computational challenges of chunking the data and of building hierarchies.

Biological communication is about affecting behavior, not pumping bits. Our final point focuses on the communicative function of language. Viewing a memory window as a communication "bottleneck" suggests that massive amounts of information must flow through the channel in question. However, the real function of a message is to influence the rich network of connotations and interconnections already present in the listener's brain (cf. Edelman 2015, sect. 2.3). Communication is about generating adaptive behavioral changes (Burghardt 1970; Green & Marler 1979) – the listener gleans from it cues relevant to decision-making. For this, a signal must be informative and reliable in the given context (Leger 1993); the amount of information is not the main issue (except as a signal of quality, as in complex courtship songs; Lachmann et al. 2001). This implies that evolutionary selection in language is for how messages fit into the information already represented by their recipient; a bottleneck may not impose significant constraints here.

NOTES

1. If verbal memory indeed evolves, language is the niche in which it does so. The target article seems to gloss over the intimate connection between cultural evolution and niche construction (Odling-Smee et al. 2003). In focusing on how "linguistic patterns, which can be processed through that bottleneck, will be strongly selected" (sect. 5, para. 3), C&C ignore the possibility of there being also selection for individuals who can better process linguistic patterns.

2. As C&C note, correctly, regarding Chunk-and-Pass, "it is entirely possible that linguistic input can simultaneously, and perhaps redundantly, be chunked in more than one way" (sect. 3.2, para. 4). This point suggests that chunking on its own, especially when carried out recursively/hierarchically, is likely to severely exacerbate the combinatorial problem faced by the learner, rather than resolve the bottleneck issue.

Memory limitations and chunking are variable and cannot explain language structure

doi:10.1017/S0140525X15000898, e84

Maryellen C. MacDonald

Department of Psychology, University of Wisconsin-Madison, Madison, WI 53706.

mcmacdonald@wisc.edu

http://lcnl.wisc.edu/people/mcm/

Abstract: Both the Now-or-Never bottleneck and the chunking mechanisms hypothesized to cope with it are more variable than Christiansen & Chater (C&C) suggest. These constructs are, therefore, too weak to support C&C's claims for the nature of language. Key aspects of the hierarchical nature of language instead arise from the nature of sequencing of subgoals during utterance planning in language production.

Christiansen & Chater (C&C) overstate both the limitations of the Now-or-Never bottleneck and the lossy character of chunking, and they are overly optimistic that memory limitations can explain the nature of language. C&C correctly note that memory limitations during planning for language production promote *incremental planning* (where planning of the utterance and its execution of action are interleaved), but the memory limitations are not as strict as they suggest. Whereas "radical incrementality" – very minimal advance planning owing to a severe memory bottleneck – once had its proponents in language production, recent studies argue for looser constraints, with more tolerance for higher memory loads and more extensive advance planning (Ferreira & Swets 2002). The extent of advance planning may even be under some degree of implicit strategic control (Ferreira & Swets 2002; Wagner et al. 2010), suggesting that, rather than the memory bottleneck controlling us, we instead can exert some control over our own memory loads during language production. The bottleneck also isn't always so severe in comprehension, and chunking isn't as uniformly eager as C&C portray. Downstream linguistic input affects interpretation of earlier material (MacDonald 1994; Warren & Sherman 1974), which shouldn't occur if chunking greedily passes off the early information to the next level. Variability in the tolerance of memory loads suggests that the Now-or-Never bottleneck is really more of a wide-mouth jar, or perhaps more of an adjustable drawstring closure, and the consequences for the nature of language will therefore need adjustment as well.

Similarly, C&C view the lossy nature of Chunk-and-Pass processing as essential to explaining the nature of language processing, but chunking is neither as lossy nor as bottom-up as they suggest. C&C argue that in speech perception, sounds are rapidly chunked into words, leaving the sounds behind, so that the just-perceived sounds do not interfere with upcoming ones. These claims create several puzzles: First, this very bottom-up characterization of chunking is inconsistent with evidence for top-down influences in perception. C&C's focus on using context only for predicting the future is misplaced, because top-down processes also allow higher-level information to elaborate earlier percepts. Examples include the word superiority effect (Cattell 1886) and the phoneme restoration effect (Warren 1970), in which word representations affect perception of their parts (letters, phonemes). If chunking is so eager and lossy, it's not clear how higher-level word information could refine the lower-level percepts that should have already been discarded by lossy chunking. Second, if the memory bottleneck is so narrow, how is there room for interference, which by definition depends on several elements being in memory at the same time? There are numerous examples of semantic and sound overlap creating memory interference over fairly long distances during both comprehension (Acheson & MacDonald 2011; Van Dyke & Johns 2012), and production (Hsiao et al. 2014; Smith & Wheeldon 2004), again suggesting that the bottleneck can't be as strict at C&C describe. Third, if lossy chunking is the solution to memory interference, why is it so easy to find interference effects? The existence of memory interference suggests that chunking may not always be so lossy after all. In at least some circumstances, there appears to be real value in non-lossy processing, such as the Levy et al. (2009) example that C&C note as well as use of prosodic information over long distances (Morrill et al. 2014). These and other examples call into question the essence of lossy, greedy, bottom-up chunking as a design feature for language.

C&C note some variability in memory limits and chunking, but they do not discuss the consequences of variability for their account. They illustrate their ideas with an individual identified as SF, who can recall vast strings of meaningless digits by chunking them into meaningful units such as dates, and using the chunks to guide production. The analogy to language is unfortunate, because SF's chunking strategies are both conscious and idiosyncratic, inviting the inference that language users' chunking units are similarly variable. In sum, if memory limitations and the lossy and eager characteristics of chunking have notable exceptions and are subject to individual differences, then it is difficult to make them the foundation of claims for the nature of human language.

More seriously, no matter how we conceive the memory bottleneck, it can explain neither the existence of a hierarchy in language representations, nor why the hierarchy has certain levels of representation across individuals and not others. Consider a non-linguistic analogy: the visual processes necessary to recognize a cup. Let's assume that these processes, also constrained by memory bottlenecks, have multiple stages of chunking and passing from low-level visual processing up to object recognition. From these perceptual stages, however, we would not want to conclude that the percept itself, the cup, has a hierarchical structure. Similarly, the memory-constrained chunking and passing for language perception, even if it works exactly as C&C describe, does not give the percept – language – its hierarchical structure.

Rather than trying to wring structure out of memory limitations, I suggest that key aspects of hierarchical structure emerge from how goals are realized in action (MacDonald 2013). Like all actions, language production must unfold over time, meaning that the various subgoals of the action must be planned and ordered in some way (Lashley 1951). For both nonlinguistic and linguistic actions, the nature of the hierarchy is constrained by the need to make decisions for some subgoals in order to plan others. To reach for a cup, the choice of which hand to use determines and must precede planning the reach. Similarly, a speaker must choose words (*cup* or *mug*?) before programming their articulation, naturally creating a hierarchy of lexical and sublexical plans. Although language and nonlinguistic action are not identical, important aspects of the hierarchical nature of language emerges from the staging of language production planning processes over time. Furthermore, although action plans are held in memory and are affected by the nature of that memory, memory limitations themselves cannot bear the explanatory burden that C&C ascribe to them.

Exploring some edges: Chunk-and-Pass processing at the very beginning, across representations, and on to action

doi:10.1017/S0140525X15000904, e85

Rose Maier and Dare Baldwin
Department of Psychology, University of Oregon, Eugene, OR 97403.
rosem@uoregon.edu **baldwin@uoregon.edu**

Abstract: We identify three "working edges" for fruitful elaboration of the Chunk-and-Pass proposal: (a) accounting for the earliest phases of language acquisition, (b) explaining diversity in the stability and plasticity of different representational types, and (c) propelling investigation of action processing.

Experience is dynamic and ephemeral, yet humans routinely generate abstract representations of their individualized experience that simultaneously achieve enough stability, plasticity, and interindividual parity to radically facilitate social and cognitive functioning. Christiansen & Chater's (C&C's) ambitious Chunk-and-Pass processing (CPP) proposal offers hope of a comprehensive and elegant account of how this can be. CPP has impressive explanatory breadth, neatly tying language acquisition to language change and language evolution, while also offering promise of a unified account of perception and cognition more generally. By C&C's own acknowledgment, however, many facets of the CPP account cry out for elaboration. In our view, three "working edges" will be (a) accounting for the earliest inception of language acquisition, (b) explaining stability and plasticity differences in learning profiles across knowledge systems (within language as well as across domains), and (c) elaborating CPP on the action processing front.

Regarding the first issue, C&C provide a workable framework for describing language acquisition once basic acoustic units have been discovered (e.g., phonemes, syllables), but do not describe how utter novices initially break into the system. Of course, there is a sizable literature investigating how infants initiate analysis of streaming speech (e.g., Vouloumanos & Werker 2007; Werker et al. 2012). One litmus test of the viability of CPP will be its ability to account for the phenomena documented in this literature within a unified Chunk-and-Pass framework. Among the complexities to be confronted here include findings indicating that infants' identification/construction of basic acoustic units may still be taking place at the same time that they are beginning to chunk longer strings of sounds together into words or morphemes. For example, infants remain quite sensitive to phonetic distributions until well into the first year; at 6 to 8 months, just 2–3 minutes of focused exposure to new distributions may be enough to temporarily rearrange infants' phonetic categories (Maye et al. 2002). And yet, by this same age, infants typically recognize at least a handful of words, including "mommy" and "daddy" (Tincoff & Jusczyk 1999), their own name (Bortfeld et al. 2005; Mandel et al. 1995), and several body part terms, such as "feet" and "hand" (Tincoff & Jusczyk 2012). Does CPP somehow build linguistic structure even without clear basic units over which to operate (in contradiction to hypotheses C&C articulate on this matter; e.g., sect. 3.2, para. 1)? Alternatively, does CPP operate on units only as they reach some criterion of availability, so that words composed of early-identified phonemes would potentially be available for chunking, whereas words with more difficult-to-identify phonemes are not? Or do processes other than Chunk-and-Pass need to be brought in to account for the earliest phases of language acquisition?

The second working edge we identify relates to stability and plasticity of representations. C&C note that stability and plasticity trade off: Learning depends on representations being updated to incorporate new content, but at the same time, some degree of stability is needed to avoid new information overwhelming previously acquired information. They argue that stability is a natural product of the compression that occurs during Chunk-and-Pass processing. The processing of linguistic content is "lossy" – the only features retained are those that are captured by a learner's current model of the language, making it difficult to dramatically alter that model since the features necessary to do so are likely the very ones lost in compression. This seems persuasive on the face of it, but leaves unclear how CPP can account for a different stability/plasticity issue: namely, the observation that representations of different types display distinct stability/plasticity profiles. In language, acquired representations of some kinds (e.g., phonetic and syntactic representations) display a strong propensity to stabilize and become markedly resistant to change (e.g., Johnson & Newport 1989; Kuhl 2004; Lenneberg 1964; Yoshida et al. 2010), whereas a variety of evidence suggests that other representational types (e.g., open-class lexical items) seem to display considerably more plasticity (e.g., Curtiss 1977; Newport 1990; Talmy 2000; Weber-Fox & Neville 1996). In question is whether these different plasticity profiles across representational types arise naturally from CPP. Are there differences in the information to be encoded across various types of representations such that the model would predict an emphasis on stability in some cases versus ongoing plasticity in others? Alternatively, will it be necessary to look to mechanisms beyond CPP to account for such differences, such as diverse neural commitment timetables?

Our third "working edge" focuses on action processing as a particularly fruitful target for broadening the scope of CPP-related investigation. Intuitively, language and action processing seem closely linked. Language can be regarded as one form of action, after all, and both language and action are subject to the Now-or-Never bottleneck, making them amenable to a CPP account, as C&C themselves note. Strikingly, however, investigation regarding action processing lags considerably behind language. One glaring example is the lack of a generally accepted inventory of basic actions, comparable to inventories of phonemes or syllables in language (cf. interesting but small-scale efforts along these lines, such as *therblig*, Gilbreth & Gilbreth 1919). Another example concerns hierarchical structure, which seems to be a fundamental organizing principle of both action and linguistic representations. To illustrate in the action context, observers typically note that an action such as getting a cup of coffee comprises embedded subgoals, such as getting a mug from a cupboard, placing it on a counter, pouring coffee into the mug, and so on. At the same time, relevant levels of that hierarchy seem not to be as crisp or well-defined as they are in language. A "learning to process" account may provide welcome guidance for continuing attempts to gain purchase on the