

# Register in computational language research

Shlomo Engelson Argamon  
Illinois Institute of Technology

Shlomo Argamon is Professor of Computer Science and Director of the Master of Data Science Program at the Illinois Institute of Technology (USA). In this article, he reflects on the current and potential relationship between register and the field of computational linguistics. He applies his expertise in computational linguistics and machine learning to a variety of problems in natural language processing. These include stylistic variation, forensic linguistics, authorship attribution, and biomedical informatics. He is particularly interested in the linguistic structures used by speakers and writers, including linguistic choices that are influenced by social variables such as age, gender, and register, as well as linguistic choices that are unique or distinctive to the style of individual authors. Argamon has been a pioneer in computational linguistics and NLP research in his efforts to account for and explore register variation. His computational linguistic research on register draws inspiration from Systemic Functional Linguistics, Biber's multi-dimensional approach to register variation, as well as his own extensive experience accounting for variation within and across text types and authors. Argamon has applied computational methods to text classification and description across registers – including blogs, academic disciplines, and news writing – as well as the interaction between register and other social variables, such as age and gender. His cutting-edge research in these areas is certain to have a lasting impact on the future of computational linguistics and NLP.

**Keywords:** computational linguistics, natural language processing, style, stylistics, text classification

## Introduction

Register, broadly construed, refers to variation in language usage within different functional contexts. This article surveys research on register specifically within

‘computational language research’. By this term I mean to include work in (a) computational linguistics, which seeks to understand language by reference to implementing computational models, (b) natural language processing, which seeks to develop automated methods for understanding, extracting information from, and generating natural language texts, and (c) corpus linguistics, which uses computational methods applied to large bodies of text to study linguistic structure. This article will consider how register variation is addressed computationally in all of these areas, not always distinguishing among them.

Historically, most computational language research has sidestepped the question of register, either by working with a ‘representative corpus’ of text, however defined, or by limiting attention to one or a few specific text types. While there is widespread acknowledgement of the obvious differences between spoken and written text, and between formal and informal registers (say, published articles versus social media), until relatively recently, the computational study of linguistic register was a niche area and received little attention in computational work on language overall.

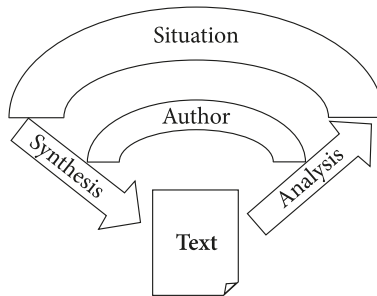


Figure 1. Register analysis and synthesis

My aim in this article is to make sense of the plethora of diverse approaches to register from a computational perspective and to suggest how future research might most productively be pursued.

Within computational language research, work on register can be divided into two broad categories, based on the research goals (cf. Figure 1). The first is ‘register analysis’, studying how to computationally classify texts according to register, as well as how to analyze register characteristics computationally. The second is ‘register synthesis’, developing techniques to generate texts in a given register, either starting from a formal semantic representation, or by ‘translating’ texts from one register to another.

Given a corpus of texts labeled according to different registers, computationally classifying new texts according to register is a well-defined task. It does,

however, require a definition of register sufficiently precise that human annotators can label texts accordingly with high inter-rater reliability, which is not always easy to achieve. Register classification can comprise a stand-alone task (Stamatatos, Fakotakis, & Kokkinakis 2000; Biber & Conrad 2001; Argamon, Koppel, Fine, & Shimoni 2003; Finn & Kushmerick 2006; Santini 2006; Herring & Paolillo 2006; Abbasi & Chen 2007; Dong, Watters, Duffy, & Shepherd 2008; Sharoff, Wu, & Markert 2010) or may be used to derive insights into larger questions related to linguistic variation (e.g., Atkinson 1992; Argamon, Dodick, & Chase 2008; Eisenstein, Smith, & Xing 2011; Teich, Degaetano-Ortlieb, Kermes, & Lapshinova-Koltunski 2013; Clarke & Grieve 2017). Register labels, either manually or automatically assigned, can also be used to control for register in research on other text analysis methods (e.g., Carroll et al. 1999; Giesbrecht & Evert 2009; Sharoff et al. 2010); differences in register between training and testing data often affect outcomes for NLP tasks such as part-of-speech tagging, parsing, or information extraction. Register categories can also be used to improve information retrieval and search, by serving as a constraint on what documents are retrieved (Karlsgren 1999; Morato, Llorens, Génova, & Moreiro 2003; Amasyalı & Diri 2006; Freund, Clarke, & Toms 2006).

In text generation, explicit parameterization of register (in terms of communication goals, domain of discourse, level of formality, etc.) can be used as constraints on the process of generating natural language text, to achieve more natural and useable texts (DiMarco & Foster 1997; Kan & McKeown 2002; Power, Scott, & Bouayad-Agha 2003; Reiter & Williams 2010). This requires, of course, an elaborated theory of how register parameters are realized in text by either restrictive or preferential constraints on word choice, syntactic structures, and so on. In research on register translation by analogy to machine translation between different languages, statistical models of the relationship between linguistic characteristics of different registers can be applied to a source text to generate a semantically equivalent text in a different register (Ficler & Goldberg 2017; Jhamtani, Gangal, Hovy, & Nyberg 2017; Fu, Tan, Peng, Zhao, & Yan 2018).

Related work on computer-mediated communication (CMC) has also been significant in the study of register, since the rise of new communication modes and media due to computational communication systems has enabled clearer and more precisely targeted research on the evolution and emergence of new registers (Nowson, Oberlander, & Gill 2005; Herring & Paolillo 2006; Crystal 2011). Much of this work falls outside our scope of computational language research, as it uses traditional (non-computational) methods of textual analysis. However, corpus linguistic methods are commonly used, and some CMC registers (e.g., blogs and Twitter) are frequent testbeds for NLP research.

## 1. How is register conceptualized in computational language research?

Unfortunately, there is a fair bit of terminological fuzziness and inconsistency in the computational literature dealing with register. For many natural language processing researchers, ‘register’ is often used as a rough synonym of the vague concept ‘text type’, by which is meant a group of texts sharing some more-or-less identifiable set of linguistic characteristics, without any clear theoretical motivation. Indeed, in much of this literature, the term ‘register’ can be used interchangeably with ‘genre’ or ‘text type’, and little attention is given to the structure of register or the nature of its categories. For the most part, this sort of informal approach is taken when studies control experiments for register/genre/text type, often using labels pre-assigned by corpus developers.

Indeed, most register classification research, which seeks to determine which of a discrete set of register categories a text belongs to, usually proceeds atheoretically. Register categories are sometimes based on reliable, if simple, external features of the text (spoken vs. written, for example), sometimes on other theoretical grounds such as the community of discourse or purpose of the text, and often based on ad hoc clusters of shared linguistic characteristics. Register categories considered in such studies tend to be chosen for convenience, e.g., based on what categories are provided by a particular corpus such as the British National Corpus (2007), or for the specific needs of a particular study, such as comparison of, e.g., peer-reviewed scientific papers vs. science popularizations. Little attention is usually placed on the functional relationship between linguistic features of register categories and the needs of their typical situations of use, though at times analysis results are interpreted in light of such functional considerations (Argamon et al. 2008; Teich et al. 2013).

For our purposes, we adopt the perspective of Biber and Conrad (2009:6) who define register as “a [language] variety associated with a particular situation of use (including particular communicative purposes).” A register is described by that situational context and the linguistic features typical of the register, along with a description of how those features function specifically within that particular context of language use. That is, the linguistic features describing a register are not arbitrary, but form a complex that is useful for particular communicative purposes in a particular context. We note that register is thus primarily defined by the situations in which it is used and only secondarily by the linguistic features by which it is recognized. This perspective is not universally recognized within computational language research, however.

Where computational language research has historically grappled seriously with the concept of register, generally one of two theoretical perspectives is taken. Both views accord broadly with Biber and Conrad, in that they take register

variation as essentially functional in nature, deriving from the fact that different communicative contexts require different linguistic resources to be brought to bear, resulting in (statistically) different language varieties. They differ, however, in how they structure register descriptions, how they differentially foreground either the situational context or the internal linguistic features of the text, and in how they formulate the relationship of register to other aspects of linguistic variation such as genre and style.

The main theoretically-motivated approach taken within register analysis is the multi-dimensional approach to register developed by Biber (e.g., 1991), which grew out of the definition above. This approach seeks to encode register character by a set of natural dimensions of variation for a mass of linguistic variables. Each dimension defines a spectrum for one aspect of register variation, such as formal-informal, narrative-informational, etc. Large numbers of linguistic features covary along each dimension, such that the set of dimensions explains much of the variation in linguistic form between registers and individual texts. With focus on the internal linguistic consistency of registers, the method remains ecumenical with respect to what external constraints (social situation, domain, function, historical context, etc.) are relevant to register. Two fundamental notions in this view are (a) that register variation is (or can be) continuous, varying along multiple fundamental dimensions of variation, and that (b) registers come at different levels of generality, with more general registers not being strictly comparable to more specific registers.

The theoretical view which predominates in work on register synthesis, on the other hand, is that of Systemic Functional Linguistics (SFL), which elaborates on the situational context by considering register to comprise distinct language varieties that serve differentiated social and communicative functions (Halliday, McIntosh, & Stevens 1968). This notion of register was elaborated by Martin (1992) to be determined by the contextual variables of 'field' (the type and domain of social discourse), 'tenor' (the relationship between the speaker and audience, and their relevant social roles), and 'mode' (parameters of textual organization, such as the communication channel and discourse goals). Each of these constrains the typical linguistic features that will be used in a text, based on functional requirements such as politeness, clarity, precision, available information bandwidth, and so forth. This articulation is particularly congenial for work on text generation systems, where values for field, tenor, and mode (or their equivalents) can be set manually, and researchers work on translating those into methods for determining the manner in which internal semantic representations are to be realized in text, for example choosing between complex or simple syntactic structures (Bateman 1997). Discourse-level generation work (e.g., Hovy 1991; Moore & Paris

1993) using Rhetorical Structure Theory (Mann & Thompson 1988) also relies on this concept of register.

## 2. How does register relate to the research goals within computational language research?

### 2.1 Aims of research in computational register analysis

Register analysis seeks to determine or characterize the register(s) of a text or of a collection of texts. A key task is to determine the character of particular registers, and their relationships to each other, in terms of what distributions of linguistic features they are realized by. Effective methods for register analysis can also be used for numerous applications including improving text analysis accuracy by controlling for register (Carroll et al. 1999; Kakkonen & Sutinen 2008; Giesbrecht & Evert 2009; Sharoff et al. 2010; Rehbein & Bildhauer 2017), information retrieval focused on particular registers and genres (Karlgrén 1999; Crowston & Kwasnik 2003; Morato et al. 2003; Amasyalı & Diri 2006; Freund et al. 2006; Vidulin, Luštrek, & Gams 2007), and analyzing the evolution of linguistic norms in different discourse communities (Biber & Finegan 2001; Spielman, Gondelaers, & Geeraerts 2006; Argamon et al. 2008; Teich & Fankhauser 2011; Degaetano-Ortlieb et al. 2016).

There are two main approaches to register analysis: classification, which seeks to divide texts into distinct register categories, and multi-dimensional analysis, which seeks to find a set of continuous dimensions, each representing a spectrum of underlying register-relevant variation.

#### 2.1.1 *Aims of register classification*

Traditionally, registers have usually been thought of as discrete categories of texts, such as spoken vs. written (at the most general level) or news stories vs. editorials (at a more specific level). This viewpoint has been adopted by much computational work on analysis of register, viewing it fundamentally as a classification or clustering problem of determining which texts fall into which register categories. As mentioned above, this work varies greatly in its theoretical approach to register, from essentially atheoretical approaches that treat register as a vague notion of ‘text type’ to those that have a stronger theoretical treatment of register emerging either from context of creation/use of the text (external) or clusters of cooccurring linguistic features (internal). In much of this work, little distinction is made between the terms ‘genre’ and ‘register’; since the computational methods for classifying texts based on genre or register differ little, we will discuss both types of

work interchangeably here. We do note that this lack of theoretical rigor in this work is problematic and discuss this issue further below.

The main operational goal of work on computational register analysis, viewing registers as discrete categories, has been to develop methods for accurately classifying texts into their proper register categories. As such, the register analysis problem is formulated as a text classification problem (Sebastiani 2002). The task is generally split into two subtasks: extracting useful features from the text and building a classification model using those features. In register classification, as in most text classification research, the focus has been on feature extraction, and generic classification machine learning methods are used. The question then is to determine a set of textual features that (a) are useful for a given register classification task and (b) can be accurately extracted using extant natural language processing methods.

Little work, if any, has examined this problem specifically from the standpoint of register, as opposed to other stylistic questions such as genre and authorship – researchers have sought ‘stylistic features’ of texts that are correlated with different styles, contrasted with ‘topical features’ that correlate with different topics or domains of discourse, which are typically used for information retrieval and related tasks. A wide variety of such features have been proposed (Stamatatos et al. 2000; Finn & Kushmerick 2006; Argamon & Koppel 2010), including relative frequencies of function words, part-of-speech n-grams, character n-grams, syntactic constructs, and systemic-functional categories, as well as type/token ratios, word and sentence length, and other textual statistics. These are discussed in more detail below.

The classification methods described above are all *supervised* methods, in that the training data are all labeled with the desired register categories. There is also some work on using *unsupervised* learning (i.e., clustering) for register analysis. These methods seek to find natural register categories in a corpus, to find the categories in a ‘bottom-up’ way, rather than relying on manual annotation from a set of predefined register labels. Results are evaluated by comparing the sets of documents in each cluster to sets derived from manually annotated data. Since the number of clusters output is a parameter, care must be taken to compare results for different numbers of output clusters, evaluating the clusters for plausibility compared to the manual annotations. Clustering, as opposed to supervised categorization, may be advantageous when seeking to characterize the different registers that may evolve with the emergence of new communications mediums such as web pages, blogs, or the like.

Gries, Newman, and Shaoul (2011) performed a systematic study of clustering over n-gram features in two corpora to study how well n-grams can be used as features to distinguish registers and sub-registers. Their analysis found considerable

structure in the corpora, with clustering giving results that matched well to fine-grained sub-registers, with 3-grams (word triples) giving overall the best results. Other clustering studies include Santini's influential studies of the genre characteristics of web pages (Santini 2005, 2006, 2008). Her results showed clear differentiation of register/genre categories emerging in an evolving computer-mediated communication medium. Clustering can also be used as an adjunct to classification analysis, to validate stability of results, as used by Teich and Fankhauser (2010) in their study of register differences between articles in different scientific fields (this study is discussed in more detail below).

When successful classification for a particular register task is achieved, examination of the features most influential in that classification can lend some insight into the linguistic difference between the registers being studied. For example, Biber and Barbieri (2007) show that different university registers use different characteristic lexical bundles (recurrent sequences of words), which they theorize realize different discourse frames and stances required by the different registers.

In this vein, register classification has also been used as to address social science research questions, by determining if two bodies of text from different communities differ linguistically, and whether such differences (if found) support theories about relevant differences in how the communities conceptualize, communicate, or organize information. For example, studies of variation in written and spoken registers of various World Englishes (Szmrecsanyi 2009; Gries & Mukherjee 2010) have shown clear differences between varieties of English in fundamental linguistic character such as analyticity/syntheticity and point towards how such analysis can help elucidate evolutionary relationships between these varieties. Classification studies of scientific writing from different disciplines have given insight into how the scientific register evolves (Atkinson 1992) and how different fields construct and communicate knowledge differently (Argamon et al. 2008; Teich et al. 2013).

### 2.1.2 *Aims of multi-dimensional analysis*

Douglas Biber (1991) pioneered another view into the nature of register and its relationship with underlying linguistic features, on the one hand, and related notions such as genre and text type, on the other. The central idea is that registers (and genres) differ along various continuous dimensions of covariation in sets of linguistic features. While the positioning of different registers at different points on each dimension critically can be related to aspects of their different contexts-of-use, the multi-dimensional approach methodologically takes a text-internal view of register, considering as primary the linguistic features characterizing different registers, irrespective of external situational variables.



Initial research in this paradigm focused on extracting and validating what the primary such dimensions are and correlating them with existing register and genre categories. The main idea is, viewing texts as points in a high-dimensional vector space (whose axes are linguistic variables), to use statistical techniques such as principal components analysis (PCA) (Jolliffe 2011) to determine the main directions of stylistic variation within that space. Subsequently, linguistic analysis of the main variables participating in each dimension, along with consideration of what kinds of texts have what scores in each dimension, can give a meaningful characterization of the spectrum that each dimension covers, such as ‘involved – informational’ (factor 1) and ‘narrative – non-narrative’ (factor 2). Different register categories reflect choices of regions in each of these dimensions; a multi-dimensional analysis of a text gives a fine-grained picture of the overall register/genre character of the text.

It is worth noting that the features considered in most applications of this approach are local lexico-grammatical features – lexical items, parts-of-speech, choices of syntactic structure – and do not include features of the overall discourse structure of the text. This focus accords with the view (Biber et al. 2007) that views register as realized through lexico-grammatical variation, as distinguished from genre, which is largely realized via variation in discourse structure. It might be useful to probe this distinction computationally, by using existing methods to extract aspects of rhetorical and discourse structure (Marcu 2000) and use them in multi-dimensional analysis to see what macrostructures might (or might not) covary with lexico-grammatical features between registers. To the best of my knowledge, this has not been done to date.

Biber’s original factors as well as his methodology have been applied to register and genre analysis in a number of domains after his original study. The factors in that study have emerged from analysis of diverse corpora, corroborating the general nature of that work (Lee 2001). Work applying multi-dimensional analysis to different corpora has been used to study the evolution of and variation in scientific writing (Atkinson 1992, 1998; Conrad 1996), the register and genre structure of blogs (Grieve et al. 2010), how textual cohesion is differentially realized in register (Louwerse & Graesser 2004), how interpersonal stance markers relate to register distinctions (Pavalanathan, Fitzpatrick, Kiesling, & Eisenstein 2017), and more.

Briefly, we note the stream of computational work that restricts the definition of ‘register’ to the single dimension of formality, defining register as the level of formality of the language. Formality distinguishes well between spoken and written registers, as well as between personal registers (such as letters and diaries) and published registers (such as fiction and non-fiction essays and books). Formality measures such as Heylighen and Dewaele’s (1999) F-score can be used to estimate

the formality of a text for corpus linguistic studies (Brooke, Wang, & Hirst 2010; Mosquera & Moreda 2012; Sheika & Inkpen 2012). This perspective can be viewed as choosing a single Biberian dimension for analysis and use.

## 2.2 Aims of research in computational register synthesis

Computational treatment of register has also been important in text generation, the problem of generating natural text from a structured semantic representation of information to be expressed (Reiter & Dale 2000; Gatt & Krahmer 2018). Text generation systems have been used in a variety of applications, from journalism to presenting healthcare information to intelligent tutoring to multimedia presentation and gaming (Di Eugenio & Green 2010). An important question in text generation for decades has been determining how to generate text that both represents the intended meaning and is also appropriate to a given situation of use. This requires the ability to generate text in different registers for the same propositional meaning (however represented). Two main approaches to this problem have been articulated (cf. Reiter & Williams 2010). The earliest approaches (e.g., Hovy 1990) to generating text in different styles formulated the problem as giving the user control over explicit stylistic features of the text to be generated and applying appropriate constraints to choices made between possible lexical items, syntactic structures, sentence complexities, punctuation, etc. during the text generation process.

A later development was corpus-based methods, that sought to generate texts matching the style of a given body of training texts (a reference corpus). The idea is to run a text generation system in such a way that statistics of linguistic features in the output text will be similar to those of the reference corpus (see, e.g., Langkilde-Geary 2002; Paiva & Evans 2005). This framework allows different styles to be generated based on examples, without explicit determination of the necessary generation parameters, which can be difficult to do well.

## 3. What are the major methodological approaches that are used to analyze or account for register in computational language research?

### 3.1 Methods of register analysis research

Register analysis, whether in the categorical or the multi-dimensional paradigm, comprises a two-stage process: feature extraction and corpus analysis.

In feature extraction, the linguistic structure of the text is analyzed to extract relevant features, and a representation of the text is constructed out of occurrence

statistics of these features, usually as a numeric vector. Occurrence statistics can be as simple as raw numbers of the features in the text or their relative frequencies or can include combined statistics like tf-idf (Aizawa 2003) or such measures as 'lexical gravity' (Gries & Mukherjee 2010).

The corpus analysis stage takes as input a collection of feature-based document representations and computes a representation of the registers being studied, whether as a classification model (in the categorical paradigm) or a set of dimensions of register variation (in the multi-dimensional paradigm). Here too, a variety of different technical approaches can be taken, giving different results, even using the same feature extraction method.

### 3.1.1 *Stylistic text features*

A simple and effective set of features (and one of the oldest) for stylistic text analysis is the relative frequencies of function words. Different function words frequencies are indicative of different grammatical choices, and yet are not expected to vary greatly with the topic of the text. For this reason, they are often used for stylistic analysis studies and have proven to be quite efficacious for a variety of such tasks, including register analysis (e.g., Biber 1991; Nowson 2006; Argamon & Levitan 2005; Herring & Paolillo 2006). Typical modern studies using function words in English use lists of a few hundred words, including pronouns, prepositions, auxiliary and modal verbs, conjunctions, and determiners. Numbers and interjections are usually included as well since they are essentially independent of topic, though they are not, strictly speaking, function words.

Another approach is to directly estimate relative frequencies of different syntactic constructions; this has been made possible by development of fast and reasonably reliable natural language parsing techniques. A number of studies have found that using syntactic features can often improve results over traditional word-based analysis alone (Stamatatos et al. 2000). Syntactic structures can be computed using full syntactic parsers (Martin & Jurafsky 2000), but such systems are often brittle when dealing with very long and complex sentences or informal and ungrammatical texts. Shallow parsing, by contrast, seeks just to identify occurrences of certain phrase types (such as noun phrases) without extracting full syntactic structures (Hammerton, Osborne, Armstrong, & Daelemans 2002), and so can be less error prone in those situations, although they provide less-detailed information. Some studies have used even simpler approximations to syntactic features to good effect, such as the frequencies of short sequences of parts-of-speech, or combinations of parts-of-speech and words (Glover & Hirst 1996; Tambouratzis, Markantonatou, Hairetakis, Vassiliou, Tambouratzis, & Carayannis 2000; Koppel, Argamon, & Shimoni 2002).

Taxonomies based on Systemic Functional Linguistics (Halliday & Matthiessen 2004) can be built to represent grammatical and semantic distinctions between classes of words, phrases, and syntactic structures at different levels of abstraction. When applied to lexical items, such taxonomies are a generalization of function word features combined with some parts-of-speech (Argamon et al. 2007). For example, (in English) a pronoun may refer either to a discourse Participant (first or second person) or Non-Participant (third person), if a Participant, it may be the Speaker (*I, me*), Speaker-Plus (*we, us*), or the Addressee (*you*); if a Non-Participant, it may be an Individual (*he, she, it*), or a Collective (*they, them*); and so on. Such categories form hierarchical taxonomies; the relative frequencies of occurrence within a text of the children of a category show the relative preferences within the text for the different options for realizing the parent category's function. Used as classification features, these relative frequencies have been shown to be effective for register analysis and are particularly helpful for relating the linguistic differences between registers to functional differences in their contexts of use (Argamon et al. 2008; Teich et al. 2013).

The frequencies of character n-grams, while lacking in direct linguistic motivation or interpretation, have proven to be useful in capturing stylistic variation in lexical, grammatical, and orthographic preferences for a variety of tasks, without the need for linguistic background knowledge (making application to different languages trivial). Relative frequencies of character n-grams have proven surprisingly effective for register and genre classification (Kanaris & Stamatatos 2007; Amasyali & Diri 2006) as well as other style classification tasks such as authorship attribution (Kjell 1994; Kešelj, Peng, Cercone, & Thomas 2003; Stamatatos 2008), document similarity (Damashek 1995), and L1 identification (Koppel et al. 2005). Since character n-grams also capture aspects of document content, however, they must be used with caution and careful experimental control to ensure that the models that are constructed represent register variation and not accidentally correlated topics instead.

### 3.1.2 *Methods of register classification*

For analyzing register in categorical terms, text classification techniques (Sebastiani 2002), usually based on machine learning, are used. The idea is straightforward: training texts are represented as numerical vectors, labeled by their register categories, and machine learning methods are used to find a function that distinguishes between the categories that minimizes some loss function over the training set. Different algorithms will produce different results, with greater or lesser ability to generalize accurately to new data (not in the training set).

A great variety of machine learning algorithms for classification have been applied to stylistic text classification at one time or another, with varying degrees

of success. Among the most straightforward are *k*-nearest neighbor (Kjell et al. 1995; Hoorn et al. 1999; Zhao & Zobel 2005) which categorizes according to the label(s) of the nearest document(s) in the training set, and Naive Bayes (Kjell 1994; Hoorn et al. 1999; Peng et al. 2004), which chooses the category label with the highest probability, assuming that features occur conditionally independently of each other given the text's category.

Excellent results for register and genre classification have been achieved using diverse machine learning algorithms including support vector machines (De Vel et al. 2001; Diederich et al. 2003; Koppel & Schler 2003; Abbasi & Chen 2007, Koppel et al. 2007; Zheng et al. 2006; Sharoff et al. 2010) and neural networks (Matthews & Merriam 1993; Merriam & Matthews 1994; Kjell 1994; Lowe & Matthews 1995; Tweedie et al. 1996; Hoorn 1999; Waugh et al. 2000). Rule-based learning (Holmes & Forsyth 1995; Holmes 1998; Argamon et al. 1998; Koppel & Schler 2003; Abbasi & Chen 2007; Zheng et al. 2006) and Bayesian regression (Genkin et al. 2006; Madigan et al. 2006) have also been applied with some success.

As noted above, the choice of feature set is more important to successful register classification than the specific learning algorithm used. In general, the learning methods that have been found to work best are those that (a) aggregate the impact of many different features together (rather than selecting a small number of 'important' features) and (b) are resistant to overfitting, being overly responsive to fine detailed distinctions in the training set (which can lead to reduced performance on new data). These include support vector machines, lasso and Bayesian regression, and some neural network methods.

Register classification is evaluated in the same way as general text classification methods (Sebastiani 2002). A corpus of documents is labeled manually for register categories, and inter-rater reliability is measured. Provided good reliability for labeling is achieved, the corpus is used for training and testing the model, ensuring that testing is done on texts not used for training. Cross-validation (Kohavi 1995) is typically used, in which the full annotated data is divided randomly into several (usually 10) equal-sized 'folds' subsets, then training the model on all but one fold and testing on the remaining fold. Results from all the folds are averaged to give an overall effectiveness measure. Effectiveness is measured by classification accuracy as well as average precision and recall over the different register categories (Goutte & Gaussier 2005). These other measures are needed since overall accuracy can be misleading if one category predominates in the data. In such a case, high overall accuracy can be achieved just by classifying all examples by the label of the majority class (i.e., smaller categories will not be accurately recognized at all). Hence, proper evaluation requires considering the precision and recall of each category on its own.

Such ‘class imbalance’ also has implications for classifier training – since most classification algorithms implicitly attempt to optimize for overall accuracy, an imbalanced training set will lead to a model with suboptimal performance. The usual ways to deal with this problem are ‘undersampling’, randomly keeping only a fraction of examples in the majority category to achieve balanced category sized in the training, or ‘oversampling’, randomly choosing examples from minority categories to be duplicated in the training set, and hence increase the cost of errors on those categories.

### 3.1.3 *Methods of multi-dimensional analysis*

In multi-dimensional analysis, the goal is to find the ‘natural’ dimensions of variation among core grammatical features of the language. Principal Components Analysis (Jolliffe 2011) or Factor Analysis (Loehlin 1998) is used to compute the sets of linguistic features that most frequently co-occur in a corpus. These are called the dimensions of variation for the corpus. Numeric weights are computed for features in each dimension, enabling computation of a score for any text in a given dimension. Analysis of which features covary in each dimension and the relationships between the dimensional scores for different texts or registers enables a linguistic interpretation of how aspects of register variation are represented by the different dimensions.

The earliest work on this approach (Biber 1989) used factor analysis to compute a set of dimensions based on a diverse corpus of English documents, constructed to represent a wide variety of spoken and written registers of British English (mainly from the LOB and LL Corpora, cf. Johansson, Leech, & Goodluck 1978; Svartvik & Quirk 1980). The work used 67 linguistic features in 16 categories including lexical, morphological, and syntactic features. Automatic syntactic analysis and tagging was performed to identify occurrences of each feature, then its frequency per 1,000 words was recorded for each text in the corpus, so that each text was represented by a 67-element numeric vector. Factor analysis was then applied to the corpus to extract the main dimensions of variation. Each factor assigns a weight between  $-1$  and  $1$  to every linguistic feature – if the feature has a positive weight, it is indicative of a positive value for the factor, if negative, for a negative value. In this way, the factor defines a spectrum for a co-varying set of linguistic features; the value of the factor for a given text represents where that text is on a particular functional linguistic dimension. By examining the commonalities of the features with highly positive and highly negative weights for a factor, we can form an interpretation of what the factor means.

Using this methodology, Biber (1991) identified seven factors of variation related to register, interpreted as:

1. Informational vs. Involved production
2. Narrative vs. Non-narrative concerns
3. Explicit vs. Situation-dependent reference
4. Overt expression of persuasion
5. Abstract vs. Non-abstract information
6. Online informational elaboration
7. Academic hedging

The generality of the result has been supported by the fact that similar studies on other corpora give substantially the same factors (Biber 2003, 2004; Xiao 2009; Clarke & Grieve 2017), though factors differ somewhat in saliency across different corpora, depending on the exact mix of registers and genres present. For example, in Clarke and Grieve's study on Twitter, one dimension primarily reflected tweet length, in that the short length of tweets means that noticeably more varied lexico-grammar appears in longer tweets.

Application of multi-dimensional analysis to new corpora can either be done by performing a new factor analysis and considering the relationship of the extracted factors to the original factors from Biber (1989, 1991), or by using the original factors and loadings to analyze the new corpus directly. Factor scores for documents in different registers or genres of texts give insight into their linguistic and functional content, showing the key distinguishing characteristics of each register.

Conrad's (1996) study of academic biology texts using multi-dimensional analysis showed that differences between academic and popular nonfiction texts can be quantified in meaningful ways and that important differences within the academic genre can be seen, specifically between textbooks and research articles. Conrad argues that these findings can be used to improve how we teach students to read and understand scientific texts by clarifying the different rhetorical modes that they use. Atkinson (1992) studied the evolution of scientific prose style by (in part) examining how dimensional scores change over time in analysis of the *Edinburgh Medical Journal* (started in 1733). He found clear change over time from a somewhat involved style to a more modern informational style and from largely narrative concerns to mostly non-narrative concerns in the most recent articles. The study shows a gradual change of rhetoric and style in scientific communication, arguing that such change is more evolutionary than due to sudden paradigm shifts.

As noted, application of Biber's original method to different sorts of texts and different feature types largely corroborates the original set of dimensions, with some interesting variations at times.

Grieve et al. (2010) used these methods to study linguistic variation in blogs, corroborating Biber (1991) by finding some of the same factors, as well as finding other factors specific to their blog corpus, such as Addressee Focus (whether the blog addresses the reader directly or not) and Thematic Variation (whether blog posts are thematically unitary or address multiple topics).

In another vein, Louwrese and Graesser (2004) used the same methodology of factor analysis but over cohesion features of texts, rather than the lexical and syntactic features used in Biber's original study. Results corroborated the original factors found by Biber, despite the use of different features that measured local and global cohesion of the text (Louwrese & Graesser 2004).

A recent application of the multi-dimensional paradigm, by Pavalanathan et al. (2017), finds dimensions of linguistic variation related to how interpersonal stances are taken in texts. Their method first generates a lexicon of stance markers, starting from a curated set of seed terms, and then applies multi-dimensional analysis to frequencies of those markers in a corpus of 530 million Reddit comments. The resulting dimensions include the involved/informational language dimension from Biber's original studies, as well as others more focused on stancetaking such as narrative/dialogue-orientation, standard/non-standard variation, and positive/negative affect.

### 3.2 Methods of register synthesis research

To synthesize a text in a particular register, choices must be made in how a text is constructed, while still realizing the intended meaning. These choices are inherently integrated within the text generation system, and so we must first discuss how text generation systems work. Reiter and Dale (2000) break down text generation into three basic tasks: (a) document planning, determining the overall structure and content of a text, (b) microplanning, determining how the needed information is going to be expressed, in linguistic terms, in sentences and clauses, and (c) realization, generating an actual text conforming to these determinations. Generating text in a particular register amounts to placing constraints (both requirements and preferences) on decisions made at the document planning and microplanning stages.

While some aspects of overall textual organization related to register may be handled at the document planning stage, most aspects of style in extant text generation systems are handled (when they are) during the microplanning stage. These aspects include (Reiter & Williams 2010): lexical choice (which words to use to express information), information aggregation (how much to express in each sentence), information ordering (what order to present information), syntactic choice (what syntactic structures to use), and punctuation choice (full



stops vs. exclamation marks, or comma-separated lists vs. numbered lists). When generating text, constraints on these aspects of text construction can be applied to influence the style of the text to be generated. For example, text in a formal register may be generated by constraining lexical choice to prefer formal and Latinate words, information to be more aggregated in fewer (longer) sentences, and syntax to prefer more complex constructions.

The importance of register-based constraints on text generation was realized early in the development of the Penman system (Hovy, Lavid, Maier, Mittal & Paris 1992) which incorporated a notion of text type, and in the related KOMET system (Bateman, Maier, Teich & Wanner 1991) which used the notion of a communicative situation. In both cases, and in much subsequent work, the SFL theoretical framework (Halliday & Hasan 1989; Martin 1992) was used to describe register as comprising field (the events, states, and participants in the communicative situation), tenor (the roles and relative statuses of participants), and mode (the kind of communication channel). Text generation systems relate these parameters to constraints on the various linguistic variables described above.

In a text generation system, control over stylistic choices can either be made *explicitly*, by specific control of constraints on document and microplanning, or *implicitly*, by specifying that the system should generate text in a style similar to that of a particular document or corpus. We will discuss each methodology in turn.

In the explicit control scenario, a user explicitly sets the values of a set of parameters that constrain style-related choices in text generation, generally at the microplanning stage. For example, in the Iconoclast system (Power, Scott, & Bouayad-Agha 2003) for generating medical information summaries, the user can set preferred values for:

- Paragraph length,
- Sentence length,
- Frequency of connectives,
- Frequency of passive voice constructions,
- Frequency of pronoun usage,
- Frequency of semicolon usage,
- Frequency of comma usage,
- Level of technical terms allowed, and
- Level of vertical (bulleted) lists allowed.

The system also had aggregated preset values giving higher-level style profiles such as ‘broadsheet’ and ‘tabloid’ that could be selected.

Each of these parameter settings introduces a constraint into the microplanning process, which collectively are expressed as a preference function that scores different microplanning choices accordingly. This will cause the system to make

choices that result in longer/shorter paragraphs or sentences, aggregate/separate information to get more/fewer semicolons and commas, and so forth.

In implicit style control, the system is given a document or corpus whose style it is to imitate. One approach to corpus-based style control is manual analysis of the corpus to extract constraints and rules that can be applied during text generation (Hovy 1990; McKeown, Kukich, & Shaw 1994; Reiter & Williams 2010). A difficulty with this approach is that developing such rules is very labor intensive, and the results tend to be of only limited generality.

An approach potentially more useful and generalizable is to measure corpus statistics and use machine learning to generate constraints and rules for text generation in a matching style. Kan and McKeown (2002) describe a text summarization system that learns content planning rules from an annotated corpus for generating summaries matching the style of the corpus. The system uses rule learning (Cohen 1995) to find content-planning rules that match well the lexical statistics (including bigrams and collocations) of the target corpus. A similar approach was taken in SumTime-Mousam (Reiter, Sripada, & Hunter 2005), which created microplanning rules for generating weather forecasts based on machine learning over a corpus of human-authored forecasts. These rules were, however, filtered by domain experts before they were used in the system. Belz (2005) describes a version of the system that used automatically learned rules without any human input – the results were not as good as those of SumTime-Mousam, but were still judged to be quite readable and understandable.

The fully corpus-based approach, by contrast, seeks to generate texts matching the style of a given reference corpus without explicit manual determination of generation parameters. A simple, but time-consuming, method is to search systematically through different parameter settings for the text generation system, comparing the various texts generated with different settings to the reference corpus, and picking the settings that produces a text with the most similar feature statistics (cf. Langkilde-Geary 2002). A more long-term efficient method is described by Paiva & Evans (2005), whose system first analyzes the reference corpus to infer the parameter settings most likely to produce a style consistent with it, and then uses those predetermined settings to generate new texts. The first step is done by generating a variety of texts for multiple parameter settings and analyzing the correlations between choices taken in generation and similarity scores between generated texts and the reference corpus. The result is a score for each possible choice in generation for how likely it is to generate texts that match the reference corpus; the most likely choices can then be used to generate texts statistically matching the reference style.

Related to this corpus-based approach for generating text in a particular register are techniques for directly ‘translating’ texts from one register or style to

another. The general goal is to take a text in one register (or style) and generate text expressing the same content, but in a different register. One of the earliest such works in Xu, Ritter, Dolan, Grishman, & Cherry (2012) work on paraphrasing modern texts in Shakespeare's style, which treated the problem as one of machine translation (MT), applying a standard phrase-based statistical MT method (Och & Ney 2003; Koehn et al. 2007). The task was formulated by taking the original versions of the plays and learning a translation model between them and translations of them into Modern English. The learning process found correspondences between words and phrases, such as 'thou' → 'you', as well as probabilities of changing word order. They applied the model to a play not used in training (*Romeo and Juliet*) and compared results to human translations.

Recently, research has grown on this 'style translation' task using neural network methods, addressing it as a variant of the 'style transfer' problem for images (Gatys et al. 2016), using sequence-to-sequence networks, whose input and output are each a sequence of tokens. Variables denoting the desired style are added as parameters either as special tokens in the input text (Sennrich et al. 2016) or as extra inputs to the neural network (Ficler & Goldberg 2017). Typically, these approaches use Long Short-Term Memory (LSTM) networks (Hochreiter & Schmidhuber 1997). An LSTM network is given each token in a text (usually a sentence) as a vector-encoded input one at a time and computes both an update to a set of memory units and a vector-encoded output token as a function of the input and the current memory state. It thus directly translates one sequence of tokens into another sequence of tokens. Training is done by comparing network outputs to gold-standard training data (what sequence should result) and adjusting weights in the network to make outputs more similar to the training for the given inputs.

For style translation, an LSTM network is trained on input/output pairs of sentences in two different styles (say, Shakespeare and a modern translation), with addition inputs encoding which pair of styles is being trained on (so the network can automatically generate the correct style after it is trained). For generation, these style parameters will influence the probability distribution of lexical items generated and the length/complexity of sentences generated by the network, so it will generate sentences in a particular style. Improvements have been made by applying adversarial network training (Goodfellow et al. 2014), in which the style translating network tries during training to fool a second network that determines whether texts are in the target style or not. This training method has been applied recently to improve accuracy at style translation in this paradigm (Prabhumoye, Tsvetkov, & Salakhutdinov 2018; Fu et al. 2018).

#### 4. What does a typical register study look like in computational language research?

To address this question, I will describe three case studies, one on register classification, one on multi-dimensional analysis, and one on register synthesis.

##### 4.1 Classification analysis: Multidisciplinary scientific texts

In an excellent exemplar of classification analysis of register characteristics, Teich and Fankhauser (2010) consider the question of how clusters of linguistic characteristics emerge in new registers as multidisciplinary scientific fields coalesce. To do this, they analyze texts from the DaSciTex corpus which comprises journal articles from nine scientific fields, including interdisciplinary ('mixed') fields. They are:

- A. computer science
- B. 'mixed' disciplines:
  - B1: computational linguistics
  - B2: bioinformatics
  - B3: computer aided design/construction in mechanical engineering
  - B4: microelectronics/VLSI
- C. 'pure' disciplines
  - C1: linguistics
  - C2: biology
  - C3: mechanical engineering
  - C4: electrical engineering

The research investigated how the registers of the 'mixed' (B) disciplines relate to those of their respective core 'pure' disciplines (C) and that of computer science research (A), which they draw from. This study provides a good case study of how classification methods are used for register analysis, and the kinds of broader conclusions that can be drawn from such analysis.

The study comprises several analyses to determine the character of the registers being examined. First, the overall characteristics of the scientific register of the DaSciTex corpus of scientific articles is analyzed by comparison with the Freiburg-LOB Corpus of British English (FLOB). The authors focused on linguistic features likely to indicate theoretically predicted characteristics of scientific text: abstractness, technicality, and informational density. These included the fractions of words that are common nouns, lexical verbs, and adverbs; the type-token ratio to indicate use of specialized terminology; and lexical density as a proxy

for informational density. Discriminability measures such as information gain (Quinlan 2014; Raileanu & Stoffel 2004) show which features individually most distinguish the two categories (scientific and general text); as expected higher abstractness, information density, and terminological diversity (lower type-token ratio) are characteristics of scientific text. The discriminating power of these features together was explored by using them as input into support vector machine classification and k-means clustering; these gave high classification accuracy indicating strong differences between the document groups. Additionally, these features do not strongly distinguish between different subcategories of the DaSciTex corpus, indicating that all these scientific registers are similar in these respects.

The authors analyze the register characteristics of the different scientific disciplines within the DaSciTex corpus by learning linear classifiers (using support vector machines) to distinguish articles from the different disciplines by occurrence frequencies of distinctive (measured by information gain) nouns and lexical verbs, respectively. Classification accuracy is high, 96%, and 87%, respectively. Examination of the confusion matrices, which show how often documents in one category were classified as each of the others, enables measuring the linguistic similarity of different document categories – categories that are more often confused with each other are more similar. Results here showed that the language of the mixed disciplines are each most similar to their respective pure disciplines (B1 and C1, B2 and C2, etc.), and somewhat to computer science (A); also electrical engineering (C4) is similar to computer science overall.

A more fine-grained analysis of the computer science (A), computational linguistics (B1), and linguistics (C1) subcorpora was done by learning linear classifiers, using as features the different verbs that appear with ‘we’ as a subject in the articles. These show what activities the authors are portrayed as engaging in by the articles. The most significant such features for computer science were formal activities (e.g., ‘prove’ and ‘define’); for computational linguistics were experimental (e.g., ‘collect’ and ‘examine’); and for linguistics were communicative (e.g., ‘argue’ and ‘read’) and cognitive (e.g., ‘see’ and ‘feel’). This gives some insight into how scientists in the different disciplines understand the key activities of their disciplines, and how they construct themselves as acting, in their research reports.

Classification analysis thus can give multiple views into the register variations under study, by a combination of univariate and multivariate results, and exploration of classification accuracy with misclassification rates for different feature sets and subsets of documents. As Teich and Fankhauser comment:

In addition to ranking features by their individual discriminatory power, we can explore their *collective* contribution to register discrimination (multivariate

analysis). Also, having available information about misclassifications in the form of the confusion matrix, we can investigate the context of typical features/terms in correctly classified and in misclassified texts, analyzing differences and commonalities between registers at class level as well as at instance level.

(Teich & Fankhauser 2010: 245; italics in original)

## 4.2 Multi-dimensional analysis: Abusive language

Multi-dimensional analysis was used recently to good effect by Clarke and Grieve (2017) in a study of racist and sexist language on Twitter. They study a corpus of 16,914 Tweets from the Waseem and Hovy (2016) Twitter corpus that have been manually coded as racist (1972 tweets), sexist (3383 tweets), or neither (11,559 tweets), with the goal of understanding register variation between abusive and non-abusive tweets and the similarities and differences between racist and sexist language on Twitter.

Due to the short length of tweets (typically under 30 words), standard multi-dimensional analysis (MDA) cannot be applied, since it relies on measuring and comparing feature frequencies within each text, and any given feature might appear once in a given tweet, if at all. The authors therefore use a variation called multiple correspondence analysis (MCA) (Husson, Lê, & Pags 2010), which reduces categorical data to underlying dimensions in a similar way to how MDA reduces quantitative data. Linguistic features in the tweets are therefore treated as categorical variables (occurrence vs. non-occurrence) and MCA used to extract underlying dimensions of variation. The analysis used 81 linguistic features that occurred in at least 1% of the tweets, including “tense and aspect markers, place and time adverbials, personal pronouns, questions, nominal forms, passives, subordination, complementation, adjectives and adverbs, modals, specialized verb classes, coordination, negation and other lexical classes, such as amplifiers, down-toners and conjunctions” (Clarke & Grieve 2017: 2). Additionally, they included medium- and genre-specific features such as “hashtags, URLs, capitalization, imperatives, comparatives, and superlatives” (Clarke & Grieve 2017: 2)

Clarke and Grieve extract four dimensions from the corpus using MCA, however the first dimension is strongly correlated with tweet length. As one might expect, since tweet length is small on average with a high standard deviation, a key dimension of variation will be text length. Hence, they remove that dimension from consideration in further analysis. The other three dimensions are interpreted as:

1. Interactivity: Variation between tweets directly interacting with others and those simply asserting something. Features of interactive tweets include question marks, question and clause-initial DO, WH-words, and first and second

- person pronouns, while those of non-interactive (informational) tweets include existential *there*, *be* as a main verb, numbers and attributive adjectives, quantifiers, prepositions, and proper nouns, as well as other features of informational text such as nominalizations and contrastive conjunctions.
2. Antagonism: Ranging from aggressive and antagonistic tweets to those of a more conciliatory style. Features of antagonistic tweets include those indicating direct confrontation and emotion, such as question and initial DO, clause-initial verbs, question marks, and second person pronouns, as well as nominalizations, possessive pronouns, emoticons, and exclamation marks. Conciliatory features included lack of second person pronouns, especially with first and third person pronouns in subject and object position, progressive verbs, and perception verbs.
  3. Attitudinality: Variation in whether a tweet explicitly expresses an attitude, or frames itself as factually propositional. Features indicating attitudinality include comparatives, predicative adjectives, first person pronouns, existential *there*, and paucity of nouns; indicating factuality include public verbs mark indirect or reported speech, marked aspect (perfect or progressive), passive voice, URLs, and numbers.

These three dimensions define a space of register variation within which the abusive (racist and sexist) tweets exist. Further analysis of the distributions of dimension scores for racist and sexist tweets shows some interesting differences. Data visualization and the Wilcoxon signed-ranked test showed sexist tweets to be more interactive and attitudinal than racist tweets, on average. The authors interpret these findings in light of Herring et al.'s (1995) study of silencing strategies in sexist discourse as the sexist tweets perhaps aiming to take control of the Twitter discussion and silence female voices, while racist tweets serve more to spread racist ideology by storytelling and persuasion, in the modes described by van Dijk (1993).

This study shows how multi-dimensional analysis applied to multiple related document categories can elucidate the linguistic structure of register variation and contribute to understanding the potential functions of such variation.

### 4.3 Text generation: Customized medical information

As a case study for text generation in variable register, consider WebbeDoc (DiMarco & Foster 1997), a system developed to produce customized information about the HealthDoc medical text generation project. The idea is to present the same information about the HealthDoc project in different ways depending on reader characteristics such as professional role (physician, computational linguist,

funder, layperson), age, and reading level. The work does not explicitly deal with the notion of register but rather with aspects of textual style that in the aggregate amount to different registers for different audiences and purposes.

WebbeDoc works from an internal structured representation of a complete ‘master document’ about the project, which can be realized in different kinds of actual texts. This internal master document (MD) contains elements which are abstract sentence specifications (written in a structured Text Specification Language, or TSL) that express elemental propositions. These are connected within a structure that represents ordering constraints, rhetorical relations, coreference links, and formatting information (such as use of vertical lists or illustrations). Each element is also annotated with user model constraints, labeling the user categories (e.g., ‘layperson’) or textual styles (e.g., ‘high technical level’ or ‘informal’) for which the element should be selected. The master document as a whole is structured into topics and subtopics, each of which can be divided into ‘version sets’ which contain the different ways to tailor the information in the topic for different audiences (per the selection annotations). In text generation, document planning will choose between versions of each subtopic, appropriate to the given parameter settings. Preferred ordering of subtopics may also vary with intended audience and style.

Document planning seeks to construct a coherent text expressing the information in the master document relevant to a given audience in an appropriate style. The system chooses to use those specific MD elements whose annotations are consistent with the parameters specified, and then seeks to satisfy consistent ordering preferences, and general rhetorical and coreference constraints to ensure completeness and coherence of the information to be presented. Constraint satisfaction is used to find such a coherent information structure to present (Marcu 1997). Aggregation of matching entities is applied to simplify syntactic structures and determine appropriate use of coreferential expressions, and then text is generated from the internal representation. A variety of repair strategies is applied to ensure the consistency and coherence of the resulting document.

Unfortunately, as with much other work in customizable text generation, no user studies were performed on WebbeDoc. However, examples show the plausibility of generated texts for different audiences.

## **5. What are the most promising areas of future research on register in computational language research?**

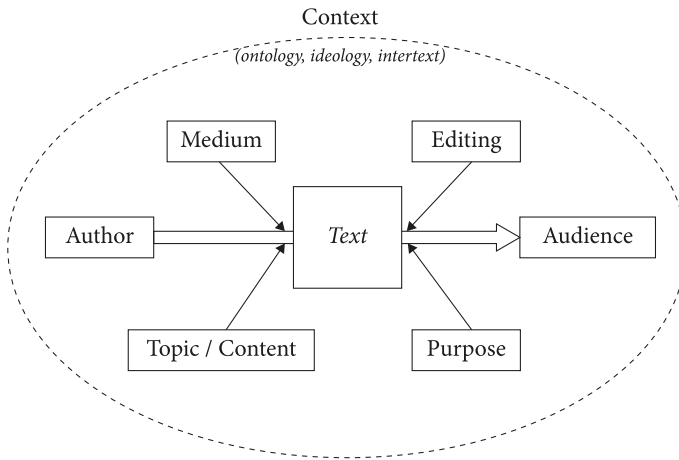
Computational research on register has the potential to both benefit from and contribute to more elaborated and precise theoretical and empirical analyses of



register in general. As we have discussed, notions of register and genre, especially as used in computational language research, are often vague and imprecise, relying on intuitive categories of ‘text types’ or small sets of register categories that are different enough to be easily distinguished, such as ‘spoken’ vs. ‘written’, or ‘news articles’ vs. ‘research papers’. The most consistent and clear approach to date has been Biber’s multi-dimensional approach, with a clear empirical basis and elaborated dimensions based on well-accepted linguistic features. However, there is as yet no general model relating these dimensions directly to aspects of the communicative situation (as in the SFL theory of register), to accepted register categories, in the categorical approach, or to the choices made by a text generation system that realizes a text in a given register. It is here that computational research into register variation and connected phenomena can make a fundamental contribution to the understanding of register and the connection between the communicative situation and linguistic style and form.

In broad outline, the research program I suggest is to use and combine existing and new computational register analysis and synthesis methods to elaborate and test detailed and empirically based models of register-based linguistic variation. (The program is closely related to Matthiessen’s [2015] registerial cartography; see further below.) If we consider the parameters and constraints that determine the form of a text (see Figure 2), they include many disparate elements, many of which feed into traditional notions of register, such as the medium of communication, the purpose of the text and intended audience, ideological roles/relationships among the author, editor(s), and audience, and the intertextual context of the text. The sort of precise and elaborated model envisioned here would describe causal relationships between parameterizations of these elements and the linguistic form of the text, likely mediated by intermediate variables such as Biber’s dimensions and register/genre categories at different levels of specificity. Such models, perhaps in the form of hierarchical probabilistic models (Koller & Friedman 2009) or structured causal models (Pearl 2009), would enable precise empirical predictions from models of how the social and communicative context influences the construction of different texts and gives rise to different registers.

An immediate research goal is to put the notion of register categories, or text types, on a firmer empirical footing. Linking work on register categorization with multi-dimensional analysis would be an excellent step – exploring how register categories might be viewed as regions within a multi-dimensional ‘style space’, and their relationships measured and compared. Rigorous work comparing the dimensions that emerge from analysis of different corpora would be needed, so that a standardized dimensional representation could be determined, or at least that regularities in how to understand any variation in the extracted dimensions. (It seems likely that one underlying set of dimensions can be derived, with varia-



**Figure 2.** High-level view of parameters and constraints on textual form and style. Adapted from Argamon & Koppel (2010)

tion attributable to corpus composition, but this will require some work to prove.) Comparison between models in different languages and cultures will be necessary, with the goal of finding typological universals; most work has been done in English, though work in other languages is consistent with multi-dimensional analyses in English (Biber 1995; Berber Sardinha 2017). This work will give a clearer basis for the descriptive analysis of register, enabling improved large-scale empirical work in the area.

A second key focus for connecting disparate research in the area is integrating register analysis and synthesis; while both active research areas use similar theoretical concepts of register and linguistic variation, there is no cross-over between the research, and so no clarity as to the relationships between register dimensions, occurrence of linguistic features, and planning decisions during text generation. Register classification systems can be used to evaluate the results of text generation systems, and correlated with user perception studies, to better validate schemes for categorization and measurement of register variation. Comparison of Biber's multi-dimensional model and the internal representations learned by style transfer neural networks may also help elaborate the ways in which linguistic representation space can usefully be structured to represent situational variation.

As clearer and more integrated representation schemes are developed for register varieties and relationships between them, register can be more easily addressed properly within work on natural language processing (NLP) and computational linguistics. Much work in these fields, as noted above, does not properly account for register, or does so superficially by controlling for some pre-labeled 'text types' or the like. Given broader articulated models for register, and

tools to implement register classification and profiling, more sophisticated methods for register-aware evaluation of NLP tools can be developed, and better experimental controls for register can be applied.

Furthermore, these same methodological and computational tools will enable qualitatively greater progress in the large-scale descriptive task termed by Matthiessen (2015) as registerial cartography – mapping the functional language varieties in a language/culture and the relationships between them. More sophisticated computational analysis tools will enable larger scale studies than have been heretofore possible, and unified representational schemes will allow comparison and integration of the results of different studies on different corpora. Historical and typological studies will also allow for the descriptive study of register evolution at the large scale, expanding greatly the scope of such studies as Atkinson's (1992) of the development of the modern scientific register and Gries and Mukherjee's (2010) work on the evolution of World Englishes. As Matthiessen notes, development of the theory of register beyond a certain point requires the detailed description and mapping of many extant register varieties and the relationships among them, to ground the theory. (See also Matthiessen & Teruya 2015 for a good example.) One grand challenge for computational research on register, therefore, will be to develop the representations and tools necessary for this task at a large scale.

The other grand challenge will be the development of the methodological and computational tools necessary for empirical verification of theories of register. This would seem to require the ability to specify clear, articulated models of the causal relationships between situational parameters (social roles and relationships, communication medium, etc.), linguistic features, and intermediate representational levels, and then to test these models empirically, either by analysis of large corpora, or by generating texts according to specification and rigorously measuring human responses to the texts.

## References

- Abbasi, A., & Chen, H. (2007). Categorization and analysis of text in computer mediated communication archives using visualization. In *Proceedings of the 7th ACM/IEEE-CS Joint Conference on Digital Libraries* (pp. 11–18). New York: ACM.
- Aizawa, A. (2003). An information-theoretic perspective of tf-idf measures. *Information Processing & Management*, 39(1), 45–65. [https://doi.org/10.1016/S0306-4573\(02\)00021-3](https://doi.org/10.1016/S0306-4573(02)00021-3)
- Amasyali, M. F., & Diri, B. (2006). Automatic Turkish text categorization in terms of author, genre and gender. In *International Conference on Application of Natural Language to Information Systems* (pp. 221–226). Berlin: Springer.

- Argamon-Engelson, S., Koppel, M., Avneri, G. (1998). Style-based text categorization: What newspaper am I reading? In *Proc. of AAAI Workshop on Learning for Text Categorization*, 1998 (pp. 1–4).
- Argamon, S., & Levitan, S. (2005). Measuring the usefulness of function words for authorship attribution. In *Proceedings of the 2005 ACH/ALLC Conference*.
- Argamon, S., & Koppel, M. (2010). The rest of the story: Finding meaning in stylistic variation. In *The structure of style* (pp. 79–112). Berlin: Springer.  
[https://doi.org/10.1007/978-3-642-12337-5\\_5](https://doi.org/10.1007/978-3-642-12337-5_5)
- Argamon, S., Dodick, J., & Chase, P. (2008). Language use reflects scientific methodology: A corpus-based study of peer-reviewed journal articles. *Scientometrics*, 75(2), 203–238.  
<https://doi.org/10.1007/s11192-007-1768-y>
- Argamon, S., Koppel, M., Fine, J., & Shimoni, A. R. (2003). Gender, genre, and writing style in formal written texts. *Text*, 23(3), 321–346. <https://doi.org/10.1515/text.2003.014>
- Argamon, S., Whitelaw, C., Chase, P., Hota, S. R., Garg, N., & Levitan, S. (2007). Stylistic text classification using functional lexical features. *Journal of the American Society for Information Science and Technology*, 58(6), 802–822. <https://doi.org/10.1002/asi.20553>
- Atkinson, D. (1992). The evolution of medical research writing from 1735 to 1985: The case of the Edinburgh Medical Journal. *Applied Linguistics*, 13(4), 337–374.  
<https://doi.org/10.1093/applin/13.4.337>
- Bateman, J. A., Maier, E. A., Teich, E., & Wanner, L. (1991). Towards an architecture for situated text generation. In *Proceedings of the ICCICL* (pp. 289–302).
- Belz, A. (2005). Statistical generation: Three methods compared and evaluated. In *Proceedings of ENLG-2005* (pp. 15–23).
- Berber Sardinha, T. (2017). Text types in Brazilian Portuguese: A multi-dimensional perspective. *Corpora*, 12(3), 483–515. <https://doi.org/10.3366/cor.2017.0129>
- Biber, D. (1989). A typology of English texts. *Linguistics*, 27(1), 3–44.  
<https://doi.org/10.1515/ling.1989.27.1.3>
- Biber, D. (1991). *Variation across speech and writing*. Cambridge: Cambridge University Press.
- Biber, D. (1995). *Dimensions of register variation: A cross-linguistic comparison*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511519871>
- Biber, D. (2003). Variation among university spoken and written registers: A new multi-dimensional analysis. *Language and Computers*, 46, 47–70.
- Biber, D. (2004). Conversation text types: A multi-dimensional analysis. In *Le poids des mots: Proc. of the 7th International Conference on the Statistical Analysis of Textual Data* (pp. 15–34). Louvain: Presses universitaires de Louvain.
- Biber, D., & Barbieri, F. (2007). Lexical bundles in university spoken and written registers. *English for Specific Purposes*, 26(3), 263–286. <https://doi.org/10.1016/j.esp.2006.08.003>
- Biber, D., & Conrad, S. (2001). Register variation: A corpus approach. In D. Schiffrin, D. Tannen, & H. E. Hamilton (Eds.), *The handbook of discourse analysis* (pp. 175–196). Malden, MA: Blackwell.
- Biber, D., & Conrad, S. (2009). *Register, genre, and style*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511814358>
- Biber, D., & Finegan, E. (2001). Diachronic relations among speech-based and written registers in English. In S. Conrad & D. Biber (Eds.), *Variation in English: Multi-dimensional studies* (pp. 66–83). Harlow: Pearson Education.

- Brooke, J., Wang, T., & Hirst, G. (2010). Automatic acquisition of lexical formality. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters* (pp. 90–98). Stroudsburg, PA: Association for Computational Linguistics.
- Carroll, J., Minnen, G., & Briscoe, T. (1999). Corpus annotation for parser evaluation. In *Proceedings of the EACL workshop on LINC*, June 1999.
- Clarke, I., & Grieve, J. (2017). Dimensions of abusive language on twitter. In *Proceedings of the First Workshop on Abusive Language Online* (pp. 1–10).
- Cohen, W. W. (1995). Fast effective rule induction. In *Proceedings 12th International Conference on Machine Learning* (pp. 115–123). Burlington MA: Morgan Kaufmann.
- Conrad, S. M. (1996). Investigating academic texts with corpus-based techniques: An example from biology. *Linguistics and Education*, 8(3), 299–326.  
[https://doi.org/10.1016/S0898-5898\(96\)90025-X](https://doi.org/10.1016/S0898-5898(96)90025-X)
- Crowston, K., & Kwasnik, B. H. (2003). Can document-genre metadata improve information access to large digital collections? *Library Trends*, 52(2), 345–361.
- Crystal, D. (2011). *Internet linguistics: A student guide*. London: Routledge.  
<https://doi.org/10.4324/9780203830901>
- Damashek, M. (1995). Gauging similarity with n-grams: Language-independent categorization of text. *Science*, 267(5199), 843–848. <https://doi.org/10.1126/science.267.5199.843>
- De Vel, O., Anderson, A., Corney, M., & Mohay, G. (2001). Mining e-mail content for author identification forensics. *ACM Sigmod Record*, 30(4), 55–64.  
<https://doi.org/10.1145/604264.604272>
- Degaetano-Ortlieb, S., Kermes, H., Khamis, A., & Teich, E. (2016). An information-theoretic approach to modeling diachronic change in scientific English. In C. Suhr, T. Nevalainen, & I. Taavitsainen (Eds.), *Selected papers from Varieng – From data to evidence (d2e)*, Helsinki, Finland. Leiden: Brill.
- Diederich, J., Kindermann, J., Leopold, E., & Paass, G. (2003). Authorship attribution with support vector machines. *Applied Intelligence*, 19(1–2), 109–123.  
<https://doi.org/10.1023/A:1023824908771>
- DiMarco, C., & Foster, M. E. (1997). The automated generation of Web documents that are tailored to the individual reader. In *Proceedings of the AAAI-97 Spring Symposium on Natural Language Processing for the World Wide Web*, Stanford, CA.
- Dong, L., Watters, C., Duffy, J., & Shepherd, M. (2008). An examination of genre attributes for web page classification. In *Proceedings of HICSS* (pp. 133). IEEE.
- Eisenstein, J., Smith, N. A., & Xing, E. P. (2011). Discovering sociolinguistic associations with structured sparsity. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1* (pp. 1365–1374). Stroudsburg, PA: Association for Computational Linguistics.
- Ficler, J., & Goldberg, Y. (2017). Controlling linguistic style aspects in neural language generation. In *Proceedings from the Conference on Empirical Methods in Natural Language Processing (EMNLP) Workshop on Stylistic Variation* (pp. 94–104).  
<https://doi.org/10.18653/v1/W17-4912>
- Finn, A., & Kushmerick, N. (2006). Learning to classify documents according to genre. *Journal of the American Society for Information Science and Technology*, 57(11), 1506–1518.  
<https://doi.org/10.1002/asi.20427>
- Freund, L., Clarke, C. L., & Toms, E. G. (2006). Towards genre classification for IR in the workplace. In *Proceedings of the 1st International Conference on Information Interaction in Context* (pp. 30–36). ACM. <https://doi.org/10.1145/1164820.1164829>

- Fu, Z., Tan, X., Peng, N., Zhao, D. & Yan, R. (2018). Style transfer in text: Exploration and evaluation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- Gatt, A., & Krahmer, E. (2018). Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61, 65–170.
- Genkin, A., Lewis, D., & Madigan, D. (2006). Large-scale Bayesian logistic regression for text categorization. *Technometrics*, 49(3), 291–304. <https://doi.org/10.1198/004017007000000245>
- Giesbrecht, E., & Evert, S. (2009). Is part-of-speech tagging a solved task? An evaluation of POS taggers for the German web as corpus. In *Proceedings of the Fifth Web as Corpus Workshop* (pp. 27–35).
- Glover, A., & Hirst, G. (1996). Detecting stylistic inconsistencies in collaborative writing. In M. Sharples & T. van der Geest (Eds.), *The New Writing Environment* (pp. 147–168). London: Springer. [https://doi.org/10.1007/978-1-4471-1482-6\\_12](https://doi.org/10.1007/978-1-4471-1482-6_12)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems 27* (pp. 2672–2680).
- Goutte, C., & Gaussier, E. (2005, March). A probabilistic interpretation of precision, recall and F-score, with implication for evaluation. In *European Conference on Information Retrieval* (pp. 345–359). Berlin: Springer.
- Gries, S. T., & Mukherjee, J. (2010). Lexical gravity across varieties of English: An ICE-based study of n-grams in Asian Englishes. *International Journal of Corpus Linguistics*, 15(4), 520–548. <https://doi.org/10.1075/ijcl.15.4.04gri>
- Gries, S. T., Newman, J., & Shaoul, C. (2011). N-grams and the clustering of registers. *Empirical Language Research Journal*, 5(11).
- Grieve, J., Biber, D., Friginal, E. & Nekrasova, T. (2010). Variation among blogs: A multi-dimensional analysis. In *Genres on the Web* (pp. 303–322). Springer, Dordrecht. [https://doi.org/10.1007/978-90-481-9178-9\\_14](https://doi.org/10.1007/978-90-481-9178-9_14)
- Grieve, J., Biber, D., Friginal, E. and Nekrasova, T. (2010). Variation among blogs: A multi-dimensional analysis. In A. Mehler, S. Sharoff, & M. Santini (Eds.), *Genres on the Web* (pp. 303–322). Dordrecht: Springer. [https://doi.org/10.1007/978-90-481-9178-9\\_14](https://doi.org/10.1007/978-90-481-9178-9_14)
- Halliday, M., & Hasan, R. (1989). *Language, Context, and text: Aspects of language in a social-semiotic perspective*, 2nd ed. Oxford: Oxford University Press.
- Halliday, M. A., McIntosh, A., & Stevens, P. (1968). The users and uses of language. In J. Fischman (Ed.), *Readings in the sociology of language* (139–169). The Hague: Mouton. <https://doi.org/10.1515/9783110805376.139>
- Halliday, M. A. K., & Matthiessen, C. (2004). *An introduction to functional grammar*. London: Routledge.
- Hammerton, J., Osborne, M., Armstrong, S., & Daelemans, W. (2002). Introduction to special issue on machine learning approaches to shallow parsing. *Journal of Machine Learning Research*, 2, 551–558.
- Herring, S., Johnson, D. A., & DiBenedetto, T. (1995). This discussion is going too far!: Male resistance to female participation on the internet. In K. Hall & M. Bucholtz (Eds.), *Gender articulated: Language and the socially constructed self* (pp. 67–96). New York: Routledge.
- Herring, S. C., & Paolillo, J. C. (2006). Gender and genre variation in weblogs. *Journal of Sociolinguistics*, 10(4), 439–459. <https://doi.org/10.1111/j.1467-9841.2006.00287.x>
- Heylighen, F., & Dewaele, J. (1999). Formality of language: definition, measurement and behavioral determinants. *Internet Bericht, Center "Leo Apostel", Vrije Universiteit Brussel*.

- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Holmes, D. (1998). The evolution of stylometry in humanities scholarship. *Literary and Linguistic Computing*, 13(3), 111–117. <https://doi.org/10.1093/lc/13.3.111>
- Hoorn, J. F., Frank, S. L., Kowalczyk, W., & van Der Ham, F. (1999). Neural network identification of poets using letter sequences. *Literary and Linguistic Computing*, 14(3), 311–338. <https://doi.org/10.1093/lc/14.3.311>
- Hovy, E., Lavid, J., Maier, E., Mittal, V., & Paris, C. (1992). Employing knowledge resources in a new text planner architecture. In *Aspects of automated natural language generation* (pp. 57–72). Berlin, Heidelberg: Springer. [https://doi.org/10.1007/3-540-55399-1\\_5](https://doi.org/10.1007/3-540-55399-1_5)
- Hovy, E. H. (1990). Pragmatics and natural language generation. *Artificial Intelligence*, 43(2), pp.153–197. [https://doi.org/10.1016/0004-3702\(90\)90084-D](https://doi.org/10.1016/0004-3702(90)90084-D)
- Hovy, E. H. (1991). Approaches to the planning of coherent text. In R. Dale, E. Hovy, D. Rösner, & O. Stock (Eds.), *Natural language generation in artificial intelligence and computational linguistics* (pp. 83–102). Boston, MA: Springer. [https://doi.org/10.1007/978-1-4757-5945-7\\_3](https://doi.org/10.1007/978-1-4757-5945-7_3)
- Husson, F., Lê, S., & Pags, J. (2010). *Exploratory multivariate analysis by example using R*. London: Chapman & Hall CRC. <https://doi.org/10.1201/b10345>
- Jhamtani, H., Gangal, V., Hovy, E., & Nyberg, E. (2017). Shakespearizing modern language using copy-enriched sequence to sequence models. In *Proceedings of the Workshop on Stylistic Variation at EMNLP 2017* (pp. 10–19). <https://doi.org/10.18653/v1/W17-4902>
- Johansson, S., Leech, G. N., & Goodluck, H. (1978). *Manual of information to accompany the Lancaster-Oslo/Bergen Corpus of British English, for use with digital computer*. Oslo: Department of English, University of Oslo.
- Jolliffe, I. (2011). Principal component analysis. In M. Lovric (Ed.), *International encyclopedia of statistical science* (pp. 1094–1096). Berlin: Springer. [https://doi.org/10.1007/978-3-642-04898-2\\_455](https://doi.org/10.1007/978-3-642-04898-2_455)
- Kakkonen, T., & Sutinen, E. (2008). Coverage-based evaluation of parser generalizability. In *Proceedings of the Third International Joint Conference on Natural Language Processing, Volume-II*.
- Kan, M. Y., & McKeown, K. R. (2002). Corpus-trained text generation for summarization. In *Proceedings of the International Natural Language Generation Conference* (pp. 1–8).
- Kanaris, I., & Stamatatos, E. (2007). Webpage genre identification using variable-length character n-grams. In *Proceedings of the 19th IEEE International Conference on Tools with Artificial Intelligence* (p. 3–10). Washington, DC. <https://doi.org/10.1109/ICTAI.2007.107>
- Karlgren, J. (1999). Stylistic experiments in information retrieval. In T. Strzalkowski (Ed.) *Natural Language Information Retrieval* (pp. 147–166). Dordrecht: Springer. [https://doi.org/10.1007/978-94-017-2388-6\\_6](https://doi.org/10.1007/978-94-017-2388-6_6)
- Kešelj, V., Peng, F., Cercone, N., & Thomas, C. (2003). N-gram-based author profiles for authorship attribution. In *Proceedings of the Conference Pacific Association for Computational Linguistics, PACLING*, 3 (pp. 255–264).
- Kjell, B. (1994a). Authorship attribution of text samples using neural networks and Bayesian classifiers. In *IEEE International Conference on Systems, Man and Cybernetics*, San Antonio, TX. <https://doi.org/10.1109/ICSMC.1994.400086>
- Kjell, B., Woods, W. A., Frieder, O. (1995). Information retrieval using letter tuples with neural network and nearest neighbor classifiers. In *IEEE International Conference on Systems, Man and Cybernetics* (Vol., pp. 1222–1225). Vancouver, BC.

- Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., & Herbst, E. (2007). Moses: open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions*.  
<https://doi.org/10.3115/1557769.1557821>
- Kohavi, R. (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. In C. S. Mellish (Ed.). *Proceedings IJCAI-95*, 14(2), 1137–1145. Montreal, Quebec.
- Koller, D., Friedman, N., & Bach, F. (2009). *Probabilistic graphical models: Principles and techniques*. Cambridge, MA: The MIT press.
- Koppel, M., Argamon, S., & Shimoni, A. R. (2002). Automatically categorizing written texts by author gender. *Literary and Linguistic Computing*, 17(4), 401–412.  
<https://doi.org/10.1093/lc/17.4.401>
- Koppel, M., & Schler, J. (2003). August. Exploiting stylistic idiosyncrasies for authorship attribution. In *Proceedings of IJCAI'03 Workshop on Computational Approaches to Style Analysis and Synthesis* (Vol. 69, pp. 72–80).
- Koppel, M., Schler, J., & Zigdon, K. (2005). Determining an author's native language by mining a text for errors. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining* (pp. 624–628). ACM.  
<https://doi.org/10.1145/1081870.1081947>
- Langkilde-Geary, I. (2002). An empirical verification of coverage and correctness for a general-purpose sentence generator. In *Proceedings of the International Natural Language Generation Conference* (pp. 17–24).
- Lee, D. (2001). Genres, registers, text types, domains, and styles: Clarifying the concepts and navigating a path through the BNC jungle. *Language Learning and Technology*, 5(3), 37–72.
- Loehlin, J. C. (1998). *Latent variable models: An introduction to factor, path, and structural analysis*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Louwerse, M. M., & Graesser, A. C. (2004). Coherence in discourse. In P. Strazny (Ed.), *Encyclopedia of linguistics*. Chicago, IL: Fitzroy Dearborn.
- Lowe, D., & Matthews, R. (1995). Shakespeare vs. Fletcher: A stylometric analysis by radial basis functions. *Computers and the Humanities*, 29, 449–461.  
<https://doi.org/10.1007/BF01829876>
- Madigan, D., Genkin, A., Lewis, D. D., Argamon, S., Fradkin, D., & Ye, L. (2006). Author identification on the large scale. In *Proc. of Classification Society of N. America*, 2005.
- Mann, W. C., & Thompson, S. A. (1988). Rhetorical structure theory: Toward a functional theory of text organization. *Text-Interdisciplinary Journal for the Study of Discourse*, 8(3), 243–281. <https://doi.org/10.1515/text.1.1988.8.3.243>
- Marcu, D. (2000). *The theory and practice of discourse parsing and summarization*. Cambridge: The MIT press.
- Marcu, D. (1997). From local to global coherence: A bottom-up approach to text planning. In *AAAI/IAAI* (pp. 629–635).
- Martin, J. H., & Jurafsky, D. (2000). *Speech and language processing*. Englewood Cliffs, NJ: Prentice-Hall.
- Martin, J. R. (1992). *English text: System and structure*. Amsterdam: John Benjamins.  
<https://doi.org/10.1075/z.59>



- Matthews, R., & Merriam, T. (1993). Neural computation in stylometry: An application to the works of Shakespeare and Fletcher. *Literary and Linguistic Computing*, 8(4), 203–209. <https://doi.org/10.1093/lc/8.4.203>
- Matthiessen, C. M. I. M. (2015). Register in the round: Registerial cartography. *Functional Linguistics*, 2(1), 9. <https://doi.org/10.1186/s40554-015-0015-8>
- Matthiessen, C. M. I. M., & Teruya, K. (2015). Grammatical realizations of rhetorical relations in different registers. *Word*, 61(3), 232–281. <https://doi.org/10.1080/00437956.2015.1071963>
- McKeown, K., Kukich, K., & Shaw, J. (1994). Practical issues in automatic documentation generation. In *Proceedings of the Fourth Conference on Applied Natural Language Processing* (pp. 7–14). Stroudsburg, PA: Association for Computational Linguistics. <https://doi.org/10.3115/974358.974361>
- Merriam, T., & Matthews, R. (1994). Neural computation in stylometry II: An application to the works of Shakespeare and Marlowe. *Literary and Linguistic Computing* 9, 1–6. <https://doi.org/10.1093/lc/9.1.1>
- Moore, J. D., & Paris, C. L. (1993). Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics*, 19(4), 651–694.
- Morato, J., Llorens, J., Génova, G., & Moreiro, J. A. (2003). Experiments in discourse analysis impact on information classification and retrieval algorithms. *Information Processing & Management*, 39(6), 825–851. [https://doi.org/10.1016/S0306-4573\(02\)00081-X](https://doi.org/10.1016/S0306-4573(02)00081-X)
- Mosquera, A., & Moreda, P. (2012). A qualitative analysis of informality levels in web 2.0 texts: The Facebook case study. In *Proceedings of the LREC workshop: @ NLP can u tag# user generated content* (pp. 23–29).
- Nowson, S. (2006). The language of weblogs: A study of genre and individual differences. Unpublished PhD dissertation. University of Edinburgh.
- Nowson, S., Oberlander, J., & Gill, A. J. (2005). Weblogs, genres and individual differences. In *Proceedings of the 27th Annual Conference of the Cognitive Science Society* (pp. 1666–1671). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Och, F. J., & Ney, H. (2003). A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1), 19–51. <https://doi.org/10.1162/089120103321337421>
- Paiva, D. S., & Evans, R. (2005). Empirically-based control of natural language generation. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL05)* (pp. 58–65).
- Pavalanathan, U., Fitzpatrick, J., Kiesling, S., & Eisenstein, J. (2017). A multi-dimensional lexicon for interpersonal stancetaking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (pp. 884–895). Vancouver, CA: Association for Computational Linguistics.
- Pearl, J. (2009). *Causality*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511803161>
- Peng, F., Schuurmans, D., & Wang, S. (2004). Augmenting naive bayes classifiers with statistical language models. *Information Retrieval*, 7(3–4), 317–345. <https://doi.org/10.1023/B:INRT.0000011209.19643.e2>
- Power, R., Scott, D., & Bouayad-Agha, N. (2003). Generating texts with style. In *International Conference on Intelligent Text Processing and Computational Linguistics* (pp. 444–452). Berlin: Springer. [https://doi.org/10.1007/3-540-36456-0\\_47](https://doi.org/10.1007/3-540-36456-0_47)
- Prabhumoye, S., Tsvetkov, Y., Salakhutdinov, R., & Black, A. W. (2018). Style transfer through back-translation. In *Proceedings of Association for Computational Linguistics Conference*. Stroudsburg, PA: ACL.

- Quinlan, J. R. (2014). *C4. 5: Programs for machine learning*. Oxford: Elsevier.
- Raileanu, L. E., & Stoffel, K. (2004). Theoretical comparison between the GINI index and information gain criteria. *Annals of Mathematics and Artificial Intelligence*, 41(1), 77–93. <https://doi.org/10.1023/B:AMAI.0000018580.96245.c6>
- Rehbein, I., & Bildhauer, F. (2017). Data point selection for genre-aware parsing. In *Proceedings of the 16th International Workshop on Treebanks and Linguistic Theories* (pp. 95–105).
- Reiter, E., & Dale, R. (2000). *Building natural language generation systems*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511519857>
- Reiter, E., Sripada, S., Hunter, J., & Yu, J. (2005). Choosing words in computer-generated weather forecasts. *Artificial Intelligence*, 167(1–2), 137–169. <https://doi.org/10.1016/j.artint.2005.06.006>
- Reiter, E., & Williams, S. (2010). Generating texts in different styles. In S. Argamon, K. Burns, & S. Dubnov (Eds.), *The structure of style. Algorithmic approaches to understanding manner and meaning* (pp. 59–75). Heidelberg: Springer.
- Santini, M. (2005). Genres in formation? An exploratory study of web pages using cluster analysis. In *Proceedings of the 8th Annual Colloquium for the UK Special Interest Group for Computational Linguistics (CLUK05)*. Manchester, UK.
- Santini, M. (2006). Some issues in automatic genre classification of web pages. In *Proceedings of JADT 2006: 8èmes Journées Internationales d'Analyse statistique des Données Textuelles*.
- Santini, M. (2008). Zero, single, or multi? Genre of web pages through the users' perspective. *Information Processing & Management*, 44(2), 702–737. <https://doi.org/10.1016/j.ipm.2007.05.011>
- Sebastiani, F. (2002). Machine learning in automated text categorization. *ACM Computing Surveys (CSUR)*, 34(1), 1–47. <https://doi.org/10.1145/505282.505283>
- Sennrich, R., Haddow, B., & Birch, A. (2016). Controlling politeness in neural machine translation via side constraints. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 35–40).
- Sharoff, S., Wu, Z., & Markert, K. (2010). The Web Library of Babel: Evaluating genre collections. In *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10)*.
- Sheika, F. A., & Inkpen, D. (2012). Learning to classify documents according to formal and informal style. *Linguistic Issues in Language Technology*, 8(1), 1–29.
- Speelman, D., Gondelaers, S., & Geeraerts, D. (2006). A profile-based calculation of region and register variation: The synchronic and diachronic status of the two main national varieties of Dutch. In A. Wilson, D. Archer, & P. Rayson (Eds.), *Corpus Linguistics Around the World* (pp. 181–194). Amsterdam: Rodopi. [https://doi.org/10.1163/9789401202213\\_015](https://doi.org/10.1163/9789401202213_015)
- Stamatatos, E. (2008). Author identification: Using text sampling to handle the class imbalance problem. *Information Processing & Management*, 44(2), 790–799. <https://doi.org/10.1016/j.ipm.2007.05.012>
- Stamatatos, E., Fakotakis, N., & Kokkinakis, G. (2000). Automatic text categorization in terms of genre and author. *Computational Linguistics*, 26(4), 471–495. <https://doi.org/10.1162/089120100750105920>
- Svartvik, J., & Quirk, R. (1980). *A corpus of English conversation*. Lund: Gleerup.
- Szmrecsanyi, B. (2009). Typological parameters of intralingual variability: Grammatical analyticity versus syntheticity in varieties of English. *Language Variation and Change*, 21(3), 319–353. <https://doi.org/10.1017/S0954394509990123>

- Tambouratzis, G., Markantonatou, S., Hairetakis, N., Vassiliou, M., Tambouratzis, D., & Carayannis, G. (2000). Discriminating the registers and styles in the Modern Greek language. In *Proceedings of the workshop on Comparing corpora* (Vol. 9, pp. 35–42). Stroudsburg, PA: Association for Computational Linguistics.
- Teich, E., & Fankhauser, P. (2010). Exploring a corpus of scientific texts using data mining. *Language & Computers*, 71(1), 233–247.
- Teich, E., Degaetano-Ortlieb, S., Kermes, H., & Lapshinova-Koltunski, E. (2013). Scientific registers and disciplinary diversification: a comparable corpus approach. In *Proceedings of the Sixth Workshop on Building and Using Comparable Corpora* (pp. 59–68).
- The British National Corpus*, version 3 (BNC XML Edition). (2007). Distributed by Bodleian Libraries, University of Oxford, on behalf of the BNC Consortium. <<http://www.natcorp.ox.ac.uk/>>
- Tweedie, S. Singh, & Holmes, D.I. (1996). Neural network applications in stylometry: The Federalist Papers. *Computers and the Humanities*, 30(1), 1–10.  
<https://doi.org/10.1007/BF00054024>
- van Dijk, T. A. (1993). Stories and racism. In D. K. Mumby (Ed.), *Narrative and social control: Critical perspectives*. Newbury Park, CA: Sage. <https://doi.org/10.4135/9781483345277.n6>
- Vidulin, V., Luštrek, M., & Gams, M. (2007). Using genres to improve search engines. In *Proceedings of the International Workshop Towards Genre-Enabled Search Engines* (pp. 45–51).
- Waseem, Z., & Hovy, D. (2016). Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proceedings of Proceedings of NAACL-HLT 2016*. (pp. 88–93). Stroudsburg, PA: The Association for Computational Linguistics <<https://aclweb.org/anthology/N/N16/N16-2.pdf>>
- Waugh, S., Adams, A., & Tweedie, F. J. (2000). Computational stylistics using Artificial Neural Networks. *Literary and Linguistic Computing*, 15(2), 187–198.  
<https://doi.org/10.1093/lc/15.2.187>
- Xiao, R. (2009). Multi-dimensional analysis and the study of world Englishes. *World Englishes*, 28(4), 421–450. <https://doi.org/10.1111/j.1467-971X.2009.01606.x>
- Xu, W., Ritter, A., Dolan, B., Grishman, R., & Cherry, C. (2012). Paraphrasing for style. In *Proceedings of COLING 2012* (pp. 2899–2914).
- Zhao, Y. & Zobel, J. (2005). Effective and scalable authorship attribution using function words. In *Asia Information Retrieval Symposium* (pp. 174–189). Heidelberg: Springer.
- Zheng, R., Li, J., Chen, H., & Huang, Z. (2006). A framework for authorship identification of online messages: Writing-style features and classification techniques. *Journal of the American Society for Information Science and Technology*, 57(3), 378–393.  
<https://doi.org/10.1002/asi.20316>

## Address for correspondence

Shlomo Engelson Argamon  
Department of Computer Science  
Illinois Institute of Technology  
10 West 31st St.  
Chicago, IL 60616  
USA  
argamon@iit.edu