



Reply to comment

Attending to the forest and the trees

Reply to comments on “Toward a computational framework for cognitive biology: Unifying approaches from cognitive neuroscience and comparative cognition”

W. Tecumseh Fitch

Received 10 July 2014; accepted 11 July 2014

Available online 30 July 2014

Communicated by L. Perlovsky

I am grateful to the commentators for their kind and generally positive comments, and the many interesting and challenging remarks and observations. While limitations of space and time preclude me pursuing each of the interesting lines of potential discussion or debate opened up by these commentaries, I will try to respond to all of the major points, roughly following the original order of the target article [1].

Beyond outworn dichotomies

I am very pleased at how many of the commentators agreed that many of the “standard dichotomies” in cognitive science have outlived their usefulness, and that the “nature versus nurture” dichotomy was singled out for particular disdain (by **Bornkessel-Schlesewsky, Dennett, Okanoya, and Perfors [2–5]**). The only exception is **Dennett’s** argument that some variant of the hardware/software distinction deserves continued use. I agree with Dennett’s point: we need to have some concise way to talk about the flexibility with which brains can change the type of computations they engage in, switching rapidly and easily from (for example) a chess-playing virtual machine to a car-driving virtual machine. But, given the dynamically changing nature of the neural wetware, I think the metaphor that such flexibility involves different programs (“software”) being loaded into a fixed neural “hardware” misleads as much as it illuminates.

Although substituting “wetware” for hardware is a step in the right direction, perhaps a metaphor that remains entirely within the information-processing domain would capture this distinction better – for example that between the operating system and (multiple) programs running within it. This “OS/app” distinction makes no commitments regarding hardware implementation (DOS or iOS can run on many different hardware platforms), but the distinction highlights the dynamic, multitasking property of modern computing that maps naturally onto parallel, multitasking aspects of brain function. From this viewpoint, one might rephrase my dendrophilia hypothesis in terms of an OS

DOI of original article: <http://dx.doi.org/10.1016/j.plrev.2014.04.005>.

DOIs of comments: <http://dx.doi.org/10.1016/j.plrev.2014.05.009>, <http://dx.doi.org/10.1016/j.plrev.2014.05.010>, <http://dx.doi.org/10.1016/j.plrev.2014.06.008>, <http://dx.doi.org/10.1016/j.plrev.2014.06.020>, <http://dx.doi.org/10.1016/j.plrev.2014.06.016>, <http://dx.doi.org/10.1016/j.plrev.2014.07.007>, <http://dx.doi.org/10.1016/j.plrev.2014.06.010>, <http://dx.doi.org/10.1016/j.plrev.2014.07.004>, <http://dx.doi.org/10.1016/j.plrev.2014.06.007>, <http://dx.doi.org/10.1016/j.plrev.2014.06.002>.

<http://dx.doi.org/10.1016/j.plrev.2014.07.008>

1571-0645/© 2014 Published by Elsevier B.V.

upgrade, which provides new system-wide subroutines, available to any app running on the human brain, for parsing, processing and storing tree structures.

On the virtues of formalization

Most commentators also agreed in general terms with the value of formalization in cognitive neuroscience, despite considerable disagreement about *which* formalizations are most promising. The one exception is the commentary of **Bornkessel-Schlesewsky** and colleagues [2], who point out that “the fact that a theory is computationally explicit does not automatically render it biologically plausible.” I agree with this statement, of course, but my point was not to pursue formalization for its own sake (or for its beauty or appeal) but as a *description language* within which specific testable hypotheses can be stated and their predictions rigorously spelled out. It is then an empirical question whether these predictions are met in (biological) reality, and it is this process of making and testing predictions that will allow us, given a set of initially plausible hypotheses, to falsify most of them and home in on and further test the winners.

I propose that the theory of computation, especially the subset termed formal language theory, provides such a description language, allowing us to clearly and unambiguously state testable hypotheses. This body of mathematics provides a notation to describe potentially infinite stringsets, a rich body of theorems that map “languages” (stringsets) onto the computational machinery that can generate or parse them, and a rigorous way to determine if two hypotheses are formally equivalent (e.g., if they make all the same predictions). Thus it provides the kind of agreed-upon modeling formalism we need to rigorously formulate and state our competing hypotheses. But after this has been done, I agree with the commentators that attention must focus on empirical validation or rejection of these hypotheses. I am under no illusions that the theory of computation, with its stacks and queues and rewrite rules, provides anything even close to a final model of biological computation. But it does provide a suitable scaffolding that will help build such a model of natural computation – an endeavor that should be a central shared goal of cognitive neuroscience in the future.

Representations, models, and Bayesian priors

Despite their agreement concerning the importance of a formal and basically Bayesian approach, both Perfors and Johnson offer important caveats. Starting with **Perfors**, she notes that a crucial prerequisite for any Bayesian approach concerns the available form(s) of *representation*, because this is what defines the hypothesis space available to the learner [5]. I agree, and this is what I intended (generally using the term “model” rather than “representation”) when stating that “after meeting certain minimal conditions about independence, statistical approaches can be applied to, and probability distributions derived for, *any* definable type of entity or model.” But the types of entities “visible” for probabilistic inference will vary from species to species and “a Bayesian organism that happens to be incapable of perceiving or remembering tree structures cannot even begin to solve this problem.” Thus the dendrophilia hypothesis specifically concerns species differences in *representations*, defined as “the type[s] of models an organism is equipped to construct,” and my proposal is that it is precisely species differences at this level – representational/model-building capacity for tree representations – that differentiate humans from most other species. The Bayesian learning machinery may nonetheless be, at least to a first approximation, identical. So there is no disagreement (except perhaps terminological) on this count.

However, **Perfors** also notes that despite the *in principle* power of a Bayesian approach, applying it *in practice* with complex model spaces and datasets can be problematic or even intractable. Here, specifically for the case of language, she is rephrasing Chomsky’s “poverty of the stimulus” argument in machine-learning terms: the search space is so vast that no guarantees of optimality or even convergence can be expected in a finite lifetime. Here I fear that she underestimates the value of a strong set of priors in reaching convergence (perhaps not optimal, and perhaps not to precisely the algorithm that generated the data, but close enough for practical purposes). Remember that the child learning language has access not just to priors on syntactic structures but also to a wealth of cognitive presumptions that greatly constrain the space of possible *meanings* of those structures (more technically, on the required *transductions* from signals to meanings). As I observed, there is no need for these cognitive constraints to be specific to either language or humans – any constraints will help. It is clear that the child acquiring language does *not* search the entire possible hypothesis space, but nonetheless *does* reliably converge on something close to the language of the previous generation. I think that a rich set of constraints – mostly cognitive and mostly shared with other species – would be of great value in this process, and I definitely agree with Perfors that “this is an area we need to investigate much more heavily.”

Johnson's commentary agrees about the importance of trees in human language, but notes that the problem of representing trees neurally remains unsolved [6]. I agree, and do not claim to have an answer to this important problem. My goal in the target article was to highlight its central importance for understanding human language and its neural implementation. A more worrisome issue is his observation of the difficulty of demonstrating that, given the many difficulties associated with firmly inferring supra-regularity from a Bayesian viewpoint, we may underestimate the degree to which supra-regularity exists in nonhuman communication systems. I am very sympathetic to this observation, since my colleagues and I have been trying for years, unsuccessfully, to clearly demonstrate non-trivial hierarchical structure in complex bird song (e.g., in the highly complex and seemingly hierarchical songs of the mockingbird *Mimus polyglottos*). Unlike for human language, where we have the rich resources of intuitive understanding and semantic meaning to help demonstrate tree-structure clearly, no such aids are available for complex nonhuman vocal systems like birdsong. I still suspect that at least some animal vocal systems are generated by context-free algorithms, but I am painfully aware that refuting or confirming this suspicion remains a major challenge. Indeed, as noted in **Honing and Zuidema's** comment [7], rigorous demonstrations of true hierarchy have proven illusive even in the cases of human language and music. I hope that experts in machine learning and statistical inference will lend a hand in this endeavor in the future. Thus, while challenging, I do not think these problems are insoluble given an adequate conceptual framework and data.

Appreciating both the trees and the forest

Both **Huber** and **Bornkessel-Schlesewsky** [2,8] noted the importance of networks or “systems level” neuroscience in understanding biological computation, and I wholeheartedly agree with them: we need to look at the forest as well as the trees. A central complaint of my target article was that the connectionist models that have served as the dominant metaphor for brains in the last thirty years of cognitive science focus *only* on the forest, and ignore the trees. My argument, based on fundamental advances in both methodology and theory, is that we need to attend to the trees (specifically, the tree-form of individual neurons) as well. As stated in the target article “the complex, tree-like form of individual neurons plays an important, and indeed central, role in determining the integration and firing behavior of cells, and thus in the computations of the networks they are part of.” Perhaps in arguing the importance of neuronal form so vigorously, I neglected to adequately emphasize that it is only *within* networks that these cellular trees can perform their individual computational tasks. Even the simplest nervous systems (of cnidarians like anemones and jellyfish) are made up of nerve nets. **Tree networks** are the key computational structure, not isolated trees.

Regarding “interneurons” I think that **Huber's** apparent disagreement is purely terminological [8]. By his definition (where “interneuron” means any neuron that is neither primary motor or primary sensory) virtually all of the neurons I discuss in my target article are “interneurons” (e.g. the pyramidal cells in Figs. 2 and 3 of the target article). There is however a more specific interpretation of this term: the local and mostly inhibitory neurons in cerebral cortex, and it is this much smaller class of cells that the article Huber cites is referring to. Although I did not specifically mention such cells in my target article, they provide excellent examples of my main point in that article: that the properties and form of individual cells are central to understanding how the brain computes. Consider, for example, two well-defined cortical inhibitory interneurons illustrated in Fig. 1 (named, not incidentally, for the form of their neurite arbors). Basket cells (1A), which make up roughly 15% of cortical cells, are neurons whose axons encircle the cell bodies of multiple pyramidal cells. Basket cells presumably play some role in normalizing the output emanating from their group of target pyramidal cells. Chandelier cells (1B), also termed axo-axonic cells, specifically target the axon hillock, the site at which axon potentials are initiated. Although both cell classes express the inhibitory neurotransmitter GABA, and both class targets pyramidal cells, their computational roles in neural circuits cannot be understood without considering both their own form, and that of the pyramidal cells that they target. Thus they clearly exemplify the overall point of section three of my paper, about the central importance of the form of individual neurons for brain computation. Huber is however correct about my ignoring glial cells in my computational account: I know of no convincing evidence that they play a fundamental *computational* role in cognition.

Are trees abstractions or cognitive realities?

The commentary by **Honing and Zuidema**, while sympathetic to the overall framework I present, takes issue with my “faith in trees” as all-encompassing models of language and music [7]. First, I certainly don't think that trees alone are adequate models of music or language: there are many other classes of structures (from simple sequences to more

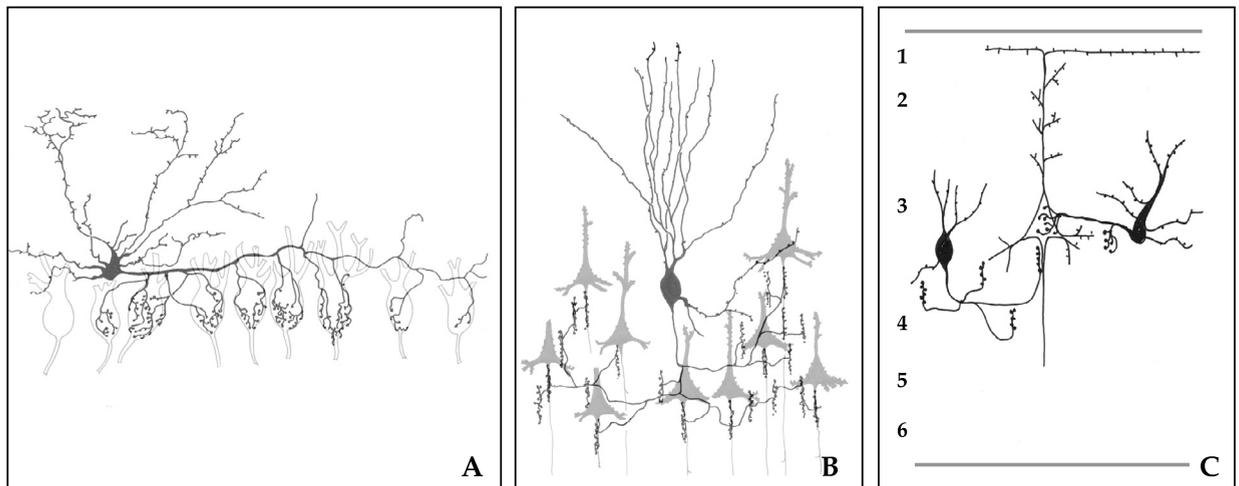


Fig. 1. **Function follows form** – Two different classes of local inhibitory interneuron are illustrated. **Basket cells** (Panel A) form dense axonal synapses surrounding the cell bodies of multiple target cells (in this case, cerebellar Purkinje cells, in gray), thus helping to normalize the firing of a group of neurons. **Chandelier cells** (Panel B) specifically target the axon hillock of multiple target cells (in this case, cortical pyramidal cells, in gray). For clarity, only a subset of the target cells are shown here. Because the axon hillock is the site of spike initiation, inhibition here can prevent cell firing entirely (or limit it to a narrow time window). Panel C illustrates a single cortical pyramidal cell, with its basket and chandelier cell inhibitor, to clarify the distinction in synaptic targeting between these two inhibitory cell types. Such specificity, which cannot be captured in a “spherical neuron” model, plays a central role in contemporary models of neural computation.

richly interconnected graphs) that must play a role. My strong focus on trees is not meant to entail ignoring these other computationally relevant structures. A more telling issue these commentators raise is whether trees are simply useful abstractions (for researchers) or cognitively and neurally real (as I suggest). While I agree that the jury is still out on this central question, I find that recent empirical studies provide grounds for optimism [9,10] that tree structures are indeed generated by the brain in the process of parsing music and language.

Regarding the possibility of circularity of tree inference, these commentators offered the example of metrical trees where the terminals – the data to be explained by the model – are already present in the signal (as notes or sonic events). This seems a poor example to illustrate circularity. First, in rhythmic cognition we often infer a pulse that is *not* unambiguously present in the stimulus – only the simplest music has perceptual events at each slot at the pulse level but none elsewhere. Indeed, the presence of rests (unsounded “events”) illustrates the *inferential process* in rhythmic parsing most clearly [cf. 11]. More generally, in parsing any sequential signal into a tree-formed representation, most of the work is done at the higher levels: grouping the perceptual events (“leaves” or “terminals”) into higher-level chunks, and further grouping those at still higher levels. Thus, even if the input signal does have each terminal represented as a perceptual event, it does not obviate the higher-level grouping analysis into a tree.

Is neural computation continuous or discrete?

A more fundamental disagreement voiced by **Honing and Zuidema** concerns whether neural computation is best considered continuous or discrete. I personally see little need to choose one over the other: I think that the correct answer is “both”. While at some levels (e.g., the local field potential) measurable neural activity is for practical purposes continuous, at others (e.g., how many ions pass through a calcium channel in a given time, or the quantized release of specific numbers of vesicles at a synapse) it is discrete. Most importantly, each individual neuron takes what is practically a fluctuating *continuous* input in its dendritic tree, but converts this to a *discrete* output: either the neuron fires an action potential or not (at least for spiking neurons – the vast majority). Each individual spiking neuron can be considered a miniature analog to digital converter. Normally, neural computation must pass many times through such A-to-D processes. Thus I do not see the question of whether the brain is “at heart” continuous or discrete as a central issue – the core computational units (neurons) are both.

However, it remains clear that *time* in the brain is a continuum, even for an all-or-none event like an action potential or an ion passing through a channel. This contrasts with most current computing technology, where a discrete clock, broadcasting to all computing elements, plays a central role. I know of no evidence for such a discrete clock in

the brain, which must therefore be considered a continuous-time, concurrent-processing system. However, there is growing agreement that one important role of the brain oscillations characteristic of the local field potential (LFP) and coarser measures such as EEG is to provide a widely broadcast, continuous proxy for a discrete clock. For example, [12] have suggested that the coupling of slow theta waves to faster gamma waves in particular brain regions may act as a time code to allow complex multispike information to be integrated and processed in an orderly way. Similar ideas are being explored by many researchers interested in the computational roles of brain oscillations [cf. 13–16]. Thus, even such prototypically continuous signals as EEG or LFP may reflect an underlying system the brain uses to essentially quantize or discretize continuous time.

Distinguishing many forms of hierarchy in the brain

Both **Bowling** and **Petkov** [17,18] use the term “neural hierarchy” in senses different from my usage in the target article, which illustrates the considerable ambiguity of the term “hierarchy” as it is used in modern neuroscience. One usage concerns the arrangement and connectivity of different regions of cortex (e.g. from V1 to V2, etc.), and is (as **Petkov** points out) a characteristic aspect of any mammalian brain. Brains, even those of much simpler organisms, for example the marine worm *Platynereis* [19], exhibit hierarchy in this limited anatomical sense, and this clearly has no implications for their ability to construct neural/cognitive representations of trees. Similar, we can speak of “temporal hierarchies” in which fast-changing processes are “nested” within more slowly changing oscillations (e.g. in the context of brain rhythms). Again these are typical features of most living organisms [13], and the possession of hierarchy in this temporal sense does not imply a capacity for hierarchical representations in my sense. Here, however, there may be some interesting connection between specific types of “nested brain rhythms” and neural processing capacity [12,20], but these were not the focus of my discussion. Thus **Petkov** is certainly correct that “it is incorrect to assume that hierarchical neural processes subservise only hierarchical representations” – but I made no such assumption. This misunderstanding illustrates a considerable danger of polysemous words like “hierarchy”, and the importance of keeping the distinctions clear is explored in more detail in [21].

Trees at different scales

Bowling correctly notes that I offer no concrete suggestions about how the tree form of individual neurons relates to the types of abstract cognitive trees that play a central role in music and language, if indeed any such connection exists [17]. I did not intend to suggest any straightforward connection between these two very different levels of description. My only assertion is that trees *exist* at both levels, are important, and cannot be ignored if we are to understand neural function in computational terms. At the level of cellular function, dendritic and axonal form play a crucial role in determining how, when and why a cell fires. At the level of cognitive function, tree representations play a central role in (at least) human cognitive specializations like music and language. If we aim to build bridging theories between these two levels of explanation, as I think we must in cognitive neuroscience, trees will play a role at both levels. Our theoretical bridges must span from “tree to shining tree” regardless of the importance of trees (or lack thereof) in the intervening spans. As mentioned in the target article, I suspect that there will be some non-trivial connections, but at present this remains nothing more than a suspicion. A model or testable hypothesis concerning such a connection remains a topic for future research, and it plays no important role in my proposed framework.

Dendrophilia = deliberative planning + theory of mind?

Embracing both formal and comparative viewpoints, **Steedman** offers a new and intriguing hypothesis about the origins of our capacity to generate and recognize meaningful tree structures in language [22]. He starts with our capacity for deliberative planning which, because of its need for tracing out multiple causal “if-then” chains, and then backtracking from failed searches, generates tree structures that would require a pushdown automaton (PDA) to implement. In this part of his hypothesis, he echoes a widespread intuition, dating back at least to [23], that complex motor action and planning provide a plausible precursor for a more generalized dendrophilia. I am very sympathetic to this viewpoint and have reviewed the evidence for it elsewhere [21], and I find the argument that such capacities were present in earlier hominins like *Homo erectus* very plausible. However, I think Steedman overstates the case that such complex planning exists in chimpanzees (see below). Kohler’s classic observations of Sultan’s “insightful” tool combinations, however intriguing, do not meet the standards of modern comparative research and should be considered fascinating anecdotes rather than clear demonstrations. (Kohler also reports groups of chimpanzees dancing and

keeping time together, an observation that no modern observer, in lab or field, has echoed, and that thus seems very doubtful.)

The big step in Steedman's hypothesis occurs when a deliberative planning capacity is augmented with a theory of mind, allowing other individuals to be seen as affordances embedded in complex plans. Although recognizing conspecifics as affordances is certainly not unique to humans (to an infant primate, there is no more reliable affordance of food, transport, and comfort than its own mother), recognizing others' *mental* states and proclivities, and incorporating these into plans, seems to be a uniquely well-developed capacity of our species. This is why cooperative "shared intentionality" is arguably one of our most important and special traits [24,25]. Again, Steedman is in good company with this proposal: the hypothesis that selection for social intelligence drove general cognitive capacities is a well-established one [26–28], as is the idea that social intelligence played a key role in language evolution [29–31].

What uniquely differentiates **Steedman's** hypothesis from these predecessors is its formal specificity. He observes that combining theory-of-mind agent-affordance models (involving one form of stack) with a deliberative-planning action-tree system (another form of stack) yields precisely the formal class of computational machinery that computational linguists find is needed to process human language: a linear embedded pushdown automaton (LEPDA). Such a system can generate and process all (and only) the forms of syntactic structures attested in the world's languages. There could be few better illustrations of the value of combining comparative, cognitive, and computational viewpoints in a way that generates testable predictions (e.g., concerning ape planning, or human brain circuitry). I thus wholeheartedly applaud both Steedman's specific hypothesis and the broader vision behind it. It makes good sense, for example, of the facts about Broca's area reviewed in the target article, and offers a clear and explicit alternative to my sensory-motor exaptation hypothesis. I do see music as a problem case for his argument, since music exhibits tree structure without the propositional semantic underpinnings that are a core component of trees in language. Humans somehow manage to absorb and generate musical trees without recourse to any rich semantics-based description language.

Do chimpanzees plan in trees?

I remain unconvinced, however, by **Steedman's** argument, based on Kohler's observations of chimpanzee planning, that "animals like chimpanzees must have a PDA" (pushdown automaton) capable of building arbitrary tree structures or parsing context-free languages [22]. This remains an interesting hypothesis, but should not be treated as a fact. The evidence against it comes from two sources: chimpanzee tool use and chimpanzee "cooperative" hunting. Despite many decades of detailed and astute observation [32–35], current data on chimpanzee tool use emphasizes the differences between chimpanzee and human planning capacities, and suggest that, as Greenfield concluded thirty years ago, the upper limit reached by chimpanzees is equivalent syntactically to that of a two-year old child. Because chimpanzees exhibit the most complex tool use among nonhuman primates [36,37], this conclusion would generalize to other primates. The tool use data thus essentially mirror conclusions based on enculturated apes using symbolic systems [e.g., 38], and do not support Steedman's contention that a full PDA, capable of parsing or generating a rich diversity of tree structures, predated our split from chimpanzees.

Another source of doubt concerning supra-regular planning in chimpanzees comes from the debate about whether the species engages in cooperative hunting. Chimpanzees often hunt, kill, and eat other primates, especially red colobus monkeys. Although early studies of chimpanzee hunting found that it was often solitary and showed few signs of intentional cooperation [39–41], later studies in one population suggested a significant level of cooperative hunting in at least some chimpanzees [42–44]. While the issue remains contested [cf. 45], many field researchers now agree that there is some evidence for cooperation during chimpanzee hunts [46]. But the question of whether this involves shared intentions, or seeing others' intentions as affordances, as required in Steedman's model, remains an open question. Thus, while I applaud Steedman's appeal to comparative data to ground his hypothesis, I believe that the existing data warrant more caution about ape abilities than evidenced in his commentary. I do hope that researchers will use his formal model to generate and test more specific predictions about the cognitive and neural underpinnings of (cooperative) planning in apes.

The instinct to learn – a Bayesian perspective

Last but not least I turn to **Okanoya's** fascinating observations about "late-spiking" neurons in cingulate cortex, and learning differences in Bengalese (domesticated) and munia (wild type) finches [4]. The first case is an excellent

example of how a clearly stated cognitive requirement (the need in associative learning to maintain a memory trace intervening between stimulus and reward) and a neural observation (an unusual class of cortical neurons whose spiking patterns show long delays from stimulation) can be combined to generate further testable hypotheses.

I am also very pleased by Okanoya's willingness to embrace Bayesian terminology to describe differences in song-learning abilities and proclivities in the finches whose songs he has studied for many years. Munias show relatively simple songs, but after several hundred years of domestication by Japanese bird fanciers, their domesticated cousins ("Bengalese finches") sing much more elaborate songs [47,48]. Written records suggest that the only conscious selection pressure was on fertility and plumage, not song, and that the increased song complexity results from (covert) sexual selection by female finches, who breed more readily when exposed to complex song. Domestication freed up the system from other constraints (e.g., predation or species identification). While, as mentioned in Okanoya's commentary, there are already differences in the innate song patterns these strains produce in isolation without tutoring, there are also significant differences in their capacity to learn complex songs: while Bengalese finches can learn either their own or munia's songs equally well, munias achieve 99% accuracy on their own song but only 80% accuracy on Bengalese song. Thus, domestication has led to changes both in the "innate template" underlying song *and* the song-learning mechanism itself [49] – a perfect illustration of the "instinct to learn" Peter Marler emphasizes as the basis of vocal learning [50].

Marr meets Tinbergen meets reverend Bayes

In summary, I was very pleased by the willingness of all commentators to embrace the general framework outlined in my target article and to amend, challenge, and extend it in various important ways. In spite of many remaining disagreements, it seems that all commentators agree about the value of formal models in helping to build new bridges between the neural and cognitive sciences, and specifically about the value of a predictive (rather than passive representational) perspective in such models. Thus – despite some cautionary notes – my optimism for a future merging of Marr's and Bayes' insights seems to be echoed by most commentators, and I can only hope the same will be true of the wider community in cognitive neuroscience.

Although fewer commentators discussed the virtues of a comparative approach, where species differing in particular cognitively relevant ways are studied to help understand the mechanisms underlying these cognitive capacities, no one criticized this component of my framework. I am happy in this case to interpret silence as assent, allowing me to introduce a third perspective, the most fundamentally biological, to the mix. This perspective is best exemplified by the great Nobel Prize-winning ethologist Niko Tinbergen, whose comparative research program and classic statement of the need to address biological questions from multiple viewpoints, including mechanism, development, phylogeny, and function, is a guiding principle for most comparative biologists [51]. Thus, a triumvirate of three great thinkers: Marr, Bayes, and Tinbergen – sit at the head of the table in the framework advanced in my target article. I hope that the future will bring further constructively critical discussion along the lines of these commentaries, along with focused empirical research, to address the fundamental topic of how the brain computes the mind.

Acknowledgements

I gratefully acknowledge the support of ERC Advanced Grant SOMACCA (#230604).

References

- [1] Fitch WT. Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition. *Phys Life Rev* 2014;11:329–64 [in this issue].
- [2] Bornkessel-Schlesewsky I, Schlewsky M, Small SL. Implementation is crucial but must be neurobiologically grounded. Comment on "Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition" by W. Tecumseh Fitch. *Phys Life Rev* 2014;11:365–6 [in this issue].
- [3] Dennett DC. The software/wetware distinction. Comment on "Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition" by W. Tecumseh Fitch. *Phys Life Rev* 2014;11:367–8 [in this issue].
- [4] Okanoya K. Learning, epigenetics, and computation: an extension on Fitch's proposal. Comment on "Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition" by W. Tecumseh Fitch. *Phys Life Rev* 2014;11:389–90 [in this issue].

- [5] Perfors A. Representations, approximations, and limitations within a computational framework for cognitive science. Comment on “Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition” by W. Tecumseh Fitch. *Phys Life Rev* 2014;11:369–70 [in this issue].
- [6] Johnson M. The forest for the trees. Comment on “Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition” by W. Tecumseh Fitch. *Phys Life Rev* 2014;11:371–2 [in this issue].
- [7] Honing H, Zuidema W. Decomposing dendrophilia. Comment on “Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition” by W. Tecumseh Fitch. *Phys Life Rev* 2014;11:375–6 [in this issue].
- [8] Huber L. Brains are not just neurons. Comment on “Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition” by Fitch. *Phys Life Rev* 2014;11:373–4 [in this issue].
- [9] Koelsch S, Rohrmeier M, Torrecuso R, Jentschke S. Processing of hierarchical syntactic structure in music. *Proc Natl Acad Sci* 2013;110:15443–8.
- [10] Pallier C, Devauchelle A-D, Dehaene S. Cortical representation of the constituent structure of sentences. *Proc Natl Acad Sci* 2011;108:2522–7.
- [11] Fitch WT. Rhythmic cognition in humans and animals: distinguishing meter and pulse perception. *Front Syst Neurosci* 2013;7:1–16.
- [12] Lisman JE, Jensen O. The theta-gamma neural code. *Neuron* 2013;77:1002–16.
- [13] Buzsáki G, Logothetis NK, Singer W. Scaling brain size, keeping timing: evolutionary preservation of brain rhythms. *Neuron* 2013;80:751–64.
- [14] Fries P. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn Sci* 2005;9:474–80.
- [15] MacKay WA. Synchronized neuronal oscillations and their role in motor processes. *Trends Cogn Sci* 1997;1:176–83.
- [16] Vinck M, Womelsdorf T, Buffalo EA, Desimone R, Fries P. Attentional modulation of cell-class-specific gamma-band synchronization in awake monkey area V4. *Neuron* 2013;80:1077–89.
- [17] Bowling D. Cognitive theory and brain fact: insights for the future of cognitive neuroscience. Comment on “Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition” by W. Tecumseh Fitch. *Phys Life Rev* 2014;11:377–9 [in this issue].
- [18] Petkov C. Mental structures and hierarchical brain processing. Comment on “Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition” by W. Tecumseh Fitch. *Phys Life Rev* 2014;11:382–8 [in this issue].
- [19] Randel N, Asadulina A, Bezares-Calderón L, Verasztó C, Williams E, Conzelmann M, et al. Neuronal connectome of a sensory-motor circuit for visual navigation. *ELife* 2014:e02730.
- [20] Buzsáki G. *Rhythms of the brain*. New York: Oxford University Press; 2006.
- [21] Fitch WT, Martins MD. Hierarchical processing in music, language and action: Lashley revisited. *Ann NY Acad Sci* 2014;1316:87–104.
- [22] Steedman M. Evolutionary basis for human language. Comment on “Toward a computational framework for cognitive biology: Unifying approaches from cognitive neuroscience and comparative cognition” by Tecumseh Fitch. *Phys Life Rev* 2014:382–8 [in this issue].
- [23] Lashley K. The problem of serial order in behavior. In: Jeffress LA, editor. *Cerebral mechanisms in behavior, the Hixon symposium*. New York: Wiley; 1951. p. 112–46.
- [24] Tomasello M, Carpenter M, Call J, Behne T, Moll H. Understanding and sharing intentions: the origins of cultural cognition. *Behav Brain Sci* 2005;28:675–735.
- [25] Herrmann E, Call J, Hernández-Lloreda MV, Hare B, Tomasello M. Humans have evolved specialized skills of social cognition: the cultural intelligence hypothesis. *Science* 2007;317:1360–6.
- [26] Byrne R. Machiavellian intelligence. *Evol Anthropol* 1997;5:172–80.
- [27] Humphrey NK. The social function of intellect. In: Bateson PPG, Hinde RA, editors. *Growing points in ethology*. Cambridge: Cambridge University Press; 1976. p. 303–17.
- [28] Jolly A. Lemur social behavior and primate intelligence. *Science* 1966;153:501–6.
- [29] Worden R. The evolution of language from social intelligence. In: Hurford JR, Studdert-Kennedy M, Knight C, editors. *Approaches to the evolution of language*. New York: Cambridge University Press; 1998. p. 148–66.
- [30] Seyfarth RM, Cheney DL, Bergman TJ. Primate social cognition and the origins of language. *Trends Cogn Sci* 2005;9:264–6.
- [31] Fitch WT, Huber L, Bugnyar T. Social cognition and the evolution of language: constructing cognitive phylogenies. *Neuron* 2010;65:795–814.
- [32] Greenfield PM. Language, tools, and brain: the ontogeny and phylogeny of hierarchically organized sequential behavior. *Behav Brain Sci* 1991;14:531–95.
- [33] McGrew WC. *The cultured chimpanzee*. Cambridge, UK: Cambridge University Press; 2004.
- [34] McGrew WC. Savanna chimpanzees dig for food. *Proc Natl Acad Sci* 2007;104:19167–8.
- [35] Sanz C, Morgan D, Gulick S. New insights into chimpanzees, tools, and termites from the Congo basin. *Am Nat* 2004;164:567–81.
- [36] McGrew WC. *Chimpanzee material culture*. Cambridge, UK: Cambridge University Press; 1992.
- [37] Pruetz JD, Bertolani P. Savanna chimpanzees, pan troglodytes verus, hunt with tools. *Curr Biol* 2007;17:412–7.
- [38] Savage-Rumbaugh ES, Murphy J, Sevcik RA, Brakke KE, Williams SL, Rumbaugh DM. Language comprehension in ape and child. *Monogr Soc Res Child Dev*, vol. 58. 1993. p. 1–221.
- [39] Goodall J. *The chimpanzees of Gombe: patterns of behavior*. Cambridge, MA: Harvard University Press; 1986.
- [40] Busse CD. Do chimpanzees hunt cooperatively? *Am Nat* 1978;112:767–70.
- [41] Stanford CB, Wallis J, Mpongo E, Goodall J. Hunting decisions in wild chimpanzees. *Behaviour* 1994;131:1–18.
- [42] Boesch C, Boesch H. Hunting behavior of wild chimpanzees in the Tai National Park. *Am J Phys Anthropol* 1989;78:547–73.
- [43] Boesch C, Boesch-Achermann H. *The chimpanzees of the Tai forest*. Oxford: Oxford University Press; 2000.
- [44] Boesch C. Cooperative hunting in wild chimpanzees. *Anim Behav* 1994;48:653–67.
- [45] Tomasello M, Carpenter M, Call J, Behne T, Moll H. Author’s response: in search of the uniquely human. *Behav Brain Sci* 2005;28:721–34.
- [46] Watts DP, Mitani JC. Hunting behavior of chimpanzees at Ngogo, Kibale National Park, Uganda. *Int J Primatol* 2002;23:1–28.

- [47] Okanoya K. Sexual display as a syntactic vehicle: the evolution of syntax in birdsong and human language through sexual selection. In: Wray A, editor. *The transition to language*. Oxford: Oxford University Press; 2002.
- [48] Okanoya K. Song syntax in Bengalese finches: proximate and ultimate analyses. *Adv Study Behav* 2004;34:297–346.
- [49] Takahashi M, Okanoya K. Song learning in wild and domesticated strains of white-rumped munia, *Lonchura striata*, compared by cross-fostering procedures: domestication increases song variability by decreasing strain-specific bias. *Ethology* 2010;116:396–405.
- [50] Marler P. The instinct to learn. In: Carey S, Gelman R, editors. *The epigenesis of mind: essays on biology and cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates; 1991. p. 37–66.
- [51] Tinbergen N. On aims and methods of ethology. *Z Tierpsychol* 1963;20:410–33.