research is not as far-fetched as one may think: Tversky and Kahneman introduced a simulation heuristic in the 1970s according to which people predict the likelihood of an upcoming event by how easy it is to simulate it. Other heuristics may well be worth exploring (in line with the affect heuristic, emotionally charged words, e.g., stupid, boring, may be predicted more rapidly). Our point is simply that predictive language processing is likely to be complex and may make use of a set of rather diverse mechanisms (with many yet unexplored).

Importance of prediction-by-simulation in development. P&G suggest that analysis of children's prediction abilities might throw light on the distinction between prediction-by-association and prediction-by-simulation and place stronger emphasis on prediction-by-association in young children: Prediction-by-association might play a more important role when listeners and speakers have little in common with each other, such as the case of children listening to adults' talking.

In a recent experiment examining 2-year-olds' prediction abilities, however, we found that, consistent with prediction-by-simulation, only toddlers in possession of a large *production* vocabulary are able to predict upcoming linguistic input in another speaker's utterance (Mani & Huettig 2012; see Melzer et al. 2012, for similar results in action perception). Furthermore, if, as P&G suggest, covert imitation is the driving force of prediction-by-simulation, then 18-month-olds are equipped with the cognitive pre-requisites for covert imitation: Covert imitation can modulate infants' eye gaze behaviour around a (linguistically relevant) visual scene (Mani & Plunkett 2010; Mani et al. 2012) similar to adults' behaviour (Huettig & McQueen 2007). Prediction-by-simulation may also be an important developmental mechanism to train the production system (Chang et al. 2006). In sum, prediction-by-simulation appears to be crucial even early in development, and hence prediction-by-association is not necessarily the simple prediction mechanism which dominates early childhood.

Mediating factors. Finally, there are many mediating factors (e.g., literacy, working memory capacity, cross-linguistic differences) involved in predictive language processing whose interaction with anticipatory mechanisms have been little explored and whose importance, we believe, has been vastly underestimated. Mishra et al. (2012), for instance, observed that Indian high literates, but not low literates, showed language-mediated anticipatory eye movements to concurrent target objects in a visual scene. Why literacy modulates anticipatory eye gaze remains to be resolved, though literacy-related differences in associations (including low-level word-to-word contingency statistics, McDonald & Shillcock, 2003), online generation of featural restrictions, and general processing speed are likely to be involved. Similarly, Federmeier et al. (2002) found that older adults are less likely to show prediction-related benefits during sentence processing with a strong suggestion that differences in working memory capacity underlie differences in predictive processing. The influence of such mediating factors may greatly depend on the situation language users find themselves in: Anticipatory eye gaze in the visual world, for instance, requires the building of online models allowing for visual objects to be linked to unfolding linguistic information, places, times, and each other. Working memory capacity may be particularly important for anticipatory processing during such language-vision interactions (Huettig & Janse 2012).

More work is also required with regard to the specific representations which are pre-activated in particular situations. Event-related potential studies have shown that even the grammatical gender (van Berkum et al. 2005), phonological form (DeLong et al. 2005), and visual form of the referents (Rommers et al. 2013) of upcoming words can be anticipated. Most of these studies, however, have used highly predictive "lead-in" sentences. It also remains to be seen to which extent these specific representations are activated in weakly and moderately predictive contexts. Last but not least, languages differ dramatically in all levels of linguistic organisation (Evans & Levinson 2009). These cross-linguistic differences are bound to have substantial impacts on the specifics (and degree) of anticipatory processing a particular language affords.

Future work could usefully explore the cognitive reality and relative importance of the potential mechanisms and mediating factors mentioned here. Even though Occam's razor may favour single-mechanism accounts, we conjecture that multiple-mechanism accounts are required to provide a complete picture of anticipatory language processing.

Toward a unified account of comprehension and production in language development

doi:10.1017/S0140525X12002658

Stewart M. McCauley and Morten H. Christiansen
Department of Psychology, Cornell University, Ithaca, NY 14853.
smm424@cornell.edu christiansen@cornell.edu
http://cnl.psych.cornell.edu

Abstract: Although Pickering & Garrod (P&G) argue convincingly for a unified system for language comprehension and production, they fail to explain how such a system might develop. Using a recent computational model of language acquisition as an example, we sketch a developmental perspective on the integration of comprehension and production. We conclude that only through development can we fully understand the intertwined nature of comprehension and production in adult processing.

Much like current approaches to language processing, contemporary accounts of language acquisition typically assume a sharp distinction between comprehension and production. This assumption is driven, in large part, by evidence for a number of asymmetries between comprehension and production in development. Comprehension is usually taken to precede production (e.g., Fraser et al. 1963), although there are certain instances in which children exhibit adult-like production of sentence types that they do not appear to comprehend correctly (cf. Grimm et al. 2011). Evidence for such asymmetries strongly constrains theories of language acquisition, challenging integrated accounts of development and, by extension, integrated accounts of adult processing. Hence, it is key to determine the plausibility of a unified framework for acquisition that is compatible with evidence for comprehension/production asymmetries.

Although Pickering & Garrod's (P&G's) target article may be construed as a useful point of departure in this respect, P&G pay scant attention to how such a unified system for comprehension and production might develop. As a result, they implicitly subscribe to a different, questionable distinction often made in the language literature: the separation of acquisition from adult processing. In light of this, and given the tendency of developmental psycholinguists to view comprehension and production as separate systems, we briefly sketch a unified developmental framework for understanding comprehension and production as a single system, instantiated by a recent usage-based computational model of acquisition (McCauley & Christiansen 2011; submitted). Importantly, our approach is consistent with evidence for comprehension/production asymmetries in development, even while uniting comprehension and production within a single framework.

Our computational model, like that of Chang et al. (2006), simulates both comprehension and production, but it goes beyond this and previous usage-based models (e.g., Borensztajn et al. 2009; Freudenthal et al. 2007) in that (a) it learns to do so incrementally using simple distributional information; (b) it offers broad, cross-linguistic coverage; and (c) it accommodates a range of developmental findings. The model learns from corpora of child and child-directed speech, acquiring item-based knowledge in a purely incremental fashion, through online learning using backward transitional probabilities (which infants track; cf. Pelucchi et al. 2009). The model uses peaks and dips in transitional probabilities to chunk words together as they are encountered, incrementally building an item-based "shallow parse" as each incoming utterance unfolds. The model stores the word sequences it groups together, gradually building up an inventory of multiword chunks – a "chunkatory" – which underlies both comprehension and production. When the model encounters a multiword utterance produced by the target child of a corpus, it attempts to generate an identical utterance using only chunks and transitional probabilities learned up to that point. Crucially, the very same chunks and distributional information used during production are used to make predictions about upcoming material during comprehension. This type of prediction-by-association facilitates the model's shallow processing of the input. The model's comprehension abilities are scored against a state-of-the-art shallow parser, and its production abilities are scored against the target child's original utterances (the model's utterances must match the child's).

The model makes close contact with P&G's approach in that it uses information employed during production to make predictions about upcoming linguistic material during comprehension (consistent with recent evidence that children's linguistic predictions are tied to production; cf. Mani & Huettig 2012). However, our approach extends P&G's account from prediction to the acquisition and use of linguistic knowledge itself; comprehension and production rely upon a single set of statistics and representations, which are reinforced in an identical manner during both processes.

Moreover, our model's design reflects recent psycholinguistic findings that have hitherto remained largely unconnected, but which, when viewed as complementary to one another, strongly support a unified framework for comprehension and production. First, the model is motivated by children's use of multiword units in production (Bannard & Matthews 2008), which cautions against models of production in which words are selected independently of one another. The model's primary reliance on the discovery and storage of useful multiword sequences follows this line of evidence. Second, the model is motivated by evidence that children, like adults, can rely on shallow processing and underspecified representations during comprehension (e.g., Gertner & Fisher 2012; Sanford & Sturt 2002). Shallow processing, supplemented by contextual information (e.g., tied to semantic and pragmatic knowledge) may often give children the appearance of comprehending grammatical constructions they have not yet mastered (and therefore cannot use effectively in production). The model exhibits this in its better comprehension performance; through chunking, the model can arrive at an item-based "shallow parse" of an utterance, which can then be used in conjunction with semantic and pragmatic information to arrive at a "good enough" interpretation of the utterance (Ferreira et al. 2002). On the production side, however, the model – like a child learning to speak – is faced with the task of retrieving and sequencing words and chunks in a particular order. Hence, asymmetries arise from differing task demands, despite the use of the very same statistics and linguistic units during both comprehension and production.

Such an abandonment of the "cognitive sandwich" approach to acquisition clearly has implications for adult processing. If, as we suggest and make explicit in our model, children learn to comprehend and produce speech by using the same distributional information and chunk-based linguistic units for both tasks, we would expect adults to continue to rely on a unified set of representations. This is corroborated by studies showing that, like children, adults not only rely on multiword units in production (Janssen & Barber 2012), but also use multiword sequences during comprehension (e.g., Arnon & Snider 2010; Reali & Christiansen 2007). This evidence further suggests that prediction-by-association may be more important for language processing than assumed by P&G, not just for children as indicated by our model, but also for adults. It is only by considering how the adult system emerges from the child's attempts to comprehend and produce linguistic utterances that we can hope to reach a complete understanding of the intertwined nature of language comprehension and production.

What does it mean to predict one's own utterances?

doi:10.1017/S0140525X12002786

Antje S. Meyer[a,b] and Peter Hagoort[a,b]

[a]*Max Planck Institute for Psycholinguistics, 6500 AH Nijmegen, The Netherlands;* [b]*Radboud University Nijmegen, 6525 HP Nijmegen, The Netherlands.*

antje.meyer@mpi.nl **peter.hagoort@mpi.nl**
www.mpi.nl

Abstract: Many authors have recently highlighted the importance of prediction for language comprehension. Pickering & Garrod (P&G) are the first to propose a central role for prediction in language production. This is an intriguing idea, but it is not clear what it means for speakers to predict their own utterances, and how prediction during production can be empirically distinguished from production proper.

Pickering & Garrod (P&G) offer an integrated framework of speech production and comprehension, highlighting the importance of predicting upcoming utterances. Given the growing evidence for commonalities between production and comprehension processes and for the importance of prediction in comprehension, we find their proposal timely and interesting.

Our comment focuses mainly on the role of prediction in language production. P&G propose that speakers predict aspects of their utterance plans and compare these predictions against the actual utterance plans. This monitoring process happens at each processing level, that is, minimally at the semantic, syntactic, and phonological level.

Given the important role of prediction in comprehension and the well-attested similarities between production and comprehension, the idea that prediction should play a role in speech production follows quite naturally. Nevertheless, to us the proposal that speakers predict their utterance plans does not have immediate appeal. This is because, in everyday parlance, prediction and the predicted event have some degree of independence. It is because of this independence that predictions may or may not be borne out. It makes sense to say a person predicts the outcomes of their hand or jaw movements, as these outcomes are not fully determined by the cognitive processes underlying the predictions, but depend, among other things, on properties of the physical environment that may not be known to the person planning the movement. Similarly, it makes sense to say that a listener predicts what a speaker will say because the speaker's utterances are not caused by the same cognitive processes as those that lead to the listener's prediction. Speaker and listener each have their own, private cognition and therefore the listener's expectations about the speaker's utterance may or may not be met.

We can predict our own utterances. For instance, based on memory of past experience, I can predict how I will greet my family. However, such predictions concern overt behavior rather than plans for behavior, and they occur offline rather than in parallel with the predicted behavior. Just like predictions about other persons, my predictions of my own utterances may or may not be borne out, depending on circumstances not known at the moment of prediction. I may, for instance, deviate from my predicted greeting if I find my family standing on their heads.

Such offline predictions of overt behavior differ from the predictions proposed by P&G. In their framework, speakers predict their utterance *plans* as they plan them, with prediction at each planning level running somewhat ahead of the actual planning. Importantly, the predictions are based on the same information as the predicted behavior, namely, the speaker's intention