The Music of Consciousness: Can Musical Form Harmonize Phenomenology and the Brain?

Dan Lloyd • Trinity College, USA • dan.lloyd/at/trincoll.edu

> Context • Neurophenomenology lies at a rich intersection of neuroscience and lived human experience, as described by phenomenology. As a new discipline, it is open to many new questions, methods, and proposals. **> Problem** • The best available scientific ontology for neurophenomenology is based in dynamical systems. However, dynamical systems afford myriad strategies for organizing and representing neurodynamics, just as phenomenology presents an array of aspects of experience to be captured. Here, the focus is on the pervasive experience of subjective time. There is a need for concepts that describe synchronic (parallel) features of experience as well as diachronic (dynamic) structures of temporal objects. **> Method** • The paper includes an illustrative discussion of the role of temporality in the construction of the awareness of objects, in the tradition of Husserl, James, and most of 20th century phenomenology. Temporality illuminates desiderata for the dynamical concepts needed for experiment and explanation in neurophenomenology. **> Results** • The structure of music – rather than language – is proposed as a source for descriptive and explanatory concepts in a neurophenomenology that encompasses the pervasive experience of duration, stability, passing time, and change. **> Implications** • The toolbox of cognitive musicology suddenly becomes available for dynamical systems approaches to the neurophenomenology of subjective time. The paper includes an illustrative empirical study of consonance and dissonance in application to an fMRI study of schizophrenia. Dissonance, in a sense strongly analogous to its acoustic musical meaning, characterizes schizophrenia at all times, while emerging in healthy brains only during distracting and demanding tasks. **> Constructivist content** • Our experience of the present is a continuous and elaborate construction of the retention of the immediate past and anticipation of the immediate future. Musical concepts are almost entirely temporal and constructivist in this temporal sense – almost every element of music is constructed from relations to non-present musical/temporal contexts. Musicology may offer many new constructivist concepts and a way of thinking about the dynamical system that is the human brain.

> Key words • **Neurophenomenology, music, ontology, temporality, fMRI, schizophrenia.**

Introduction: Time and consciousness

It is a truism that time is important to human consciousness. But as with many truisms, obviousness leads to a kind of neglect. In this case, one acknowledges time as a dimension of experience measured by the clock, while failing to recognize fully how pervasive the experience of time actually is. For example, consider an ordinary object, a teacup. The teacup has a history and a future, of course, but at first glance we might suppose that we see it just as a momentary presence, a three-dimensional form without noticed temporal properties. Our first glance presupposition is tacitly abetted by the schoolbook picture of visual perception, in which the eye works like a camera and its images are translated and broadcast among various brain areas. The job of perception in the schoolbook story is to build a stable three dimensional model of the scene here and now, a model in which various curves and shadows are parsed as a teacup, itself a stable concept without important temporal properties. Conscious life, according to this view, is a series of such snapshots. Each is conceptually rich in itself, but in the normal case, where we are simply seeing the scene and recognizing the objects arrayed there, time is not particularly invoked.

When we step outside of the schoolbook scheme, however, we quickly observe that real encounters with teacups are not nearly so simple. To begin with the most obvious, our perceptual system cannot take in (and understand) a complex scene in a single glance. Observing a teacup, we fixate on the details of the handle, then the rim, then the saucer, then something the background, then back to the handle again. From this flowing filmstrip a percept emerges, in which we see the glimpses as a cup on a table. Already at this point, time has made its stealthy entrance. The multiple glances are not simultaneous. To construct the teacup, the succession of visual fragments must be retained long enough to synthesize an interpreted object in a coherent setting. This achievement is made more complicated by the elementary consideration that we observers are in continuous motion. As our eyes make their saccadic leaps, we turn our head, and possibly move in other ways. Assembling a visual scene requires coordinating the jumpy images with the equally jumpy traces of efferent commands to muscles of the eye and the rest of the body. We must somehow keep each image paired with concurrent position information and keep the parade in order, and from this build a teacup. To the extent that I perceive a teacup here and now, and not a teacup fragment, I need to hold in mind events in succession that have occurred in the immediate past.

In short, every teacup has a history that is folded into its presence here and now, if it is to be recognized as a cup at all. It has a future too, as is apparent if we should desire a sip of tea. Now we must anticipate the responses of the cup to our grasp and lift, a set of expectations entailing a fluid grasp of the physics of our bodies and environment. How hard must the handle be squeezed? How much upward force will be required to lift the cup? How must these forces be modulated to keep the tea inside the cup? What path will bring cup to lip? These, among other expectations, shape how the cup is seen. Although these pasts and futures are not part of the sensory field, they are among the contents of ordinary consciousness (Husserl 1974). They can be foregrounded through scenarios where the sensory presentation is fixed but the temporal properties altered. For example, suppose we discover that the teacup in question is a very detailed hologram? The difference between a real teacup and a hologram is not the look of it but the clutch. As one's hand passes unimpeded through the ghostly handle, the contents of consciousness shift dramatically. What shifts, however, is not the sensory presentation but the expectations one holds.

These observations of the ubiquitous presence of subjective time may perhaps begin to hint at the phenomenological complexity of ordinary states of awareness. Even the simple cases in their most simplistic descriptions rapidly spiral towards Proustian elaboration. (Note how different the teacup appears when you realize that it and its contents have been sitting on the desk overnight. Cheers!) In 1907, Edmund Husserl devoted an entire book to the simple encounter with objects (*Thing and Space*, Husserl 1974), in which pure phenomenology led to ideas that would be repeated by phenomenologists (especially Merleau-Ponty 1962), and rediscovered in James J. Gibson's ecological psychology and in the recent philosophical fascination with embodiment (Gibson 1979; Clark 2003; Noë 2004). Husserl also opened the enduring theme of temporality in phenomenology (Husserl 1966); in this, he was influenced by William James (1890). The ubiquitous immediate comet tail of recollection he called "retention." The equally ubiquitous branches of immediate possibility he called Protention. These flank the immediately given, which Husserl called the "primal impression." The Husserlian sandwich of retention, primal impression, and protention characterizes every state of consciousness, regardless of intentional object. Importantly, the three facets of temporality are co-present in the subjective now. In perceiving the teacup at 12:00:00, I am also aware of its presentation at 11:59:59, and its likely appearance at 12:00:01 – all of this packed into the experience at 12:00:00.

These phenomenological observations, however elementary, set a high bar for scientific theories of consciousness. Consider, for example, the binding problem, as presented in a standard reference work:

> "Information processing in the human brain is highly parallel. This means that different features of an object are processed in different parts of the brain. For example, the color and the shape of a red square are coded by different neurons in the visual system (visual field). However, we do not perceive 'red' and 'square shaped' separately but a 'red square.' The binding problem deals with the question of how features that are processed in parallel are bound to the one unique percept."
> (Herzog 2009: 388)

Temporality multiplies one red square into many, as even a static percept continuously updates a lengthening history. Its shifting retentions and protentions must also be bound together, but in a binding that preserves the structural subjective distinctions between past, present, and future. For every object of consciousness, in short, the brain must build a timeline, store it, and (crucially) keep it continuously present to mind as a non-sensory dimension of every act of perception. Consciousness is a synchronic structure in which a formidable sea of diachronic information is represented and continuously updated.

Temporality is a feature of not just teacups, of course. While we have not argued here that explaining temporality would be sufficient for a scientific neurophenomenology, temporality is certainly a prominent and ubiquitous dimension of conscious life. Accordingly, the overall project of neurophenomenology is at least in large part the project of neurotemporality. Time, as experienced, will be a necessary component.

The gears

Cognitive science circa 2013 draws on two broad paradigms for its explanations, which comprise the tools for a science of consciousness: dynamical systems theory and computational modeling. "Dynamical system" and "computer" are each labels broad enough to cover anything and everything, but as it happens the discourse of cognitive science has tended toward somewhat more specific and concrete subtypes. Neural networks have been the reigning dynamical system, a natural choice for bridging brain and cognition. Computational theories have tended to stress language-like computational symbol systems, most notably championed as the "language of thought" proposals of Fodor and others (Fodor 1975; Field 1978; Fodor 2008). Neither of these is intended primarily to apply to consciousness. Nonetheless, their different explanatory powers immediately invite application in neurophenomenology. How do these two approaches fare in application to the problem of consciousness? In particular, how are they equipped to address the complexities of subjective time?

To begin with language, decades of work in linguistics and computation theory have illuminated the possible gears of the computational mind. We have the representational theory of mind and the language of thought in which the mind is a kind of software implemented in the neural networks of the brain. Can it stretch usefully toward neurophenomenology?

The powers of language are obvious and well-known. After all, the moderately rich internal content of the teacup example above is conveyed entirely in words, and I can reference it with just the phrase "the teacup example." That is efficient, as well as vastly flexible (Did I mention that the teacup is blue?). Syntax and pragmatics allow endless molecular combinations of atomic symbols, and of course time is included, through tense and many other explicit and implicit time markers. Syntax affords language much of its power, but semantics lends a hand. For example, the atomic symbols (and hence their molecular expansions) are arbitrary. Nothing about the symbols, each taken in isolation, constrains the reference or meaning that the symbol might convey, once it is deployed as part of a system of symbols.

While language can describe just about anything (consciousness included), it is not so clear that conscious experience is itself a stream of symbols. (Extensive critiques include Dreyfus 1972; Varela, Thompson & Rosch 1991; Wheeler 2005; Thompson 2007.) Here, we can mention just the mismatch of serial symbols and synchronous temporal experience. The descriptive phrase, "the teacup example," did not get filled with content in a stroke. Language is serial, and pearls of rich content are strung along with many less luminous beads. If we sample the stream of description in the slim narrative of teatime, most of the symbols we find are denoting bare and local facets of the scene, if not still less. Consider any sentence word-by-word with special attention to the content at each time point – what is usually denoted is very limited. Conscious experience, in contrast, seems to exceed this trickle. Perhaps the language of thought admits many parallel streams, but then the binding problem returns. There is no reason that the streams will bear any special relationship to one another. We could also broaden the window of consideration, considering the content of the text sentence-by-sentence rather than word-by-word. This still seems to fall short of the phenomenology, but also faces the very problem of temporality itself, as a serial stream of tokens must be framed in a single synchronous representation.

Synchronicity comes naturally in parallel distributed processing, the most popular exemplar of dynamical system. Even at the dawn of the neural network era, the capacity for accommodating simultaneous constraints was celebrated (Rumelhart, McClelland et al. 1986). Early connectionist models demonstrated that networks of neuron-like nodes could be adjusted (through neural network learning) to solve reaching and grasping problems (Hinton 1984). In this case, a network maintained and updated vectors whose elements represented some of the interplay of sensory information evident in phenomenology. Neural networks can be exquisitely sensitive to their immediate antecedent conditions, inviting theoretical discussions of dynamical systems and time (van Gelder 1999; Varela 1999a). Concretely, models that include recurrent feedback loops in networks can capture explicit temporal information as well (Elman 1990).

These are very promising starting points. However, between the simple models and the theoretical "in principle" discussions lies a gulf. Simple models can be understood as functioning either via the interaction of neuron-like computational units or as rudimentary functional divisions (layers or other discrete groupings of basic units). But tractable explanations of real brains (and real consciousness) are assumed to comprise functional architectures and their dynamical interactions that leave units and layers behind. The assumed complexity is abbreviated in the idea of high-dimensional state spaces, within which complex network states appear as points. The concept of a high-dimensional state space partitioned into regions that can be assigned to various cognitive states is an appealing simplification of otherwise unintelligible processes, but it remains vague in application. By what principles should the state space regions be demarcated? Are there general rules that apply across tasks and modalities? I suggest that there is a missing level in dynamical systems theorizing, something usefully situated between units and layers on the one hand, and regions of state space on the other. This missing element is not a distinct process, but rather distinct entities – we need an enriched scientific ontology to frame hypotheses about the gears of conscious experience. Phenomenology illuminates the features that this intermediate layer should capture; here, we have focused on temporality, for which synchronous and sequential relations are essential. This desideratum distinguishes the missing entities from language-like serial symbol systems. However, the middle entities should nonetheless capture the some or all of the generative (syntactic) properties of language since the distinctions language draws are generally distinctions available to awareness.

In short, both dynamical systems (such as neural networks, typically) and linguistic systems (typically understood as computations over serial symbol strings) face similar lacunae: both are illustrated through numerous simple examples and analogies that establish that each approach is plausible and sufficient in principle to account for a broad swath of cognition (and by extension, consciousness); moreover, both are buttressed by sophisticated theoretical discussions of their capacities (with occasional nods to phenomenology, at least on the dynamical systems side). Between these low-level demonstrations and the high-level theoretical proposals is a missing middle ground, where our teacup is found. What is missing are accounts of conscious capacities that approach the richness of human consciousness without leaving behind its fundamental phenomenology

The music of thought

To supplement dynamical systems and (partly) displace the language of thought, it is proposed that neurophenomenology appropriate the practical ontology of music. Specifically, we now consider the possibility that the entities created and deployed by the brain in the processes of consciousness itself are analogous to musical entities. "Music" here is meant broadly as a set of systematic arrangements and constraints within cultures for creating extended structures of sound. Music is thus the analogue of natural language, and will inform theorizing much as language does. Accordingly, we might hypothesize a "music of thought." In some ways, the music of thought hypothesis is a small revision in the language of thought idea, but in others it makes a large and useful difference.

As a prelude, however, note that music is not a frivolous sidebar to human evolution. For at least 42,000 years humans have been making musical instruments (Higham et al. 2012). Speech and song use the same anatomy, so it is reasonable to assume that singing and talking mixed in the most ancient social formations. Indeed, several researchers have proposed that music precedes language in human prehistory (Froese, Ikegami & Beaton 2012). All known cultures have music, and most create and share musical memes with enthusiasm. If one looks for a uniquely human production that could be the model of cognition, music shares the initial appeal of language. Indeed, it may seem that music is just a subtype of language (Patel 2007). Most conspicuously, it shares a generative syntax (Lerdahl & Jackendoff 1983). Like language syntax, musical syntax informs ordinary listeners (and informal musicians) with strong anticipations of the next sounds in a piece and sensitizes listeners to violations. As with language, there are experts for whom the syntax is explicit and invites theoretical elaborations. But even non-experts are adept: just one French horn player who replaces a B with a B-flat will stand out in the middle of a symphony, even to non-expert listeners.

Music has much in common with language, but two conspicuous differences make music an appealing metaphor for consciousness. The first is that musical symbol systems define most, if not all, of their symbols using time. The elements of a musical system generally include pitch, timbre, harmony, rhythm, and melody. Pitch is arguably the atom of music from which the other structures are built. Pitch itself is dependent on frequency, a time-dependent concept; but more importantly, the syntax governing well-formed symbol complexes in music is entirely temporal. Musical systems include elaborate specifications of simultaneous interactions among pitches – these are the rules of harmony. Rhythm and melody are governed by diachronic laws. Duration defines the elemental beats of music, and sequences of durations specify rhythm. Sequences of pitches together with their durations define melodies. The rules of harmony, then, define and interpret simultaneous interactions of rhythms and melodies. Finally, timbre, like harmony, is specified by the variety of overtones of an instrument and their temporal envelope (Patel 2007).

Music thus falls intriguingly in the space between computation and dynamical systems. Like a language system, a musical system comprises discrete atomic symbols governed by a rich combinatorial syntax, capturing the main appeal of the language of thought hypothesis. Like a dynamical system, musical entities emerge from simultaneous parallel interactions. Musical syntax constrains these interactions, allowing expressions that are both distributed and syntactically correct. In this respect, a musical structure and its dynamical possibilities seems better suited to phenomenology, which demands an infrastructure of temporality and simultaneity.

Language affords a sharp distinction between syntax and semantics, a distinction absent in musical systems. Music does not have language-like semantics. Occasionally, music theorists have attempted to specify how music can represent scenes and actions, but these ideas are striking in their failure. No orchestra can ever communicate "The cat is on the mat." Musical representation thus requires blurry denotations of indefinite conditions such as "feeling" or "movement." Even with this hedge, music seems to be unable to represent even the simplest subject-predicate propositions. Even if a piece somehow depicts sadness, it is unable to say who or what is sad (Kivy 2002; Scruton 2009).

This sketchy semantics is wildly inconsistent with the highly specific syntax of all musical systems, which are exacting in their constraints on composition and performance. To some music theorists and philosophers, this implies that music is nothing but syntax, pure form without content (for example, Hanslick 1854, Scruton 1976). For others, music carries rich content, but is entirely self-referential (e.g., Kivy 1990). Every note in a piece represents other notes and ultimately the whole piece and even the musical tradition in which it falls. According to this view, while musical symbols have an absolute constitution of pitch and duration, their interpreted meaning is always in reference to their musical context. Every tone also specifies interval, harmony, tempo, meter, and melody. Indeed, a core "utterance" in musical discourse is repetition, and its primal modification is variation from repetition. Unlike language, musical repetition is never redundant. At every scale, every element of a piece denotes by resemblance (or variation) the elements around it. The famous four note motif that opens Beethoven's Fifth is already a repetition-and-variation structure. Its immediate repetition, one tone lower, is of course a reference to its first occurrence. And so it continues, ultimately reappearing in the third movement as a rhythmic motif on a single chord. A large part (if not all) of the artistry and pleasure of the symphony rides on these symmetrical formal interconnections. For present purposes, it is not necessary to take a position with respect to formalism and self-reference. We can refer to their common ground as "internalism," the view that musical syntax does not travel with a semantics of external, extra-musical reference.

It was noted above that musical entities are largely constituted through temporal relations, analogous to the constitution of states of awareness. Internalism, as defined here, also has an intriguing analogy with phenomenology. Internalism as a phenomenological position would imply that a state of consciousness is constituted through relations among other states of consciousness, and that the objects of awareness are manifolds of states of consciousness, unified through temporal relations. Phenomenology reveals the importance of temporality in all awareness – at every moment, the teacup of the example above is coordinated with its past and potential future appearances. Although there is a sensory presence (Husserl's "primal impression"), this is like the absolute identity of a chord. The edifice of meaning attendant to the percept is entirely contextual. With this discussion in mind, Husserl's metaphor for consciousness takes on a new resonance.

> "When, for example, a melody sounds, the individual notes do not completely disappear when the stimulus or the action of the nerve excited by them comes to an end. When the new note sounds, the one just preceding it does not disappear without a trace; otherwise, we should be incapable of observing the relations between the notes which follow one another.… On the other hand, it is not merely a matter of presentations of the tones simply persisting in consciousness. Were they to remain unmodified, then instead of a melody we should have a chord of simultaneous notes or rather a disharmonious jumble of sounds…" (Husserl 1991: 30)

The musical analogy may be especially apt.

None of the above contradicts the importance of the reality of biological brains in their ecological context. The music of thought does not sound in a vat, but rather emerges from myriad internal and external causes. I suggest, however, that those causes not be confused with content or the subjective "feel" of experience ("what it is like," Nagel 1974).

From metaphor to model

Music offers the prospect of a temporal grammar for dynamical states. The discussion above highlights the fit of musical structure with phenomenology. For the moment, however, this is merely a metaphor. Metaphors can usefully spotlight aspects of their targets, but for a science of consciousness we would like to reconstrue metaphorical features as measureable observations. How will this proceed? With a hypothesis as broad as the music of thought, many avenues are open. For example, in another study I demonstrate that a statistical measure of sparseness within symbol systems distinguishes natural languages from music (using several hundred examples from several languages and several world music traditions). In effect, sparseness measures a feature of syntax of these systems, which this study examined as single symbol tokens, dyads, and triplets. The sparseness measures used sharply distinguished language from music. Then, the same analytic technique was applied to nearly one hundred fMRI image series from three experimental groups to show that these sequences of brain states were indistinguishable from musical sequences (Lloyd 2011a).

In this paper, I will examine another empirical application, taking up the musical property of consonance. Consonance and dissonance are psychological percepts arising from the degree of interference between tones (Plomp & Levelt 1965; Sethares 2005). When frequencies are similar but not identical, their oscillations rapidly alternate between constructive and destructive interference. Over a certain critical interval, this creates an audible beating that listeners experience as roughness. When a tone includes overtones (as do most musical tones), then dissonance can arise between neighboring overtones. Although this is a perceptual feature of heard tones, like many musical features it correlates strongly with physical properties of the sound. Thus there can be a straightforward physical calculation of consonance for any pair of complex signals. Physical consonance is not psychological (experiential) consonance, but their close psychophysical correlation warrants the application of the term as an objective descriptor of frequency relationships.

But consonance is a feature of audible sounds, and whatever else we may claim about the brain, we know its myriad computations unfold in silence. Nonetheless, the analysis generalizes to any phenomena that can be analyzed as oscillations. Brains are included, particularly considering the alpha, theta, and gamma oscillations detected by EEG. Functional MRI yields time series of images, and each pixel ("voxel" in fMRI parlance) is an oscillator amenable to analysis as a fundamental frequency plus many overtones – although the sampling rate of fMRI affords oscillations at much lower frequencies than those observed with EEG. These voxel time series may or may not be consonant with each other, using the same sense of consonance that applies to complex tones. This will be a useful measurement if we can use it to draw distinctions between brains in different conditions.

Here, I will examine consonance as a global dimension of the brains of healthy subjects and schizophrenia patients. Data originally collected by Abigail Garrity et al. comprised 15 schizophrenia patients and 18 healthy controls performing the same auditory oddball task: each subject listened to a sequence of tones, and pressed a button for a different "target" tone appearing at random; the sequence was interrupted by odd distracting noises ("surprises"), also occurring at random. Each run was 372 seconds in length, with images acquired every 1.5 seconds (for details, see Garrity et al. 2007). Images were preprocessed via standard methods (with motion-correction, normalization to a standard anatomical template, and spatial "smoothing" to reduce isolated spikes of noise). Then, the image series were further processed using independent component analysis (ICA). ICA identifies temporally coherent networks, that is, sets of voxels that vary together over time, activating and deactivating in unison (Calhoun et al. 2002). In effect, ICA discovers networks of correlated activity that can then be analyzed as "super-voxels." A rather small set of independent components (~20) generally captures 90% or more of the variance in the data. Here, ICA identified twenty temporally coherent networks (brain regions activating in unison) for each subject, and subsequent analysis was based on these components.

At this point, each component time series was analyzed with the fast Fourier transform to yield a spectrum of frequencies for each, as if each oscillating component were a tone with a complex timbre. Each of these signals can then be compared pair-wise to all of the others, with a consonance (dissonance) value calculated for each pairing (Sethares 2005). That is, each component tone (with its overtones) can be compared to all the others, one by one, and the consonance measured for each two-tone chord. Thus, among twenty independent components there will be 190 unique pairs with an equal number of consonance measures. We can calculate mean global dissonance for each subject, in effect taking the brain as a keyboard of "tones" and assessing its overall harmoniousness. This measure separates the two groups ($p < .01$, two-tailed t-test). Mean dissonance is greater in the patient group; brains affected by schizophrenia have more tones and overtones whose frequencies interact roughly. The sharpness of the divide is more apparent when the dissonance values are not averaged. That is, we compare all of the pair-wise dissonance measurements for all subjects. In this analysis, components from the brains of schizophrenia patients differed from healthy controls with a significance of $p < 0.000001$. Overall, by both the aggregate and mean measures, healthy subject brains are more internally consonant than the brains of schizophrenia patients.

Do these differences reflect a steady state in the subject brains, or do the specific components of the experimental task provoke different transient responses? We can assess this by calculating the instantaneous consonance of each subject through the experiment, as follows. In the two calculations above, dissonance was calculated over the entire series of 248 images. But at each time point, component intensity varies. Thus, if two mutually dissonant components are both very active at a particular time, we can consider the global dissonance of the system to be higher at that time. In this way, the global magnitude of dissonance at any time is modulated by the instantaneous intensity of the more or less dissonant components. Dissonance thus becomes a varying global parameter for each subject.

With this continuous measurement, the circumstances of sudden dissonance can be discerned. In healthy subjects, a transient dissonance spike, compared to the baseline levels of dissonance in each subject, occurs during two experimental events: at the sounding of a target tone, requiring a button press; and at the random interruptions with a varying distracting noise. (The target dissonance increase is significant ($p < .001$) as is the increase in the presence of a distractor ($p < .02$).) The two events share several properties. At the most general level, they are arousing and capture attention, and both require a decision to act (or withhold action). In contrast, schizophrenia patients show no significant difference in momentary dissonance in response to any event in the experiment.

If we consider the mean dissonance levels in the two groups, a possible interpretation emerges. Patient brain components are consistently more dissonant than those of healthy controls, and their level of dissonance is statistically uniform for the entire experiment. In contrast, during target and distractor events, healthy subjects exhibit transient spikes of dissonance. That is, during arousing stimuli that capture attention, healthy brains become temporarily more like the brains of patients. Conversely, the brains of patients are continuously like those of healthy subjects in aroused, attention-demanding situations.

The feeling of dissonance

As is generally the case in consciousness science, we have arrived at the moment of phenomenological interpretation: something in an experiment has produced a measurable alteration in the brain. The experimental events are conscious experiences. The brain alteration is thus correlated with the state of mind. This is the reasoning of "neural correlates of consciousness" (NCCs) (Chalmers 2000; Koch 2004). The present study is unusual, however, in that it examines a global property of the brain, and, moreover, that the property is derived by analogy from the analysis of musical sound. Both of these differentia are potentially a good fit with phenomenology. Consciousness science, being young, is open to new methods for characterizing global brain dynamics appropriate to a holistic, structured, interdependent, temporal flow of consciousness. Language offers formal structures for characterizing cognition overall, but language does not offer an adequate model of the rich tapestry of conscious life. In the earlier sections of this paper, it was argued that musicology has developed analytical concepts appropriate to phenomenology, especially in its temporal foundations. Consonance is one such analytical concept. Although it ultimately derives from listeners' response to sound, like many other musical concepts it has a physical analogue, which enables it to be applied to signals of all types, including the varying activity of regions of the brain, as detected by fMRI. Data from 33 subjects confirm that dissonance does indeed distinguish experimental conditions and subject groups.

What then is the neurophenomenological meaning of consonance and dissonance? Consonance between signals increases when their component frequencies are harmonic (identical, or integer multiples of an implied fundamental). Consonant signals are more often synchronized; they are more coherent. Synchrony has been widely discussed as a mechanism for binding percepts (Engel & Singer 2001; Varela et al. 2001; Melloni et al. 2007). Though the consonance discussed here is at low frequencies, it could be well suited to organizing experiences over seconds and minutes as in Francisco Varela's "1 scale" (Varela 1999b). Interacting frequencies have also been proposed as a mechanism for subjective awareness of duration, another structural aspect of consciousness (Matell & Meck 2004). Conversely, schizophrenia patients experience deficits in these capacities (Lloyd 2011b), and other signs of disrupted temporality (Gallagher & Varela 2003). Intrusion by targets and distractors temporarily desynchronizes brain components in situations that phenomenologically are also intrusive and attention-demanding. The healthy brain quickly reestablishes synchrony and returns to its baseline levels of consonance. The conditions of arousal and exogenous distraction subside. Consciousness continues, of course, but in the less aroused, "default" mode, with its mix of internally and externally directed awareness. Disordered brains, in contrast, do not consistently establish a similar degree of consonance. Schizophrenia may be a disorder of harmony.

Conclusion

Music, like language, is culturally universal, developmentally early, uniquely rich in humans, and evolutionarily adaptive (for sexual selection and possibly communication, see Patel 2007). If language is embedded in our biology, then it is likely that music is as well, possibly preceding language in evolutionary development (Froese, Ikegami & Beaton 2012). Music shares with neural networks a capacity for parallel distributed representation, arguably a necessary condition for consciousness. Musical systems are as diverse as human culture, but share in posing constraints on interactions among musical elements. These constraints and the concepts used to describe them carve out a subspace of dynamical systems: musical dynamical systems. A musical dynamical system thus has the ability to represent temporal properties that are essential to full-blown human consciousness.

Music, again like language, has its formal methods of study. Its synchronic and diachronic form and temporal dynamics have been described and modeled in cognitive musicology (e.g., Sethares 2005; Huron 2006; Sethares 2007). Although these formal properties are based on sound, they are readily adapted to the study of any complex system. For any system, the fit of formal musical properties can be probed empirically. Some systems will be "music-like," and some will not – perhaps systems that are "language-like" will exclude musical descriptions. The cornucopia of musicology is waiting to be explored. In this paper, a single musical property showed potential as a tool for understanding consciousness in health and schizophrenia, but this is only a tiny corner of an uncharted domain of exploration. Music offers tools waiting for adoption in the study of the conscious brain.

Acknowledgments

The author would like to thank Tom Froese and an anonymous referee for the journal, *Constructivist Foundations*, for their extensive comments, and Susan Stuart, Bryony Pierce, and Mike Beaton for their generous editorial guidance throughout the writing and revision of this paper.

DAN LLOYD
is the Thomas C. Brownell Professor of Philosophy at Trinity College, Connecticut. He is the author of *Radiant Cool: A Novel Theory of Consciousness* (2004), a book joining noir fiction with a theory of consciousness, and *Simple Minds* (1989). His current projects include *Ghosts in the Machine* (forthcoming), a philosophical dialogue about minds, brains, and computers, and *Subjective Time* (in press), an anthology on the philosophy, psychology, and neuroscience of experienced temporality. His work at the intersection of brain and music is the subject of a documentary in progress, *The Music of the Hemispheres*, directed by Elisa da Prato.

References

Calhoun V. D., Adali T., Pearlson G. D., van Zijl P. C. & Pekar J. J. (2002) Independent component analysis of fMRI data in the complex domain. Magnetic Resonance in Medicine 48(1): 180–192.

Chalmers D. (2000) What is a neural correlate of consciousness? In: Metzinger T. (ed.) Neural correlates of consciousness: Empirical and conceptual questions. MIT Press: Cambridge MA: 17–40.

Clark A. (2003) Natural-born cyborgs: Minds, technologies, and the future of human intelligence. Oxford University Press, Oxford.

Dreyfus H. L. (1972) What computers can't do: A critique of artificial reason. Harper and Row, New York NY.

Elman J. (1990) Finding structure in time. Cognitive Science 14: 179–211.

Engel A. K. & Singer W. (2001) Temporal binding and the neural correlates of sensory awareness. Trends in Cognitive Sciences 5(1): 16–25.

Field H. H. (1978) Mental representation. Erkenntnis 13(1): 9–61.

Fodor J. A. (1975) Language of thought. New York, Crowell.

Fodor J. A. (2008) LOT 2: The language of thought revisited. Oxford University Press, Oxford.

Froese T., Ikegami T. & Beaton M. (2012) Nonhuman primates cannot decontextualize and objectify the actions of their conspecifics. In: Scott-Phillips T. C., Tamariz M., Cartmill E. A. & Hurford J. R. (eds.) The evolution of language: Proceedings of the 9th international conference (EvoLang9). World Scientific Publishing, Singapore: 126–133.

Gallagher S. & Varela F. J. (2003) Redrawing the map and setting the time: Phenomenology and the cognitive sciences. Canadian Journal of Philosophy 29: 93–132.

Garrity A. G., Pearlson G. D., McKiernan K., Lloyd D., Kiehl K. A. & Calhoun V. D. (2007) Aberrant "default mode" functional connectivity in schizophrenia. American Journal of Psychiatry 164(3): 450–457.Gibson J. J. (1979) The ecological approach to visual perception. Houghton Mifflin, Boston.

Hanslick E. (1854) Vom musikalisch Schönen [The Beautiful in Music]. R. Weigel, Leipzig.

Herzog M. (2009) Binding problem. In: Binder M., Hirokawa N. & Windhorst U. (eds.) Encyclopedia of neuroscience 1 (A–C) Springer, Berlin: 388–391. Available at http://www.springerreference.com/docs/html/chapterdbid/114175.html

Higham T., Basell L., Jacobi R., Wood R., Ramsey C. B. & Conard N. J. (2012) Testing models for the beginnings of the Aurignacian and the advent of figurative art and music: The radiocarbon chronology of Geissenklosterle. Journal of Human Evolution 62(6): 664–676.

Hinton G. E. (1984) Parallel computations for controlling an arm. The Journal of Motor Behavior 16: 171–194.

Huron D. B. (2006) Sweet anticipation: Music and the psychology of expectation. MIT Press, Cambridge MA.

Husserl E. (1966) Zur Phänomenologie des inneren Zeitbewusstseins [Phenomenology of inner time consciousness]. Martinus Nijhoff, The Hague. Originally published in 1928.

Husserl E. (1974) Ding und Raum [Thing and space]. Lectures of 1907. Martinus Hijhoff, The Hague.

Husserl E. (1991) On the phenomenology of the consciousness of internal time (1893–1917). Translated by J. B. Brough. Kluwer, Dordrecht.

James W. (1890) The principles of psychology. H. Holt, New York.

Kivy P. (1990) Music alone: Philosophical reflections on the purely musical experience. Cornell University Press, Ithaca NY.

Kivy P. (2002) Introduction to a philosophy of music. Clarendon Press, Oxford.

Koch C. (2004) The quest for consciousness: A neurobiological approach. Roberts & Company, Englewood CO.

Lerdahl F. & Jackendoff R. (1983) A generative theory of tonal music. MIT Press. Cambridge MA.

Lloyd D. (2011a) Mind as music. Frontiers in psychology 2: 63.

Lloyd D. (2011b) Through a glass darkly: Schizophrenia and functional brain imaging. Philosophy, Psychiatry, and Psychology 18(4): 18.

Matell M. S. & Meck W. H. (2004) Corticostriatal circuits and interval timing: Coincidence detection of oscillatory processes. Cognitive Brain Research 21(2): 139–170.

Melloni L., Molina C., Pena M., Torres D., Singer W. & Rodriguez E. (2007) Synchronization of neural activity across cortical areas correlates with conscious perception. The Journal of Neuroscience 27(11): 2858–2865.

Merleau-Ponty M. (1962) Phenomenology of perception. Routledge and Kegan Paul, London. Originally published in French as: Merleau-Ponty M. (1945) Phénoménologie de la perception. Callimard, Paris.

Nagel T. (1974) What is it like to be a bat? Philosophical Review 83: 435–450.

Noë A. (2004) Action in perception. MIT Press, Cambridge MA.

Patel A. D. (2007) Music, language, and the brain. Oxford University Press, Oxford.

Plomp R. & Levelt W. J. (1965) Tonal consonance and critical bandwidth. The Journal of the Acoustical Society of America 38(4): 548–560.

Rumelhart D. E., McClelland J. L. & the PDP research group (1986) Parallel distributed processing: Explorations in the microstructure of cognition. Volume I. MIT Press, Cambridge MA.

Scruton R. (1976) Representation in music. Philosophy 51(197): 273–287.

Scruton R. (2009) Understanding music: Philosophy and interpretation. Continuum, London.

Sethares W. A. (2005) Tuning, timbre, spectrum, scale. Springer, Berlin.

Sethares W. A. (2007) Rhythm and transforms. Springer, Berlin.

Thompson E. (2007) Mind in life: Biology, phenomenology, and the sciences of mind. The Belknap Press of Harvard University Press, Cambridge MA.

van Gelder T. (1999) Wooden Iron? Husserlian phenomenology meets cognitive science. In: Petitot J., Varela F. J., Pachoud B. & Roy J.-M. (eds.) Naturalizing phenomenology: Issues in contemporary phenomenology and cognitive science. Stanford University Press, Stanford CA: 245–265.

Varela F. J. (1999a) Present-time consciousness. Journal of Consciousness Studies 6(2–3): 111–140.

Varela F. J. (1999b) The specious present: A neurophenomenology of time consciousness. In: Petitot J., Varela F. J., Pachoud B. & Roy J.-M. (eds.) Naturalizing phenomenology: Issues in contemporary phenomenology and cognitive science. Stanford University Press, Stanford CA: 266–317.

Varela F. J., Lachaux J. P., Rodriguez E. & Martinerie J. (2001) The brainweb: Phase synchronization and large-scale integration. Nature reviews. Neuroscience 2(4): 229–239.

Varela F. J., Thompson E. & Rosch E. (1991) The embodied mind: Cognitive science and human experience. MIT Press, Cambridge MA.

Wheeler M. (2005) Reconstructing the cognitive world: The next step. The MIT Press, Cambridge MA.

RECEIVED: 18 FEBRUARY 2013
ACCEPTED: 27 JUNE 2013