Body-Knowing and Neural Networks:

Is a computer's ability to learn human skills a victory for reductionism?

Charles W. Lowney

Abstract: Harry Collins argues that while Collective Tacit Knowledge is irreducible, Somatic Tacit Knowledge is explicable and reducible. He provides evidence by showing how a machine can learn and perform tasks that were once thought to be irreducibly human (or animal). Reductionists would agree that if a machine can perform a task, it is in principle explicable. I show how Polanyi's structure of tacit knowing and learning provides strong parallels to connectionist (Neural Network) architectures and dynamical systems-- the very sort of computational processes and strategies that have achieved the most success in machine learning and hold the most promise for AI. In contrast to Collins, I argue that connectionist systems display an irreducibility. Ontically, machines and living beings are what Polanyi calls "dual control systems" that display emergence. Epistemically, explanations in terms of component parts/inputs and their laws/relations are insufficient (Polanyi), and explanations via higher-order conceptions wholes/outputs manifest a "dimension shift" (Paul Smolensky) that creates a mismatch with lower level processes.

Machines are learning better and better to perform skills that were once thought to be uniquely human—such as driving a car or learning a language. We have made these advances through connectionist (also called "neural network") parallel processing systems. Even skills that Michael Polanyi believed to be beyond a computer's ability, such as recognizing a face, can now be successfully performed by machines. Does this mean that all our skills can be reduced down to matter in motion? Are these technological advances a victory for material reductionism?

In order to distinguish ourselves from dead matter and our free actions from causal necessity, we humans have attempted to draw lines that distinguish us from our animal bodies and from machines. Polanyi suggests a different tack. To find our freedom and creativity, we should instead look at what we have in common with animals and machines.

Following Polanyi, I recommend here that we rethink the way we draw the lines between the reducible and irreducible. First, I will show strong similarities between structures of tacit knowing and connectionist architectures, both in how they process information and how they learn behaviors, but I also show how these computational methods mimic the irreducibility (irreversibility, unspecifiability, and inexhaustibility) that Polanyi sees in tacit knowing and skill learning. Next, I will examine whether or not the ability for such a machine to model tacit learning and knowing makes that knowledge reducible. I will examine Harry Collins' reasons for answering "yes," (they are reducible to material causal strings) and then provide Polanyian reasons for saying "no" (the material causal laws underdetermine the emergent functionality). In looking to our bodies, and to computing systems, we will find a tacit learning and knowing that is irreducible to dead matter and its immutable laws. In conclusion, I will briefly reflect on what this means for the process of evolution, the value of animals, and the prospects of Artificial Intelligence, but in order to emphasize the importance of this issue for human beings, I will first briefly reflect on what is at stake.

Are we mere aggregates of matter in motion? Slaves to the determinate laws of physical causality? Or are we irreducibly special in a way that makes us free? Over the centuries we've drawn different lines demarcating the ordinary from extraordinary and the material from the free. It seems, however, the advancement of knowledge has continually forced us to give up ground to deterministic causes and material reduction, and retreat behind a smaller and smaller perimeter, but this, I argue, is an illusion.

To be special, we first primarily differentiated ourselves from dead matter. *Living beings* were special and different. Then, as we came to see the workings of animal bodies as the workings of machines, we came to see living bodies as reducible to mere matter as well. So, we drew a new line. We said that what makes us special is something else: our *reason* and ability to use *language*. But with the advancement of computer technology, that line, too, was in jeopardy. As mechanical procedures and computers began to do more and more of what we called "thinking," phenomenologists like Maurice Merleau-Ponty and Hubert Dreyfus showed us a different way to draw the line, between body-knowing and explicit propositional knowing, or as Gilbert Ryle would say, between "knowing how" and "knowing that." This body-knowing also displayed a creative human way of being in the world that resisted analysis and reduction. We see it in improvisational dance, jazz music, and martial arts.

But while embodied knowing seems on the right tack, it can also draw a line that is misconceived and, as such, is subject to erosion. Dreyfus, for example, in attempting to protect our "knowing how," from reducibility frequently used the example of driving a car. A computer, he claimed, could *never* drive a car (*What Computers Can't Do*). It was not that the task was too complex, but that it relied on irreducible human skills that could not be laid out in terms of the sort of explicit instructions a computer would require.

Harry Collins draws different a line: *Collective* Tacit Knowledge (CTK) cannot (or is very unlikely to) be reduced. One of his examples, in *Tacit and Explicit Knowledge*, was that a computer would never be able to drive a car in traffic, on a complex roadway, to a destination. He, too, saw this as not simply a matter of the complexity of instructions—computers can do complex things—but as a matter of linguistic and social knowledge that cannot be reduced to any a causal string or algorithm.

Both Dreyfus and Collins argue that if a machine can perform a task, it is reducible. But Collins concedes more to the material reductionists. Collins acknowledges that—although it may mix with

CTK—body-knowing or Somatic Tacit Knowledge (STK) can in principle be made explicit. He argues that the body is like a machine, and, ultimately, if a machine can do it, it is reducible. So, he creates another line that pits our creativity and freedom against the rising tide of scientific achievement.

Well, now we do have self-driving cars that are able to navigate through traffic to a preprogrammed destination. Were Dreyfus and Collins wrong? Not really. The problem with these lines of defense against reducibility is not that they do not demarcate important emergent skills and understanding. The problem is that their characterization of irreducibility often makes it seem that we are quixotically fighting against the tide of science, and that we are being forced constantly into retreat as scientific knowledge advances.

Polanyi takes a different approach; he isolates a process of *tacit knowing* that replicates at various evolutionary and epistemic levels. He finds this tacit meaning integration at the physiological, cognitive, linguistic and social levels, and so for him we have an irreducible *process* of tacit integration taking place even at sublinguistic and subsocial levels. Polanyi also links tacit knowing and emergence together, so we might very well have some form of irreducible tacit knowing even at the level of primitive living organisms. The lines Dreyfus and Collins point out are important demarcations. They show emergent skills at highly developed levels that display more degrees of freedom, but they also tend to perpetuate misdirection, since we can observe the process of tacit knowing at all emergent levels.

If Polanyi is right about the limits of reduction in science, we do not have to retreat. We can see each credibly proposed line—life, language, reason, skills, society—as an achievement providing more degrees of freedom, and we can locate innovation and irreducibility in the *process* of tacit knowing and the *performance* of emergent being.

I. Can a machine –or Artificial Intelligence—*do* or *perform* what we consider to be achievements of tacit knowing?

1. Modeling Body-Knowing and Learning with Connectionist Networks

The AI machines that are the most promising for simulating human behavior and skills work with connectionist architectures. They are the sort of computers that train on "big data" and can "learn" the appropriate responses to solve problems. Connectionist systems perform best in pattern recognition, and they are being used to do things like recognize and mimic human speech to give us virtual assistants like Siri and Alexa. While Polanyi himself did not believe that computers could model or imitate tacit knowing processes, he primarily had an ordinary digital computer in mind (see Bocharova 22). In mid-twentieth century discussions, he rejected early ideas about how computers could produce "knowledge."

He insisted, much like Dreyfus, that informal skillful elements (which he later would call "tacit") performed by humans could not be replicated (Polanyi, 1952). Today he might think differently, seeing that many of the skills that he used as examples of tacit knowing are capable of being effectively simulated by connectionist programs. For instance, our ability to take two two-dimensional images and construct a three-dimensional map is now mimicked by cars with "eyesight" technology; and the ability to recognize one face out of a million (LTI 4) can now be accomplished with facial recognition software. Polanyi would likely now notice the similarities in tacit knowing and connectionist architectures, but he also would undoubtedly notice how these networks also mimic inexplicability.

a. Similarities in Architectural, Holistic, and Hierarchical Features

The architecture of tacit knowing is strongly analogous to the architecture of connectionist systems. Polanyi discusses how various "clues" integrate together in a way that provides the focal output of a joint meaning or "joint comprehension."

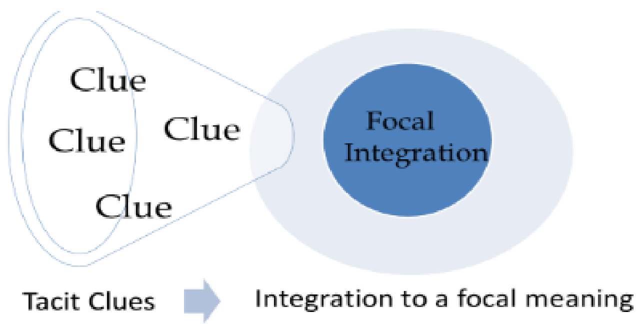

Similarly, in connectionist systems we have many "input units" that connect and organize into individual nodes, and these integrate again in a complex process that produces resultant "output" nodes.

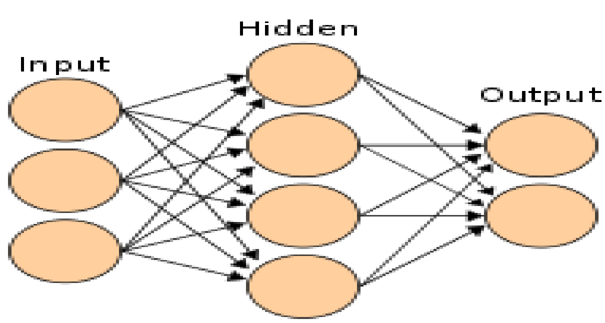

The non-linear vectors from assorted inputs to an output also models the difference in meaning between the clues (inputs) in isolation and their joint meaning (output). Connectionists who model language, for instance, discuss "distributed representations," by which the inputs are sub-symbolic or proto-representations and the resultant nodes are the symbolic representations proper. Different experiences involving the sound "coffee" do not individually amount to the symbol COFFEE (Smolensky

1991). The inputs in isolation are just sounds with some similarities in pattern and provide no such representation (like pixels in a photo), but the output gains a different, holistic significance. The inputs, i.e., iterations of the word "coffee," do not function like a symbol or meaning, while the node COFFEE does perform that function.

Also, some of the distributed inputs that go into the symbol COFFEE might be part of the network that goes into the symbol, DRINK, when a different activation pattern is triggered. This is similar to Polanyi's description of the exercise of linguistic skills, e.g., the way some of the same letters might go into different words, or the same words might go into different sentences, each combination being a joint comprehension of overlapping clues into different meanings. In the exercise of a physical skill, e.g., martial arts, this is similar way some of the clues (muscles in coordination) that go into a punch (which requires balance) might also go into a kick (which also requires balance).

Just as a joint comprehension of the clues/input can provide a higher-level meaning/output, the way tacit knowing can be hierarchized for Polanyi is also mimicked in connectionist systems. As well as operating in parallel to form different nodes, a next layer of nodes themselves can feed forward into another level of processing in a larger system. COFFEE and DRINK, for instance, can be nodes that act as some of the inputs (clues) feeding forward into a node representing DRINKING COFFEE (joint meaning).

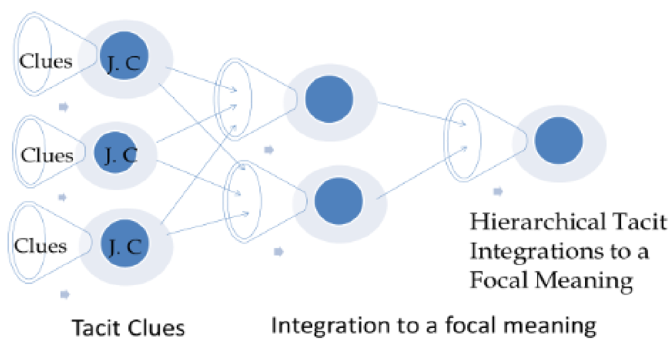

Going back to examples from somatic processes, the layered hierarchization of clues to joint foci can be modeled in connectionist networks, as one node (say the performance of a punch with a step) together with other nodes (stepping forward with a kick) feeds forward into another node (stepping forward with a kick and transitioning to a punch).

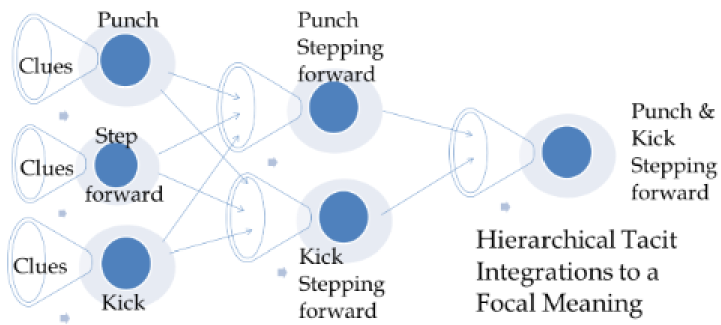

But those input processes also need to be adjusted in terms of output goals at each level for them to work effectively together in response to new environmental conditions (so, for instance, the step has to be adjusted so one can maintain balance and transfer weight on the right vector while striking one's opponent). In order to perform the subordinate and ordinate task properly, each meaning at each level must adjust in relation to each other. The adjusting of the weights between the connections, like the adjustment of the movements and positioning of the muscles, nerves, and bones, is part of a learning process.

b. Similarities in Training the Systems

Connectionist networks, like the martial artist, are programmed to learn how to perform tasks with the help of some coaching, and one task builds upon others. To help the networks learn a specific task, the programmer can consciously modify the weights between the connections. This is comparable to the instructor providing correction to the student, or the student consciously adjusting his or her own bodily movements too align with their current conscious understanding of the goal. The programmer can also provide the network with repeated examples, just as the practitioner can break down individual techniques or movements and train the body through repetition. The practitioner thus engrains or habituates the new associations and activation patterns— like a musician might practice a new chord, so that the output is more effective or appropriate when played in a song.

Polanyi discusses skillful production, specifically in the arts, in terms of "intuition" and "technical invention" (*Meaning* 96-98). When we have a problem to solve and we do not know the end goal or conception and the imagination seeks out a possible solution. An "intuition" is a discovery that is a spontaneous integration of subliminal clues to a joint comprehension that acts as a solution to the problem. "Technical invention," in contrast, is a sort of engineering. We do know the end goal and the discovery comes with the body or mind's ability to organize the clues towards that outcome. When one consciously isolates particular movements in a technique and trains them toward a fixed result, this is similar to the engineering part of the artistic process. While this somatic adjustment in tacit knowing often happens with conscious oversight, it can also happen without it.

In order to learn, the network might not need a programmer to train it and manually adjust weights of certain connections. Networks can be set up with a recurrent system that feeds back information into earlier input units. This backpropagation is the technique that can automatically make corrections towards some given end. In these feedback processes the weights or strengths of connections in the network are changed and activation patterns are altered towards a desired goal without the constant guidance of a trainer.[1] These automatic feedback systems are comparable to the *body itself* working towards the goal of getting the motion right. In developing a skill, the body uses its own intentionality, and plays with possibilities, to work towards and discover the right end goal.

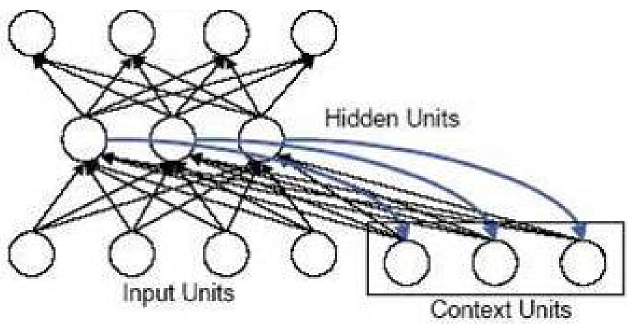

This unconscious somatic fine-tuning works in the way information backpropagates from final nodes or from "hidden layers" within the network. This information makes a difference to the next round of input weights and activation patterns.[2] The connectionist model can thus catch how the body gains experience and improves in executing a technique, or series of techniques, not just by repeating the same but by zeroing in on the right moves over time "spontaneously." For instance, after months of attempting to properly follow instructions, the Brown Belt may finally "realize" a how to put hip power into a punch correctly. This unconscious play with possibilities, and narrowing towards a joint comprehension, is similar to discovery by "intuition" in the artistic process.

c. Similarities in Intentional and Innovative Tendencies.

In reaching a new level of skill, the body itself works back and forth between intuition and invention, mutually adjusting its performance. In both directions, we move intentionally in anticipation of a solution to a problem. The intentionality that we can see in the body's discovery, and the freedom and creativity it can exhibit in an ability to diverge from internal or externally set goals, may be modeled in versions of connectionist networks that use a dynamical systems approach. A network can be set up so

---

[1] We see this in Simple Recurrent Networks as in Jeffrey Elman's (1990) or Michael Jordan's (1986) work; see Bechtel and Abrahamsen,180.

[2] Backpropagation is not necessary for effective learning; a better designed system can learn from fewer examples or training cycles.

that it develops its own goals and learns how to learn in a given environment. The intentionality can be mimicked with preset or emergent "attractors" that increase the probability of one outcome rather than another as the system develops. A dynamic system interacting with an environment, or another dynamic system, can achieve a stabilization whereby it can reinforce a particular pattern, or it can move down a gradient of increasing probability to a new pattern or outcome. The body works with an external environment and not just the feedback it produces in isolation. The body could develop its own attractors and repulsions, and then—when working with a trainer in a wider system—the body's own intentional trajectory can be affected, putting it on a path to a new stable set of actions and reaction. For Polanyi, living systems, and not just conscious human beings, respond to a "telic field" that helps them move forward in ontogenetic and in evolutionary history (PK 403). This need not be the work of an intelligent being, nor a force at the end of the history of time, but it can be an emergent environment that coaxes an entity along a path toward richer meanings or ways of being.

2. **Connectionist Irreducibility: Can we make explicit what is happening in connectionist systems?**

So it seems that connectionist networks paired with a dynamic systems approach might actually do a fairly good job modeling somatic tacit knowing in the martial arts, or language learning and use, but these systems also seem to model or mirror the irreducibility of tacit knowing. Jeffrey Elman (2014), for instance, presents a connectionist "words as cues dynamical model" (137) inspired by the notion that "words do not have meaning but rather are cues to meaning" (138), which seems very friendly to Polanyi's structure of tacit knowing.

Tacit knowing cannot be made fully explicit for Polanyi at least in part because of the *irreversibility* of the gestalt, i.e., the holistic nature of tacit integration. Firstly, in the example of distributed representations, we see an irreducibility in the non-linear gestalt change of meaning from inputs to outputs. Distributed representations are not yet symbols. We do not get to the concept COFFEE by a mere aggregation of iterations of "coffee."

Secondly, we see inexplicability in the tangle of "hidden layers" that intervene in connectionist systems between the input units and the output. Hidden layers of nodes preprocess information to help obtain the solution to the problem or task that the network is being developed to solve or achieve (Bechtel & Abrahamsen 70). These are layers of connections and nodes that are not programmed ahead of time, but which assist the network in attuning itself to the desired outputs. There is an *intractability*, as the process of clue (or input) integration passes through hidden layers whose functionings are not directly accessible to the programmer and which are often too complex to explicate. A recent issue of Scientific American discusses this problem of decomposition in networks using big data (Musser, 2019).[3]

Decomposition is helpful because when a neural net machine tells us what our next move is—for instance perhaps, in stabilizing an economy—we’d often like to know *why* this is the right next move and not wish to blindly adhere to the dictates of an algorithm. The gestalt of tacit integration into new meaning and the inability to turn attention around on the clues to make them all fully explicit is thus modeled in the complexity of the inputs into output nodes that manifest a different quality, and in the intractability of the connections and their weights.

In connectionist learning we see another dimension of the inexplicably tacit connected to both the holistic shift in meaning and its intractability. This is a mismatch between the system as it functions and any efforts to explicate that functioning. This coordinates with the unspecifiability and inexhaustibility of clues that Polanyi discussed when one is no longer "dwelling in" clues/particulars but is attempting to "turn around" on them to make them explicit. First we notice that although connectionist systems are trained to the parameters of rules, they are not operating by those rules. The rules we apply will be higher-order descriptions using a conceptual system that does not quite capture what is actually going on at the processing level of clue/input connection and integration.

In connectionist networks, the organizing rules provided by the higher level are constraints but do not describe the causal processes. This is similar to the way grammatical rules of syntax, for connectionists, are not the true processing motors of linguistic composition. Connectionist systems display rule-like behavior without explicit rules.[4] Smolensky, for instance, describes learning how to play chess. One operates by the explicit rules until a connectionist network is built up that allows one to "intuitively process" the rules and, so to speak, "see" the possible moves (See Bechtel & Abrahamsen 193, citing Smolensky 1988). Connectionist "intuitive processors" are gradually built, but there is no one to one relation between the explicit rules *we* use and the processing rules.

So there are two senses here in which our attempts to make connectionist systems, or martial arts learning and performance, fully explicit will miss the mark. First of all, explication of the tacit by conceptualization is inadequate because the lower-level processes are operating by subsidiary somatic processes, e.g. muscle, nerves, their signals, mechanics and chemistry, and that is not completely captured in the higher-order conceptual descriptions (in, e.g., "Step forward and extend your arm to punch" or

[3] The article discusses how looking back upon an integration (as Polanyi might say), through meta-learning and a disentanglement of representations, not only helps a machine learn, but might present some of the "baby steps" toward make a machine conscious (Musser 63).

[4] See the section “Raven’s Matrices: Rule-Following without Rules” in Lowney, et al. 2020. Regularities without rules goes back to Rumelhart and McClelland's 1986 connectionist approach to the production of past tenses, as mentioned in Bechtel and Abrahamsen (2002, 153).

"Combine a subject with a verb to make a sentence"). There is a processing level at work that may require a different sort of description. Secondly, our descriptions of tacit knowledge (the rules we see consciously as applying) are top-down approximations and are therefore incomplete and inexhaustible (as we need more and more post hoc refinements as exceptions crop up). Both of these senses come together in learning and performing. The artist is tacitly learning the right way to move with the help of explicit rules—but is not strictly following the rules when artfully producing in accord with their training.[5] The rules provide a higher-level description and do not adequately describe the actual lower level processes involved in the network's performance (i.e., the tacit processes).

That the higher-level conceptual explanation is "mismatched" in its efforts to make the tacit knowledge explicit is described well by Martin Davies in "Connectionism, Modularity and Tacit Knowledge" (1989). Davies understands tacit knowledge in terms of the "causal-explanatory structure which underlies, or is antecedent to, the pieces of knowledge that [e.g.] the speaker has concerning complete sentences" (542). Linguistic competence, like a skillful performance, requires a mastery rooted in tacit knowledge, since we are not typically consciously aware of the rules by which we are speaking/acting. The explicit rules that we come up with, i.e., the "tacit knowledge descriptions" (547) are interpretations of what is really going on causally.[6] But the level at which we provide rules is not only a coarser—as opposed to a more fine-grained—explication, there is also a "mismatch between the tacit knowledge description and the network description" (550).[7] Looking back from the higher conceptual level, something like grammatical rules with something like symbols works well as an approximation of the knowledge working tacitly, but there is a "semantic *dimension shift*... [Davies uses Smolensky's phrase, 551]...between the conceptual and the subconceptual [interpreted] description ...of the operation of the network" (551-552).

While this connectionist insight affirms Polanyi's notions of the irreversibility, unspecifiability, and inexhaustibility of tacit knowledge, it adds another layer of inexplicability that Polanyi didn't quite see. Polanyi thought we internalized or "interiorized" the rules and dwelt in them (*Meaning*, 62; *KB*, 214) but the explicit rules for connectionists are higher level approximations that attempt to describe the "intuitive processor," as Smolensky calls it, that develops in response to training. In "internalizing" the rules we are developing rule-like processes, but the explicit rules themselves are not internalized nor are they the actual rules by which we abide when we act. The mismatch comes, according to Davies, from

[5] See my "From Science to Morality: A Polanyian Perspective on the Letter and the Spirit of the Law" (2009).
[6] In modeling language, Davies considers the level for tacit description below the level of our theoretical characterizations, but above the processing or algorithmic level, and, in turn, the processing level is also above the physiological level (Davies 547).
[7] Davies notes here that tacit knowledge itself need not involve conceptualization (551), he thus likes the word "cognize" for tacit knowing, so as to gloss over the pseudo-problem of how the tacit can be knowledge (542).

using the concepts of the higher-level theory (its explicit concepts) to try to describe or untangle the process relevant rules or algorithms that function tacitly. In spite of his misapprehension, this fully affirms Polanyi's notion that a tacit integration of clues can form a higher-level meaning that cannot be a mere aggregate of the lower-level clues and their rules.

II. Does a *machine's ability* to perform what we consider acts of tacit knowing mean that *that knowledge* is now fully explicable?

Can the skillful performance and creativity of the body in action be fully reduced to an explicit formulation? As we have just seen, it is possible that the machine model itself is modeling the inexplicability of the tacit knowing process. But if the somatic processes and productions can be modeled by *any* machine, this, for Collins, is an indication that this tacit is explicable.

Collins has done remarkable work exploring the conception of tacit and explicit knowledge. He has made helpful distinctions, and has provided many examples and insights, that catalogue various sorts of tacit knowing and help to expand Polanyi's conception. Collins, however, makes the most significant cut between tacit and explicit knowledge at the level of the social and the linguistic, e.g., in being able to understand a joke, or when a word is intentionally misspelled. This is indeed an important and distinct level that manifests new achievements and possibilities, but for Collins CTK is the *only* sort that is not in principle (as yet) subject to an explication that would make the tacit fully explicit. He puts the bodily skills of animals and humans on par with deterministic mechanical processes. "[C]ats and dogs and sieves and trees cannot be said to 'know' any explicit knowledge, they shouldn't be said to know any tacit knowledge either. In fact, they don't 'know' anything; they just transform strings" (78).

For Polanyi, as Collins notes by quoting Polanyi, "We may say in general that by acquiring a skill, whether muscular or intellectual, we achieve an understanding which we cannot put into words and which is continuous with the inarticulate faculties of animals" (Collins, 76 quoting Polanyi, PK 90), but this continuity features the unformalizable and unspecifiable aspects of tacit knowing that confound a strict determinism.

1. Collins on Making the Tacit Explicit

Collins' Relational Tacit Knowledge (RTK) is certainly involved in learning a somatic skill; RTK takes place when a master relays knowledge tacitly to an apprentice. In the Martial Arts, there is usually a teacher or coach involved, who demonstrates and explains techniques—though insufficiently until the student grasps the idea with his body. We also see somatic tacit knowledge (STK) in the body's skillful performance, but both RTK and STK are tacit knowledge in a qualified sense for Collins. They are ultimately *fully explicable* for Collins, and only our human limitations mask that explicability from us. The complexity of the network is thus masking an "in principle" reducibility and even connectionist AIs can be made fully explicit. There may be a conceptual dimensional shift, but the processing machine shows there to be a causal string, which is a hallmark of explicability for Collins.

For Collins, only collective tacit knowledge (CTK) is irreducible— "for now " he might add, since it would require techniques for explication that we don't yet have but might eventually develop. In *Artifictional Intelligence* (2018), Collins provides conditions that would have to be met for AI to achieve human tacit abilities, such as passing a rigorous Turing Test that demonstrates "ubiquitous" common sense (See chapter 10). Such an achievement would show human abilities to be explicable and as fully reducible as STK is for Collins. But, as yet, Collins believes CTK to be a special irreducible human ability.

There is indeed CTK involved in the "programming" of the martial artist or the baker; there are social rules and context sensitivity at work guiding when it is appropriate to do what. As Collins says of the martial artist, "one way to win would be to smash the opponent's head with a baseball bat, but that would not be a *right* move, the criteria are collective" (Collins, 2011-2012, 41). My claim is not that productive skills and arts are not nested in higher order social contexts which act as constraints on lower-order meanings. My claim is that there are dual control systems at lower levels as well that function in a similar way to that of the *individual human* to *social-collective* level of functioning, and that those levels also display forms of the creativity we see in higher-level discovery.

For Collins, the actual performance of the technique by the body is a sort of "prosthetic"— a machine response that is tied into social practices at either end, i.e., between natural, pre-social functioning and the functioning within the social sphere. Higher CTK is thus involved in both learning an art and how and when it is performed. That's true, of course, but I don't want our conscious experience of this level to obscure the possibility that—also—in the course of learning and performing an art, the body is also displaying the tacit knowing (particulars *to* whole; individual *to* complex) dynamical process that is both innovative and intentional/telic. And so the performance of less social and less mentally adept animals also displays an irreducible tacit structure—even before language and society— that is also both innovative and intentional/telic. Although it might seem like animals' productive skills are fully determinate and explicit causal processes, I believe this impression comes from a *post hoc* and a flattened view, linked together with a sophisticated yet limited understanding of what makes the tacit explicit, i.e., linked with a reductionist understanding of causal strings.

Collins contrasts tacit knowledge with explicit knowledge. Explicit knowledge can be gained if we can elaborate or transform a process, here a somatic process, into an interpretable "string" (TEK 9,10).

Strings are material things that engage in causal relations. For Collins, human bodily skills and craft knowledge reduce down to physical processes that can thus be explicated in terms of material and causal strings. So although the body often *seems* to demonstrate tacit abilities we cannot explicate, those processes can be made explicit. "*Elaboration*" is Collins' first sense of explicability (81), it basically just means we need more details about the process, i.e., we need a longer string. This is generally how RTK becomes explicit; we write out instructions that were glossed over and passed from master to apprentice unthematized. The details we have, however, may need to be transformed into a different string that provides a better affordance for us to (focally) grasp; this string transformation is Collins' sense two of explication (81). This happens, for instance, when there is *translation* of a text from a language we don't understand into our native language.

According to Collins, if a *machine can be made to do a task* that a human performs, this also acts as the transformation of the string into an explicitly interpretable form. This is Collins' third sense of explicability (81). Although a productive process might work tacitly relative to our contingent psychology, if we can make a machine that can produce the same results, we now have an interpretable string that makes the process explicit. Through this transformation we now have the ability to analyze the machine to its explicit parts and state clearly how the parts function together to produce the desired result. This form of making the tacit explicit is similar to the ability to transform an informal verbal argument into a formal system of logic. In describing the aptness of Collins' conception of interpretable strings to capture the paradigmatically explicit, I noted that mechanical modeling "is an explication because we interpret that mechanical model as a formal language in which the process has been successfully embedded, and the *working* mechanical model acts as a syntactical proof" (Lowney 2011-2012, 24).[8]

2. Machines, Animals and Martial Artists: Operant conditioning and reducible complexity?

[8] In my (2011-2012) paper, "Ineffable, Tacit, Explicable, Explicit: Qualifying Tacit Knowledge in the Age of 'Intelligent' Machines," I show how Collins is indeed on the right tack in connecting the paradigmatically explicit with causal strings that are interpretable by looking at the development of formal systems and their understanding in the history of analytic philosophy. While I provide Collins' definition and note the need for an interpretation, e.g., on p. 21, Collins in his reply to my essay, believes I reduced his definition to the mere presence of a transformation from one string into another without an interpretation (Collins 2011-2012, 41). I agree that transformation is a necessary but not sufficient condition for explication. My point was that it is easy to forget that we are *always* relying on tacit knowledge, as Polanyi claims, and I point to an instance where Collins seemed to slip. The most common reason for such a slip comes from the vectorial quality of tacit inference— is *so easy* to forget that we are always attending from tacit clues when we are looking at something focally, as explicit. *Any* interpretation relies on tacit clues. We see the explicit paradigmatically come to life in formal languages and proofs and forget that these have their significance only against a tacit background context of meaning. This is something Frege saw that the mathematical formalist missed (Lowney 2011-2012, 24).

Looking at connectionism, somatic tacit knowing, and martial arts training from Collins' perspective, it looks like it all might ultimately be the engineering of strings. It seems a matter of building the right string, or programming the computer-machine or body in the right way. Perhaps then, as Collins claims, all the innovation and context sensitivity is happening at the collective level, where the programmer, or my instructor and I, understand and apply the concepts and strategies of the social practice? Looking down from the top it can seem that way. But I believe we would be ignoring the creative contribution of the body and its role in the discovery process. We would ignore the variety of possibilities that the body has chosen from and the numerous variables that needed to be balanced in order to initiate a proper response. We might, in retrospect, flatten out the productive process by selecting one channel to one end result, and then transform that isolated chain into a physically causal machine process. From this flattened perspective it can then seem like there was nothing but rote learning and determinate processes involved.[9] So for Collins, connectionist systems or neural networks and bodies are no more intentional and inexplicable than animals: "cats and dogs—or trees and sieves for that matter" (TEK 77).

Although, as we have seen, the tacit knowing process, as modeled on connectionist networks, has features which appear irreducible, Collins cites several good reasons for why we should see reducibility and a determinate causal chain at work: (1) First, the training of an AI Machine or an animal or a martial artist seems to work like operant conditioning, which in Collins' mind betrays a fully deterministic origin (TEK 75). (2) Second, Collins also points out, quite correctly, that connectionist programs can be run on digital computers. Hence there is a physical transformation from a nonlinear complex connectionist string into a linear digital one (TEK 75). This, to Collins' mind, shows that processes modeled by connectionist networks are intractable only as a matter of complexity; apart from human limitations, they are then in principle reducible and explicable.

Collins is not alone in his assessment of somatic tacit processes as fully determined from below. There are connectionists who would agree. Davies**,** for instance, although he supports a Polanyian view with respect to *knowing* the tacit (epistemologically) would side with Collins (ontologically) with regard to the causal structure of tacit knowing. Even in presenting the mismatch of our concepts to connectionist networks, Davies presents tacit knowledge as being part of a "strict causal systematicity" (553), and Collins would apparently agree.[10]

---

[9] Collins (2018) seems to recognize the illegitimacy of this flattening when it comes to CTK, but not for STK. For example, he says a computer cannot tell when a word is misspelled on purpose, because this requires irreducible CTK (4). He acknowledges that we might then attempt to repair this glitch by adding some lines of code to correct the isolated incident (11), i.e., what I am calling a post hoc flattening, but that wouldn't belie the fact that CTK was involved. I hope to show a similar context sensitivity and flexibility for the pre-CTK body.

[10] Davies says: "Tacit knowledge does not have to be explicitly represented; it can be realized by the presence of a processor. But tacit knowledge is a matter of strict causal systematicity in the transitions mediated by that processor—causal systematicity mirroring the derivational systematicity in the theory that is tacitly known" (553).

A strict determinism here is anathema to Polanyi's conception of tacit knowing. As Jean Bocharova points out, determinism pushes against the unformalizability and unspecifiability that Polanyi saw in the tacit (2014-2015, 22). Bocharova describes how Polanyi rejected an "automatic neurological model" that might be formalized by a computer because he rejected that "the nervous system—functioning automatically according to the known laws of physics and chemistry—determines the workings which we normally attribute to the mind of an individual" (Polanyi, PK 262; quoted in Bocharova 23). Bocharova goes on to point out how the determinism (that we might associate with Davies' conception here) is not endemic to all descriptions of connectionist networks, for example, those developed by Jeffrey Elman (Bocharova 23).[11]

When connectionist systems are combined with a dynamic systems understanding, it is possible to model the ability to learn, e.g., to hone in on a result, and to do so in a way that allows degrees of freedom and even what we call creativity. Dynamic systems can develop attractors that provide a propensity towards a certain outcome, but there are also degrees of freedom that emerge. While Collins identifies operant conditioning with causal determinism, it is not quite that. You may train your dog to stay on command, but if she sees a rabbit, she may be torn between staying and chasing. Similarly, connectionist systems display a flexibility and they display an intentionality that often diverges from the will of the programmer. In this, they show emergent features similar to the kind Polanyi described, by which lower-level entities and processes "enable but do not direct—the emergence of new structures and processes" (Bocharova 23).

Living bodies themselves are emergent for Polanyi. The DNA does not act like a determinate cause in physics or even chemistry. And as Bocharova points out, like Polanyi, some

> connectionists also reject the claim that genes 'code' for higher-order cognition or complex neural structures. Instead of likening genes to a 'computer program' or 'blueprint,' they use the metaphor of genes as 'catalysts' that enable—but do not direct—the emergence of new structures and processes (Elman et al. 1999, 350-351). In this sense, connectionism agrees with Polanyi's view that DNA '*evokes* the ontogenesis of higher levels, rather than *determining* them' (KB, 235; emphasis original). (Bacharova 23)

Although connectionist learning and operant conditioning seem deterministic, they are probabilistic, and their possibilities may be constrained by emergent potentialities that open up in the complex system of a wider or higher environment. We train animals —and our bodies—and

[11] Bocharova cites Elman, et al. 1999, 350-351; and we have seen the same above in Elman 2014.

connectionist systems with something like the reinforcement of operant conditioning, but operant conditioning it is not causally deterministic and isn't the whole story. Since it increases the probability of a certain response it does not act quite like a physical force. One might say it is the implementation of a telic cause or formal cause, rather than an efficient or material cause. It operates in the way a dynamic system can create an attractor when it comes into contact with another dynamic system. Or in the way an open system, for Polanyi can take part in generating a "telic field" that moves it forward towards a particular configuration, e.g., the way the DNA of a developing embryo operates together with the epigenetic landscape (KB 219, 232). Operant conditioning also has its explanatory limitations. It is like taking some of the implicit tacit clues that we can explicitly identify from a complex integrative system and working just with those and peripheral associations (Polanyi, LTI 14). It hardly identifies a causal mechanism— let alone the mechanism as it is indwelt by a living being.

But what of the ability for a connectionist system to run on a digital computer? Does that not prove that all that complexity is ultimately, in principle, tractable and determined? Well, it could be that the more intentional and creative features of somatic learning have not yet been properly modeled by even connectionist systems—and that is likely. A machine that mimics a human ability, such as a connectionist machine, is like a working model. The working model is analogous but not identical to the target. For example, one can model the movement of the earth around the sun by swinging a rock around with a string. Some information will be made explicit, but all the relevant features will not be modeled. Collins claims that AI machines cannot yet pass the Turing Test, but that might be a matter of time and technological innovation.[12] Turing Tests, etc., are ways to judge whether enough of the relevant features are modeled for an AI to be considered to have our sort of thinking abilities—but even here we still go only by the clues that we can make explicit.

The ability to run an apparently irreducible connectionist program on a digital computer could also be another example of a dual control system, and so it would still harbor an irreducibility.[13] A connectionist system (upper level control) is constraining features of a digital computers (lower level control) in order to display connectionist functions; one "machine" is acting upon the boundary conditions left open by another "machine"—and this seems plausible. But even if a connectionist network running on a digital computer is just one *one level* machine—a string—that, for Polanyi, could still indicate that

[12] It is still an open question whether or not the creativity and telic intention we have can be adequately modeled by a machine—one may have to have the affordances of a living system in an environment in order to mimic that sort of ability. It does seem that computers can get somewhat better at "guessing right," which is a tacit ability for Polanyi that a scientist, fully apprenticed and now expert in her field, could better display than a novice. Creative innovations are different than totally random changes. So a computer programmed to mimic creativity through randomness is not quite catching on, though both are generally stochastic processes.

[13] For more see my "Rethinking the Machine Metaphor since Descartes: The Irreducibility of Bodies, Minds and Meanings" (2011).

we are dealing with something emergent and not fully determined by its enabling conditions, because it is still a machine.

3. The Irreducibility of Tacit Knowing and Dual Control Systems

We don't need to answer the more difficult question of whether or not a machine can fully model STK in order to answer the question of whether STK has irreducibly tacit features. Even assuming that a machine can successfully model riding a bicycle, hammering a nail, or competing in a mixed martial arts match, we can question whether this would reduce the tacit knowledge involved to an explicitly interpretable "string" that fully explicates the tacit clues and their connections/integrations into a new meaning or entity.

Elaborating on, transforming, or untangling the material string to make it explicitly interpretable is not the main issue for Polanyi, and this can be seen in how even simple machines are irreducible for him. Even if we can eventually untangle the complexity of a neural network by transforming it into a more manageable string, it would not make the tacit integration—that movement from parts to whole—explicit for Polanyi. The clues/conditions that we describe would not be functioning in the same way they did to bring us the emergent quality and, even if successful, a working model would mirror the same irreducible process of tacit integration/emergence: There is still a semantic dimensional shift here that is indicative of an ontic emergence.

In an emergent system (rather than a mere aggregate), there can be no fully determinate formal explication in terms of the lower level principles or meanings. Just as there is an integration of clues to focal meaning, there is an integration of parts to whole. Here we see how the structure of tacit knowing can come together in tandem with the structure of emergent being for Polanyi, and how both blend into one at a rudimentary level of description when we are talking about living systems and machines.[14] This comes across vividly when we see computers—tangible entities—mimicking tacit knowing processes. The change of meaning at the higher level of integration is displayed in the need to use higher-order principles to account for the system's operation. Hence for Polanyi, a machine represents a different ontological entity than its parts and has a different meaning than they do.

To identify a machine requires a higher-level description in terms of its function and its engineering principles; it's very existence cannot be recognized by the lower-level principles delineated by physics and chemistry. A complete physical description of a steam engine, for instance, could not tell

---

[14] See my (2011-2012) paper, page 30, and my "From Epistemology to Ontology to *Epistemontology*" (2013-2014). The ontological and epistemological diverge, of course, when we develop concepts and attempt to use them to capture ontological processes.

you whether it was working or broken (KB 176). Lower level descriptions in terms of physical or chemical strings are thus insufficient for a full explication. Any machine is thus an example of a dual control system. No physical laws will be violated, but the possible configuration of the physical elements will be constrained by a higher-order emergent whole and its principles. The emergent entity will be able to control, to some extent, the boundary conditions left open by the lower level. So while the lower level has and maintains its controls, the higher level exercises its constraints as well.

I maintain, with Polanyi, that the body itself is a system with many dual level controls, and its *knowing how* has an irreducibility even before the emergence of higher order entities, such as conscious minds or social systems that it both makes possible and subserves. If this is so, then there is an irreducibility of STK similar in structure to that of CTK operative before the higher-order social constraints we see displayed in CTK's influence upon the martial artist or our driving of a car. The body's operations display the structure of tacit knowing, and its irreducibility.

The irreducibility that comes from the tacit knowing process, linked together with emergent being, displays the problem of trying to understand any lower-level, subsidiary process in terms of higher-level concepts, and any higher-level process in terms of the lower. For instance, for Polanyi we have a hierarchy from lower to higher: physical, chemical, biological, psychological, social and then personal. The hierarchization is a convenient way of simplifying more complex dependency relationships in a manageable way. In the emergentist schema, the emergent level is ordinate and in between the subordinate (parts) and the superordinate (environment it is a part of). When we look at activities at our own level, we have subordinate and superordinate systems acting as tacit clues. For example, in a portrait of Queen Elizabeth, both the colored pixels (subordinate) and our knowledge of human faces and the queen (superordinate) are in play. We understand and make particulars explicit in terms of their higher context (it is a picture of Queen Elizabeth) or in terms of lower level contexts (colored pixels make lines and shapes), which won't be enough on their own to tell us who or what we are talking about.

We have the irreducibility that Collins sees only in STK at many levels, it is just harder for us to see that irreducibility. CTK irreducibility is easier to see because we cannot grasp the uppermost social superordinate level explicitly, and that's what we'd need to fit lower-level pieces together in a coherent explicit picture. STK irreducibility is harder to see because it is easy to think we have made something fully explicit when it is several layers below us on the hierarchy of being. The lower level on its own, however, is insufficient for the identification of real emergent entities or new joint comprehensions of meaning.

Looking at the underdetermined nature of lower level laws we can now see why Collins fourth sort of explication, *scientific explanation* (81), does not make somatic tacit knowledge fully explicit. Polanyi, as a scientist, is fully in favor of finding the best explanations we can. We can make somatic processes and even the behavior of dynamical systems explicit in terms in a scientific explanation. Polanyi would not be averse to this, but just as there are hierarchies in being, there are different explanatory levels, and each has its own limitations. It is the difference in emergent levels detectible through a semantic "dimension shift" that primarily shows machines, bodies and other emergent entities irreducible.

A scientific explanation can take several forms. An explanation can take the form of showing how the particulars fit into a higher-order theory, e.g., using engineering principles to describe a machine, or it can show how the material parts function together according to lower level physical laws. But just as there is an inevitable mismatch in making tacit knowledge explicit from the top-down, there is also a mismatch from the bottom up: Looking down from above, i.e., with emergent higher-level conceptual descriptions, we only approximate the tacit causal process and indirectly circumscribe them. And when we attempt to use lower level descriptions, e.g., describing the parts and the causal process with physical and chemical laws, we miss the higher-level of meaning and inadequately describe the emergent properties of the real activity or entity, i.e., the proper way to understand a higher-order context. For reasons both of the approximation, unspecifiability and inexhaustibility of a top-down approach, and of the underdetermination of a bottom-up approach, one might say that *knowing how* cannot in principle be completely captured by *knowing that*.

For Polanyi, ontic systems, such as bodies and machines—though they can die or be broken— are both irreducible, and the epistemic processes of perception and skillful activities follow that same structure. So the possibility that a machine—some sort of AI computer—could mimic human activities and human cognitive capacities would mean different things for Polanyi than for Collins. For Collins, it would mean that we have found transformations of those strings that would make what was always explicable explicit. For Polanyi, that a machine could perform the action could still be a mark of the irreducibly tacit.

III. Concluding Remarks

1. Proto-Structures

For Collins, if we can make a machine to perform a task that —for humans—usually requires apprenticeship, experience or tacit skill— then that performance does not really display tacit knowing in its strictest sense: it is still "mimeomorphic" and mechanical rather than "polimorphic" (TEK 55) and it does not truly display what is irreducibly tacit, i.e., that which *cannot* (as yet) be explicated (by known techniques). There is still plenty of collective tacit knowledge shaping such activities and what they mean for us, but the performance itself is not strictly an irreducibly tacit achievement.

For Polanyi, the tacit structure pre-exists linguistic and social knowledge. It is an intentional structure that forms the basis of simple organic activities such as those involved in perception. So one can say of even rudimentary bodily processes that tacit knowing is taking place and there are hierarchies of these structures that build on one another in the course of evolution. The skills animals use in hunting are the basis of the skills that we use as we develop crafts and arts, whereas, for example, animal hunting is not an example of tacit knowing for Collins (TEK 78). Bodily tacit processes have proto-meanings (one might say) that linguistic and social meanings are built on. As Charles Taylor and Hubert Dreyfus claim, following Merleau-Ponty, motor intentionality is the structural basis for representational intentionality (2015, 50).

An advantage of Polanyi's approach is that it shows a knowing structure that manifests differently at different levels. There is a similarity in structure, though differences in types and orders of meaning. So there is indeed the disjunction Collins sees between the social and the somatic, and between humans with language and animals that do not have our language, but these operate like layered subsidiary-focal achievements and layered sets of dual control system that support them. There is not just one big dividing line between the explicit (and the explicable) and the irreducibly tacit located at the CTK, we instead see the emergence of tacit structures in the course of evolution, reformed, retooled or repurposed to meet new challenges. As William Bechtel notes, "evolution often works most effectively by taking components that were previously employed for one purpose and using them for other purposes. This kind of evolution occurs at the expense of decomposability, since it depends on building up additional connections within the system to build a more integrated system" (146).

Polanyi sees emergent intentionality and degrees of freedom even in the most basic forms of life and the simplest of machines. The tacit process, in a rudimentary form is already at work. As evolution builds to more complex and developed forms, we see layers of dual control systems. Higher level meanings emerge, but we still have the building of a system upon a system. There are emergent orders, but just because we can experience most vividly our inability to express human-level actions in terms of subsidiary causes and conditions, we shouldn't be blind to the layering of similar systems in other animals and in our evolutionary history, and their own form of irreducibility.

Seeing the link between the tacit and the emergent may also have implications for understanding the evolutionary process. Both how evolution retools tacit intentional structures as species develop and how a dynamic attractor works like a telic field might be part of a properly understood evolutionary picture.

2. Telic Fields

The notion of an intentional component to tacit knowing goes together with Polanyi's notion that there is some sort of telic field drawing evolution forward towards more complex and free unities. The idea of tacit knowing and emergence may go down as far as the simplest forms of life, which show a unity and a meaningful response to its environment geared toward self-sustenance. Polanyi would side with Thomas Nagel in saying that, left only to chance variation and ability to survive, "the materialist Neo-Darwinian conception of nature is almost certainly false" (2012).[15] But if there is a telic field that unlocks or establishes various potentialities that an entity has in particular environments, it can itself be emergent. Considered this way the telic does not need to be a mysterious force, but can work in the way dynamic systems do. The material conditions or environment can coordinate in a way that sets up dynamic attractors that increase the probability of the expression of certain properties or traits. So we have something like Polanyi's "maturation" (PK 395), and the notion of a telic, rather than a pre-existent or final teleological, principle at work (see Gelwick, 2005). For Polanyi, life, and then consciousness, are emergent. But they are also drawn forward into existence, by telic forces or principles that were, according to Polanyi, there from the beginning. In particular and in general development, "this field of forces would also be the gradient of a potentiality: a gradient arising from the proximity of a possible achievement" (PK 398).

Walter Gulick, sides with the Neo-Darwinians here and says that Polanyi was wrong to talk about telic fields in the way that he did, i.e., as principles that were existent from the very beginning (See Gulick, 2010-2011 and Fennell, 2013-2014). But the idea of dynamic systems seems to bridge the difference here. One can say the Neo-Darwinians are not wrong. They just need to incorporate the notion of emergent dynamic systems that can act as telic fields. Or one can say that Polanyi was not wrong. He just did not provide a properly emergent mechanism, i.e. successive dynamic attractors, to account for the possibility of what he called a telic field.

3. Intrinsic Value of Animals

The existence of tacit intentionality in non-human animals may also affect our understanding of their intrinsic value. Here there might be an important distinction between living things and machines. One might say—although they are indeed emergent things—that we put the meaning into machines, i.e., they exist as real but we recognize them only in relation to our purposes and intentions. They thus model intentionality, but do not experience or indwell it. Living organisms in contrast do have a form of intentionality that they indwell. There are differences that make a difference to them. Without us, they have their own form of meaning and purpose. Showing the intentional and creative structure of somatic

[15] Compare also Gregory Bateson on the telic as a feature of self-corrective systems (117, 118) and larger systems with some degree of autonomy, such as ecological systems.

knowing thus reinforces Holmes Rolston III's conception of emergent value in natural entities and emergent structures; we can see levels of intentional meaning and value in at least sentient beings and higher zoology, and even in lower zoology and botany. So there are moral implications here when we tie together indwelt meaning with some sort of intrinsic value that would need to be spelled out in an emergentist framework.

Dreyfus' exposition of body-knowing and Collin's exposition of CTK both show important emergent jumps in knowing, being and value. We can agree that most animals do not share collective tacit knowledge (*TEK* 76), but we do not have to defend the line between animals and humans to preserve human irreducibility against material reductionism. Polanyi, like Dreyfus, can afford to be more generous to our cousins and ancestors in other species than Collins, but would Polanyi be as generous with AIs?

4. The Future of Consciousness

As we've seen, connectionist networks are not ordinary machines. They model an inextricability with "hidden layers" and exhibit a semantic "dimension shift." They can also function like complex dynamic systems, that develop "attractors" as they interact with their environment and reinforce their own intentional trends (See Juarrero, 1999). If connectionist networks can model STK, then they might still be intractable and even inexplicable. They might be modeling dynamic systems with an emergent degree of freedom from deterministic physical causality (though they would not violate the laws of that causality). Our sort of freedom and context sensitivity may only emerge with self-consciousness within society, but degrees of freedom or proto-freedom can be seen to various degrees in various living organisms/bodies.

Can computers do what people do when they are using collective tacit knowledge? Here I agree with Collins: they (as yet) cannot. In an emergentist picture—given the notion of *multiple realizability*—there is an acknowledgement that consciousness might form from different material subsidiaries. But experience matters, and so do —at least to some extent—the affordances of the medium (e.g., biological v. silicon), since some experiences are endemic to particular mediums in particular environments. To get a conscious AI, we'd have to provide them with the right sort of bodily functions and clues, and they'd have to develop the right sort of social interactions and become autonomous to some extent. They would also need to have some EI, emotional intelligence, if we are really looking for entities like humans. As it is, AIs only have representations, and manipulate those representations; the meaning in and meaning out is still overwhelmingly provided by human interpretation.[16] AIs mimic tacit knowing, but it doesn't seem like they are yet close to the sort of systematic structure that can *exercise* tacit knowing *by dwelling in* their subsidiaries. Animals can properly be said to have their own sense of intentions. Just as

[16] For Polanyian reasons for why (current) autonomous robots are not conscious for Polanyi, see Heder and Paski (2012).

machines/artifacts are often too superficial a model of tacit processes to catch the emergent nuances—just as the AI machine is an explicit formulation of a tacit process that itself is inexhaustible and indeterminate—we seem a long way off from creating a machine with the right affordances and experiences that could give it its own sense of intention. So, while there is likely something that it is like to be a bat—as Nagel would say—it is unlikely that there is something that it is like to be Siri, or even a much more complex and well-trained AI. At this point, advanced processing machines are not emergent comprehensive entities that dwell in their parts.

What about the future of AI? Will we be able to develop conscious, sentient, human-like machines? Maybe. But if we do, we don’t have to worry about similarities with animals, martial artists or computers reducing all that is meaningful and valuable about humans—or this new AI—to rocks and strings.

Works Cited

Bateson, Gregory, 1979. *Mind and Nature: A Necessary Unity*. New York: Bantam Books.

Bocharova, Jean. 2014-2015. "The Emergence of Mind: *Personal Knowledge* and Connectionism" Tradition and Discovery 41:3, 20-31.

Bechtel, William and Abrahamsen, Adele. 2002. *Connectionism and the Mind: Parallel Processing, Dynamics and Evolution in Networks*, Oxford: Blackwell Publishers.

Collins, Harry. 2010. (TEK) *Tacit and Explicit Knowledge*, Chicago: University of Chicago Press.

Collins, H. 2011-2012. "Analysing Tacit Knowledge: Response to Henry and Lowney" *Tradition and Discovery*, 38:1, 38-42.

Collins, H. 2018. *Artifictional Intelligence*. Cambridge: Polity Press.

Davies, Martin. 1989. "Connectionism, Modularity, and Tacit knowing" *British Journal of Philosophy of Science* 40:4, 541-555.

Elman, Jeffrey L. 1990. “Finding Structure in Time.” *Cognitive Science* 14:2, 179-211.

Elman, J. L. 2014. "Systematicity in the Lexicon: On Having Your Cake and Eating It Too" in Symons, J. and Calvo, P. (eds.) *The Architecture of Cognition*, MIT Press, 115-146.

Elman, J. L., Elizabeth A. Bates, Mark H. Johnson, Annette Karmiloff-Smith, Domenico Parisi, and Kim Plunkett. 1999. *Rethinking Innateness: A Connectionist Perspective on Development*. Cambridge, MA: The MIT Press.

Fennell, Jon. 2013-2014. “Plausibility and Common Sense: *Mind and Cosmos* by Thomas Nagel,” *Tradition & Discovery*, 40:1, 45-52.

Fennell, J. 2017. “Is Polanyi’s Emergence Reductive?,” *Appraisal*, 11:2, 22-35.

Gelwick, Richard. 2005. "Michael Polanyi's Daring Epistemology and the Hunger for Teleology" *Zygon* 40:1, 63–76.

Gulick, Walter. 2010-2011. "On the Adequacy of Neo-Darwinism: A Reply to Daniel Paksi, *TAD* 38:2, 56-60.

Heder, Mihaly and Paski, Daniel, Oct 2012. "Autonomous Robots and Tacit Knowledge" *Appraisal* 9:2, 8-14.

Juarrero, Alicia. 1999. *Dynamics in Action: Intentional behavior as a Complex System*. Cambridge, MA: MIT Press.

Lowney, Charles. Fall 2009. "From Science to Morality: A Polanyian Perspective on the Letter and the Spirit of the Law" *Tradition and Discovery,* 36:1, 42-54.

Lowney, C. 2011. "Rethinking the Machine Metaphor since Descartes: The Irreducibility of Bodies, Minds and Meanings" *Bulletin of Science, Technology and Society*, 31:3, 179-192.

Lowney, C. 2011-2012. "Ineffable, Tacit, Explicable, Explicit: Qualifying Tacit Knowledge in the Age of 'Intelligent' Machines," *Tradition and Discovery*, 38:1, 18-37.

Lowney, C. 2013-2014. "From Epistemology to Ontology to *Epistemontology*" *Tradition and Discovery*, 40:1, 16-29.

Lowney, C., Levy, S., Meroney, W. and Gayler, R. 2020. "Connecting Twenty-First Century Connectionism and Wittgenstein" *Philosophia*, 48:2, 643-671.

Musser, George. May 2019. "Artificial Imagination: How machines learn creativity and common sense, among other human qualities" *Scientific American*, 59-63.

Nagel, Thomas. 1974. "What is it Like to be a Bat?" Philosophical Review, 83:4, 435–450.

Nagel, T. 2012. *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature is Almost Certainly False*. Oxford: Oxford University Press.

Polanyi, Michael. Feb. 1952. "The Hypothesis of Cybernetics" *British Journal for the Philosophy of Science*, 2:8, 312-315.

Polanyi, M. 1962. (PK) *Personal Knowledge: Towards a Post-Critical Philosophy*. Chicago: The University of Chicago Press.

Polanyi, M. Jan. 1966. (LTI) "The Logic of Tacit Inference" Philosophy. 41:155, 1-18.

Polanyi, M. 1969. (KB) *Knowing and Being Essays by Michael Polanyi*, edited by Marjorie Grene. Chicago: The University of Chicago Press.

Polanyi, M. and Prosch, Harry. 1975. *Meaning*. Chicago: The University of Chicago Press.

Rolston III, Holmes. 2017. "Naturalizing Values: Organisms and Species" in Louis Pojman, et al., *Environmental Ethics: Readings in Theory and Application*, 7th ed., Boston: Wadsworth, 107-119.

Rumelhart, David E. and McClelland, Jay. 1986. "On learning the past tense of English verbs" in Rumelhart, D. and McClelland, J., eds. 1986. *Parallel Distributed Processing: Explorations in the microstructure of cognition, vol.1 Foundations*. Cambridge, Ma: MIT Press.

Smolensky, Paul. 1988. "On the Proper Treatment of Connectionism" *Behavioural and Brain Sciences*. 11: 1-74.

Smolensky, P. 1991. "Connectionism, Constituency and the Language of Thought" in MacDonald, C. & MacDonald, G. 1995. *Connectionism: Debates on Psychological Explanation, Vol 2*. Oxford: Blackwell, 164-198.

Taylor, Charles and Dreyfus, Hubert. 2015. *Retrieving Realism*. Cambridge: Harvard University Press.