feedforward loops work continuously to produce meaning; meaning-making at a lower level does not stop when the feedforward loop sends its messages up the hierarchy. As I suggested in another context in *My Mother Was a Computer*, this process is perhaps best described as a heterarchy for that reason.

As you know, biosemiotics did not originate with information theory but with von Uexküll's *umwelt* theory. It therefore addresses the process of meaning-making by considering the meaning-maker as a subject with a specific world-view, that is, its *umwelt*. This makes it fundamentally different from the purely quantitative (and subjectless) information postulated by Shannon and Wiener. What I especially like about biosemiotics is precisely this subjective orientation, which it combines in a very convincing way with empirical research on biological processes.

As for the temporal structure of these processes (in Whiteheadian terms, its eventness), anticipation is shown to be woven into all biological processes, in the sense that phenotypic experiences give a specific life-history that affects how meaning-making subsequently occurs. A good example here would be the production of antibodies by the immune system, an anticipatory response based on past experiences.

As for moving away from an information model, the best way to describe information as it appears in biosemiotics is probably Bateson's "difference that makes a difference." I think the ocean metaphor (waves, intensities) has potential, for example in the "time-wave" crystals newly discovered in physics.[4] But for my taste, its utility so far has been somewhat obscured by a Deleuzian rhetoric, which is very difficult to reconcile with empirical research. I would not discount its possibilities, but the biosemiotic approach has been far more developed, and its potential is therefore much clearer in my view.

TS: Thanks Kate. I found your criticism of Deleuzian rhetoric in *Unthought* a timely challenge to new materialism. I understand your concerns. My main concern is, however, that some of the more dizzying rhetoric, which we all, to varying degrees, partake in, might obscure Deleuze's valuable articulation of science through philosophical ideas. More precisely, his work with Guattari in *What is Philosophy?* helps us to think through these contemporary mixtures of science and philosophical concepts in new ways. I think this is important for two reasons.

On the one hand, at a time when the neurosciences are laying down some big challenges to, and in some cases even rubbishing, philosophy, the Deleuzian non-scientific reading of science becomes ever more essential. After all, science is, for the most part, speculative; especially when it comes to figuring out consciousness. It operates very much in the virtual realm. Indeed, the desire for concrete empirical evidence to support, for example, a philosophical concept of consciousness is, in my opinion, problematic since scientific ideas are often drawn from moments of paradigmatic epistemological ignorance. Perhaps these time crystals you mention will upend the theory of time according to physics?

On the other hand though, I think that a more rigorous reading of Deleuze and more thorough grasp of his sources (Tarde, Bergson, Whitehead, Simondon) reveals a series of concepts that engage with science in ways that many in the humanities have failed to. I was first attracted to Deleuze because of my interest in science and technology. His work introduced me to Whitehead who is a major league mathematician with a deep interest in early quantum physics. In the contemporary work of philosopher-mathematicians, like Brian Rotman, for example, we find a fascinating alternative to a rather stale anti-scientific idealism in the humanities. So, I agree with you that, at one extreme, we find a rhetorical tendency in the overuse of terms like deterritorialization, lines of flight and frequent naïve rantings about rhizomes. But, at the other extreme, there's been a more

rigorous engagement with these concepts and where they originate from. To be sure, we cannot talk of deterritorialization without recourse to territorialization, reterritorialization and the refrain. This has an important bearing on how we conceptualize consciousness and nonconscious through events and assemblages.

Perhaps this is a perfect moment to move on to our next topic, which is consciousness, or more precisely those emergent conscious slices of cognition, which you say are ignored in a lot of new materialist work on the nonconscious.

To begin with, it's important to note that while other critics have contested the version of the neurosciences deployed by some new materialists (see e.g. Wetherell), we find that you similarly draw on Libet and Damasio as a starting point for grasping the nonconscious. Indeed, I'm struck by your initial enthusiasm for new materialism as an alternative to the linguistic turn—comparing it to a "burst of oxygen to a fatigued brain" (2017, 65). This concurrence does not, however, last for long. The challenging question that you pose for new materialism instead concerns why emergent consciousness is often missing from discussions on the nonconscious. As follows, I'd like to make two points.

Firstly, although I agree with you that many writers do overly focus on the nonconscious, I think this is for the reasons that you admit to; that is, new materialism is initially driven by a need to readdress the bias toward the anthropocentric subject. Again this is a shared point of interest. However, in many cases, the idea of emergent consciousness is not, I suggest, *missing*, but is instead *repositioned*, and to a great extent, *weakened*. This has clear implications for the use of a cognitive frame modeled on human subjective experience to explain the nonconscious of, say, a nonhuman.

By way of example, I know that elsewhere you used Thrift's technological unconscious to great effect (reconceiving it as a technological nonconscious), so it's interesting to quote, at length, Thrift's (2007) backpedaling response to the question of consciousness in his work. Here he presents a weakened, repositioned conscious cognition (supported by Damasio's thesis) and notes the importance of *precognition*. He says:

> [C]onsciousness can be depicted as though it hardly existed, as an emergent derivative of an unconscious. Yet it is clearly dangerous to make too little of cognition, as I perhaps did in some of my early papers. Because it is so weak (though hardly as weak as some commentators have depicted it), it has enrolled

> powerful allies which can focus and extend conscious awareness—various
> configurations of bodies and things which, knitted together as routinized
> environments, enable a range of different technologies for more thinking to
> be constructed (6-7).

We can return to the technological nonconscious later on, but here I note how, similarly drawing on Damasio and Thrift, Grusin (2010) offers a theory of affect in relation to the premediated human encounters with digital media, following, in part, a neuropsychology approach that insists upon "the inseparability of cognition from affect or emotion, often on the priority of affect and emotion to cognition and rational judgment" (78). There are many other examples of where cognition is not necessarily ignored, but weakened and repositioned in this way.

My second point refers to a notion of "unthought" I develop in *The Assemblage Brain* based on Deleuze's Whitehead-inspired *The Fold*. This conceives of a kind of unmediated non-subjective experience that is well explicated by Steven Shaviro in his book *Discognition* (2015, 17-18). Similarly influenced by Whitehead, Shaviro begins by acknowledging the kind of point you make about those new materialists who ignore consciousness having to admit that as they write about nonconceptual experience they do so through the conceptual experience of consciousness (as I just did above). There is, as such, no avoidance of, from a human perspective, cognitive consciousness. This is an unassailable fact, it would seem. Unreflective experience must itself be reflected on. There is no raw unmediated experience *without* concepts.

Fair enough, because, yes, if we limit ourselves to a human perspective of conceptual consciousness I think anthropocentrism slips back into the debate. In terms of human sentience, we ignore an opportunity to think beyond subjective thought or consider how to conceptualize sensation in ways that do not necessarily lead to meaning (as a human process of reflective interpretation and conceptual work). Of course, it seems that a concept of sensation can never escape the concept. I would therefore argue that we need to take on board Stengers' observation of Whitehead's concept of nature (nonbifurcated experience); that is, consciousness is not a "command post," but a mere "foothold" in the events of the world (Whitehead 1964, 46). As affect theory posits, conscious cognition is weak. Evidently, given the slight foothold consciousness has in the world, humans can reflect on, to some extent, their own sentience—to know who it is that feels. This is what Whitehead calls an example of the extreme plasticity of nature. But why should that mean

that we discount sensations that are not available to consciousness? The point I make in *The Assemblage Brain*, following Bergson, is that consciousness seems to block access to nonconceptual content. It filters out the unmediated experiences of the nonconscious. Are these not the same lessons we learn from neuroscience too?

Likewise, in terms of the nonhuman, Shaviro draws attention to *nonintentional sentience*. Plants, for example, "feel" or prehend (in the Whiteheadian sense) the world they encounter not in the manner that humans experience it. There is no self-concept in plants, I assume. We do not have to enter into the complexities of panpsychism here to see how nonintentional and nonconceptual sentience might be distributed throughout the world in ways very different from cognition (conscious or nonconscious).

So overall I would say that rather than ignore consciousness, new materialism repositions it and tries to reconcile the ungraspability of a subjectless experience (human and nonhuman) through the theory of affect.

KH: Tony, thanks for your comments. I want to point out, first, that in my view cognition does not exclude affect but considers affective responses to be forms of cognition. Similarly, I would not say that plants, which are minimally cognitive by my definition, are not "nonintentional" but rather have intentions (ultimately, to survive and reproduce) and are capable of creating meanings within their contexts. These issues raise questions about meta-strategies, the positioning aspects that you foreground in your response. One way to go is to jettison meaning and intention and to emphasize drives, mobilities, intensities, etc., the route that many new materialists choose. Another way to go is to extend cognition, intention and meaning to nonhuman subjects and computational networks, which is my preferred route.

Why do I choose this route, and what are its advantages from my point of view? First is a desire I share with new materialists, namely to avoid, as much as possible, the constraints and biases of anthropocentrism. Also important for me is the possibility of building bridges between biological lifeforms and computational media. Ever since Searle's Chinese room thought experiment, folks interested in computational networks have faced the challenge of asserting that computers can create meaning, that they do more than just matching or processing numbers. This is a pressing concern in the contemporary world, where human-computer assemblages are now indispensable for much of the work that gets done in developed societies. Yet there is a scarcity of approaches that can talk about these assemblages in ways that go beyond the HCI vocabularies or the programming focus of computer scientists.

To build bridges in ways that do justice both to human and computational capacities, it is necessary to find common terms that can be used to describe both and yet that are attentive to the huge differences in embodiment between the two. I do not find the Deleuzian approaches helpful in this regard, but rather obfuscating rhetorics that create a gap rather than a bridge. Perhaps this is because the primary agents driving change are drives, intensities, deterritorializations, etc., for which it is difficult or impossible to find corollaries in computational media, at least in terms that anyone working

Neuron inspired by Santiago Ramón y Cajal, Dorota Piekorz, 2017

in computer science would recognize. My approach is thus to broaden and re-define the key terms-- cognition, intention and meaning-- in ways that recognize the importance of biological embodiment and yet can also extend to designed and purposeful devices in computational networks. I can see your point about the usefulness of Whitehead and Deleuze for thinking about biological organisms, and I understand why you and many others may choose this route. Years ago, I talked with a speedboat designer who made the point that all the parameters are known—speed vs. stability, for example—and it is a matter of choosing optimal configurations for different purposes. I think somewhat the same is true of critical exploration: what path one follows depends on the ultimate goal one has in mind. I am reading further into "The Assemblage Brain" and will have more to say about your approach in subsequent installments. Thanks for bearing with me in the meantime.

TS: Once again, thanks Kate. So, it seems that we have more convergences and divergences that need fleshing out here. Let's stay with affect and cognition for now. As I proposed earlier, new materialism does not totally ignore cognition, but repositions it. I can now see how you differently reposition affect as part of cognition. Interesting! Let me put forward my position on this.

One reason I think new materialism regards affect as inseparable, yet not part of cognition, is due to the differing relations it [affect] establishes with feelings and emotion. These are distinct terms that are often confused in a lot of literature associated with broader affective and emotional turns. There is, however, a specific process put forward in affect theory wherein the nonconceptual experiences of the nonconscious (registered in the intensities of affect and immediate experience) feed through to autobiographical registers of previously felt experiences and eventually emerge as a kind of emotional cognition. So in this account, affect is not cognition, but rather emotion is regarded as a cognitive aspect of emergent affect and feelings. Emotional cognition is, as such, the capture (and some say the closure) of affect. In other words, emotion is how affect becomes conscious. There is, funnily enough, a kind of Deleuzian "levelling up" process here, whereby what we think is presupposed by what we feel at some deeper level. This is what I think Shaviro means by nonintentionality: noncognitive affect.

That aside, I'm very interested in your ideas about building bridges between humans and computation, and how these bridges might be extended more broadly to other nonhuman worlds too. I couldn't agree more. The various links currently being made between heavy social media use and mental health issues like addiction and compulsive behaviors, for example, require urgent interdisciplinary attention. I therefore get what you mean about the need to effectively communicate so that those working in computing can understand the points being made by psychologists, for example. There's also an urgent need for us all to address the Anthropocene, of course. Assemblages are crucial for this task. I'm similarly interested in how the humanities can take a less aloof position and more closely engage with the sciences at the front end of a project rather than at a point later down the line when it's too late to make a difference.

There's another divergence here, however, with your choice of analogical thinking. In my opinion, there are too many weak parallels in the analogy between information machines and brains. On one hand, I can see how a computer can be regarded as uniquely cognitive. There are, evidently, many high level cognitive processes at work in computers (calculation, data interpretation and decisions). They are certainly cognizers, as you describe them, and in this respect can often outperform human cognition. The computer is, after all, a very successful product

of information science. On the other hand though, doesn't the analogy break down when we try to describe human brains in the same context? I would say yes, since unlike computers, human brains are more than mere data processors. As Damasio argues, the deciding brain is awash with affect, feelings and emotion.

For me, the technological nonconscious is all about the relationalities of the assemblage, not an analogy between computational and biological information processors. I think this is important to our understanding of human-computer assemblages since it is not the computer itself, but instead it is increasingly the *relation* between the human and the computer that is becoming nonconscious. Indeed, I've already noted the role of precognition in the technological nonconscious (see Richard Grusin's use of it in his post 9/11 *Premediation* book). I'm also interested in developments in emotional and affective computing were progress is, it seems, inhibited by the information machines' inability to feel. As Shaviro again points out, computers *read* emotion; they do not feel. Sentiment analysis and emotional AI performs like this. This is what facial recognition software and EEG also do; they read states associated with emotional cognition. Likewise, GSR (galvanic skin response) is supposed to get closer to the so-called affective valence, but similarly this technology simply *reads* a state of arousal. Of course, computers can respond to these kinds of input (they can learn, infer and anticipate), and that's where I think the danger lies, but they are like actors in the sense that they can only *express* emotions. Even if conscious emotions did emerge, and I don't see much evidence of that right now, we wouldn't know what kind of feeling was being felt.[5] All affective computing can do is process and act as a vector for the expression of human emotion as a data input/output. In terms of affect theory, then, we might say that computers pass on affect (in a way then they can be affected and they can also affect), but they cannot feel it.

This relational aspect of affect theory isn't solely attributable to Deleuzian rhetoric either. I work with psychologists here in London who are running digital media and mental health projects. They take a more nuanced position, for sure, but, nonetheless, refer to a very similar kind of affect theory. I also don't think affect theory is alien to others working with computers. In HCI, for example, affect plays a major role in what has been called third paradigm research (Harrison et al, 2007). A good example of this is Donald Norman's (2004) *Emotional Design*. For Norman experience is processed in the brain through three interconnected levels: reflective (cognitive), behavioral (use), and visceral (affective). He explicitly references Damasio in this book. It's an interesting account and one that has

not surprisingly been met with challenges from within HCI, and particularly those using a phenomenological approach. Some of these people criticize Norman for equating emotion with information (similar to the criticism of Damasio by Bennet and Hacker I mentioned earlier). He certainly talks about emotion in terms of information flows, which I find very problematic. Others criticize his counterpoising of cognition and emotion. They argue that emotions are not the opposite of cognition, but like cognition, they are made in social and cultural interactions. Again, for me, affect theory takes this all a step further without the baggage of the overly subject–mind–centered appeal we find in phenomenology.

KH: Thanks for your comments on affect theory and your view of the kind of work they enable. The terms can be confusing; what Shaviro calls affect, Damasio calls emotion; what Shaviro calls emotion, Damasio calls feelings. Nevertheless, it makes perfect sense to me that the body processes sensory and contextual information before consciousness becomes aware of it, and that the amount of information reaching consciousness is always less than is incoming through sensory channels and interior processes. Whether these processes count as cognitive or not depends on how one defines cognition. In my definition, they are cognitive, as indicated earlier in my comments about biosemiotics. The expansion of cognition beyond consciousness/unconsciousness into nonconconscious lifeforms and into bodily processes for humans makes it possible to think about cognition as a broad spectrum encompassing all lifeforms. It also makes it possible to distinguish lifeforms from the nonliving, which for me is an important point, in contradistinction to others who want to see the boundary as highly permeable and in fact disappearing altogether. This is something of a nuance, because I also see the boundary as permeable and fluctuating, but I still want to preserve cognition for the living in the biological realm, and for computational media in the technical.

Which brings me to the issue that you raise about computers not having emotions, and beyond that, to the related issue of the profound differences in embodiment between computers and humans. I am in complete agreement with the point you make about computers simulating emotions but not feeling them. When we think about biological lifeforms and computational media in broad strokes, one of the prominent distinguishing features is that the living are formed by evolutionary dynamics where survival and reproduction take center stage, with functional, morphological, and behavioral adaptations to the environment emerging as epiphenomena from the primary dynamic of natural selection. With computational media it is the inverse. They are designed rather emerging through evolutionary forces, and designed for specific purposes they fulfill in the world.

But on further thought, we might see evolution and design as background and foreground to one another. Computational media also experience evolution of a sort as their fitness for specific tasks evolves. Usually this means rendering one platform obsolete, for example, and going to another one, something that natural evolution can rarely if ever afford to do. So computers evolve artificially through different instantiations that humans design for them. There are fitness criteria of sorts, but they come from humans who imagine, construct, and implement the purposes they want computers to fulfill in specific environments. Hence the common terminology of computational media as different "generations" of devices.

You mention that of special interest to you is the relation between computers and humans. I too think relationality is key, but the problem here is that for many situations, that relation is multilayered and infrastructural, and thus largely invisible to most humans who are in fact engaged with it, although they may not consciously realize it. A good example is the control tower at a busy airport, where there is intense engagement between the humans and the computational media; aboard the aircraft, the pilots are engaged with the actors in the tower as well as with the onboard computers. A lot of information is flowing very quickly through this cognitive assemblage to make sure everything goes smoothly. On board the plane, however, the passengers see the landing strip appear and may hear the pilot's announcement, but the rest of the assemblage is largely invisible to them and probably outside their awareness altogether. So how can we think about "relation" in these terms? The passengers are certainly "in relation to" what is happening in the control tower in some sense, but this relation is indirect and highly mediated for them. If they think about it at all, they probably vastly underestimate the importance and complexity of the computational media involved. The same kind of situation obtains in most complex assemblages that make everyday life go (more or less) smoothly in developed societies, from stop light timing to water delivery to the electrical grid to millions of other goods and services depending on computational media. The net result is a kind of blindness of most people to the extent, pervasiveness, and criticality of computational media to their daily lives, of which their laptop and cell phone are only the most visible tip of the iceberg. These they understand themselves as "in relation to," but what about all the other infrastructural mediations? These are the kinds of "relations" that I hope to address through the idea of cognitive assemblages.

This is why I am experimenting with the idea of biosemiotics and overlapping (never entirely coinciding) *umwelten* of humans and computational media. Yes, there are profound differences in embodiment, but there are also functional homologies. A homology differs from an analogy in being far more constrained, specifically in terms of the similar functionalities that constitute a homologous series. An example of a homologous series could take the form of comparing forelimbs on a human, a lion, and a whale. The different morphologies notwithstanding, what makes the comparison work are the similar functionalities that limbs possess across different phyla. Similarly, there are functional correspondences between the kinds of cognitive activities that computers carry out and those that humans do. This is not to say that brains operate like computers as posited in the computationalist model; we know this is not the case. Nevertheless, computers have intentions, make selections, and perform interpretations on flows of information. They also have a view of the "world" as it is constituted through their designs. With an in-depth understanding of how the "world" looks to them (their *umwelten*), we can arrive at a much more precise understanding of how "relations" are constituted with humans in specific cases.

With the development of neural net architectures and deep learning algorithms, the kinds of self-learning that computers do come much closer to human processing of information, with astonishing results in machine translation, competitive play, circuit design, etc. The linchpin for me that holds all this together is cognition, defined broadly as the contextual processing of information that involves interpretations and choices that lead to meaning. I think it is important to recognize that computational media do produce and process meanings, both for themselves and for other devices, and of course for humans. That is the underlying homology that makes a cognitive assemblage work.

I welcome your thoughts on all of this.

TS: Thanks Kate! Again, I would say that your efforts to get to grips with the overlapping of nature, culture, technology and biology present us with another point of convergence in our discussion. This orientation toward homology, however, draws attention to two distinct assemblage theories. As you say, this is an academic matter of what course one decides to follow, and I can see how we ultimately end up in a fairly similar place, but before we move on I think it important to distinguish between these two theories.

The assemblages I'm drawn to cautiously approach the kind of resemblances established through analogy or homology. Indeed, rather than look to similarities between function, form, or structure to explain how an assemblage comes together, the focus crucially shifts to affective capacities and differential relations between bodies. We need to go back to Spinoza to fully understand how this works,[6] but in short, assemblages (or abstract machines) are about distinctiveness rather than similitude. It is the relational capacity of a body to affect (and be affected) that takes precedence over comparisons between bodies. This is ethology, as opposed to homology or analogy, in which it is the imbrication of relations rather than comparative mapping of forms or functions that matter. I therefore appreciate what you mean when you say that your homologies never entirely coincide. Assemblages are certainly not a jigsaw puzzle.

Interference as Method, Francesco Tacchini, 2016

There are a number of advantages to this approach, I think. Firstly, I doubt that comparisons can entirely predict what a body can do. We cannot know, for example, the cognitive intentions of a plant, and even if it were possible for a computer to develop intentions of its own, we would not recognize them; certainly not by way of analogical or homological comparisons with our own sense of intentionality.

Secondly, this approach is not constrained to cognizers alone, since what you call noncognizers have affective capacities too. As a result, I'm more interested in inclusive assemblages that comprise the relations established between humans and nonhumans; that is to say, human bodies, technology, geology, climate, and so on. These might be considered as nonintentional decisions made in relation to events—a hurricane responding to climate change, for example, or more specifically, a storm changing direction due to sea temperatures that are affected by human technologies. So, humans are not cut out of these assemblages. This is a kind of deciding that exists outside of the resemblances of the cognitive frame, but nonetheless implicates human cognition in relational processes of sense making.

Thirdly, through the focus on affective relational encounters we can at the very least point to the transformational interactions between bodies. In terms of power then we can see which body has the most potent capacity to affect. I'm not suggesting, however, that this is the definitive route without disadvantages. I grasp some of the problems with regard to science in particular, where I see that analogy and homology are fairly ingrained. I've attended a number of conferences where Deleuzians and scientists have attempted to dialogue. Some have been more successful than others. For my part though, I'm more interested in the cultural and political contexts in which potent capacities are assembled.

This leads us to our final discussion point (utopia/dystopia). There's been much public debate in the UK and US (post Brexit, Trump) about what kind of dystopia we are currently in. Much of this seems to be rooted in digital cultures of fake news on social media, Trump bots, and various "outside" interferences with the "democratic" process. I've followed Neil Postman, to some extent, insofar as I have compared current dystopias with Orwell and Huxley as a way to *fabulate* digital culture. I've recently read a nice piece in the Boston Review that argues that Philip K Dick provides a much more accurate dystopic model of what's going on.[7] I suppose it's my miserable English disposition, but I am openly dystopian in my take on digital culture. It's very refreshing therefore to find that that *Unthought* has a final utopian message.

This focus on dystopia/utopia draws attention to another point of convergence in our work in a similar recognition of the "dangerous," and I would add, dystopian legacy of cybernetic control. Moreover, I welcome your effort to look beyond this kind of control to what lays outside of what you describe as the failed project of computable cybernetics: the incomputable, the undecidable, and the unknowable. Along with Jussi Parikka, I've similarly been interested in the accident and anomalies of digital culture. Our subsequent work on digital contagion and virality starts by theorizing the accident. *The Assemblage Brain* likewise follows this trajectory by looking at the incompleteness of control through Burroughs's influence on Deleuze's Control Society thesis.

So I approached your final chapter "The Utopian Potential of Cognitive Assemblages" with a lot of interest. Indeed, there is a lot in it that I agree with. Those in the humanities, for example, should certainly commit themselves to "ethical responsibilities and positive futures" in digital cultures, as well as making ethical interventions that fully understand how the operations of the computational media work (2017, 204). I agree that the humanities has indeed felt threatened by the pace and complexity of technological change. Not least because being technophobic often seems to lead to a submissive ignorance of how things actually work, but also since there's been a violent devaluation of the humanities and arts in terms of cuts to funding. Digital humanities seems like an understandable response to this attack and needs to dialogue with the humanities, as you say.

My point of departure is, evidently, an insistence that that we should also look beyond the cognitive frame to this differently oriented nonconsciousness I've tried to describe. I agree that most people are indeed "blind" to their relation to the operations of digital technology. As follows, the cognitive "reading" of humans, nonhumans, and their environments by, for example machine learning programs, is part of what I have similarly attributed to *experience capitalism*.[8] This is not so much about the cutting out of the human mind from the assemblage though, as Hansen seems to contend, but rather it is the exploitation of the mere foothold consciousness has in these technological systems. For me, it's not about rescuing human cognition from an invisible operational media (has there ever been a time when the human mind had a command post in media systems?), but instead alerting users and educating them about the ways in which their sensory experiences are operated on in ways similar to R.D. Lang's politics of experience.

The politics of [user] experience increasingly happens in the affective realm through appeals to feelings and emotions. This is why I've been interested in working with social psychologists like Darren Ellis and Ian Tucker in the UK who variously explore experience via affect, feelings, emotions, mental health, social media, and bodies in relation to technology rather than measuring discrete bodies according to normative conditions of health. There seems to be a bigger project needed here that demands that we not only get the humanities talking to the digital humanities and computer scientists; we also need to reach out to psychologists, industrial designers, HCI and digital marketing people. This is necessary since we are seeing efforts to produce habit forming experiences with social media designed to trigger intrinsic negative emotional responses.[9] This so-called dark UX is a part of what we might also call the nasty side of the affective turn wherein technology is designed to exploit the nonconscious through joyful and negative feelings as a mode of control.

There are a number of commercial products emerging from MIT's affective computing programme, for example, that begin with ethically motivated research into autism, but end up with applications in digital marketing or workplace surveillance that are a cause for concern.[10] As we have agreed in this discussion, these applications are limited to merely "reading" affect, but it's the subsequent priming of experience that I think needs attention here; how, that is, a certain kind of subjectivity emerges in the production of user experiences. So the difference is not so much in the decline of human experience of technology, but rather acknowledging how computational media experiences the human.

KH: In your useful clarification of the kinds of assemblages to which you are drawn, you point out these are assemblages connected not by homologies of form or function but the relational capacity of bodies to affect one another. This works well for your project of critiquing neurocapitalism and affective capitalism, showing how the affective capacities of humans are targeted for marketing purposes. However, one of the goals of my project, as you know, is to create a framework in which humans, nonhumans, and computational media interact with one another through what I call "cognitive assemblages," assemblages through which information, interpretations, and meanings circulate. The problem for me in emphasizing affective capacities over cognitive ones is that computers do not have emotions. Even the field of emotional computing only simulates emotions but does not actually create them in computers, as you point out. So any framework that leaves cognition out of account or underplays it does not work well for the integration of computational media into hybrid human networks; nor does it have much explanatory power about how human-computer cognitions interact

to create and extend the infrastructures on which contemporary life in developed countries increasingly depends. In fact, in your account, computers frequently are positioned as agents of exploitation (which of course they can be) and as reductive machines that threaten to reduce humans to the engineering terms that explain them. This is a valuable and necessary project, but it does little to illuminate how computational media play essential *positive* roles in creating the world as we experience it. It also does not allow us to see the extent to which human and mechanical cognitions are increasingly entwined with one another in everyday transactions other than those directly connected with capitalistic marketing.

I think this fundamental difference in goals explains a lot about our different approaches. You say, for example, that "even if it were possible for a computer to develop intentions of its own, we would not recognize them." But computers do develop intentions of their own all the time, and the people who design, program and maintain them know perfectly well in what senses these intentions are manifested within the computers and how these intentions shape the kinds of communications that take place within human–computational interactions. For an example, consider the dictation program called Dragon Dictate. The program is designed to solicit user feedback that will allow the program increasingly to refine its sense of a user's distinctive pronunciation and vocabulary. The user does this by repeating words that the computer gets wrong and typing in corrections for the computer-generated text. The program's intention is to arrive at a textual representation that accurately reflects a particular user's vocabulary, pronunciations, and other speech idiosyncrasies. Intentions are often associated with "aboutness," and here "aboutness" includes the program's ability to detect the modulations of air that create sound for humans. The program does not hear sound as humans do, but it has sensors and actuators to create digital representations of that sound within its memory and databanks.

This example can be multiplied thousands or millions of times, as computational devices are increasingly interfaced with a huge variety of different kinds of sensors and a similar multitude of actuators. Of course, in analyzing how these interactions take place, it is crucially important to include the affective capacities of humans, and that is why I define cognition in such a way that it includes affect.

This focus on cognition also brings up another difference in our approaches, again related to the different kinds of goals we have in mind. Your focus on relationality and the potential of bodies to affect one another tends to blur the line

between living and nonliving forces. If the question is how the forces of living and nonliving events interact, then the tendency will be to consider hurricanes, tornadoes, or even something as simple as water percolating through a rock as all instances of bodies affecting each other through the forces they exert. But in this approach, what tends to drop out of sight is the flexibility (or plasticity, as you discuss it, following Malabou) of living entities to respond to their environments. Rocks do not make choices, perform interpretations, or exhibit the kinds of plasticity that living organisms routinely demonstrate, even plants, even biological entities as small as a cell; (I would say, incidentally, that plants do have intentions, and that their intentions have been studied extensively by plant biologists). By contrast to the plasticity of living systems, the behaviors of nonliving entities can be explained as the resultant of all the forces involved in the interaction without needing to take choice and interpretation into account. Indeed, this is precisely the goal of fields like materials science and stress engineering, which have developed sophisticated methods to account even for critical phenomena so fickle they cannot be accurately predicted but can be successfully modeled using simulations.

My comments [above] are not intended to imply that my approach is better than yours, but rather as an observation that we each have certain goals in mind and have devised approaches that we consider appropriate for those goals. These differences notwithstanding, it is interesting that we both arrive at similar endpoints, although by fairly different routes. I have enjoyed our discussions and want to thank you for your generous engagements with my lines of thought (lines of flight?). Speaking of utopia/dystopia, I will conclude by noting that the implicit utopian hope performed through our discussions is that we can have reasonable and illuminating conversations with others whose agendas and interests partially overlap and also significantly diverge from our own. Thanks for making this possible.

TS: Well, it's very agreeable that we end here on such a fine utopian note. I agree—open dialogue is essential. That is to say, we need a dialogue that is not just limited to the humanities and the digital humanities, but moves outside of these subject lines to the so-called interdisciplinary nexus. We began our discussion here in such a fashion by noting the influence of the neurosciences, and the idea of the nonconscious in particular, on the humanities. I hope you can join us in London in the near future to carry this important discussion forward, and further consider the role of the nonconscious in our encounter with the brain sciences, computer science, HCI, industrial design and so much more.

Endnotes

1. Bennett, M. R and Hacker, P. M. S. (2013). *History of Cognitive Neuroscience*. Chichester, U.K.: Wiley- Blackwell; Harrison, Steve, Tatar, Deborah and Sengers, Phoebe. (2007). The Three Paradigms of HCI. Paper presented at Conference on Human Factors in Computing Systems, San Jose, California.

2. "By 'material' I mean matter, energy, and information, not only matter in the narrow sense." Hayles, Katherine N. (2017). *Unthought: The Power of the Cognitive Unconscious*.Chicago: University of Chicago, 218.

3. See recent discussion in Debaise, Didier (2017). *Nature as Event: The Lure of the Possible*. Durham, London: Duke University Press.

4. Wolchover, Natalie. 'Time Crystals' Could Upend Physicists' Theory of Time. [Online] Wired. Available at: https://www.wired.com/2013/04/time-crystals/

5. See Shaviro's "Thinking Like a Computer" chapter in *Discognition*, (2015). London: Repeater Books.

6. There's a more considered effort on this in Cameron Duff's *Assemblages of Health: Deleuze's Empiricism and the Ethology of Life* (2014). Dordrecht, Heidelberg, New York, London: Springer.

7. Farrell, Henry. Philip K. Dick and the Fake Humans. [online] Boston Review: A Political and Literary Forum. Available at: https://bostonreview.net/literature-culture/henry-farrell-philip-k-dick-and-fake-humans

8. Sampson, Tony David. (2018). Transitions in Human Computer Interaction: From Data Embodiment to Experience Capitalism. [online] *AI and Society*. Available at: https://doi.org/10.1007/s00146-018-0822-z

9. See Eyal, Nir. (2014). *Hooked: How to Build Habit-Forming Products*. London: Portfolio/ Penguin.

10. See for example Affectiva or Empatica, at Schwab, Katherine. This MIT Startup is Developing a Fitness Tracker for Your Brain. [online] Co.Design. Available at: https://www.fastcodesign.com/90160775/this-mit-startup-is-developing-a-fitness-tracker-for-your-brain``

References

Damasio, A. (2000). *The Feeling of What Happens: Body, Emotion, and the Making of Consciousness*. New York: Vintage Books.

Damasio, A. (2012). *Self Comes to Mind: Constructing the Conscious Brain*. New York: Vintage Books.

Debaise, Didier. (2017). *Nature as Event: The Lure of the Possible*. Translated by Michael Halewood. Durham: Duke University Press.

Ellis, D. and Tucker, I. (2015). *The Social Psychology of Emotion*. London: Sage.

Grusin, R. (2010). *Premediation: Affect and Mediality After 9/11.* Hampshire, UK: Palgrave Macmillan.

Hayles, N. Katherine. (2017). *Unthought: The Power of the Cognitive Unconscious.* Chicago: University of Chicago.

Hayles, N. Katherine. (2005). *My Mother Was a Computer: Digital Subjects and Literary Texts* Chicago: University of Chicago.

Hansen, M. (2015). *Feed-Forward: On the Future of Twenty-First Century Media.* Chicago: University of Chicago Press.

Hoffmeyer, J. (2009). *Biosemiotics: An Exploration into the Signs of Life and the Life of Signs.* Scranton, PA: University of Scranton Press.

LeDoux, J. (2003). *The Synaptic Self: How Our Brains Become Who We Are.* New York: Penguin Books.

Malabou, C. (2008). *What Should We Do With Our Brains?* New York: Fordham University Press.

Norman, Donald. (2004). *Emotional Design: Why We Love (or Hate) Everyday Things.* New York: Basic Books.

Parikka, J. and Sampson, T. (2009). *The Spam Book: On Viruses, Porn, and Other Anomalies From the Dark Side of Digital Culture.* New York: Hampton Press.

Postman, N. (1985). *Amusing Ourselves to Death: Public Discourse in the Age of Show Business.* London: Penguin Books.

Sampson, Tony D. (2016). *The Assemblage Brain: Sense Making in Neuroculture.* Minnesota:University of Minnesota.

Shaviro, S. (2015). *Discognition.* London: Repeater Books.

Stengers, I. (2014). *Thinking with Whitehead: A Free and Wild Creation of Concepts.* Cambridge, MA: Harvard University Press.

Thrift, N. (2007). *Non-Representational Theory: Space / Politics / Affect.* New York: Routledge.

Whitehead, Alfred N. (1964). *The Concept of Nature.* Cambridge: Cambridge University Press.