

Claim 3: Acceptability judgments are susceptible to differences in instructions (sect. 1.2, para. 3). Claim 3 has been directly investigated by Cowart (1997), who reports that the systematic manipulation of instructions does not change the pattern of acceptability judgments for factorial designs.

Claim 4: Acceptability judgments are impacted by sentence processing effects (sect. 1.2, para. 5). Claim 4 is technically true, but B&P exaggerate its consequences. First, many classic lexical and sentence processing effects have relatively small or negligible effects on acceptability (e.g., Featherston 2009; Phillips 2009; Sprouse 2008; Sprouse et al. 2012). Second, very few syntactic phenomena have been proposed to be fully reducible to sentence processing effects. The lone exceptions to this appear to be constraints on long-distance dependencies (e.g., Kluender & Kutas 1993; Hofmeister & Sag 2010), but in that case, a number of experimental studies have disproven the reductionist predictions (Phillips 2006; Sprouse et al. 2012; Yoshida et al. 2014). Thus, to the extent that AJs are impacted by sentence processing, it appears as though the effects can be dealt with like any other source of noise in an experimental setting.

Claim 5: Acceptability judgments reveal only set membership (sect. 1.2, para. 7). Claim 5 is confusing. It is false in the sense that, if one is interested in *set membership*, this property still needs to be *inferred* from acceptability data, using a logic that maps that data type back to the relevant cognitive computations. In this, AJs are like any other data type in cognitive science: No data types, including priming, directly reveal the underlying computations of the human brain, and all data types require a linking hypothesis between the observable data and the unobservable cognitive process.

Claim 6: Acceptability judgments have yielded no consensus theory among linguists (sect. 1.2, para. 9). Claim 6 is a strange criticism to make of any data type, especially AJs. First, the beliefs of scientists are a subjective issue based on how they weigh different kinds of evidence. Second, AJs are, by all accounts, a robust and replicable data type. Whatever disagreements there are in linguistics literature, they appear to obtain mostly at the level of *interpreting*, not *establishing*, the data (e.g., Phillips 2009).

In conclusion, we support B&P's desire to bring new evidence to bear on questions about linguistic representation. We caution, however, that advocacy for one method should not be bolstered by misleading comparisons, especially with methods such as AJs, which yield data that are demonstrably robust, highly replicable, and comparatively convenient and inexpensive to collect.

Priming is swell, but it's far from simple

doi:10.1017/S0140525X17000607, e312

Jayden Ziegler,^a Jesse Snedeker,^a and Eva Wittenberg^b

^aDepartment of Psychology, Harvard University, Cambridge, MA 02138;

^bDepartment of Linguistics, University of California, La Jolla, CA 92093.

ziegler@g.harvard.edu <http://www.jaydenziegler.com>

snedeker@wjh.harvard.edu <https://software.rc.fas.harvard.edu/lds/research/snedeker/jesse-snedeker/>

ewittenberg@ucsd.edu

<http://www.evawittenberg.com/i/start.html>

Abstract: Clearly, structural priming is a valuable tool for probing linguistic representation. But we don't think that the existing results provide strong support for Branigan & Pickering's (B&P's) model, largely because the priming effects are more confusing and diverse than their theory would suggest. Fortunately, there are a number of other experimental tools available, and linguists are increasingly making use of them.

Branigan & Pickering (B&P) tell a straightforward tale. Linguists rely on grammaticality judgments to uncover representations.

Judgments have limitations, but no other psycholinguistic methods systematically reveal linguistic structure. Fortunately, priming offers a direct window onto representation, providing evidence for two distinct levels: a surface syntactic form, independent of meaning and void of lexical content, and a semantic form that includes information about thematic roles, quantifier scope, and information structure.

We are fond of priming ourselves, but this elegant story is misleading in several ways. First, the priming literature does not strongly support the theory that B&P propose. As they dive deeper, the loose ends and contradictions emerge, but their final conclusions bypass this complexity. If we rearrange the evidence a bit, the theoretical ambiguity becomes clearer.

The primary evidence for syntactic representations comes from studies of argument alternations (dative or active-passive) that perfectly confound *surface syntax* with *thematic mappings*. B&P note that a few foundational studies demonstrated that syntax *can* be primed independent of thematic mappings (sect. 2.1). Thus, they privilege syntax in their theory. But there is now an equally robust literature demonstrating that thematic mappings can be primed independent of syntax (e.g., Cai et al. 2012; Chang et al. 2003; Cho-Reyes et al. 2016; Hare & Goldberg 1999; Salamoura & Williams 2007; Ziegler & Snedeker 2016b). B&P acknowledge this work (sect. 2.4) but treat it as a secondary, interface phenomenon: Thematic information remains separate from syntax (Fig. 1).

Similarly, the observation that priming *can* occur in the absence of lexical overlap motivates a theory in which the syntactic skeleton is separate from the lexical content. To account for the lexical boost, B&P must complicate their story, by linking lemmas to structures (sect. 2.3). But perhaps we should revisit the claim that the syntactic structure lacks lexical nodes. Indeed, function words *can* be a locus of priming (Bencini et al. 2002; Ferreira 2003). We know that only partial overlap in the syntactic skeleton is needed for structural priming (sect. 2.1), but we don't assume that the unnecessary pieces are removed from the syntactic representation. Lexical content may be similar: always present and sometimes contributing to priming via overlap.

The evidence for their semantic level is also sparse. We know: (1) Quantifier scope can be primed, (2) this priming is isolated to the particular quantifier used (e.g., *each* does not prime *every*), and (3) it abstracts away from the nouns and verbs in a sentence (Feiman & Snedeker 2016; Raffray & Pickering 2010). However, B&P's claim that scopal priming is bound to thematic roles and cannot be captured by an LF representation is controversial (Chemla & Bott 2015). It rests on a single null result with prime stimuli (*A boy climbed every tree*) that have not been shown to produce priming when thematic roles are the same. Furthermore, the manipulation used confounds verb-specific roles, thematic roles, and the notion of deep subject/object. It's just too early to conclude that scope and thematic roles are tightly coupled, or that LF isn't the locus of scopal priming.

It seems that, under the right conditions, almost any linguistic representation, mapping, or process can be primed. Consequently, evidence for priming is always interpretable to some degree (it demonstrates a commonality between prime and target). But the absence (or magnitude) of an effect is often less constraining, because there is so much variability across tasks and stimuli. In some comprehension tasks, there is no priming in the absence of verb overlap (Arai et al. 2007), while in others, abstract priming is robust (Tothathiri & Snedeker 2008a; 2008b). This problem isn't unique to comprehension. The pattern of effects in production can depend on how the sentences are elicited (stem completion vs. full sentence generation; Ziegler & Snedeker 2016a).

Understanding this instability is critical; we suspect that the answer lies in thinking through the processes involved in each task and how they engage both stored representations and representations that are constructed on the fly. To do this, we will have to move beyond the notion of priming as a static, atemporal

phenomenon that targets stable representations independent of the process of production (or comprehension).

While priming is not the transparent window that B&P promise, psycholinguists do have a much wider range of tools than the paper suggests. Some are behavioral. Novel word generalization studies going back to the 1950s have revealed structural regularities in the representation of linguistic form and meaning (Berko 1958; Fisher 2002; Gropen et al. 1989; Naigles 1990; Prasada & Pinker 1993). Similarly, artificial language learning sheds light on the representations that learners extract from linguistic data and use to guide generalization (Pothos 2007; Reber 1967; Saffran et al. 2008; for review, see Erickson & Thiessen 2015).

New methods for analyzing imaging data also provide greater constraint on representational theories. For example, multi-voxel pattern analysis, a class of machine-learning algorithms that examine patterns of neural activity (Haxby et al. 2001), has revealed regions of the left temporal cortex that appear to bind arguments to something roughly like thematic roles (Frankland & Greene 2015). Structural priming is therefore only one useful tool of many.

Finally, we think that B&P are too pessimistic in their assessment of linguists and their tools. There are longstanding traditions of experimental work in phonetics, phonology, and language acquisition. In recent years, experimental work has also become common in syntax, semantics, and pragmatics (e.g., Arunachalam 2013; Cowart 1997; Myers 2009; Sorace & Keller 2005). Indeed, the question of how armchair judgments translate into generalizable conclusions has received considerable attention (Sprouse & Almeida 2012; Sprouse et al. 2013). From our perspective, the remaining disputes do not reflect an over-reliance on grammaticality judgments or a dearth of appropriate methodologies; they stem from: (1) the close parallels between the theories that are still standing (similar operations assigned to different theoretical levels), (2) the lack of falsifiability for the contrasting features, and (3) the danger we all face of letting our “affection for [our] intellectual child[ren]” (Chamberlin 1897) guide our interpretation of the data.

Authors' Response

Structural priming and the representation of language

doi:10.1017/S0140525X17001212, e313

Holly P. Branigan and Martin J. Pickering

Department of Psychology, University of Edinburgh, Edinburgh, EH8 9JZ, United Kingdom.

holly.branigan@ed.ac.uk <http://www.ed.ac.uk/profile/holly-branigan>

martin.pickering@ed.ac.uk

<http://www.ed.ac.uk/profile/martin-pickering>

Abstract: Structural priming offers a powerful method for experimentally investigating the mental representation of linguistic structure. We clarify the nature of our proposal, justify the versatility of priming, consider alternative approaches, and discuss how our specific account can be extended to new questions as part of an interdisciplinary programme integrating linguistics and psychology as part of the cognitive sciences of language.

In our target article, we argue for an experimental approach to linguistic representation and demonstrate how structural priming can be used to develop a psychologically motivated account of

how people represent linguistic structure. If one utterance primes another, then we argue that they share structure. Patterns of priming are therefore informative about the way in which people represent language. While acknowledging the continuing value and importance of acceptability judgements in developing representational hypotheses, we proposed that priming is in many ways superior: It is an implicit behavioral measure that taps representation. It also allows us to study representation in all groups, including those such as young children who cannot make acceptability judgments.

In recent years, many experimental studies using structural priming have helped us understand language users' representations. These studies suggest that semantic information is represented separately from syntactic information. The single semantic level encodes information about scope relations, information structure, and thematic structure. The single syntactic level, which draws on well-formedness constraints concerning local linear and hierarchical relations, includes syntactic category information and some missing elements (i.e., elements that are not uttered) but does not involve syntactic movement. At present, linguists propose incompatible theories, and there is no sign that the nearly exclusive use of acceptability judgments is ever going to determine which type of theory is correct. In contrast, the findings from priming are compatible with some theories and incompatible with others.

The commentators take a unified view that researchers from different disciplines should cooperate in investigating linguistic representation, as a single integrated programme of research. They unanimously agree that an experimental approach to linguistic representation is valuable and largely accept that structural priming is informative in this regard. In other words, they do not feel that it is relevant only to the study of how language is produced and comprehended. They differ, however, about the versatility of priming, how directly it taps into representation, and its advantages over other methods (in particular, acceptability judgments). They also take different views about our conclusions regarding aspects of linguistic representation and make specific proposals for further research.

We are pleased that our proposal has met with such interest and hope that it will encourage a future programme of cooperative interdisciplinary research on linguistic representation. In what follows, we have grouped our responses to their insightful comments under 11 headings.

R1. How can linguistic representation be investigated?

We begin by clarifying the nature of our proposal. Although commentators agree about the importance of experimental methods for investigating linguistic representation, some of them seem to assume that we advocate entirely renouncing the use of acceptability judgments. But as we make clear throughout the article, our argument is that researchers interested in linguistic representation should not be restricted to using *only* acceptability judgments. Specifically, we contend that “the representations underlying language use need not and, in fact, should not be investigated *only* via [acceptability] judgments” (sect. 1, para. 2). We further argue that “[our] goal is to consider alternative (experimental) methods to acceptability judgments that potentially address the linguistic representations implicated in language processing” (sect. 1.1, para. 5), that “acceptability judgments are not *enough*,” (sect. 1.2) and that “researchers concerned with linguistic representations should not rely *solely* on such judgments, and should call on additional methodologies” (sect. 1.2, para. 10). Moreover, we argue that “acceptability judgments can be used (with appropriate controls) alongside structural priming (and perhaps other experimental methods; see sect. 1.3) as a means