Published in final edited form as:

Commun Methods Meas. 2015 ; 9(1-2): 55–77. doi:10.1080/19312458.2014.999751.

Linking Neuroimaging with Functional Linguistic Analysis to Understand Processes of Successful Communication

Matthew Brook O'Donnell and Emily B. Falk

Communication Neuroscience Lab, Annenberg School for Communication, University of Pennsylvania

Abstract

Functional linguistic models posit a systematic link between language FORM and the FUNCTIONS for which language is used. This is a systematic (and therefore quantifiable) relationship. Yet, many open questions remain about the mechanisms that link form, function and communication relevant outcomes. Neuroimaging methods can provide insight into such processes that are not apparent from other methods. We argue that the combination of neural and linguistic measures will allow insight into both individual and population-level communication processes that would not be possible using either method in isolation. We present examples illustrating this methodological integration and notes regarding the most amenable linguistic tools. We summarize a framework in which language presented to and produced by participants undergoing neuroimaging is correlated with the resulting neural data and other proximal communication outcomes allowing the triangulation of individual experimental with population level outcomes, thereby linking between micro and macro levels of analysis.

Keywords

quantitative language analysis; fMRI; neuroimaging; sentiment analysis; message propagation

The study of successful communication can take place at both micro and macro levels of analysis. For instance, positive effects of a physical activity health campaign might be measured in terms of individual attitude or behavior change or larger-scale shifts in population-level sentiment toward physical activity. At the macro- or population-level this could involve the analysis of what types of messages generate the greatest reductions in chronic disease across states, or an analysis of which messages spread most rapidly and broadly through social networks. At the micro- or individual-level this might involve the experimental testing of what message factors are found to be persuasive or what kinds of psychological processes are involved in an individual's motivation to share a message.

Different methods offer different strengths, and are applicable at different levels of analysis. Neuroimaging, for example, offers particularly rich data regarding the neurocognitive mechanisms of communication, often beyond what can be obtained using self-report

measures on small samples (Berkman & Falk, 2013; Cascio, Dal Cin, & Falk, 2013)(see Weber, this issue and Falk et al., this issue for more detailed treatment of the strengths of neuroimaging methods for theory testing). Neuroimaging on its own, however, is currently limited to application at the micro-level (c.f., Falk, Hyde, Mitchell, et al., 2013). The quantitative analysis of language samples, in contrast, does not tap as directly into the psychological and neurocognitive processes that may unfold in real time during cognitive and emotional tasks, but language samples can be obtained at individual, group, and population levels, and used as markers of individual differences and cognitive states (Pennebaker, 2011). In this way, language can be applied from the micro-level through to the macro-level. Language can also be a vehicle that links multiple levels of analysis, since it is a vehicle that spreads ideas from individuals through populations (Gruhl, Guha, Liben-Nowell, & Tomkins, 2004).

We offer two main arguments in this paper. First, we highlight the ways in which language tools can triangulate and link mechanisms that support communication phenomena at multiple levels of analysis. In this way, linguistic analysis can bring together theories about the psychology of individual level responses to interpersonal and mediated stimuli with large-scale responses observed at the population level. Language is relatively easy to collect at both the individual level and increasingly easy to aggregate at the large scale ('big data' collected from social media and other online channels). Quantitative linguistic tools, such as those discussed below, can be efficiently applied across levels and the resulting quantitative measures can be interpreted within sociopsychological frameworks. Second, we note that ideas and messages spread from person to person through relational and social networks where language serves as one primary mode for this transmission. As such, a second way to link our understanding of individual and population-level phenomena in communication science is to follow the linguistic pathways through which the ideas are initially processed in the brains of message recipients, and then to track flow beyond the individual. This would involve applying quantitative linguistic analysis to the language trace left by instances of person-to-person idea propagation.

The paper is organized as follows. First, we provide an overview of some of the key quantitative approaches to language data, specifically categorical word counting approaches and supervised machine learning techniques, such as those used for sentiment analysis. For corresponding reviews of approaches to neuroimaging, we refer interested readers to other papers in this special issue. Then we present an exploration of the two main arguments described above concerning the use of quantitative linguistic measures as a link between observed neural mechanisms and population level outcomes. We then present examples of studies that have begun to combine linguistic analysis and neuroimaging and briefly discuss theoretical insights that have emerged from these studies. Finally, we offer practical notes regarding types of analysis that would be most appropriate to achieve the goals of linking communicative processes at individual and population levels and of tracing the pathways that lead to the spread of ideas and describe avenues for future development of these approaches.

Types of quantitative linguistic measure

A broad range of approaches exist for the quantification of linguistic data, developed within fields such as computational and corpus linguistics (McEnery & Hardie, 2011), natural language processing (Jurafsky & Martin, 2008; Manning & Schütze, 1999) and information retrieval (Manning, Raghavan, & Schütze, 2008). Some approaches are rich in terms of linguistic theory while others leverage statistical modeling and learning techniques. Grimmer and Stewart (2013) present a state of the art overview of the use of techniques for automated content analysis with specific application to political texts and provide guidance as to how to select the most appropriate method and measures for a particular research question.

Quantitative Linguistic Analysis of Large-Scale Population Level Data

The use of quantitative analysis of language has become increasingly popular as more and more large-scale data sets (e.g., from social media and other online sources) have been used to analyze social movements and cultural phenomena and to make links with personality traits and emotional characteristics and behaviors (Gruzd, Doiron, & Mai, 2011; Schwartz et al., 2013; Yarkoni, 2010). The Google Books ngram corpus, for example, has been utilized to chart the emergence and decline of trends, ideologies and changes in cultural norms based on the frequency of words and phrases found in a very large sample of books published between 1800 and the present (the term 'culturomics' has been coined for this approach; also Aiden & Michel, 2013). Automated sentiment analysis of politically focused tweets has been linked with opinions as measured by traditional polling techniques and thereby demonstrated as a tool for predicting trends and polling outcomes (Balasubramanian, Routledge, & Smith, 2010; Gayo-Avello, 2013; Tumasjan, Sprenger, Sandner, & Welpe, 2011). Both the volume and sentiment (measured using linguistic analysis) of weblog discussion of movies has been shown to predict box office performance (Asur & Huberman, 2013; Mishne, 2006). Message valence, measured by sentiment analysis, has been linked with the spread of messages through social media, specifically that positive tweets are more contagious than negative ones (Gruzd, Doiron, & Mai, 2011; Gruzd, 2013). Large-scale linguistic analysis of blog posts has established the link between language usage patterns and personality among blog authors (Yarkoni, 2010).

Building on this work, linguistic analysis of status updates made by 75000 Facebook users correlated with psychological trait and demographic data for all individuals revealing systematic links between language use and personality (Schwartz et al., 2013). Crowdsourced scoring (i.e., 'how happy does this word make you feel') of a large number of highly frequent words has been used to score social media language (from Twitter) to link language with geographical and temporal factors and corroborating data from large surveys (Dodds, Harris, Kloumann, Bliss, & Danforth, 2011; Mitchell, Frank, Harris, Dodds, & Danforth, 2013). Other studies have used large online language samples to describe the general and comparative emotional state of certain groups and various demographics (Godbole, Srinivasaiah, & Skiena, 2007; Ritter, Preston, & Hernandez, 2013) and proposed methods to detect growing community tensions (Burnap et al., 2013).

In sum, it is clear that across multiple contexts, aggregated language samples have been linked to important civic and psychological outcomes at the large-scale population level. However, the mechanisms that produce these outcomes are not fully understood. As demonstrated by other contributions to this special issue (e.g., Falk et al., this issue; Weber, this issue) neuroimaging tools and techniques are well suited to the investigation of the mechanisms involved in the production and reception of communicative processes and can be linked to population level outcomes. We argue that the integration of certain types of linguistic analysis tools with neural measures can help unpack psychological and neurocognitive mechanisms behind successful communication and propagation. This is made possible because these tools are applicable to language collected across levels of analysis, namely both at the individual/experimental level and the population/observational level.

Quantitative Linguistic Tools Suitable for Linking Neural and Linguistic Measures

Here, we provide a brief summary of linguistic analysis methods that we view as most promising for linking neuroimaging and linguistic data collected in an experimental context to understand psychological and social processes involved in communication (see also Grimmer & Stewart, 2013 for an overview of quantitative linguistics methods applicable to the study of political communication in particular but of relevance to a range of areas of communication science). More specifically, we narrow our discussion to two kinds of analysis aimed at quantifying the affective dimensions of language. These are: (1) categorical word counting techniques and (2) supervised text classification using linguistic features.

Categorical word counting approaches—Much of the research on links between psychology and language use has relied on categorical word counting approaches. These can score and compare language produced by individuals in both experimental (e.g., writing, thought-listing or open-ended response tasks) and naturalist contexts (e.g., diary language, Facebook posts, transcribed conversation, other print or online published language, etc.). The underlying idea is that a spectrum of individual differences are reflected in differences of word choice and frequency of usage (Bradley & Lang, 1999; Gottschalk & Gleser, 1969; Kahn, Tobin, Massey, & Anderson, 2007; Pennebaker, Chung, Ireland, Gonzales, & Booth, 2007; Pennebaker, Mehl, & Niederhoffer, 2003; Pennebaker, 2011; Tausczik & Pennebaker, 2010). These approaches have also been applied to the study of effective communication and persuasion, both in terms of methods such as automated content analysis (Gottschalk, 1995; Grimmer & Stewart, 2013) and experimentally where language form is manipulated and the persuasive effect measured (Blankenship & Craig, 2006; Craig & Blankenship, 2011; Holtgraves & Lasky, 1999; Hosman & Siltanen, 2011; Kramer, Guillory, & Hancock, 2014). They have also been applied to the question of language synchrony or communicative accommodation during in-person and mediated conversation (Danescu-Niculescu-Mizil, Gamon, & Dumais, 2011; Gonzales, Hancock, & Pennebaker, 2010; Goode & Robinson, 2013; Ireland et al., 2011).

Categorical word counting approaches are methodologically simple but have demonstrated remarkable results in usage across a range of applications. They rely upon predefined

dictionaries (or lexicons) in which words, phrases and linguistic features are grouped into specific categories (e.g., conceptual groupings such as social, positive, negative, sensing words, or formal linguistic categories such as pronouns, verbs, nouns etc.). Then texts are scored on each of these categories by counting the number of instances of items in a category and normalizing the count according to text length. When two categories can be defined in contrast, for example, positive and negative, counts for one can be taken away from counts for the other (i.e., +1 for each positive item and -1 for each negative resulting in a final positive-negative score) (Baek, Cappella, & Bindman, 2011). Examples of categorical word counting approaches include General Inquirer (Stone, Dunphry, Smith, & Ogilvie, 1966), LIWC (Linguistic Inquiry and Word Count - Pennebaker et al., 2007; Tausczik & Pennebaker, 2010) and other affective lexicons such as ANEW (Affective Norms for English Words - Bradley & Lang, 1999) and SentiWordNet (Baccianella, Esuli, & Sebastiani, 2010). The affective lexicons noted above have been experimentally validated by collecting responses to words from the dictionary using scales such as arousal and positivity. In addition, crowd sourcing validation (e.g. Amazon MTurk) is now increasingly utilized (Dodds et al., 2011). Word category approaches have been broadly applied to a range of psychological, sociological and communicative areas. For instance, correlations between self-reported personality, behavioral data and LIWC personality categories have been demonstrated (Fast & Funder, 2008) and the construct validity of these categories established (Kahn et al., 2007).

Limitations of the category word count approach include the cost involved in producing a categorical dictionary and the assumption that words in isolation can be assigned to domain general affective meaning categories—that is, the notion that a word will be used with the same sense, and therefore affect, across different contexts, genres and speakers. A conceptual limit of word count approaches is that they adopt a naïve model of language in which a text is seen as an unordered list of words, where each word is processed independently of its context. This ‘bag of words’ of approach, common in many natural language processing (NLP) and information retrieval (IR) applications, is computationally efficient and generally produces good results. However, it fails to capture the role of word order and syntactic structure in making meaning, both at the micro-level, for example, modification ‘it's not bad but actually very good’ versus ‘it's actually not good but very bad’, and at the sentence and paragraph level, for example, a single positive sentence followed by an overall negative review ‘I did enjoy reading this book but overall I wouldn't recommend it because...’. Much of how meaning is created in language is a result of combinations of words (e.g., n-grams, collocations and constructions) that are not entirely semantically compositional (Bybee, 2010; Ellis & O'Donnell, 2012; Hoey, 2005; O'Donnell, 2011; O'Donnell, Römer, & Ellis, 2013; Sinclair, 2004). That is, the meaning and particularly the connotation or pragmatic function of a series of words is more than just the sum of the meanings of the words themselves.

However, the simplicity of implementation and methodological transparency can be seen as advantages for initial investigations combining linguistic measures and fMRI data (O'Donnell, Falk, & Lieberman, In press; Saxbe, Yang, Borofsky, & Immordino-Yang, 2013). On the whole, this is a method that we recommend for investigators who have strong theoretical predictions about the category of mental process most engaged during the

communicative process being investigated. For instance, using LIWC one could decide to use scores from the cognitive processes category if it is hypothesized that people are likely to engage in deliberative processing during stimulus exposure. And that this cognitive process is reflected in the language they procedure in response to these stimuli. In contrast, the affective processes category (or specific subcategories, e.g., positive or negative emotion) would better fit an investigation of emotional sharing. Example 1, below, demonstrates some initial success in combining category word count measures, specifically words from the LIWC social processes category, with neuroimaging data in the context of understanding how communicators decide which products to recommend to others and the language they use in doing so.

Supervised machine learning—Supervised text classification combines human coding of a small set of training texts and automatic classification of a larger set of similar texts. This is accomplished by creating a statistical model built from the analysis of features in the training set and applying it to the test set. It thereby incorporates the strengths of human communicative competence and big-data scalable statistical pattern analysis. That is, human coders excel at identifying communicative intent and effect but are relatively poor at pointing out the low-level language features that give rise to them. In contrast, high-level pragmatic and contextual understanding has proved to be a particularly challenging goal for natural language processing but computers excel at discovering low-level associations between words, text and context. Supervised text classification begins with texts annotated into categories, for example, positive and negative reviews, descriptive and evaluative depictions, interactional and informational messages, etc. Then a set of linguistic features are identified (e.g., words, phrases, part-of-speech and semantic categories, contextual cues) and their distribution over the training set extracted. A statistical modeling procedure is then selected that can provide accurate and maximally distinguishing predictions of a text's category based on the features it contains. Figure 1 summarizes this process for one form of supervised machine learning—automated sentiment analysis (Godbole et al., 2007; Pang & Lee, 2008; Pang, Lee, & Vaithyanathan, 2002). The accuracy of the resulting classifier algorithm is tested against various subsets of the annotated training data and can be tuned by changing the linguistic features used and by changing the model parameters or applying alternative statistical models.

Limitations of the use of supervised classifiers include the need for sufficient independent training data that matches your textual corpus and the need for labeled data produced by human coding or derived from a relevant outcome (e.g., product rating). Classifier performance is strongly associated with domain (i.e., text-type and genre). So a classifier that accurately classifies product reviews into positive and negative—say from a training corpus from [Amazon.com](https://www.amazon.com) of one and two-star ratings for negative and four and five-star for positive (Jindal & Liu, 2008)—will not perform well if applied to a corpus of New York Times articles without tuning or retraining.

For the methodological combination of quantitative linguistic and fMRI analysis proposed here, supervised text classification has significant advantages. It captures implicit patterns of natural language that are associated with specific communicative events (e.g., two-word phrases ‘and so’ and ‘but I’ are associated with negative product reviews), which may not be

apparent to human coders. A single outcome variable, either categorical (e.g., positive/negative, subjective/objective) or probabilistic (e.g., 90% positive) produced by a classifier can be easily incorporated into basic statistical models for fMRI data, such as the general linear model; GLM (Poldrack, Mumford, & Nichols, 2011). With supervised machine learning, the single value produced by the linguistic classifier captures a multivariate combination of motivated linguistic features (i.e., words, n-grams, frames, part-of-speech and semantic categories) that best distinguish functional language uses. Example 2 below illustrates the use of the supervised machine learning classification of language produced by participants in an fMRI experiment when they were asked to describe the stimuli they saw during the imaging session to another person. The measures produced by the classifiers can be correlated with neural activity to explore the brain-to-language link.

Steps Involved in Using Quantitative Linguistic Measures

It is beyond the scope of this paper to provide a detailed guide to using quantitative linguistic measures and methods in communication research (we refer readers to useful guides in (Bird, Klein, & Loper, 2009; Gries, 2009; Grimmer & Stewart, 2013; Pustejovsky & Stubbs, 2012). However, we outline broad steps involved in the collection and analysis of language samples for questions combining viewpoints of individual communication processes with larger population-level communication outcomes (see Figure 3).

1. Textual data collection—Across levels of analysis, the first step is to collect and transcribe or extract the text corpus that is relevant to your research question.

At the individual level (e.g., in a neuroimaging study), this can take the form of manipulating stimuli to have specific linguistic properties, or recording linguistic responses to experimental stimuli (as described in Examples 1 and 2). An example of the latter, study participants can be asked to speak or write about specific study-relevant stimuli, which are then mapped onto the mental processes recorded using neuroimaging as those stimuli were initially presented. It is also possible to record participant language recalling specific events, attitudes or values that can then be used as an implicit individual difference measure.

From a practical standpoint, collecting writing samples is more efficient from the data analysis perspective but participants may produce less text (i.e., write less than they would speak). Capturing audio or video responses has many advantages, including capturing both verbal (e.g., hesitations and reframes) and non-verbal (e.g., facial expressions) cues but involves much more work to extract text that can be used as input for quantitative linguistic analysis (i.e., transcription). More specific to the combination with neuroimaging designs that require large numbers of trials, and hence where participants typically see a large number of stimuli (e.g., novel product or TV show ideas, sets of images or PSAs) researchers may want to limit the size/amount of response for each item to reduce participant burden. For instance, participants could be asked to compose a Tweet (140 characters max) about each of the stimuli in turn or asked to describe (spoken) their memory and opinion of each item in 15 to 30 seconds.

At the population level, collection of textual data will involve identifying the source or sources of social interaction relating to the research question (e.g., a specific idea or

campaign that produces language). For instance, following the launch of a new anti-smoking PSA, these sources could include comments posted about the video on YouTube, links and response to the video on Twitter and other social media platforms and engagement with news articles describing the new campaign, and so on. Once these sources are identified, the text data can be downloaded using web-scraping tools and where available using an API (Application Programming Interface) provided by the site, which allows for direct programmatic querying and retrieval of site data in machine-readable form) (Bird et al., 2009; Munzert, Rubba, Meissner, & Nyhuis, 2014).

2. Preprocessing—The next step once your textual corpus has been collected is to clean and preprocess the data so that it is standardized (often referred to as normalization) for quantification and use with a particular linguistic tool (e.g., LIWC (Tausczik & Pennebaker, 2010), WordSmith Tools (Scott, 2012)) or with software for specific NLP tasks such as sentiment analysis. Preprocessing steps will vary depending on the NLP task undertaken, but will usually include: 1. the recognition of units or segments (i.e., words, sentences or conversational turns, paragraphs and sections, etc.), 2. the standardization of text (i.e., transformation to lower case, expansion of abbreviations, removal of links, etc.) and 3. the addition of interpretive metadata or tags (i.e., addition of grammatical class or part-of-speech tagging, marking of entities—such as persons, places, institutions—and the analysis or parsing of larger structures—including clause elements such as subject, verb and object) (Bird et al., 2009; Pustejovsky & Stubbs, 2012). Detailed discussion of preprocessing is beyond the scope of this article, but it is a well-developed area in NLP and a broad range of highly accurate automatic tools and techniques have been developed and are available. Readers are referred to the following overviews of the issues (Cardie & Wilkerson, 2008; Jurafsky & Martin, 2008; Manning et al., 2008; Manning & Schütze, 1999; Pustejovsky & Stubbs, 2012) and to software solutions for preprocessing (Bird et al., 2009; Cunningham et al., 2011; Manning et al., 2014).

3. Derive linguistic measure(s)—The next step is to select and calculate the linguistic measure or measures that capture the psychological or sociological variable of interest. This will be determined by the research questions under investigation. For instance, in the context of social media if you were interested in emotional sharing—Do people tend to share more when they are feeling positive or when they are feeling negative?—an appropriate linguistic measure would be one that captures the ratio of positive to negative word usage in a post (Kramer et al., 2014). A categorical word counting approach using a validated affective lexicon would be a strong candidate in this case. And measures from tools such as LIWC can be used as both independent and dependent or outcome variables in such instances. Where the process in question is likely a combination of multiple complex motivational and psychological processes a supervised machine learning classifier will be a better candidate. An example of such a question might be: How do people subsequently describe ideas that they intend to share when they first encounter them? Texts can be grouped according to whether a participant indicated an intention to share the idea or not and these categories can be used to train a classifier. A limitation of this approach is the need for enough data to allow for separate training data. In a neuroimaging experiment this training data can be collected through running a separate group of participants behaviorally (i.e., without the neuroimaging

component) to reduce cost. In other words, the non-scanned participants could be exposed to the same stimuli and write about them without undergoing a scanning session, and their written language used to train the classifier which would then be applied to language produced by the scanned participants.

The discussion above of quantitative linguistic measures and the two examples presented below are confined to categorical word counting and supervised machine learning. However, there are many other approaches that should also be considered. For example, unsupervised language classification approaches are well suited for more exploratory analysis of linguistic data when grouping categories and/or prominent features are not known a priori. Methods including latent semantic analysis (Babcock, Ta, & Ickes, 2014; Dumais, 2004) and latent dirichlet allocation (used for topic modeling) (Blei, Ng, & Jordan, 2003; Landauer & Dumais, 1997; Schwartz et al., 2013) can be used for this purpose, but are beyond the scope of this paper.

4. Test models using linguistic measure(s) as DV and/or IVs—The final step is to incorporate the derived linguistic measures into predictive models formulated in response to the research questions under investigation. For instance, in the analysis described in the first example below, LIWC social process scores calculated from the in-scanner stimuli (novel product descriptions) serve as an independent variable that predict the degree of neural activity in the temporoparietal junction (TPJ) across people when they are exposed to the ideas (when their self-reported intention to share is controlled for) (O'Donnell et al., In press). Equally, it is possible to treat the linguistic measure(s) as dependent variables as illustrated in example 2 where neural activity in the right TPJ predicts degree of positive evaluative language (Falk, O'Donnell, & Lieberman, 2012). Readers are referred to other contributions in this special issue that discuss and demonstrate the use of neuroimaging data within a predictive framework (Falk et al., this issue; Weber, this issue).

Integrating linguistic data with neuroimaging data

There are at least three ways that we can examine the interplay between language and the neurocognitive processes underlying communication. These include considering (1) language as input to individual-level communication processes (2) language as output to individual-level communication processes (Figure 2), and (3) language as a carrier of ideas between people (i.e., a tool to understand how ideas spread).

First, language-based stimuli are frequently used in fMRI tasks allowing for the interaction of task focus (e.g., perspective taking, valuation, etc.) and the features of language in the stimuli to be examined in relation to the resulting neural activity. For example, researchers might be interested in examining the effects of being exposed to language that contains large numbers of first person pronouns. They could examine whether the neural activity in self-related processing regions observed in response to language containing large proportions of first person pronouns mediates the relationship between the stimuli and the degree of self-focus observed in measures collected after the scan. Further, the linguistic features of stimuli can be manipulated to address questions relevant to social phenomena of communicative relevance such as priming, valence framing, hesitation, and so on (Bergen & Wheeler, 2010;

Bergen & Binsted, 2010; Blankenship & Craig, 2006; Craig & Blankenship, 2011; Hosman, Huebner, & Siltanen, 2002; Hosman & Siltanen, 2006, 2011).

Secondly, after undergoing scans participants can be asked to produce language samples (either spoken or written) reflecting upon the stimuli seen. The language of these samples can be analyzed and correlated with neural activity during exposure. (In a variant of this approach, Saxbe, Yang, Borofsky, & Immordino-Yang, 2013 collect language responses to stimuli prior to exposure to the same stimuli in scanner).

With both approaches, it is possible to take an exploratory approach mapping neural regions throughout the brain that are associated with a particular sociopsychological-linguistic correlate, or testing relationships with more targeted, a priori hypothesized regions of interest for more specific theory testing. Importantly, stimuli presented to individuals in neuroimaging experiments can be manipulated by the research team (e.g., to map brain regions associated with theoretically constructed language patterns) or in a data driven manner using linguistic stimuli gathered at the individual or population level (e.g., samples of language collected in response to stimuli post-scan, or equivalent language collected in response to similar stimuli at large scales, such as ideas or tweets that spread successfully and those that did not).¹

Third, language can be treated as a mediator of the relationship between neural responses to a set of ideas and effective transmission of those ideas. This framing, in particular, suggests that there is a high degree of applied potential for each approach. The identification of how specific combinations of language presented as input during a social/psychological task are associated with increased neural activity in specific neural regions of interest could lead to significant improvements in the creation of persuasive measures in domains such as health or political communication. Further, when language is treated as an outcome variable predicted by activity in specific neural regions, the identification of particular language patterns may be a proximal indication of certain broad types of neural activity.² Likewise, establishing neural patterns that are predictive of successful transmission of ideas may lend insight to what makes certain communicators effective and aid in developing interventions that would increase this capability in others. This would then be highly significant because language samples are easy to collect and analyze at scale while neuroimaging data are not.

¹It should be noted that it is also possible to treat linguistic data at either level of analysis as an outcome variable. Linguistic output (again, either categorized along theoretical or practical lines) can be predicted by a combination of neural responses collected during exposure in a relatively small group of participants. Such data can also be combined with other predictor variables, such as, self-report ratings relating to the stimuli (e.g. interest, intention, argument strength, etc.) and other individual difference measures (e.g. personality traits, social network variables, level of media exposure and usage, etc.).

²Such predictions would, of course, be contingent upon establishing robust associations between neural activation and specific linguistic patterns. Future work is required to establish such baselines of neural activity associated with a range of linguistic constructions independent of specific neuroimaging tasks designed to examine specific cognitive processes. Some of this work has been carried out in neurolinguistic and psycholinguistic studies, which have for instance shown consistent activation of specific areas associated with various linguistic levels (e.g. phonetic, lexical, syntactic, semantic and pragmatic) in both the production and reception of language (Awad, Warren, Scott, Turkheimer, & Wise, 2007; Menenti, Gierhan, Segaert, & Hagoort, 2011; Menenti, Petersson, & Hagoort, 2012; Silbert, Honey, Simony, Poeppel, & Hasson, 2014).

Triangulation of Mechanisms at Individual and Population Levels

Neuroimaging and linguistic data both represent implicit signals of individuals' psychological responses to communicative events and processes, but measure different points in the process and are amenable to different levels of analysis. Neuroimaging data can provide access to in-the-moment psychological responses to stimuli without the interruption of self-report measures (e.g., checking whether an individual has an intention to share an idea) and can provide additional explanatory power beyond these measures (Berkman, Falk, & Lieberman, 2011; Cascio et al., 2013; Falk, Berkman, Mann, Harrison, & Lieberman, 2010). By contrast, linguistic analysis focuses on the social inputs (e.g., language contained in stimuli) and outputs (e.g., language produced in response to stimuli). Both neural data and patterns of language use have been linked to psychological states, leading to the development of a range of analytical approaches and tools in each sub area. Linking neurocognitive mechanisms and linguistic correlates affords a number of advantages, including triangulating the underlying mechanisms of successful communication at multiple levels of analysis (Figure 3). Two starting points for creating such linkages are (see Figure 2):

- 1) Mapping the neurocognitive mechanisms that underlie responses to different types of language (e.g., in mass media or interpersonal communication). Example 1 illustrates this approach.
- 2) Mapping the psychological mechanisms that precede specific linguistic output. Example 2 provides the foundation for investigating and establishing systematic links between neural activity in response to social and communicative stimuli and proximal linguistic output (micro-level samples of language collected in an experimental setting). This method also offers promising implications as a method for further linking individual and population levels (where macro-level language samples can be acquired).

In both of these approaches, mapping neural processes associated with different language patterns can provide convergent evidence for psychological mechanisms supporting effective communication at multiple levels of analysis if comparable language samples can be collected at micro and macro scales. The integration of neuroimaging with language tools is critical because despite strong links between individual traits (e.g., personality types, gender and age), and associated functional linguistic outcomes (e.g., speaking positively or negatively about an idea), linguistic tools alone cannot directly observe and measure the psychological mechanisms that lead people to process and produce language with these methods. By contrast, a substantial body of neuroimaging research has mapped the neural correlates of a wide range of social, cognitive and affective processes, but neuroimaging tools alone cannot move beyond the confines of the laboratory.

Examining the Mechanisms Underlying the Spread of Ideas

Beyond uncovering mechanisms that may be common to different levels of analysis, bridging neural and linguistic analysis also stands to link micro and macro levels of analysis since large-scale trends often spread between individuals. Language is traceable/recordable and can be turned into a written record of human communication, making it a prime target

for communication scientists and neuroscientists interested in how ideas and other cultural units spread. In parallel, with the growing availability of large samples of language data (i.e., social media texts in response to a particular event, e.g., the Arab Spring) it is possible to follow ideas as they go viral and to examine linguistic patterns associated with such phenomena. Examining neural responses to a theoretically chosen subset of the larger-scale linguistic data in a relevant smaller sample may help illuminate why some ideas spread and others do not; whether the neural responses to these language samples are the same as those identified as preceding successful communication at the individual level; and how different types of language interact with psychological, cultural or demographic individual differences, etc. to shape the speed and depth of cascades, and how individuals decide what to pass on. It is thereby possible to create deeper intellectual links between behavior observed at a large scale and the underlying mechanisms associated with those behaviors. To this end, initial research has begun to map neural processes and linguistic features associated with the successful spread of ideas (Falk, Morelli, Welborn, Dambacher, & Lieberman, 2013; Falk et al., 2012). In turn, through an iterative research approach, once those mechanisms are mapped, the maps can be leveraged to subsequently use neural activity collected at the individual level to predict large-scale linguistic outcomes.

Examples of Integrating Quantitative Linguistic and Functional Neuroimaging

The following examples illustrate components from each of the approaches described above. Potential applications that would link individual level neural and linguistic responses with large-scale population level linguistic data are briefly discussed in the final section of this paper.

Example 1. Categorical Analysis of Input Stimuli and Associated Neural Activity

In a larger study investigating word-of-mouth processes (Falk et al., 2013, Falk, O'Donnell and Lieberman 2012 report on a different task from this dataset) participants were shown descriptions of 24 products during an fMRI session and asked to indicate whether they would recommend each product to a friend. The product descriptions consisted of positive recommendations for the product they described but varied in terms of their use of certain linguistic features, such as first person versus third person pronouns. Compare the two descriptions in the table below for the Flowbee (A) and the Periodic Element Rings (B) shown in Table 1.

Both descriptions achieve their purpose of recommending the product to the reader in a positive and enthusiastic manner, and provide enough information about the product to allow the reader to form an opinion. Notice, however, the complete lack of first person reference in the Flowbee description in comparison to the dominant use through out the Periodic Element Rings description. The recommendation of the Periodic Elements Rings is strongly implied but not directly stated, namely it does not say 'These will look great on your wall and help you ace your chemistry test!'. Similarly, the Flowbee description is explicit in the recommendation of the product without clearly stating the individual's personal experience with the product. These differences are seen quantitatively comparing the scores produced

by LIWC (Pennebaker et al., 2007) under the Self References (I, me, my) and Social word categories (you, us, friend, talk). Overall descriptions scoring high on the Self References category tend to score low on the Social category ($r=-0.37$ $p=0.078$).

The resulting neural patterns associated with each of these types of language (self references and social processes) showed a subset of neural regions associated with mentalizing or social cognition. Greater use of first person references in the product description was associated with increased activation in the dorsomedial prefrontal cortex (DMPFC) (Figure 4). This suggests that more explicit opinions, framed in the first person, appear to most robustly activate neural systems that are associated with perspective taking ('theory of mind') in the brain of the listener. It is possible that such first person statements more clearly reveal the mental states of the speaker and activate neural systems that process social knowledge. Increased use of words from the LIWC Social Processes category in the product recommendations were also associated with more activity throughout a network of neural regions associated with thinking about the mental states of others (Saxe & Powell, 2006), specifically the bilateral TPJ (but less in the DMPFC than the self references category). The use of categorical word counting in this way (specifically words from the LIWC Social Processes category) suggests a link between the use of social language and processes of social cognition as individuals evaluated ideas for subsequent sharing (see O'Donnell et al., In press for more details).

These data illustrate the combination of the linguistic quantification of stimuli (language input) with recorded neural activity from subjects during exposure to these stimuli. Specifically they provide initial evidence about a discoverable link between functional patterns of language used to frame ideas in a specific social context—to the extent it can be simulated in a controlled experimental task—and the range of resulting neural and psychological responses.

Example 2. Classification of Post-scan Language Associated with Neural Activity

In a separate task, scanned participants in the same study were directed to imagine they were acting as interns in a TV production company (Falk et al., 2012). After viewing show ideas in the scanner they were videotaped describing each show and their language was transcribed. The transcriptions were classified using a sentiment analysis (SA) algorithm trained on texts from a movie review corpus (see Falk et al., 2012 for details of the SA classifiers). The SA classifiers return both category labels (e.g., neutral, for descriptive texts, or positive/negative for evaluative texts) and classification probabilities.

Table 2 shows two different post-scan reviews of one of the shows called Beauty Queens. Both texts were classified as highly evaluative (with a 0.99 probability by the evaluative language classifier). Text A is clearly positive about the show and contains an explicit recommendation, whereas B is very negative about the idea. In this study, an association was observed between higher positivity scores and increased neural activity in two clusters in the medial prefrontal cortex (MPFC) and the posterior cingulate cortex (PCC) and precuneus (PC) (Figure 5A). These are regions that are consistently associated with self-related processes and positive value (Lieberman, 2010). Although brain regions support multiple psychological functions, and hence this explanation is one of many, the neural activity

observed is consistent with the hypothesis that for show ideas where participants experienced increased liking and self-relevance (e.g., 'I like this idea' or 'this idea is relevant to me')—they tended to use language patterns in their subsequent description of those shows that is associated with recommendation reviews, that is, highly positive.

Next, the two scores from the classifiers were combined into a single score (evaluative*positivity) to capture reviews that are strongly evaluative and high on polarity (such as those in the Table 2). Figure 5B shows the large cluster in the right TPJ, a region associated with processing of mentalizing or 'theory of mind' (Saxe & Powell, 2006), resulting from this analysis. Again with caveats related to reverse inference (see Falk et al., this volume; Weber, this volume), the study team hypothesized that when subjects were evaluating show ideas as they viewed them, a tendency to consider the social value of the idea (i.e., 'will this idea appeal to others?' or 'I think this would appeal to lots of people/these types of people') may have positioned the participants to later use language in their subsequent description of these shows that was both highly evaluative and highly positive (i.e., a strong recommendation of the idea).

These data show how we can combine neural activity from subjects during exposure to ideas with subsequent free-form reflection about these ideas. They provide initial evidence about a discoverable link between functional patterns of language and broad socio-psychological processes that can be examined using neuroimaging and other neuroscience methods. Although the combination of language and neuroimaging data is in its infancy, the examples here lay the foundation for future work that achieves the more ambitious goals outlined above. For example, a parallel experiment to that described in example two could be conducted on actual TV show pilots, with parallel data collection of large-scale audience reception as indexed by social media response. This would include the full spectrum of analysis from individual to population level illustrated in Figure 3.

Summary

Functional linguistic models posit a systematic link between language form and the functions for which language is used (Croft & Cruse, 2004; Halliday & Matthiessen, 2012). That is, language form and patterns of language usage—which can be quantified—vary systematically based on the communicative task in which the language user is engaged (Biber, 1991; Biber & Conrad, 2009). Such models assume that language is a social, communicative tool and that there is a systematic (and therefore quantifiable) relationship between linguistic FORM and communicative FUNCTION. However, while a consistent link between linguistic patterns and psychological states and traits has been empirically demonstrated and rigorously applied using tools such as LIWC, this cannot be validated without the use of other methods.

In this paper, we argue that the combination of neural and linguistic measures will allow insight into both individual and population-level psychology that would not be possible using either method in isolation. We have suggested three possible ways in which linguistic analysis can contribute to the creation of experimental stimuli, processing of experimental outputs in individual study participants, and creation of shared understanding between dyads

or in larger groups. We have described examples illustrating key points and provided additional notes regarding specific linguistic tools that may be most amenable to the advocated approach. We have also explicated some ways in which the approach has the potential of linking experimental outputs at the individual level with larger scale phenomena that have been the focus of much recent attention in large-scale text mining (Figure 2). An example of this would be a study that identifies tweets that spread and go viral in a particular context compared with similar ones that do not and explores the neural correlates of the linguistic features of each type of message in order to understand psychological and neurocognitive mechanisms related to the spread of ideas.

As with any approach, however, there are limitations to the methods presented that should be noted. First the idea that even the shortest piece of text, produced in a rich social context to accomplish a particular function, can be reduced to a series of numbers is clearly incomplete. For longer texts of more than a single sentence, the notion that a single score captures the sense of each sentence in the text is also limited. Consider a product review, from a website like [Amazon.com](https://www.amazon.com) for instance. Frequently positive facts and observations will be interspersed with things the reviewer did not like or sees as limitations. For example, overall a reviewer may like their mobile phone and give it a 4-star rating but might spend a quarter of the review discussing disappointing battery life and a lack of screen responsiveness in landscape mode in certain apps. Category word count approaches capture this variance through a mean or cumulative score. Classification approaches that produce a continuous score may be more sensitive to these factors as far as they are found in the training set. Recent work in opinion mining and sentiment analysis focuses on the local, clause by clause recognition of emotive and opinion features leading to a time series like distribution of values (Pang & Lee, 2008).

In addition, a full discussion of linguistic theory supportive of the proposed framework is beyond the scope of this paper (interested readers are referred to Biber & Conrad, 2009; Bybee, 2010; Ibbotson, 2013). As such, we have outlined only a few basic approaches to the quantification of language. Language usage and the resulting form-functional patterns, however, vary at multiple, interacting levels. For instance, it is possible to examine and quantify the difference between spoken and written language, between broad types or genres of language (i.e., academic language, literary and fictional language, the language used in news reporting and magazines, and so on), between different speakers in different regions of a country (i.e., dialect), of different demographics (i.e., gender, social class, subculture) and even between different individual speakers (i.e., idiolect). These levels of language and their quantification have been explored within a number of fields such as sociolinguistics, corpus linguistics and stylistics. We envision that once basic relationships between simpler forms of language analysis and neural function have been established, these more intricate relationships can be explored and integrated. More broadly, future development of the proposed methodology will require that researchers who choose these methods are prepared to engage with the considerable literature and body of studies within the fields of neurolinguistics and psycholinguistics, and hence we note the benefits of interdisciplinary research (i.e., collaborations between neuroscientists and linguists).

Finally, the data discussed in the examples make use of fMRI which although powerful has significant limitations for both the presentation of language—because the noise of the machine limits hearing clarity—and the production of language—because head movement must be restricted and it is challenging to capture speech. Further, the fMRI data used in the studies discussed in the examples here reduces a large number of temporally and spatially distributed points to a single measure of mean activation at each voxel or across the voxels in a region of interest. Further developments should make use of time series approaches that could link neural activity across exposure to text, aligning it with the sentiment at each point. There are limitations to the temporal resolution with fMRI (however cf. Hasson, Nir, Levy, Fuhrmann, & Malach, 2004; Hasson, Ghazanfar, Galantucci, Garrod, & Keysers, 2012; Stephens, Silbert, & Hasson, 2010) but techniques like Functional Near Infrared Spectroscopy (fNIRS) and EEG capture a much faster signal (Amodio, Bartholow, & Ito, 2013; Cutini & Brigadoi, 2014; Hoshi, 2007; Huppert, Diamond, Franceschini, & Boas, 2009). It should thereby be possible to align a neural and linguistic time series. As stated before, the dangers of reverse inference are a central concern in the linking of neuroimaging data and experimental outcomes and should receive especial consideration as neuroimaging establishes itself as a more prominent role within communication science (see Weber, this issue for more discussion).

Quantitative linguistic measures provide communication scholars access to scalable approaches in analyzing the mechanisms underlying successful communication at both the micro- (individual) and macro- (population) levels. Considerable application of such tools has already been seen in studies of automated content analysis, information diffusion and social media. In this paper we have argued for the combination of such approaches with those from neuroimaging. This allows the triangulation of individual level outcomes with population level outcomes beyond the lab, thus increasing both the internal and external validity of a given study. A core element of this linkage is the ability to apply the same methods of language quantification at all points along the individual to population level continuum or to link the levels through the spread of ideas from one level to the next. Although this work is in its infancy, this leaves much room for creative collaborations between communication scholars, linguists and neuroscientists to expand the horizons of all three disciplines.

Acknowledgements

The authors are grateful to Jason Coronel, Lynda Lin and Jiaying Liu for helpful feedback on drafts of this manuscript. This work was supported in part by NIH-1 DP2 DA035156-01 (PI: Falk).

References

- Aiden E, , Michel J-B. *Uncharted: Big Data as a Lens on Human Culture* Riverhead; New York: 2013
- Amodio DM, Bartholow BD, Ito TA. Tracking the dynamics of the social brain: ERP approaches for. *Social Cognitive and Affective Neuroscience*. 2013;nst177. doi:10.1093/scan/nst177.
- Asur S, Huberman BA. Predicting the Future with Social Media. *Applied Energy*. 2013; 112:1536–1543. doi:10.1016/j.apenergy.2013.03.027.
- Awad M, Warren JE, Scott SK, Turkheimer FE, Wise RJS. A Common System for the Comprehension and Production of Narrative Speech. *The Journal of Neuroscience*. 2007; 27(43):11455–11464. doi: 10.1523/JNEUROSCI.5257-06.2007. [PubMed: 17959788]

- Babcock MJ, Ta VP, Ickes W. Latent Semantic Similarity and Language Style Matching in Initial Dyadic Interactions. *Journal of Language and Social Psychology*. 2014; 33(1):78–88. doi: 10.1177/0261927X13499331.
- Baccianella S, Esuli A, Sebastiani F. SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining.. Presented at the Proceedings of the Seventh Conference on International Language Resources and Evaluation (LREC'10); European Language Resources Association (ELRA); 2010 Retrieved from http://www.lrec-conf.org/proceedings/lrec2010/pdf/769_Paper.pdf
- Baek YM, Cappella JN, Bindman A. Automating Content Analysis of Open-Ended Responses: Wordscores and Affective Intonation. *Communication Methods and Measures*. 2011; 5(4):275–296. doi:10.1080/19312458.2011.624489. [PubMed: 25558294]
- Balasubramanyan R, Routledge BR, Smith NA. From Tweets to Polls : Linking Text Sentiment to Public Opinion Time Series. 2010
- Bergen B, Binsted K. Embodied Grammar and Humor. *Cognitive Linguistics and Humor*. 2010
- Bergen B, Wheeler K. Grammatical aspect and mental simulation. *Brain and Language*. 2010; 112(3): 150–158. doi:10.1016/j.bandl.2009.07.002. [PubMed: 19656554]
- Berkman ET, Falk EB. Beyond brain mapping: Using the brain to predict real-world outcomes. *Current Directions in Psychological Science*. 2013; 22:45–55. [PubMed: 24478540]
- Berkman ET, Falk EB, Lieberman MD. In the trenches of real-world self-control: neural correlates of breaking the link between craving and smoking. *Psychol Sci*. 2011; 22:498–506. doi: 10.1177/0956797611400918. [PubMed: 21378368]
- Biber D. Variation Across Speech and Writing Cambridge University Press; 1991
- Biber D, Conrad S. Register, Genre, and Style Cambridge University Press; 2009
- Bird S, Klein E, Loper E. Natural Language Processing with Python O'Reilly Media; 2009
- Blankenship KL, Craig TY. Rhetorical Question Use and Resistance to Persuasion: An Attitude Strength Analysis. *Journal of Language and Social Psychology*. 2006; 25(2):111–128. doi: 10.1177/0261927X06286380.
- Blei DM, Ng AY, Jordan MI. Latent Dirichlet Allocation. *J. Mach. Learn. Res*. 2003; 3:993–1022.
- Bradley M, Lang P. Affective norms for english words (anew): Stimuli, instruction manual and affective ratings. (Technical report c-1,) University of Florida; Gainesville, FL: 1999
- Burnap P, Rana OF, Avis N, Williams M, Housley W, Edwards A, Sloan L. Detecting tension in online communities with computational Twitter analysis. *Technological Forecasting and Social Change*. 2013 doi:10.1016/j.techfore.2013.04.013.
- Bybee. Language, Usage and Cognition Cambridge University Press; 2010
- Cardie C, Wilkerson J. Text Annotation for Political Science Research. *Journal of Information Technology & Politics*. 2008; 5(1):1–6. doi:10.1080/19331680802149590.
- Cascio CN, Dal Cin S, Falk EB. Health communications: Predicting behavior change from the brain.. In: Hall P, editor *Social Neuroscience and Public Health: Foundations of an Emerging Discipline* Springer; 2013 5761
- Craig TY, Blankenship KL. Language and Persuasion: Linguistic Extremity Influences Message Processing and Behavioral Intentions. *Journal of Language and Social Psychology*. 2011; 30(3): 290–310. doi:10.1177/0261927X11407167.
- Croft W, Cruse DA. *Cognitive Linguistics* Cambridge University Press; 2004
- Cunningham H, Maynard D, Bontcheva K, Tablan V, Aswani N, Roberts I, Peters W. Text Processing with GATE (Version 6) 2011 Retrieved from <http://tinyurl.com/gatebook>
- Cutini S, Brigadoi S. Unleashing the future potential of functional near-infrared spectroscopy in brain sciences. *Journal of Neuroscience Methods*. 2014; 232:152–156. doi:10.1016/j.jneumeth.2014.05.024. [PubMed: 24880046]
- Danescu-Niculescu-Mizil C, Gamon M, Dumais ST. Mark my words!: linguistic style accommodation in social media. *WWW'11*. 2011:745–754.
- Dodds PS, Harris KD, Kloumann IM, Bliss CA, Danforth CM. Temporal Patterns of Happiness and Information in a Global Social Network: Hedonometrics and Twitter. *PLoS ONE*. 2011; 6(12):e26752. doi:10.1371/journal.pone.0026752. [PubMed: 22163266]

- Dumais ST. Latent semantic analysis. *Annual Review of Information Science and Technology*. 2004; 38(1):188–230. doi:10.1002/aris.1440380105.
- Ellis NC, O'Donnell MB. Statistical construction learning: Does a Zipfian problem space ensure robust language learning?. In: Rebuschat P, Williams J, editors *Statistical Learning and Language Acquisition* Walter de Gruyter; 2012
- Falk EB, Berkman ET, Mann T, Harrison B, Lieberman MD. Predicting persuasion-induced behavior change from the brain. *Journal of Neuroscience*. 2010; 30:8421–4. doi:10.1523/JNEUROSCI.0063-10.2010. [PubMed: 20573889]
- Falk EB, Hyde LW, Mitchell C, Faul J, Gonzalez R, Heitzeg MM, Schulenberg J. What is a representative brain? Neuroscience meets population science. *Proceedings of the National Academy of Sciences*. 2013; 110(44):17615–17622. doi:10.1073/pnas.1310134110.
- Falk EB, Morelli SA, Welborn BL, Dambacher K, Lieberman MD. Creating buzz: the neural correlates of effective message propagation. *Psychological Science*. 2013; 24(7):1234–1242. doi:10.1177/0956797612474670. [PubMed: 23722983]
- Falk EB, O'Donnell MB, Lieberman MD. Getting the word out: neural correlates of enthusiastic message propagation. *Frontiers in Human Neuroscience*. 2012; 6:313. doi:10.3389/fnhum.2012.00313. [PubMed: 23189049]
- Fast LA, Funder DC. Personality as manifest in word use: Correlations with self-report, acquaintance report, and behavior. *Journal of Personality and Social Psychology*. 2008; 94(2):334–346. doi:10.1037/0022-3514.94.2.334. [PubMed: 18211181]
- Gayo-Avello D. A Meta-Analysis of State-of-the-Art Electoral Prediction From Twitter Data. *Social Science Computer Review*. 2013; 31(6):649–679. doi:10.1177/0894439313493979.
- Godbole N, Srinivasaiah M, Skiena S. Large-scale sentiment analysis for news and blogs. *ICWSM'07*. 2007
- Gonzales AL, Hancock JT, Pennebaker JW. Language Style Matching as a Predictor of Social Dynamics in Small Groups. *Communication Research*. 2010; 37(1):3–19. doi:10.1177/0093650209351468.
- Goode J, Robinson JD. Linguistic Synchrony in Parasocial Interaction. *Communication Studies*. 2013; 64(4):453–466. doi:10.1080/10510974.2013.773923.
- Gottschalk L, A.. *Content Analysis of Verbal Behavior: New Findings and Clinical Applications* Lawrence Erlbaum Associates; Hillsdale, NJ: 1995
- Gottschalk LA, Gleser GC. *The measurement of psychological states through the content analysis of verbal behavior* University of California Press; Berkeley: 1969
- Gries ST. *Quantitative corpus linguistics with R: a practical introduction* Routledge; 2009
- Grimmer J, Stewart BM. Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts. *Political Analysis*. 2013; 21(3):267–297. doi:10.1093/pan/mps028.
- Gruhl D, Guha R, Liben-Nowell D, Tomkins A. *Proceedings of the 13th International Conference on World Wide Web ACM*; New York, NY, USA: 2004 *Information Diffusion Through Blogspace.*; 491501 doi:10.1145/988672.988739
- Gruzd A. Emotions in the Twitterverse and Implications for User Interface Design. *AIS Transactions on Human-Computer Interaction*. 2013; 5(1):42–56.
- Gruzd A, Doiron S, Mai P. Is Happiness Contagious Online? A Case of Twitter and the 2010 Winter Olympics. 2011 44th Hawaii International Conference on System Sciences (HICSS). 2011:1–9. doi:10.1109/HICSS.2011.259.
- Halliday MAK, Matthiessen C. *Halliday's Introduction to Functional Grammar* Taylor & Francis Limited; 2012
- Hasson U, Ghazanfar AA, Galantucci B, Garrod S, Keysers C. Brain-to-brain coupling: a mechanism for creating and sharing a social world. *Trends in Cognitive Sciences*. 2012; 16(2):114–121. doi:10.1016/j.tics.2011.12.007. [PubMed: 22221820]
- Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R. Intersubject synchronization of cortical activity during natural vision. *Science*. 2004; 303:1634–40. doi:10.1126/science.1089506. [PubMed: 15016991]
- Hoey M. *Lexical Priming: A New Theory of Words and Language* Routledge/Taylor and Francis Group; London, UK: 2005

- Holtgraves T, Lasky B. Linguistic Power and Persuasion. *Journal of Language and Social Psychology*. 1999; 18(2):196–205. doi:10.1177/0261927X99018002004.
- Hoshi Y. Functional near-infrared spectroscopy: current status and future prospects. *Journal of Biomedical Optics*. 2007; 12(6):062106. doi:10.1117/1.2804911. [PubMed: 18163809]
- Hosman LA, Huebner TM, Siltanen SA. The impact of power-of-speech style, arguments, and need for cognition on impression formation, cognitive responses and persuasion. *Journal of Language and Social Psychology*. 2002; 21:361–379.
- Hosman LA, Siltanen SA. Powerful and Powerless Language Forms: Their Consequences for Impression Formation, Attributions of Control of Self and Control of Others, Cognitive Responses, and Message Memory. *Journal of Language and Social Psychology*. 2006; 25:33–46.
- Hosman LA, Siltanen SA. Hedges, Tag Questions, Message Processing, and Persuasion. *Journal of Language and Social Psychology*. 2011; 30(3):341–349. doi:10.1177/0261927X11407169.
- Huppert TJ, Diamond SG, Franceschini MA, Boas DA. HomER: a review of time-series analysis methods for near-infrared spectroscopy of the brain. *Applied Optics*. 2009; 48(10):D280–D298. [PubMed: 19340120]
- Ibbotson P. The scope of usage-based theory. *Frontiers in Language Sciences*. 2013; 4:255. doi: 10.3389/fpsyg.2013.00255.
- Ireland ME, Slatcher RB, Eastwick PW, Scissors LE, Finkel EJ, Pennebaker JW. Language Style Matching Predicts Relationship Initiation and Stability. *Psychological Science*. 2011; 22(1):39–44. doi:10.1177/0956797610392928. [PubMed: 21149854]
- Jindal N, , Liu B. Proceedings of the 2008 International Conference on Web Search and Data Mining ACM; New York, NY, USA: 2008 Opinion Spam and Analysis.; 219230 doi: 10.1145/1341531.1341560
- Jurafsky D, , Martin JH. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics and Speech Recognition* 2nd ed.. Prentice Hall; Englewood Cliffs, NJ: 2008
- Kahn JH, Tobin RM, Massey AE, Anderson JA. Measuring emotional expression with the Linguistic Inquiry and Word Count. *The American Journal of Psychology*. 2007; 120(2):263–286. [PubMed: 17650921]
- Kramer ADI, Guillory JE, Hancock JT. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*. 2014; 111(24):8788–8790. doi:10.1073/pnas.1320040111.
- Landauer TK, Dumais ST. A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*. 1997; 104(2):211–240. doi:10.1037/0033-295X.104.2.211.
- Lieberman MD. Social cognitive neuroscience.. In: Fiske S, Gilbert D, , Lindzey G, editors *Handbook of Social Psychology* 5th ed.. McGraw-Hill; New York, NY: 2010 143193
- Manning CD, , Raghavan P, , Schütze H. *Introduction to Information Retrieval* Cambridge University Press; New York, NY, USA: 2008
- Manning CD, , Schütze H. *Foundations of Statistical Natural Language Processing* MIT Press; Cambridge, MA, USA: 1999
- Manning CD, , Surdeanu M, , Bauer J, , Finkel J, , Bethard SJ, , McClosky D. The Stanford CoreNLP Natural Language Processing Toolkit.; Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations 2014 5560 Retrieved from <http://www.aclweb.org/anthology/P/P14/P14-5010>
- McEnery T, , Hardie A. *Corpus Linguistics: Method, Theory and Practice* Cambridge University Press; Cambridge: 2011
- Menenti L, Gierhan SME, Segaert K, Hagoort P. Shared language: overlap and segregation of the neuronal infrastructure for speaking and listening revealed by functional MRI. *Psychological Science*. 2011; 22(9):1173–1182. doi:10.1177/0956797611418347. [PubMed: 21841148]
- Menenti L, Petersson KM, Hagoort P. From Reference to Sense: How the Brain Encodes Meaning for Speaking. *Frontiers in Psychology*. 2012; 2 doi:10.3389/fpsyg.2011.00384.
- Mishne G. Predicting movie sales from blogger sentiment. In *AAAI 2006 Spring Symposium on Computational Approaches to Analysing Weblogs (AAAI-CAAW. 2006*

- Mitchell L, Frank MR, Harris KD, Dodds PS, Danforth CM. The Geography of Happiness: Connecting Twitter Sentiment and Expression, Demographics, and Objective Characteristics of Place. *PLoS ONE*. 2013; 8(5):e64417. doi:10.1371/journal.pone.0064417. [PubMed: 23734200]
- Munzert S, , Rubba C, , Meissner P, , Nyhuis D. Automated Data Collection with R: A Practical Guide to Web Scraping and Text Mining Wiley; New York: 2014
- O'Donnell MB. The adjusted frequency list: A method to produce cluster-sensitive frequency lists. 2011; 35:135–169.
- O'Donnell MB, Falk EB, Lieberman MD. Social in, social out: How the brain responds to social language with more social language. *Communication Monographs*. In press.
- O'Donnell MB, Römer U, Ellis NC. The development of formulaic sequences in first and second language writing: Investigating effects of frequency, association, and native norm. *International Journal of Corpus Linguistics*. 2013; 18(1):83–108. doi:10.1075/ijcl.18.1.07odo.
- Pang B, Lee L. Opinion Mining and Sentiment Analysis. 2008
- Pang B, Lee L, Vaithyanathan S. Thumbs up? Sentiment Classification using Machine Learning Techniques. IN *PROCEEDINGS OF EMNLP*. 2002:79–86.
- Pennebaker JW. The secret life of pronouns: What our words say about us Bloomsbury Press; 2011
- Pennebaker JW, , Chung CK, , Ireland M, , Gonzales A, , Booth RJ. The development and psychometric properties of LIWC2007 LIWC. net; Austin, TX: 2007 [February 1, 2010]
- Pennebaker JW, Mehl MR, Niederhoffer KG. Psychological aspects of natural language. use: our words, our selves. *Annual Review of Psychology*. 2003; 54:547–577. doi:10.1146/annurev.psych.54.101601.145041.
- Poldrack RA, , Mumford JA, , Nichols TE. Handbook of Functional MRI Data Analysis Cambridge University Press; 2011
- Pustejovsky J, , Stubbs A. Natural Language Annotation for Machine Learning O'Reilly Media; 2012
- Ritter RS, Preston JL, Hernandez I. Happy Tweets: Christians Are Happier, More Socially Connected, and Less Analytical Than Atheists on Twitter. *Social Psychological and Personality Science*. 2013:1948550613492345. doi:10.1177/1948550613492345.
- Saxe DE, Yang X-F, Borofsky LA, Immordino-Yang MH. The embodiment of emotion: language use during the feeling of social emotions predicts cortical somatosensory activity. *Social Cognitive and Affective Neuroscience*. 2013; 8(7):806–812. doi:10.1093/scan/nss075. [PubMed: 22798396]
- Saxe R, Powell LJ. It's the Thought That Counts Specific Brain Regions for One Component of Theory of Mind. *Psychological Science*. 2006; 17(8):692–699. doi:10.1111/j.1467-9280.2006.01768.x. [PubMed: 16913952]
- Schwartz HA, Eichstaedt JC, Kern ML, Dziurzynski L, Ramones SM, Agrawal M, Ungar LH. Personality, Gender, and Age in the Language of Social Media: The Open-Vocabulary Approach. *PLoS ONE*. 2013; 8(9):e73791. doi:10.1371/journal.pone.0073791. [PubMed: 24086296]
- Scott M. WordSmith Tools (Version 6.0) Lexical Analysis Software; Liverpool: 2012
- Silbert LJ, Honey CJ, Simony E, Poeppel D, Hasson U. Coupled neural systems underlie the production and comprehension of naturalistic narrative speech. *Proceedings of the National Academy of Sciences*. 2014; 111(43):E4687–E4696. doi:10.1073/pnas.1323812111.
- Sinclair JM. Trust the Text: Language, Corpus and Discourse Routledge; 2004
- Stephens GJ, Silbert LJ, Hasson U. Speaker–listener neural coupling underlies successful communication. *Proceedings of the National Academy of Sciences*. 2010; 107(32):14425–14430. doi:10.1073/pnas.1008662107.
- Stone PJ, , Dunphy DC, , Smith MS, , Ogilvie DM. The General Inquirer: A Computer Approach to Content Analysis MIT Press; Cambridge, MA: 1966
- Tausczik YR, Pennebaker JW. The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. *Journal of Language and Social Psychology*. 2010; 29(1):24–54. doi:10.1177/0261927X09351676.
- Tumasjan A, Sprenger TO, Sandner PG, Welpe IM. Election Forecasts With Twitter How 140 Characters Reflect the Political Landscape. *Social Science Computer Review*. 2011; 29(4):402–418. doi:10.1177/0894439310386557.

Yarkoni T. Personality in 100,000 Words: A large-scale analysis of personality and word use among bloggers. *Journal of Research in Personality*. 2010; 44(3):363–373. doi:10.1016/j.jrp.2010.04.001. [PubMed: 20563301]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Figure 1.
Overview of the steps in a supervised machine learning approach to sentiment classification

Figure 3.
Examining communication processes from individual to population level using linguistic analysis as linkage

Figure 4. Neural activity associated with higher self-reference scores from LIWC word dictionary ($p < 0.005$, $k = 37$).

Figure 5.

A. Neural activity associated with higher positivity scores from automatic SA classification ($p < 0.005$, $k = 37$). B. Neural activity associated with higher combined evaluative*positivity scores from automatic SA classification ($p < 0.005$, $k = 37$)

Table 1

Example of two text stimuli used as input for an fMRI study of message propagation. Texts are scored using LIWC (Pennebaker et al., 2007).

(A) Flowbee	(B) Periodic Element Rings
Flowbee is awesome, you can stay at home and cut your hair exactly the way you want. It's simple, easy, and precise. Flowbee uses a vacuum to suction hair up and uses a spacer to cut the desired length. The spacer makes it impossible to cut the hair shorter than the length it is set for. So you can't mess up or get a bad hair cut. They even have a specially designed spacer for a tapered cut. So now everyone in the family can cut their hair at home, even the family pet.	I'm sort of a geek, but I like to think of myself as a fashionable geek. I recently got one of the Periodic Rings, which are literal rings of chemical elements from the periodic table. I got the "Ag" one, which is silver. It looks exactly like the element box from the periodic table. These rings are also made of exactly what they say — so the platinum version is really expensive, but the silver and gold ones are a little more reasonable. I'm hoping someday that I can collect all three to become some sort of science teacher superhero.
LWIC Self References: 0 – Social: 9.38	LWIC Self References: 4.85 – Social: 1.34

Table 2

Example of transcripts of language produced by participants in response to stimuli viewed during an fMRI session. Texts are scored using a supervised sentiment analysis classifier.

(A) Example of positive recommendation	(B) Example of negative non-recommendation
<p>beauty queens i thought looked pretty hilarious um it was about moms who were former beauty queens who raised their daughters to be beauty queens it was about the stress of um their the daughters trying to be beauty queens and it was also about the mothers a little too um and that actually looked really funny um so i would definitely recommend moving forward with that one</p>	<p>beauty queens i don't know if its i am biased cause i am a guy but beauty queens would really not appeal to me that much cause the mere fact that i don't want to see a little girl putting hair and make up on for four hours go walking across the stage getting off the stage and doing it again two or three times and then losing or winning i really don't care um so that really doesn't appeal to me or jujust the moms pushing the kids jus they do that anyway why beauty</p>
Positivity: 0.49 Evaluative: 0.99	Positivity: -0.8 Evaluative: 0.99