Individual differences in syntactic processing: Is there evidence for reader-text interactions?

Ariel N. James[a,*], Scott H. Fraundorf[b], Eun-Kyung Lee[c], Duane G. Watson[d]

[a] University of Illinois at Urbana-Champaign, Department of Psychology, 603 East Daniel Street, Champaign, IL 61820, United States
[b] University of Pittsburgh, Department of Psychology and Learning Research and Development Center, 608 Learning Research and Development Center, 3939 O'Hara Street, Pittsburgh, PA 15260, United States
[c] Yonsei University, Department of English Language and Literature, 50 Yonsei-ro, Seodaemun-gu, Seoul 03722, South Korea
[d] Vanderbilt University, Department of Psychology and Human Development, Peabody College #552, 230 Appleton Place, Nashville, TN 37203-5721, United States

ARTICLE INFO

ABSTRACT

There remains little consensus about whether there exist meaningful individual differences in syntactic processing and, if so, what explains them. We argue that this partially reflects the fact that few psycholinguistic studies of individual differences include multiple constructs, multiple measures per construct, or tests for reliable measures. Here, we replicated three major syntactic phenomena in the psycholinguistic literature: use of verb distributional statistics, difficulty of object- versus subject-extracted relative clauses, and resolution of relative clause attachment ambiguities. We examine whether any individual differences in these phenomena could be predicted by language experience or general cognitive abilities (phonological ability, verbal working memory capacity, inhibitory control, and perceptual speed). We find correlations between individual differences and offline, but not online, syntactic phenomena. Condition effects on online reading time were not consistent *within individuals,* limiting their ability to correlate with other measures. We suggest that this might explain controversy over individual differences in language processing.

In Lee Cronbach's famous presidential address to the American Psychological Association Annual Convention in 1957, he described an optimistic vision of the future of psychology in which the best of the correlational and experimental traditions joined forces as the *united discipline.* A complete theory of human behavior, he argued, requires the modeling of individual variability along with the prediction of an individual's response to varying conditions. The usefulness of such a united approach is especially clear in the domains of applied psychology: It would be best to provide an intervention that is tuned to the particular needs of each individual (Pellegrino, Baxter, & Glaser, 1999). Since that 1957 address, psychologists have taken up the challenge of the united discipline. In their 1999 review, Pellegrino, Baxter, and Glaser chart the progress of the field, focusing on the intersections of cognitive psychology and psychometrics that follow directly from Cronbach's initial interests, focusing first on "aptitude-treatment interactions", or the relationship between a student's intellectual abilities and expertise on one hand, and educational materials and instructional methods on the other.

A specific case of this type of investigation is what we will call "reader-text interactions". Substantial prior work has revealed that the time required to read a sentence or text is a function of both the individual reader and the text being read: Researchers in individual differences and educational psychology have identified important sources of variation in reading and comprehension skill (e.g., Kuperman & Van Dyke, 2011; Perfetti & Hart, 2002), and work in cognitive psychology and psycholinguistics has identified the types of words, sentences, and texts that are more difficult for comprehenders (e.g., some syntactic structures are more difficult to process; Gibson, 1998; Just & Carpenter, 1992; Waters & Caplan, 1996). What is less clear is whether and how reader and text characteristics *interact*: Are difficult sentences equally challenging for all readers? And, conversely, does variation in reading skill affect the comprehension of *all* linguistic materials, or just especially difficult ones?

Here, we investigate reader-text interactions in the domain of syntactic processing. We first consider whether there is evidence for such interactions—that is, are some types of sentences consistently more difficult for specific types of readers? Then, to the extent that we observe such differences, we consider what might explain them.

The relevance of reader-text interactions

The potential for reader-text interactions in syntactic processing is

* Corresponding author at: Macalester College, 1600 Grand Avenue, St. Paul, MN 55105, United States.
  E-mail address: ajames2@macalester.edu (A.N. James).

relevant to several broader issues in psychology. First, if variability in readers interacts with properties of texts, that can provide insights into the underlying mechanisms of language processing (for a discussion of how individual differences contribute to more general theoretical development, see Vogel & Awh, 2008). As we review below, theories of language processing make different claims about why some texts are more difficult to process. Consequently, they also imply different hypotheses about which individual differences are likely to modulate syntactic processing. For instance, theories that attribute the difficulty of some syntactic structures to comprehenders' relative inexperience with them predict that individual differences in language experience might drive differences in syntactic processing. By contrast, in accounts in which some syntactic structures are difficult because of the demands they place on memory, it is individual differences in memory capabilities that are most likely to relate to individual differences in syntactic processing.

Second, whether and which individual differences exist in syntactic processing speak to broader, fundamental questions about the architecture of the mind and language processing system. For example, as we review in greater detail below, some theories (e.g., Waters & Caplan, 2003) propose that language processing is divided into initial, automatic stages and later, interpretive stages, with only the latter subject to individual differences in working memory and other cognitive abilities. Studying individual differences in both online and offline processing allows us to test this theoretical claim. Similarly, another central question in psychology is the extent to which cognitive systems are modular rather than driven by domain-general systems (see Fodor, 1983). By understanding whether variability in the capacity of domain general systems like working memory and executive function is associated with syntactic processing, we can better understand the overall architecture of the mind: To what degree is language (and other motor and perceptual systems) modular, and to what degree does it recruit domain-general systems? It also provides an opportunity to understand *why* characteristics like high working memory are associated with positive outcomes in more complex domains like reading comprehension.

Finally, and most broadly, reader-text interactions exemplify one of the central questions of the united discipline envisioned by Cronbach: How do the skills and abilities identified by psychometricians intersect with the cognitive-processing effects discovered by experimentalists? Aptitude-treatment interactions have been reported in some educational domains. For instance, learners with greater prior knowledge learn better from different types of texts (McNamara, Kintsch, Songer, & Kintsch, 1996) and feedback (Hausmann et al., 2013) than do low-knowledge learners. Indeed, several reader-text interactions have been reported within the language processing literature. For instance, slower overall readers show larger effects of word frequency (Seidenberg, 1985), and readers with greater linguistic experience may be less sensitive to word difficulty and correspondingly more sensitive to discourse-level factors (e.g., the introduction of new concepts; Stine-Morrow, Soederberg Miller, Gagne, & Hertzog, 2008). Most relevant for the present paper, readers with greater linguistic knowledge are also more efficient at resolving syntactic ambiguity (Traxler & Tooley, 2007). On the other hand, a review of the learning-styles hypothesis—that certain learners do best under one instructional method and other learners do best with a different method—has found little evidence to date in favor of such an interaction; instead, the most well-established mnemonic effects appear to apply across learners (Pashler, McDaniel, Rohrer, & Bjork, 2008). Thus, there is a need to investigate in other domains whether the cognitive-processing effects discovered by experimentalists are consistent across individuals, and whether the important skills and abilities identified in psychometrics apply across tasks and materials.

Assessing reader-text interactions

Given the applied and theoretical relevance of potential reader-text

interactions, it is not surprising that they have been studied by both educational psychologists and cognitive psychologists. While educational psychologists have investigated reader-text interactions with the goal of promoting learning in young readers (e.g. Coté, Goldman, & Saul, 1998) and comprehension among students (e.g., McNamara et al., 1996), a complementary literature grew in cognitive psychology as theories of reading began to include ideas about individual differences in cognitive abilities. An influential example is Just and Carpenter (1992), who proposed, and reviewed evidence, that differences in capacity between individuals correlate with differences in reading ability. Since then, psycholinguists have employed individual differences to promote both memory-capacity theories of language comprehension (e.g., Fedorenko, Gibson, & Rohde, 2006, 2007; Gibson, 1998, 2000), competing experience-based theories (MacDonald & Christiansen, 2002, discussed in greater detail below), and a number of other explanations that combine language-specific and domain-general mechanisms (e.g., Engelhardt, Nigg, & Ferreira, 2017; Farmer, Fine, Misyak, & Christiansen, 2017; Novick, Trueswell, & Thompson-Schill, 2010; Payne et al., 2014; Swets, Desmet, Hambrick, & Ferreira, 2007; Van Dyke, Johns, & Kukona, 2014).

As the individual differences approach in psycholinguistics has continued to grow in popularity in recent years, it is important to take a step back and assess its progress toward the *united discipline*. These psycholinguistic investigations are nested within the experimental approach, investigating language-processing effects that have been previously shown across subjects using controlled linguistic stimuli. So, the question is whether these investigations live up to the ideals of the *correlational* approach. Here, we describe several methodological demands identified by the correlational approach and discuss how these constraints may have contributed to a lack of consensus regarding individual differences in syntactic processing.

First, a critical insight from measurement theory is that two variables can be observed to correlate only to the degree that there is meaningful variation in those individual variables and to the degree that such variation is reliably measured (Spearman, 1904). If there are genuine, stable individual differences in syntactic processing, those individuals who show large syntactic-processing effects on one subset of items should also show large effects on another, similar subset. By contrast, a failure to observe such correlations would suggest that either (a) there are not consistent individual differences in syntactic processing or (b) such differences exist, but our methods cannot reliably detect them.

For instance, consider a scenario in which all readers read a syntactically complex sentence 300 ms more slowly than a syntactically simple sentence. In this case, there is clearly a *text* effect—one sentence is more difficult than another—but there is no reader-text *interaction* because *all* readers found the complex sentence more difficult than the simple sentence to the same degree. In this scenario, it would be impossible for any other construct (such as verbal working memory) to explain individual differences in syntactic processing because such variation was not observed to begin with.

Unfortunately, while past investigations of individual differences in syntactic processing have sometimes used measures of working memory and other cognitive abilities that have been normed for their reliability, researchers have only rarely assessed whether we observe meaningful variation across individuals in the syntactic processing effects themselves (but see Swets et al., 2007 for one application of psychometric principles to syntactic processing). Thus, before we ask *why* individuals might differ in syntactic processing, it is first necessary to establish that such individual differences *exist* at all. If we cannot observe consistent individual differences in syntactic processing to begin with, differences in online syntactic processing cannot be expected to relate to any other measure.

Second, individual differences are best assessed with multiple measures. "Perhaps the most valuable trading of goods the correlator can offer," Cronbach (1957) states, "…is his multivariate conception of