

Natural Language Analysis and the Psychology of Verbal Behavior: The Past, Present, and Future States of the Field

Journal of Language and Social Psychology
2021, Vol. 40(1) 21–41
© The Author(s) 2020



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0261927X20967028
journals.sagepub.com/home/jls



Ryan L. Boyd¹  and H. Andrew Schwartz²

Abstract

Throughout history, scholars and laypeople alike have believed that our words contain subtle clues about what we are like as people, psychologically speaking. However, the ways in which language has been used to infer psychological processes has seen dramatic shifts over time and, with modern computational technologies and digital data sources, we are on the verge of a massive revolution in language analysis research. In this article, we discuss the past and current states of research at the intersection of language analysis and psychology, summarizing the central successes and shortcomings of psychological text analysis to date. We additionally outline and discuss a critical need for language analysis practitioners in the social sciences to expand their view of verbal behavior. Lastly, we discuss the trajectory of interdisciplinary research on language and the challenges of integrating analysis methods across paradigms, recommending promising future directions for the field along the way.

Keywords

language analysis, computational social science, natural language processing, attention

Human language is truly fascinating. When we speak, our words provide structure to an otherwise intangible relationship between the self and the world. When we write, we give shape to the very substance of our innermost, subjective experiences. Language

¹Lancaster University, UK

²Stony Brook University, NY, USA

Corresponding Author:

Ryan L. Boyd, Department of Psychology, Lancaster University, D04, Fylde College, Lancaster LA1 4YF, UK.

Email: r.boyd@lancaster.ac.uk

is quintessential human potential: it has relatively fixed rules and forms, yet we can string together a combination of words that are completely unique, formulating ideas that have never been shared before—not ever—in the entire history of humankind. Even the most dull, trivial text becomes interesting when approached as a scientific question: what motivated the creation of such a boring document?

Language is at once obvious and perplexing. We can write epic novels, sing songs, and communicate intimate thoughts and feelings through language, all without explicit awareness of the psycholinguistic machinery at play. The seemingly paradoxical nature of language has made it a recurring topic of psychological study, inspiring countless theories on verbal behavior that range from the empirical to the fantastical. Theories that connect language to our thoughts, feelings, and behaviors have largely co-evolved with *psychological language analysis* methods—the practice of extracting psychological insights from a person’s verbal behavioral data.¹ Over the past century, language analysis methods have grown increasingly objective, systematic, and empirical, allowing for the deployment of sophisticated text analytic frameworks for psychological research in the digital age.

Today, however, language analysis theory rests at something of a standstill, despite the unprecedented availability of natural language datasets to scholars and language analysis practitioners (e.g., social media posts, SMS, e-mail, etc.). With the advent of new natural language processing frameworks, increasingly complex and nuanced approaches to psychological language analysis are appearing rapidly, and we believe that the field is due to expand both its text analysis toolkit as well as its theories about *how* human psychology is reflected in language. In this article, we provide a broad discussion of the past, present, and possible future of language analysis research in the social sciences. Specifically, we discuss:

- (1) How psychological language analysis arrived at its current state;
- (2) Today’s prevailing theory of language analysis and its shortcomings;
- (3) The immediate horizon of language analysis theory;
- (4) Key developments in this new age of personalized language data and the necessity for interdisciplinary collaboration.

The Secrets of Verbal Behavior: From the Mystical to the Statistical

In the early days of psychology, language was ascribed a vague, almost supernatural power to reveal the bubbling cauldron of the preconscious (Freud, 1915; Laffal, 1964). As the field increasingly turned to scientific methods, however, researchers grew skeptical of psychodynamic theories that often made murky or circular claims about the relationship between language and the rest of mental life. Allport (1942), for example, aired significant concerns about the need for more rigorous theory-testing in language analysis studies, rejecting much of his predecessors’ musings on language as unscientific and methodologically flawed. Allport saw massive potential in

verbal behavior as a source of psychological data, however, he also recognized the sheer difficulty of unifying diverse approaches to language analysis frameworks under a single roof.

In the 1960's, the so-called General Inquirer showed incredible promise as one of the first serious efforts toward a generalized and universal approach to psychological language analysis. The General Inquirer was, in many ways, a reflection of the widespread technological optimism of early Space Age social sciences (e.g., Kranz, 1970; Stone et al., 1962): a massive computer with software that could “read” texts and deduce the inner workings of the human mind was, quite literally, the stuff of science fiction (e.g., Clarke & Kubrick, 1968). By having a computer score texts for psychosocial phenomena, it was hoped that the General Inquirer would remove subjectivity, bias, and human error from the measurement process. Put another way, the General Inquirer aspired to become a standardized, scientific tool: something of a language analysis analogue to the microscope or the Erlenmeyer flask.

Early demonstrations showcased the system's sheer potential: research ranging from the scientific study of suicide notes to political documents to folktales highlighted the range of social scientific topics that could be studied computationally (see, e.g., Colby et al., 1963; Holsti, 1965). Ultimately, however, the General Inquirer system never really caught on: it was complicated, difficult to maintain, and costly beyond most academics' means (Psathas, 1969). Despite years of tweaking and adjusting the underlying dictionaries, researchers struggled to get satisfactory results that aligned with the “gold standard” of hand-scored texts (Smith, 1968). The gradual decline of the General Inquirer allowed language analysis methods of questionable validity to remain prominent. Other approaches, however, absorbed important lessons from the General Inquirer, and began favoring greater statistical analyses of words over subjective readings (see, for example, Martindale & Martindale, 1988; Thomas & Duszynski, 1985; Wilson, 2006). The General Inquirer marked a slow turning of the tide that would bring with it a newer, more scientific rationale for language analysis as a psychological research method.

LIWC: Little Words, Big Secrets

Historically, most language analysis research attempted to infer a person's psychology by mapping out the *content* of what they said: for example, charting a client's trajectory through treatment by counting how often the concepts of “hope” or “social alienation” are mentioned (Gottschalk & Gleser, 1969). By skimming over the small, “meaningless” parts of language—words like “the” and “than” and “not”—it was believed that the psychological significance of a text could be discovered through counting how often different *meanings* were conveyed (e.g., How often does the text talk about *gaining* versus *losing* power? Does the phrasing imply an action occurred *actively* or *passively*? Is the author conveying *certainty* or *caution*?).

Following the rise and fall of the General Inquirer, the idea that language may reveal hidden, esoteric clues about human psychology began taking on new forms.

By the late 20th century, researchers were learning that the much-ignored “particles” of language could provide remarkably powerful insights into a multitude of psychosocial phenomena. The small, mundane parts of speech that language analysts had largely ignored—things like pronouns (I, you, we), articles (the, a, an), and negations (no, not, never)—were, in fact, reliable indicators of a *how* a person thinks, regardless of *what* they were talking about (Pennebaker, 2011; Weintraub, 1989). Some of the most important verbal behavioral markers of a person’s psychology, it seemed, were hidden in plain sight.

At the same time, some of the first truly user-friendly text analysis software—in particular, Linguistic Inquiry and Word Count, or LIWC (Pennebaker & Francis, 1999)—made the prospect of objective, automated, and transparent psychological text analysis a reality. The rise of LIWC coincided with the dawn of widespread personal computing and, for the first time in scientific history, the typical social scientist could now perform serious, objective, automated computations on language data. LIWC was nothing short of revolutionary for psychologists, and the widespread enthusiasm for language analysis was reborn.

The new wave of language analysis research was founded on a simple principle—the same principle, in fact, that had guided the General Inquirer: *word frequencies represent attentional habits* (Stone et al., 1966; Tausczik & Pennebaker, 2010). People who use high rates of articles (the, a, an) and prepositions (next, above), for example, tend to focus on formal or concrete concepts and their inter-relations (Pennebaker et al., 2014). People with higher social status and confidence are more focused on the external social environment than themselves, using more “you” and “royal we” words than “I” words (Kacewicz et al., 2014).

LIWC’s rapid rise to prominence helped to cement its underlying logic as the dominant rationale in the language analysis world. Indeed, much of our current understanding about the relationship between verbal behavior and psychology has been scaffolded, either directly or indirectly, by the relatively simplistic method of LIWC that is “word counting.” At its core, the word counting approach is one of scanning texts and tallying up the frequency of words from predefined categories—categories that are informed by psychological theory: emotion words (happy, upset, angry), agentic words (do, able, try), thinking words (think, understand, guess), self-referential words (I, me, my), and so on. The relative frequencies of each category, then, are typically interpreted as reflecting a person’s relative focus on each domain—something of an eye tracker, but for a person’s words.²

From the word counting perspective, then, the “secrets” that could be extracted from a person’s language were fundamentally different from those of classical psychoanalysis. Now, the mysteries contained within people’s words were effable and calculable, but more subtle and statistical: too small to see with the naked eye, but discernible using computerized analyses of large numbers of documents. By using words to measure microscopic differences in what people pay attention to, we now had a direct pipeline into their minds and, by extension, their personalities (Boyd & Pennebaker, 2017), life experiences (Chung & Pennebaker, 2018), cultures (Michel et al., 2011), and societies (Iliev et al., 2016).

The Attentional View of Verbal Behavior

The early success and widespread popularity of LIWC helped to cement the word counting method and its basic psychometric logic—that words correspond to attentional focus—as the norm. Today, the majority of language analysis theory assumes that word frequencies reflect the relative amount of attention a person pays to a given domain, even if not stated explicitly. Pronouns are not pronouns: they are markers of who we pay attention to. “Emotion” words are not some reflection of the *experience* of emotions themselves: they are merely diagnostic of one’s *attention* to affective states. The “words as attention” approach to language analysis has been a simplistic, albeit necessary, foothold for helping the field get as far as it has.

The Merits of “Words as Attention”

For decades, the basic mantra of “words reflect attention, and nothing more” has been an intuitive, straight-forward, and reliable framework for interpreting and contextualizing language analysis research findings. The principle that verbal behavior chiefly reflects attention has been the anvil against which modern language analysis has been forged and validated.³ Research on personality, for example, tells us that social attention is the core, defining feature of extraversion (Ashton et al., 2002). If words reflect attention, then we might reasonably expect extraverts to use more “social” words—this prediction turns out to be true (Mairesse et al., 2007).

Under the “words as attention” model, language analysis research has been able to stretch across an incredible number of topics. Over the years, a strong nomological network—one that describes the correlations between various psychological processes and attentional patterns—has been woven from the fabric of “words as attention.” The words that people use when describing their homes tell us whether they are focused on the relaxing or stressful properties of their abodes (Saxbe & Repetti, 2010). Computerized language analysis reveals that psychologists who focus on their social connections live longer (Pressman & Cohen, 2007), and that students who focus (perhaps a bit too much) on eating, drinking, and sex tend to fare worse in university (Robinson et al., 2013). Word frequencies tell us how much attention two people are paying to each other (Ireland & Pennebaker, 2010) and one’s attention to various facets of the self (Rodríguez-Arauz et al., 2017). Nearly everywhere that we find psychological language analysis today, we find the implicit assumption that verbal behavior reflects attentional processes above all else.

One of the greatest strengths of the “words as attention” interpretation has been its extensibility into otherwise difficult areas of study. Once it has been established that “cognitive” words are related to things like problem-solving and stress, for example (see, e.g., D’Andrea et al., 2012), we can indirectly estimate a person’s stress levels simply by counting their words. This basic “validation–extension” research cycle has expanded language analysis research across virtually every domain of scientific research, including public health (Eichstaedt et al., 2015), politics (Jordan et al., 2018), neuropsychology (Lumma et al., 2018), affect (Garcia & Rimé, 2019), personality

(Schwartz et al., 2013), and genetics (Mehl et al., 2017). Even in the absence of other psychological measures, the existing nomological network for language analysis facilitates meaningful theory-testing, ranging from the study of daylight cycles on mood (Golder & Macy, 2011) to personality-driven business outcomes (Block et al., 2019).

The Pitfalls of “Words as Attention”

There is no question that the “words as attention” view of verbal behavior has been invaluable for advancing the science of psychological language analysis. Nevertheless, this attentional view has gradually become something of a “rocking chair theory” for language analysis scholars: like a rocking chair, “words as attention” keeps us moving, but without actually taking us very far. Put another way: “words as attention” permits the use of word frequencies as estimates of other, more interesting psychological phenomena, but it does little to further our understanding of *how* verbal behavior actually reflects psychological processes.

There are several growing reasons for why the field must expand beyond “words as attention” and, here, we highlight two specifically. First, “words as attention” gives us limited material with which to meaningfully expand our nomological network which, in turn, constrains the very types of empirical questions that can be probed using word counting approaches. Second, at its very core, “words as attention” is too simple a model of language, one that was never designed (or intended) to accommodate the full complexity and functionality of verbal behavior.

Overwhelmingly Descriptive Research

Under the “words as attention” model, the contribution of natural language analysis to research findings is generally interpreted through the pinhole of “paying attention to X is correlated with Y .” This correlational approach to language analysis contributes most obviously to nomological research, yet constricts research findings to conclusions like “extraverts focus on people and positive emotions” and “depressed people are self-focused relative to non-depressed people.” Psychologists working with language data have bent over backwards to demonstrate the validity of language-based psychological inferences, gaining mainstream credibility in the process. The majority of validation efforts have been geared toward demonstrating that language-based measures of psychological process show convergent validity with other conceptualizations of the same measures (e.g., self-reports, behaviors, etc.) in domains like well-being (Vine et al., 2020), social status (Kacewicz et al., 2014), values (Bardi et al., 2008), and so on. However, the “over-validation” of word counting has resulted in something of a language analysis “crud factor”—an overly complicated and meaningless network of intercorrelations with little-to-no theoretical value (Meehl, 1990).

When we unquestioningly follow the “words as attention” model of language analysis, we have stopped trying to understand *how* verbal behavior reflects a person’s psychology. Perhaps the most pressing issue with today’s overly-descriptive language analysis research is that verbal behavior is seldom used to unpack the very things that

are purportedly being measured. Much like the perennial debate in personality psychology, one could argue that much—perhaps most—of today’s language analysis research is purely *descriptive* of psychological phenomena, but not *explanatory* of those same phenomena (see, e.g., McCrae & Costa, 1995). Rather than using natural language analysis to tease apart the antecedents, facets, and consequences of anger, for example, most studies simply use verbal behavior as a presumed measure of anger itself, or one’s attention to anger, or something very similar.

Fortunately, there is considerable room to expand and enrich our understanding of how verbal behavior reveals psychological processes. First, we recommend continued efforts toward experimentation at every stage of research, from study formulation to design to execution to analytic decisions. Today, with the widespread availability of *in situ* language data, language analysis research has become increasingly correlational, often by necessity. Data from Twitter, for example, does not come pre-packaged with Big 5 scores or fMRI data: psychometric specificity is often the price we pay for cheap, readily accessible data at an unimaginable scale. While valuable in its own right, we encourage supplementing psychometrically “thin” data with theory-testing through subsequent experimental work and additional, carefully selected follow-up samples (see Boyd et al., 2020). We also recommend avoiding “language luggage” studies: data collection schemes where language is collected without express purpose, but instead under the rationale that it can just be “unpacked” later. Such study designs ultimately serve to create verbal behavioral data that seldom reflects constructs of interest: it is difficult to simply “create” a linguistic measure of empathy, for example, from language samples generated in settings where empathy is unlikely to be a relevant construct (e.g., students reading a PowerPoint presentation for a class project are unlikely to suddenly become overwhelmed with feelings of a shared emotional experience with their audience).

Second, we suggest the adoption of a broader view of verbal behavioral data: one that permits words to tell us about more than attentional focus, and one that uses language to better understand *how* psychological processes operate. In the study of anger or aggression, for example, a typical LIWC-driven text analysis might involve counting up “anger” words to predict future aggressive behaviors. An alternative approach, however, might include an experiment wherein the links between affect and cognition more deeply probed and tested. For example, the emotional state of “anger” has several hypothesized appraisal tendencies associated with it: externalization, control, and certainty (Lerner & Keltner, 2000). An experimental condition that induces anger, then, might explore verbal behavior to determine the degree to which a person uses more “you” and “they” words (reflecting greater externalization), more agentic words (reflecting greater control), and more absolutist words (reflecting greater certainty). In this way, the deeper meaning, purpose, and utility of language is preserved and informs broader psychological theory rather than being reduced to “angry people use more ‘anger words.’” By decoupling a literal, one-dimensional interpretation of the relationship between words and anger, we return to a more intuitive place for language analysis—one that capitalizes upon (rather than discards) its complexity and nuance.

Isn't Language Supposed to be. . . You Know. . . Social?

Ironically, treating words as attention has biased psychological language analysis—particularly computerized language analyses—away from the inherently social components of verbal behavior.⁴ If we treat a person's language as the verbal equivalent of an eye tracker, we overlook the fundamental essence of language itself: to transmit and coordinate complex information (broadly defined) with others. Scholars who are unfamiliar with the above-discussed historical rationale, traditions, and context of today's automated language analysis methods often find themselves deeply bothered over the fact that word counting is so ubiquitous, yet does not seem to reflect even the most basic functions of language. It is widely rumored that, every night, one unlucky discourse analyst is tormented by nightmares of being chased down an endless hallway by an axe-wielding, anthropomorphic copy of LIWC.

The social aspects of verbal behavior have traditionally been downplayed in computerized language analysis primarily due to the sheer difficulty of collecting and analyzing language data from two or more people; until very recently, capturing the multitude of variables at play during social interactions was an absolutely colossal task. Often, researchers would “split the difference,” so to speak, by analyzing the language of mass media such as newspapers, speeches by political leaders, or other sociocultural artifacts as a proxy for group psychology (e.g., Hogenraad, 2003; Robinson et al., 2017). However, these studies tell us little about good, old fashioned *social interactions*: the type of exchange one might have with a coworker or new acquaintance at a dinner party. Even in the context of studying rich online spaces designed around social interaction (e.g., social networks, internet forums), most language analysis studies profile each group member separately, treating them largely as agents operating within a complete social vacuum.

Promisingly, the past decade has seen a steadily growing number of studies that directly explore language as complex, bidirectional social processes. Language analysis techniques such as Language Style Matching (LSM) and Latent Semantic Similarity (LSS) are currently thought to give us a “broad strokes” picture of interpersonal attention, for instance (e.g., Ta et al., 2017). While still primarily rooted in the idea of “words as attention,” it stands to reason that the more we know about *when* and *how much* two people are paying attention to each other, the better we can understand how social processes operate (see, e.g., Bayram & Ta, 2019; Doorn & Müller-Frommeyer, 2020; Heuer et al., 2020). After all, social interaction implies interpersonal attention, and the degree to which people are focused on each other is the foundation upon which social interactions are built. Nevertheless, the attentional processes captured by LSM (see Gasiorek et al., 2021) and LSS are but a small piece of a much, much larger puzzle. Attentional patterns are necessary, but not sufficient, to explain the myriad of underlying motivations and goals at play during social interactions (Berger, 2014; Horowitz et al., 2006).

Relatedly, there has been a recent increase in language analysis research that explores *social cognitive* domains, such as the “why's” and “how's” of deception, persuasion, and person perception as they occur through verbal behavior. In particular,

research on persuasion has begun to thrive, with scholars starting to leverage language analysis methods in the exploration of how social appraisals are shaped by subtle linguistic cues (Althoff et al., 2014; Markowitz, 2019). Similarly, investigations of *how* one's personality is inferred by others through natural language has seen recent and steady growth (Berry et al., 1997; Tong et al., 2020).

Lastly, and most recently, there have been several exciting developments in observing more elaborate social phenomena in natural language, such as group identification and interpersonal motives (Danescu-Niculescu-Mizil et al., 2013; Niculae et al., 2015). Though it is still too early to say what the impact of these developments will be for psychological and communication sciences, we urge readers to pay close attention: this work signals some of the most important paths forward for all social sciences.

The Future of Language Analysis in Psychology

The traditional word counting approach to computerized language analysis has repeatedly shown its value for psychological research. While psychologists have largely adopted a “stick with what works” mentality to text analysis methods, computer scientists and computational linguists have accelerated into an all-out arms race for better models of language, with the most rapid innovations coming from the fields of Natural Language Processing and Natural Language Learning. Combined with the accessibility of language data afforded by digital technology, the potential for these new, more computational models of language is enormous for psychological science. Given that our words reflect who we are as psychological beings, increasingly accurate models of human language will necessarily produce an increasingly clear reflection of human thought, feelings, and behavior. In the near future, computation will be used not just to model language, but to more fully realize the potential of verbal behavior as a window into the deepest secrets of human psychology.

Increasingly Faithful Models of Human Psychology Through Verbal Behavior

Ask any social psychologist, discourse analyst, communication scientist, anthropologist, historian, or sociolinguist and they will gladly tell you: *context matters*. Perhaps the most common critique of the word counting approach is that many/most words can take on radically different meanings depending on the context in which they appear (Grimmer & Stewart, 2013). If each instance of a given word were diagnostic of where the author's attention is focused, one might expect the word would (in theory) map to a specific attentional target with some consistency. However, traditional word counting treats all instances of each word as equal: “play” in “the children *played* the game” is treated as psychometrically equivalent to “they went to the *play*,” “I got *played* by the scammer,” or “the children *played* the math video.” Of course, we intuitively know this to be untrue: some uses of “play” are more positive than others, whereas others have no emotional connotation whatsoever.

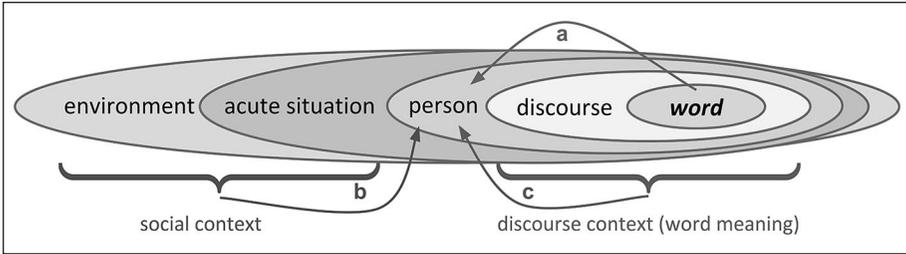


Figure 1. The various layers of context within which a word is observed. (a) The word count approach attempts to map words onto a person’s psychology, bypassing important contextual factors, (b) in more advanced language analysis approaches, social context is brought into the mix (e.g., as additional covariates or through experimental treatment/control groups), and (c) at the same time in more advanced natural language processing techniques, discourse context is leveraged to extract word meaning and, thus, accurate psychological meaning.

Variations in the meaning of the word “play,” then, should tell us very different things about a person’s psychology; indeed, research on word context finds this to be the case. Rouhizadeh et al. (2016), for example, showed that the psychological meaning of pronouns can reverse depending on the discourse context. In this work, the authors showed that well-established relationships between first-person pronouns and gender, age, and depression changed substantially depending on how they were used. For example, while females used somewhat more first-person singulars (I, me, my) than males, males were more likely to use first-persons in the context of physical attack verbs (hit, kick, strike). Despite the fact that self-referential language is one of the more robust linguistic correlates of depression (Holtzman et al., 2010), Rouhizadeh et al. (2016) compellingly showed that not all self-references are equal.

In the social sciences, we typically think of “context” as that which exists beyond the person, often conceptualized as the physical or social environment, or perhaps the specific situation within which an individual is being studied. However, we can also think of words as occurring within *nested* contexts: verbal behaviors occur within a discourse, which occurs at a specific time and place and involving specific individuals, and so on (e.g., Gray & Biber, 2011). Ultimately, the assumptions about how a given word reflects some underlying psychological process is tentative—completely functional, but still tentative—without a basic consideration of the discourse in which the word occurs (see Figure 1).⁵

Put another way: if our “model” of language fails to account for the context in which it occurs, we lose clarity on what individual words tells us about a person’s psychology in the first place. Looking at individual words under the “words as attention” paradigm, the psychological meaning attributed to the person comes from a decontextualized understanding of the word, thus yielding an ambiguous psychosocial sense (Markowitz & Hancock, 2019). A more ideal attentional framework would be a “words *in context* as attention” approach, whereby features of the individual

themselves are implicitly folded into the quantification of their words and what they are likely to reflect, psychologically speaking (Garten et al., 2019).

Embeddings and Beyond

The skeptical reader may be thinking “Big deal! We already account for context when hand-coding texts!” True as this may be: just as automation of the word counting method has removed the subjective, inconsistent “judgement calls” of manual scoring, so too will we benefit from computerized models of language that objectively and consistently appreciate the complex message-altering power of context (see Bursleson, 2010). As our ability to computationally model all aspects of language improves, including context, our ability to objectively and reliably probe the human mind in a standardized, reproducible, and scientific fashion grows.

An important step toward the inclusion of context in computerized language analysis has been the introduction of word embeddings. Briefly described, word embeddings represent each word as a vector of values that reflect its position within a high-dimensional geometric space (much like X, Y, and Z coordinates, but *far* more than 3 dimensions). The principle behind how these vectors are statistically “learned” is the assumption that word meaning is derived from word context—the so-called “distributional hypothesis” best summarized by Wittgenstein (1953): “the meaning of a word is its use in language.” One of the first computational realizations of word embeddings was *latent semantic indexing* (Deerwester et al., 1990), which is still in use today.⁶ Using such models, it becomes possible to statistically calculate something of a fuzzy representation of an entire text’s meaning (e.g., a diary entry, a tweet, a speech), then calculate the similarity of a document’s meaning to psychologically important concepts, such as “happiness” or “fairness” (Garten et al., 2018).

Nevertheless, most widely used word embeddings still treat each individual word as having a singular meaning or “sense.” Perhaps the most promising technology that approaches a more accurate, “true to psychology” model of language is the so-called “contextual word embeddings,” introduced only within the last few years in computational linguistics. These approaches to natural language processing, which utilize a deep learning technique known as transformers, produce a unique vector representing each word, as described above. Importantly, however, the precise details of each word’s vector is dependent on its exact context (Devlin et al., 2019). Thus, the word “play” in “got played by the scammer” and “played the game” will have totally different representations. Such approaches have led to unprecedented increases in performance in natural language applications ranging machine translation to speech recognition by digital personal assistants (Yang et al., 2019).

Due to their novelty, contextual embeddings have yet to see serious action in psychological research aside from classification benchmarking (e.g., Matero et al., 2019). However, we believe that the entire discipline of psychological language analysis is poised for a quantum shift. For the first time ever, words can be automatically, objectively, and reproducibly quantified in a way that inherently accounts for at least one level of context: discourse. In essence, contextual embeddings give a more stable

approach to move beyond function words, mapping a nearly endless supply of open-ended content words to a finite number of dimensions quantifying their meaning.

Most importantly, transcending the word counting approach into one that computationally “understands” social and environmental contexts allows us to use language as more than simply a window into a person’s attentional focus. As we create models of language that are increasingly “true to life” in terms of how humans generate, understand, and use natural language, we will gain new opportunities to probe social psychological questions. We can already see the stirrings of some such work within the field, wherein language models are being probed for validity insofar as they reflect known psychological phenomena. For example, well-established social biases that can be found in studies of “implicit associations” have recently been shown to have analogues in how social concepts are represented in word embedding models (Bolukbasi et al., 2016; Caliskan et al., 2017). Similarly, the changing social landscape can be seen in the rapid increase of research into stereotypes, prejudices, and biases as they are formulated, embedded, and surreptitiously transmitted in our language (Carpenter et al., 2017; Sneffjella et al., 2018). We are excited to see what can be learned about ourselves through the study of increasingly veridical language models created by highly collaborative teams of scholars spanning the sciences.

Greater Interdisciplinarity

We are entering an exciting time for language analysis in general, and for social psychologists in particular. Teams of computer and cognitive scientists have generated large, precise models of language going well beyond word counting. Such models are beginning to show up in the social sciences proper, often precipitating first through interdisciplinary conference proceedings and public whitepapers. However, the empirical paradigms under which newer language models are generated are often difficult to transplant directly into psychological research.

Critically, the increasing complexity of natural language processing methods directly necessitates interdisciplinary cooperation not just in the conduct of research, but also in the creation of new paradigms that facilitate the rapid and meaningful integration of new language analysis methods into the social sciences. The merits of interdisciplinary collaboration are often-discussed (e.g., Cummings & Kiesler, 2016), however, we acknowledge that lip service to interdisciplinarity is cheap. We believe that interdisciplinarity will be central to the future of language analysis research, however, it is worth taking a moment to underline the need for *true* collaboration across all types of boundaries.

The Bird in Borrowed Feathers

With the rise of digital technology, social scientists have gradually become more computationally savvy, on average, with programming languages like *R* and Python gaining traction. Similarly, computer scientists have begun to recognize the need for better theoretical models of human behavior, moving beyond a purely “informatic” view of

social psychology (Wang et al., 2007). Accordingly, there has been an explosion of research at the intersection of computer science and the social sciences; most of this work falls under the umbrella term of “computational social science.”

In practice, however, a majority of research branded as “computational social science” consists of little more than intelligently engineered data pipelines that use one type of human-generated data (e.g., text or videos) to predict another (e.g., scores on questionnaires; behaviors), often providing no real test of social theories or insights into human psychology. In more extreme cases, research in this vein can be caricatured as machine learning exercises masquerading as social science. Equally problematic is the flip side of the same coin: social scientists haphazardly wielding powerful machine learning algorithms and poorly-validated language models, deriving statistically flawed interpretations of their data. For every social scientist who rolls their eyes at *yet another* computer scientist butchering basic theories of human behavior, there is a computer scientist wincing in pain at *yet another* psychologist excitedly proclaiming that “big data” and “machine learning” solved their research problems with *unbelievable* accuracy. In the end, each field has its own scientific norms (i.e., *paradigms*, in the Kuhnian sense): fitting psychological theory into the norms of machine learning or vice versa rings of the classic “square peg into a round hole” problem.

The way forward may require shedding field-specific norms in favor of developing *new* paradigms in computational social science that suit the unique needs of the field. Initial steps toward this goal are visible in the form of practical guides written within specific disciplines that outline other fields and norms. For example, political science papers describing computational text analysis (Grimmer & Stewart, 2013); psychology papers describing machine learning (Stachl et al., 2020), and machine learning papers describing social sciences (O’Connor et al., 2011). We see this trend moving toward the development of a new field at the intersection of several social and computational sciences, where such survey works are no longer “translated” from one domain to another but, rather, an interdisciplinary space that can be free to develop its own norms. We encourage readers to keep a highly open-minded approach to new developments in today’s interdisciplinary gaps—innovations will likely emerge rapidly and take on unfamiliar but powerful forms.

Conclusion

Language as a vehicle for studying the human condition has existed since the birth of psychology and has gone through several evolutions over the past century. Early psychology subscribed to language being a powerful window into the human psyche, but its study was ill-defined and pre-scientific. Later, an empirical thrust came about, leading to approaches where statistical inference could be applied, but the methods for quantifying psychological meaning were extremely weak. We are only now beginning to realize the true power of objective, computational methods of psychological language analysis. As our computational models of language grow increasingly true-to-life in form and function, we expect to be able to peer into the inner workings of the mind in ways that are currently unimaginable. The secrets of language are being

unlocked in new and exciting ways, and we sit at the cusp of an absolutely revolutionary shift in how we conduct social scientific research.

Acknowledgments

The authors would like to thank Dave Markowitz, Gabriella Harari, Howie Giles, and Karolina Hansen for their helpful and insightful feedback on earlier drafts of this manuscript.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Preparation of this manuscript was funded, in part, by grants from the Swiss National Science Foundation (#196255) and the US National Institutes of Health—NIAAA (r01 aa028032). The views, opinions, and findings contained in this chapter are those of the author and should not be construed as position, policy, or decision of the aforementioned organizations unless so designated by other documents.

ORCID iD

Ryan L. Boyd  <https://orcid.org/0000-0002-1876-6050>

Notes

1. In this article, we use terms like “language analysis,” “psychological language analysis,” and “natural language analysis” interchangeably, referring to the idea of extracting psychological insights through the study of a person’s verbal behaviors. We also use terms such as “computerized text analysis,” “automated text analysis,” and “natural language processing” to refer to computerized methods designed to accomplish the same tasks in an automated fashion.
2. The simplicity and transparency of the word counting approach ranks among the most celebrated features of LIWC within psychology. Elsewhere, LIWC’s simplicity is often bemoaned as its greatest shortcoming—something of a lightning rod for ire in the more linguistic and computationally savvy circles of scholarship. What is commonly misunderstood, however, is that the methodological logic of LIWC and its spiritual predecessors was never intended as a model of language itself. Rather, the idea of counting psychologically “telling” words is grounded in psychometric theory and the statistical analysis of human behavior—empirical principles upon which the entirety of psychological science is based. No researcher involved in LIWC’s many iterations has ever believed that its list of “positive emotion” words, for example, is comprehensive, semantically cohesive, or even particularly well-suited to the task of estimating “sentiment.” Instead, the LIWC dictionaries have always been conceptualized more as fuzzy indicators of other, more interesting latent psychological phenomena.
3. We note that while the “words as attention” interpretation of verbal behavior is implicit to most mainstream language analysis research today, it is not the only interpretation (and not even necessarily the only *correct* interpretation). Other conceptualizations of language

approach verbal behavior as tapping into the degree or orientation of dimensional psychological processes, typologies that reflect psychodynamic “modes” of thinking, or even stages of psychosocial development (see, e.g., Bucci & Maskit, 2006; Conway et al., 2020; Lanning et al., 2018). Nevertheless, the majority of mainstream language analysis frameworks can be seen to rely on the implicit assumption of “words as attention” if we dig deep enough.

4. Of course, we do not mean to suggest that *nobody* studies language through the lens of social process: the social facets of verbal behavior are quite literally the foundation for dozens of academic disciplines. Rather, we note a paucity of work in *computerized, psychological* language analysis literature that explicitly treats language as a rich stream of ongoing social processes.
5. In our case, “discourse” is considered to be a sequence of words (or local context) in which a particular word appears, forming a coherent unit of meaning (Prasad et al., 2008).
6. Similar (but more computationally intricate) versions of the same idea include the currently popular, Google-developed *word2vec* (Mikolov et al., 2013).

References

- Allport, G. W. (1942). *The use of personal documents in psychological science*. Social Science Research Council.
- Althoff, T., Danescu-Niculescu-Mizil, C., & Jurafsky, D. (2014). How to ask for a favor: A case study on the success of altruistic requests. In E. Adar, P. Resnick, M. D. Choudhury, B. Hogan, & A. H. Oh (Eds.), *Proceedings of the Eighth International Conference on Weblogs and Social Media* (pp. 12–21). The AAAI Press.
- Ashton, M. C., Lee, K., & Paunonen, S. V. (2002). What is the central feature of extraversion? Social attention versus reward sensitivity. *Journal of Personality and Social Psychology*, *83*(1), 245–252. <https://doi.org/10.1037/0022-3514.83.1.245>
- Bardi, A., Calogero, R. M., & Mullen, B. (2008). A new archival approach to the study of values and value–behavior relations: Validation of the value lexicon. *Journal of Applied Psychology*, *93*(3), 483–497. <https://doi.org/10.1037/0021-9010.93.3.483>
- Bayram, A. B., & Ta, V. P. (2019). Diplomatic chameleons: Language style matching and agreement in international diplomatic negotiations. *Negotiation and Conflict Management Research*, *12*(1), 23–40. <https://doi.org/10.1111/ncmr.12142>
- Berger, J. (2014). Word of mouth and interpersonal communication: A review and directions for future research. *Journal of Consumer Psychology*, *24*(4), 586–607. <https://doi.org/10.1016/j.jcps.2014.05.002>
- Berry, D. S., Pennebaker, J. W., Mueller, J. S., & Hiller, W. S. (1997). Linguistic bases of social perception. *Personality and Social Psychology Bulletin*, *23*(5), 526–537. <https://doi.org/10.1177/0146167297235008>
- Block, J. H., Fisch, C. O., Obschonka, M., & Sandner, P. G. (2019). A personality perspective on business angel syndication. *Journal of Banking & Finance*, *100*, 306–327. <https://doi.org/10.1016/j.jbankfin.2018.10.006>
- Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., & Kalai, A. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Proceedings of the 30th International Conference on Neural Information Processing Systems* (pp. 4356–4364). Curran Associates Inc.

- Boyd, R. L., Pasca, P., & Lanning, K. (2020). The personality panorama: Conceptualizing personality through big behavioural data. *European Journal of Personality, 34*(5), 599–612. <https://doi.org/10.1002/per.2254>
- Boyd, R. L., & Pennebaker, J. W. (2017). Language-based personality: A new approach to personality in a digital world. *Current Opinion in Behavioral Sciences, 18*, 63–68. <https://doi.org/10.1016/j.cobeha.2017.07.017>
- Bucci, W., & Maskit, B. (2006). A weighted dictionary for referential activity. In J. G. Shanahan, Y. Qu, & J. Wiebe (Eds.), *Computing attitude and affect in text* (pp. 49–60). Springer.
- Burleson, B. R. (2010). The nature of interpersonal communication: A message-centered approach. In C. R. Berger, M. E. Roloff, & D. R. Roskos-Ewoldsen (Eds.), *The handbook of communication science* (pp. 145–164). Sage. <https://dx.doi.org/10.4135/9781412982818.n9>
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science, 356*(6334), 183–186. <https://doi.org/10.1126/science.aal4230>
- Carpenter, J., Preotiuc-Pietro, D., Flekova, L., Giorgi, S., Hagan, C., Kern, M. L., Buffone, A. E. K., Ungar, L., & Seligman, M. E. P. (2017). Real men don't say "cute": Using automatic language analysis to isolate inaccurate aspects of stereotypes. *Social Psychological and Personality Science, 8*(3), 310–322. <https://doi.org/10.1177/1948550616671998>
- Chung, C. K., & Pennebaker, J. W. (2018). What do we know when we LIWC a person? Text analysis as an assessment tool for traits, personal concerns and life stories. In *The SAGE handbook of personality and individual differences: The science of personality and individual differences* (pp. 341–360). Sage. <https://doi.org/10.4135/9781526451163.n16>
- Clarke, A. C., & Kubrick, S. (1968). *2001: A space odyssey*. The New American Library.
- Colby, B. N., Collier, G. A., & Postal, S. K. (1963). Comparison of themes in folktales by the General Inquirer system. *The Journal of American Folklore, 76*(302), 318–323. <https://doi.org/10.2307/537928>
- Conway, L. G., Conway, K. R., & Houck, S. C. (2020). Validating automated integrative complexity: Natural language processing and the Donald Trump Test. *Journal of Social and Political Psychology, 8*(2), 504–524. <https://doi.org/10.5964/jssp.v8i2.1307>
- Cummings, J. N., & Kiesler, S. (2016). Collaborative research across disciplinary and organizational boundaries. *Social Studies of Science, 35*(5), 703–722. <https://doi.org/10.1177/0306312705055535>
- D'Andrea, W., Chiu, P. H., Casas, B. R., & Deldin, P. (2012). Linguistic predictors of post-traumatic stress disorder symptoms following 11 September 2001. *Applied Cognitive Psychology, 26*(2), 316–323. <https://doi.org/10.1002/acp.1830>
- Danescu-Niculescu-Mizil, C., West, R., Jurafsky, D., Leskovec, J., & Potts, C. (2013). No country for old members: User lifecycle and linguistic change in online communities. In *Proceedings of the 22nd International Conference on World Wide Web* (pp. 307–318). <https://doi.org/10.1145/2488388.2488416>
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science, 41*(6), 391–407. [https://doi.org/10.1002/\(SICI\)1097-4571\(199009\)41:6<391::AID-ASII>3.0.CO;2-9](https://doi.org/10.1002/(SICI)1097-4571(199009)41:6<391::AID-ASII>3.0.CO;2-9)
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (pp. 4171–4186). <https://doi.org/10.18653/v1/N19-1423>

- Doorn, K. A., & Müller-Frommeyer, L. (2020). Reciprocal language style matching in psychotherapy research. *Counselling and Psychotherapy Research, 20*(3), 449–455. <https://doi.org/10.1002/capr.12298>
- Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., Sap, M., Weeg, C., Larson, E. E., Ungar, L. H., & Seligman, M. E. P. (2015). Psychological language on Twitter predicts county-level heart disease mortality. *Psychological Science, 26*(2), 159–169. <https://doi.org/10.1177/0956797614557867>
- Freud, S. (1915). *The unconscious*. Penguin Classics.
- Garcia, D., & Rimé, B. (2019). Collective emotions and social resilience in the digital traces after a terrorist attack. *Psychological Science, 30*(4), 617–628. <https://doi.org/10.1177/0956797619831964>
- Garten, J., Hoover, J., Johnson, K. M., Boghrati, R., Iskiwitch, C., & Dehghani, M. (2018). Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior Research Methods, 50*(1), 344–361. <https://doi.org/10.3758/s13428-017-0875-9>
- Garten, J., Kennedy, B., Hoover, J., Sagae, K., & Dehghani, M. (2019). Incorporating demographic embeddings into language understanding. *Cognitive Science, 43*(1), e12701. <https://doi.org/10.1111/cogs.12701>
- Gasiorek, J., Weatherall, A., & Watson, B. (2021). Interactional adjustment: Three approaches in language and social psychology. *Journal of Language and Social Psychology, 40*(1), 102–119. <https://doi.org/10.1177/0261927X20965652>
- Golder, S. A., & Macy, M. W. (2011). Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science, 333*(6051), 1878–1881. <https://doi.org/10.1126/science.1202775>
- Gottschalk, L. A., & Gleser, G. C. (1969). *The measurement of psychological states through the content analysis of verbal behavior*. University of California Press.
- Gray, B., & Biber, D. (2011). Corpus approaches to the study of discourse. In K. Hyland & B. Paltridge (Eds.), *The continuum companion to discourse analysis* (pp. 138–152). Continuum.
- Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis, 21*(3), 267–297. <https://doi.org/10.1093/pan/mps028>
- Heuer, K., Müller-Frommeyer, L. C., & Kauffeld, S. (2020). Language matters: The double-edged role of linguistic style matching in work groups. *Small Group Research, 51*(2), 208–228. <https://doi.org/10.1177/1046496419874498>
- Hogenraad, R. (2003). The words that predict the outbreak of wars. *Empirical Studies of the Arts, 21*(1), 5–20. <https://doi.org/10.2190/HJWQ-QRBX-0C2E-VJYA>
- Holsti, O. R. (1965). East-West conflict and Sino-Soviet relations. *The Journal of Applied Behavioral Science, 1*(2), 115–130. <https://doi.org/10.1177/002188636500100202>
- Holtzman, N. S., Vazire, S., & Mehl, M. R. (2010). Sounds like a narcissist: Behavioral manifestations of narcissism in everyday life. *Journal of Research in Personality, 44*(4), 478–484. <https://doi.org/10.1016/j.jrp.2010.06.001>
- Horowitz, L. M., Wilson, K. R., Turan, B., Zolotsev, P., Constantino, M. J., & Henderson, L. (2006). How interpersonal motives clarify the meaning of interpersonal behavior: A revised circumplex model. *Personality and Social Psychology Review, 10*(1), 67–86. https://doi.org/10.1207/s15327957pspr1001_4

- Iliev, R., Hoover, J., Dehghani, M., & Axelrod, R. (2016). Linguistic positivity in historical texts reflects dynamic environmental and psychological factors. *Proceedings of the National Academy of Sciences*, 113(49), 7871–7879. <https://doi.org/10.1073/pnas.1612058113>
- Ireland, M. E., & Pennebaker, J. W. (2010). Language style matching in writing: Synchrony in essays, correspondence, and poetry. *Journal of Personality and Social Psychology*, 99(3), 549–571. <https://doi.org/10.1037/a0020386>
- Jordan, K. N., Pennebaker, J. W., & Ehrig, C. (2018). The 2016 U.S. presidential candidates and how people tweeted about them. *SAGE Open*, 8(3), 1–8. <https://doi.org/10.1177/2158244018791218>
- Kacewicz, E., Pennebaker, J. W., Davis, M., Jeon, M., & Graesser, A. C. (2014). Pronoun use reflects standings in social hierarchies. *Journal of Language and Social Psychology*, 33(2), 125–143. <https://doi.org/10.1177/0261927X13502654>
- Kranz, P. (1970). What do people do all day? *Behavioral Science*, 15(3), 286–291. <https://doi.org/10.1002/bs.3830150311>
- Laffal, J. (1964). Freud's theory of language. *The Psychoanalytic Quarterly*, 33(2), 157–175. <https://doi.org/10.1080/21674086.1964.11926307>
- Lanning, K., Pauletti, R. E., King, L. A., & McAdams, D. P. (2018). Personality development through natural language. *Nature Human Behaviour*, 2(5), 327–334. <https://doi.org/10.1038/s41562-018-0329-0>
- Lerner, J. S., & Keltner, D. (2000). Beyond valence: Toward a model of emotion-specific influences on judgement and choice. *Cognition and Emotion*, 14(4), 473–493. <https://doi.org/10.1080/026999300402763>
- Lumma, A.-L., Valk, S. L., Böckler, A., Vrtička, P., & Singer, T. (2018). Change in emotional self-concept following socio-cognitive training relates to structural plasticity of the prefrontal cortex. *Brain and Behavior*, 8(4), e00940. <https://doi.org/10.1002/brb3.940>
- Mairesse, F., Walker, M. A., Mehl, M. R., & Moore, R. K. (2007). Using linguistic cues for the automatic recognition of personality in conversation and text. *Journal of Artificial Intelligence Research*, 30(1), 457–500. <https://doi.org/10.1613/jair.2349>
- Markowitz, D. M. (2019). What words are worth: National Science Foundation grant abstracts indicate award funding. *Journal of Language and Social Psychology*, 38(3), 264–282. <https://doi.org/10.1177/0261927X18824859>
- Markowitz, D. M., & Hancock, J. T. (2019). Deception and language: The contextual organization of language and deception (COLD) framework. In T. Docan-Morgan (Ed.), *The Palgrave handbook of deceptive communication* (pp. 193–212). Springer International Publishing. https://doi.org/10.1007/978-3-319-96334-1_10
- Martindale, A. E., & Martindale, C. (1988). Metaphorical equivalence of elements and temperaments: Empirical studies of Bachelard's theory of imagination. *Journal of Personality and Social Psychology*, 55(5), 836–848. <https://doi.org/10.1037/0022-3514.55.5.836>
- Matero, M., Idnani, A., Son, Y., Giorgi, S., Vu, H., Zamani, M., Limbachiya, P., Guntuku, S. C., & Schwartz, H. A. (2019). Suicide risk assessment with multi-level dual-context language and BERT. In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology* (pp. 39–44). <https://doi.org/10.18653/v1/W19-3005>
- McCrae, R. R., & Costa, P. T. (1995). Trait explanations in personality psychology. *European Journal of Personality*, 9(4), 231–252. <https://doi.org/10.1002/per.2410090402>
- Meehl, P. E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports*, 66(1), 195–244. <https://doi.org/10.2466/pr0.1990.66.1.195>
- Mehl, M. R., Raison, C. L., Pace, T. W. W., Arevalo, J. M. G., & Cole, S. W. (2017). Natural language indicators of differential gene regulation in the human immune system. *Proceedings*

- of the *National Academy of Sciences*, 114(47), 12554–12559. <https://doi.org/10.1073/pnas.1707373114>
- Michel, J.-B., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Pickett, J. P., Hoiberg, D., Clancy, D., Norvig, P., Orwant, J., Pinker, S., Nowak, M. A., & Aiden, E. L. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, 331(6014), 176–182. <https://doi.org/10.1126/science.1199644>
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 26* (pp. 3111–3119). Curran Associates. <http://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf>
- Niculae, V., Kumar, S., Boyd-Graber, J., & Danescu-Niculescu-Mizil, C. (2015). Linguistic harbingers of betrayal: A case study on an online strategy game. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (pp. 1650–1659). <http://www.aclweb.org/anthology/P15-1159>
- O'Connor, B., Bamman, D., & Smith, N. A. (2011). Computational text analysis for social science: Model assumptions and complexity. In W. Mason, J. W. Vaughan, & H. Wallach (Eds.), *Second Workshop on Computational Social Science and Wisdom of the Crowds* (pp. 1–7). Curran Associates.
- Pennebaker, J. W. (2011). *The secret life of pronouns: What our words say about us*. Bloomsbury Press. <https://doi.org/10.1093/llc/ftq006>
- Pennebaker, J. W., Chung, C. K., Frazee, J., Lavergne, G. M., & Beaver, D. I. (2014). When small words foretell academic success: The case of college admissions essays. *PLoS ONE*, 9(12), e115844.
- Pennebaker, J. W., & Francis, M. E. (1999). *Linguistic Inquiry and Word Count (LIWC): A computer-based text analysis program*. Lawrence Erlbaum.
- Prasad, R., Dinesh, N., Lee, A., Miltsakaki, E., Robaldo, L., Joshi, A., & Webber, B. (2008, May). The Penn Discourse TreeBank 2.0. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*. LREC 2008, Marrakech, Morocco. http://www.lrec-conf.org/proceedings/lrec2008/pdf/754_paper.pdf
- Pressman, S. D., & Cohen, S. (2007). Use of social words in autobiographies and longevity. *Psychosomatic Medicine*, 69(3), 262–269.
- Psathas, G. (1969). The General Inquirer: Useful or not? *Computers and the Humanities*, 3(3), 163–174. <https://doi.org/10.1007/BF02401609>
- Robinson, M. D., Boyd, R. L., Fetterman, A. K., & Persich, M. R. (2017). The mind versus the body in political (and nonpolitical) discourse: Linguistic evidence for an ideological signature in U.S. politics. *Journal of Language and Social Psychology*, 36(4), 438–461. <https://doi.org/10.1177/0261927X16668376>
- Robinson, R. L., Navea, R., & Ickes, W. (2013). Predicting final course performance from students' written self-introductions. *Journal of Language and Social Psychology*, 32(4), 469–479. <https://doi.org/10.1177/0261927X13476869>
- Rodríguez-Arauz, G., Ramírez-Esparza, N., Pérez-Brena, N., & Boyd, R. L. (2017). Hablo Inglés y Español: Cultural self-schemas as a function of language. *Frontiers in Psychology*, 8, 1–15. <https://doi.org/10.3389/fpsyg.2017.00885>
- Rouhizadeh, M., Ungar, L., Buffone, A., & Schwartz, H. A. (2016). Using syntactic and semantic context to explore psychodemographic differences in self-reference. In *Proceedings of*

- the 2016 Conference on Empirical Methods in Natural Language Processing (pp. 2054–2059). <https://doi.org/10.18653/v1/D16-1219>
- Saxbe, D. E., & Repetti, R. (2010). No place like home: Home tours correlate with daily patterns of mood and cortisol. *Personality and Social Psychology Bulletin*, 36(1), 71–81. <https://doi.org/10.1177/0146167209352864>
- Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., Shah, A., Kosinski, M., Stillwell, D., Seligman, M. E. P., & Ungar, L. H. (2013). Personality, gender, and age in the language of social media: The open-vocabulary approach. *PLoS ONE*, 8(9), e73791. <https://doi.org/10.1371/journal.pone.0073791>
- Smith, M. S. (1968). The computer and the TAT. *Journal of School Psychology*, 6(3), 206–214. [https://doi.org/10.1016/0022-4405\(68\)90017-4](https://doi.org/10.1016/0022-4405(68)90017-4)
- Sneffjella, B., Schmidtke, D., & Kuperman, V. (2018). National character stereotypes mirror language use: A study of Canadian and American tweets. *PLoS ONE*, 13(11), e0206188. <https://doi.org/10.1371/journal.pone.0206188>
- Stachl, C., Pargent, F., Hilbert, S., Harari, G. M., Schoedel, R., Vaid, S., Gosling, S. D., & Bühner, M. (2020). Personality research and assessment in the era of machine learning. *European Journal of Personality*. Advance online publication. <https://doi.org/10.1002/per.2257>
- Stone, P. J., Bales, R. F., Namenwirth, J. Z., & Ogilvie, D. M. (1962). The General Inquirer: A computer system for content analysis and retrieval based on the sentence as a unit of information. *Behavioral Science*, 7(4), 484–498. <https://doi.org/10.1002/bs.3830070412>
- Stone, P. J., Dunphy, D. C., Smith, M. S., & Ogilvie, D. M. (1966). *The General Inquirer: A computer approach to content analysis*. MIT Press.
- Ta, V. P., Babcock, M. J., & Ickes, W. (2017). Developing Latent Semantic Similarity in initial, unstructured interactions: The words may be all you need. *Journal of Language and Social Psychology*, 36(2), 143–166. <https://doi.org/10.1177/0261927X16638386>
- Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, 29(1), 24–54. <https://doi.org/10.1177/0261927X09351676>
- Thomas, C. B., & Duszynski, K. R. (1985). Are words of the Rorschach predictors of disease and death? The case of “whirling.” *Psychosomatic Medicine*, 47(2), 201–211.
- Tong, S. T., Corriero, E. F., Wibowo, K. A., Makki, T. W., & Slatcher, R. B. (2020). Self-presentation and impressions of personality through text-based online dating profiles: A lens model analysis. *New Media & Society*, 22(5), 875–895. <https://doi.org/10.1177/1461444819872678>
- Vine, V., Boyd, R. L., & Pennebaker, J. W. (2020). Natural emotion vocabularies as windows on distress and well-being. *Nature Communications*, 11, Article 4525. <https://doi.org/10.1038/s41467-020-18349-0>
- Wang, F.-Y., Carley, K. M., Zeng, D., & Mao, W. (2007). Social computing: From social informatics to social intelligence. *IEEE Intelligent Systems*, 22(2), 79–83. <https://doi.org/10.1109/MIS.2007.41>
- Weintraub, W. (1989). *Verbal behavior in everyday life*. Springer.
- Wilson, A. (2006). Development and application of a content analysis dictionary for body boundary research. *Literary and Linguistic Computing*, 21(1), 105–110. <https://doi.org/10.1093/lc/fqi014>
- Wittgenstein, L. (1953). *Philosophical investigations*. Basil Blackwell.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R. R., & Le, Q. V. (2019). XLNet: Generalized autoregressive pretraining for language understanding. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, & R. Garnett (Eds.), *Proceedings*

of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019) (pp. 5753–5763). Curran Associates. <http://papers.nips.cc/paper/8812-xlnet-generalized-autoregressive-pretraining-for-language-understanding.pdf>

Author Biographies

Ryan L. Boyd is an assistant professor of Behavioural Analytics in the Department of Psychology at Lancaster University, United Kingdom. His research uses language analysis methods to study psychological phenomena that stretch across social, personality, and health psychology, with occasional work in forensic psychology and the digital humanities. He is the author of over a dozen free, open source text analysis applications and is a co-developer of Linguistic Inquiry and Word Count.

H. Andrew Schwartz is an assistant professor of Computer Science at Stony Brook University, in New York. His interdisciplinary research focuses on developing large-scale language analyses for psychological and health sciences, as well as integrating human factors into natural language processing approaches to improve results for diverse populations. He is a PI of the World Well-Being Project, a multidisciplinary consortium of Computer Scientists and Psychologists studying physical and psychological well-being based on language in social media.