Identifying Personality Types Using Document Classification Methods

Mike Komisin and Curry Guinn

Department of Computer Science, University of North Carolina Wilmington
mkomisin@gmail.com, guinnc@uncw.edu

Abstract

Are the words that people use indicative of their personality type preferences? In this paper, it is hypothesized that word-usage is not independent of personality type, as measured by the Myers-Briggs Type Indicator (MBTI) personality assessment tool. In-class writing samples were taken from 40 graduate students along with the MBTI. The experiment utilizes naïve Bayes classifiers and Support Vector Machines (SVMs) in an attempt to guess an individual's personality type based on their word-choice. Classification is also attempted using emotional, social, cognitive, and psychological dimensions elicited by the analysis software, Linguistic Inquiry and Word Count (LIWC). The classifiers are evaluated with 40 distinct trials (leave-one-out cross validation), and parameters are chosen using leave-one-out cross validation of each trial's training set. The experiment showed that the naïve Bayes classifiers (word-based and LIWC-based) outperformed the SVMs when guessing Sensing-Intuition (S-N) and Thinking-Feeling (T-F).[1]

Introduction

The Myers-Briggs Type Indicator (MBTI) is the most widely used personality assessment tool in the world. According to Myers (1998), an individual has a natural preference in each dichotomy. The Myers-Briggs typology contains four functional dichotomies: (1) the Thinking-Feeling (T-F) dichotomy describes whether someone is logical in their judgments, or whether they base their decisions in personal or social values. (2) Judging-Perceiving (J-P) describes how an individual reveals themselves to the outside world. If an individual prefers Judgment, then they will reveal their Thinking or Feeling nature. If they prefer Perception, then they will reveal their Sensing or Intuitive nature. (3) Sensing-Intuition (S-N) reflects the two ways in which people are Perceiving --- a Sensing type will rely on the 5 senses and concrete observation while an Intuitive type will draw upon conceptual relationships or possibilities when gathering information. (4) Lastly, what Jung referred to as attitude, Extraversion-Introversion (E-I), deals with how a person focuses their energy and attention—whether outwardly focusing their perception or judgment on other people or inwardly focusing upon concepts and ideas, respectively.[1] Myers and Briggs work outlines 16 unique personality types using different combinations of the four bipolar continuums, or dichotomies (Center for Applications of Psychological Type [CAPT], 2010).

In our study, we hypothesize that word choice in personal essays is not independent of Myers-Briggs personality type. We will explore this hypothesis in two ways: (1) by examining the specific word choices used in corpus of essays of subjects whose MBTI scores are known; and (2) by examining the semantic categories of these words. Standard document classification techniques will be employed to differentiate essays based on these features. In this paper, we will contrast two classifiers: a probabilistic Naïve Bayes classifier and a Support Vector Machine (SVM).

Related Work

Pennebaker and King (1999) examined stream-of-consciousness (SOC) writings in terms of linguistic dimensions and personality trait. For their experiment, they utilized the Five Factor model, i.e. personality traits, as opposed to the MBTI personality types. The NEO-PI was utilized for correlation analysis with the MBTI (Furnham, 1996). From the study, correlations were observed between four of the five factors: Agreeableness with Thinking-Feeling; Conscientiousness with Judging-Perceiving;

[1] Note that this Jungian definition of the Extraversion/Introversion dichotomy differs from that given by the Five Factor model which tends to focus on sociability and outgoingness.

Extraversion with Extraversion-Introversion; and Openness with Sensing-Intuition. The dimension Neuroticism is not correlated with any of the Myers-Briggs categories, causing some criticism of the MBTI by McCrae and Costa (1989). These relationships also add to the credibility of the Myers-Briggs type theory in that both assessments are reliable indicators of personal dimensions that do not change much at all over time (Myers, 1998).

Pennebaker and King's experiment shows statistically significant correlation between four linguistic dimensions and the Five Factor scores of the authors. The four linguistic dimensions were derived by Pennebaker and King (1999) using principal component analysis on the Linguistic Inquiry and Word Count (LIWC) dimensions from 838 stream-of consciousness (SOC) writing samples. The four dimensions derived from the study were labeled: Immediacy, Making Distinctions, The Social Past, and Rationalization. Making Distinctions, for example, is a dimension comprised of four LIWC text categories: tentative (e.g. depends, guess, hopeful, luck), exclusive (e.g. but, either, or), negation (e.g. can't, hasn't, neither, not), and inclusive (e.g. and, with, both). Their work highlights correlations between the LIWC categories and the Five Factor scores for individuals. For example, three categories in the Making Distinctions dimension (tentative, exclusive, and negations) correlate negatively with Extraversion on the Five Factor scores.

Chung and Pennebaker (2008) used the LIWC to analyze self-descriptive essays. Through principle component analysis using varimax rotation, they were able to show that factor analysis on adjectives in the essays produced 7 factors which they found to be statistically significant indicators of the Big Five traits. Interestingly, some of the factors were unipolar and some exhibited bipolarity. Factors included 7 broadly labeled categories: Social, Evaluation, Self-Acceptance, Negativity, Fitting In, Psychological, Stability, and Maturity. The highest factor, Sociability, included self-descriptive adjectives like quiet, shy, outgoing, reserved, comfortable, open, friendly, and insecure. One interesting point is that participants that used words like shy, or quiet, were actually more likely to show positive correlation with Extraversion in the Five Factor scores. However, Pennebaker & King's (1999) analysis of stream-of-consciousness (SOC) writings suggested the statistically significant positive correlation between the LIWC Social category and Extraversion.

A Korean version of the Linguistic Inquiry and Word Count (KLIWC) was used to analyze eighty stream-of-consciousness writings with respect to Myers-Briggs and the Five Factor model (Lee et al., 2007). Lee et al. introduce correlations between the KLIWC and Myers-Briggs types, but the focus of the study is primarily on linguistic categories and does not provide a means of comparison across the same 64 psychological and contextual dimensions as the English LIWC.

Mairesse and Walker (2006) analyzed the same recorded speech as Pennebaker and Mehl using the machine-learning algorithm, RankBoost. Their findings showed they were able to successfully predict many of the Five Factor traits of individuals using several feature sets including LIWC, prosody (pitch, intensity, and speech rate), and with features gathered using the MRC psycholinguistic database.

Methodology

Data Collection

For our study, data was collected as part of a graduate course on conflict management in which students took the Myers-Briggs Type Indicator and completed a Best Possible Future Self (BPFS) exercise. Over 3 semesters, data was collected from 40 subjects.

Best Possible Future Self Writing (BPFS) Exercise

The Best Possible Future Self essay contains elements of self-description, present and future, as well as various contexts (e.g. work, school, family, finances). King's (2001) BPFS exercise is as follows:

> *Think about your life in the future. Imagine that everything has gone as well as it possibly could. You have worked hard and succeeded at accomplishing all of your life goals. Think of this as the realization of all of your life dreams. Now, write about what you imagined (p. 801).*

Word stemming

Word stemming is a simple way to reduce the feature set of a corpus, and, in doing so, reduce the sparseness of the data set. The simplest methods of word stemming use rule-based processes like dropping suffixes -e, -es, -ed, and -ing. In computer literature, the most commonly used stemming algorithm is Porter stemming (Porter, 1980). It is a rule-based approach to suffix stripping. The Natural Language Toolkit (Bird et al., 2009) provided an implementation of the Porter stemmer used in our trials.

Stop-words

Stop-words are words which generally act as syntactic sugar—for example, articles such as the, a, or an give little insight into the content of a document but make the meaning of content words more clear. Stop-word filtering has become commonplace in natural language processing, often improving the accuracy of word-based classifiers by eliminating common words which offer less contextual meaning. An English corpus of stop-words is included in the Natural Language Toolkit (Bird et al. 2009).

Smoothing

Data smoothing techniques account for terms in which a previously unseen word is encountered in the test case. In this study we tried a number of smoothing techniques including Lidstone smoothing, Good-Turing smoothing, and Witten and Bell Smoothing (Chen and Goodman, 1999).

Naïve Bayes

The naïve Bayes model utilizes a joint probability word distribution with priors calculated from the training set. Naïve Bayes accounts for prior probabilities of a class when attempting to make a classification decision. It uses what is referred to as the maximum a priori (MAP) decision rule to classify unseen texts using the word frequencies for each class's bag-of-words. In the study, each competing preference, e.g. Introversion vs. Extraversion, has its own bag-of-words to help classify the test documents. Leave-one-out cross validation is used to test the precision and recall of the classifier.

Support Vector Machines

Support Vector Machines have their roots in binary linear classifiers. Geometrically-based, they transform data into a higher-dimensional space using a kernel function, then, use a separating hyperplane to make a decision boundary. The decision boundary allows for previously unseen samples to be classified based on the hyperplane which separates attributes according to their associated training labels (Cortes & Vapnik, 1995).

Linguistic Inquiry and Word Count (LIWC)

The Linguistic Inquiry and Word Count program was used to provide an alternative feature set to the word-frequency set. LIWC processes multiple text files and returns a word-frequency distribution based on well-defined categories, e.g. money, social, positive, and negative. The LIWC utilizes sixty-four psychological and social dimensions which may overlap or aggregate. The hierarchy of categories begins with four dimensions: linguistic, psychological, relativity, and current concerns. These dimensions are comprised of multiple categories, and words may belong to more than one category.

Experiments

The data was collected over three semesters in 2010 and 2011 as part of a course on conflict management offered to graduate students. The data consists of two parts–the Myers-Briggs Type Indicator Step II (MBTI) results and the Best Possible Future Self (BPFS) essays. The MBTI Step II gives a more detailed classification than the MBTI, but only the primary dichotomies were utilized for classification. The BPFS exercise was given first.

Experimental Goals

In total, there are four MBTI personality type dichotomies: Extroversion or Introversion (E-I), Sensing or Intuition (S-N), Thinking or Feeling, Judging or Perceiving. Our classification problem is a binary decision for each dichotomy. Leave-one-out cross-validation is used as an unbiased approach to model selection (Elisseeff & Pontil, 2003). Forty students were given the MBTI Step II with a summary of these results given in Table 1. The MBTI reports include a clarity index for each dichotomy. The scale of these clarity scores range from 0 to 30. The clarity score reflects the consistency of an individual to convey a given preference within the questionnaire (Myers, 1998). Thus, a low clarity score means that the assessment is less clear, and a high clarity score denotes that the questionnaire's decision is very clear for a given preference (see Table 2 for a representative sample). In one experiment, clarity scores were used to select a subset of participants for text classification. In another, we used the entire set of subjects.

Myers-Briggs Preferences	Est. Population Distribution	Sample Distribution
Extraversion	49.0%	65.0% (26)
Introversion	51.0%	35.0% (14)
Sensing	70.0%	52.5% (21)
Intuition	30.0%	47.5% (19)
Thinking	45.0%	47.5% (19)
Feeling	55.0%	52.5% (21)
Judging	58. 0%	60.0% (24)
Perceiving	43.0%	40.0% (16)

Table 1: Population and Sample Distributions by Preference (sample counts in parentheses).

Personality Type	E	I	S	N	T	F	J	P
ESFJ	24		9			17	28	
ISTJ		5	2		5		14	
ISTJ		14	24		1		22	
ENFP	24			21		24		18
ISFJ		26	4			8	1	
ENFJ	19			23		4	17	
ESFJ	12		8			4	21	
ENFP	25			10		10		18
ESTP	26		5		11			30
ISTP		1	13		18			6

Table 2: Sample MBTI scores of some participants.

The Best Possible Future Self essays were categorized based on the corresponding MBTI score and a summary of the token and unique word count is presented in Table 3. WPD is words tokens per document; WTD is unique word types per document.

Myers-Briggs Dichotomy	Word Tokens	Unique Words	Avg. WPD	Avg. WTD
Extraversion	10428	1859	401	72
Introversion	5275	1140	377	81
Sensing	7913	1455	377	69
Intuition	7790	1594	410	84
Thinking	6879	1348	362	71
Feeling	8824	1685	420	80
Judging	6210	1389	388	87
Perceiving	9493	1649	396	69

Table 3: Text features of BPFS essays.

Myers-Briggs Dichotomy	Word Tokens	Unique Words	Avg. WPD	Avg. WTD
Extraversion	5631	1376	217	53
Introversion	2834	846	202	60
Sensing	4335	1067	206	51
Intuition	4130	1178	217	62
Thinking	3718	1015	196	53
Feeling	4747	1224	226	58
Judging	3312	1030	207	64
Perceiving	5153	1207	215	50

Table 4: Text features of BPFS essays after Porter stemming and stop-word filtering.

Word-based Naïve Bayes Classification

The naïve Bayes classifier utilizes a bag-of-words input space, i.e. a count of the token occurrences for each given training set. Each preference is tested and represented using a bag-of-words. For example, Extraversion and Introversion will each be assigned a bag-of-words and be pitted against one another using the MAP decision rule. The term occurrences of the training set data are used to calculate the conditional probabilities of the word types, given their labeling. The distribution of prior probabilities (i.e. the probability that an arbitrary document belongs to a given class) will be calculated from the training data, as well. Preliminary analysis revealed there were no significant differences in classification when using the data set of words prior to Porter stemming and stop-word filtering versus the data set after those filters. Therefore, to reduce the feature space we did employ both filters. Preliminary analysis also revealed that for our dataset, Lidstone smoothing performed better than either Witten-Bell or Good-Turing smoothing. Two trials are conducted for each classifier, one using the entire dataset and one using the subset based on highest clarity scores (the top 75th percentile of samples ordered by clarity score for each dichotomy).

Word-based Naïve Bayes Classification Results

First, an analysis was conducted using simple Bayesian statistics gathered for the words present in the training corpus and tested using leave-one-out cross validation. The results are presented in Figure 1. Precision and recall scores were near 70% for Sensing (S) and Intuitive (N) types. We then sampled from our corpus only data whose MBTI clarity scores were in the top 75th percentile of the sample. Figure 2 is a presentation of the results of using a Naïve Bayes classifier on that reduced dataset. By eliminating the data with low clarity, the results are much better with higher precision and recall scores for three dichotomies: Sensing/Intuitive, Thinking/Feeling, and Judging/Perceiving.

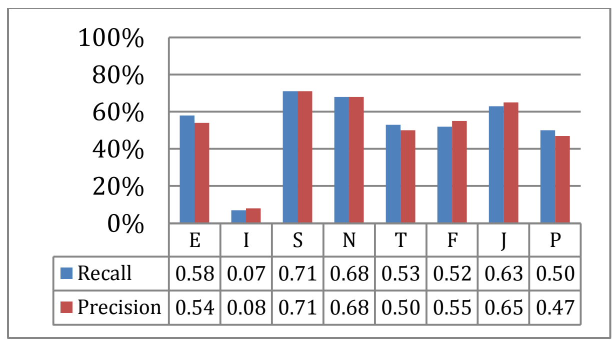

Figure 1: Word-based classification results using naïve Bayes with leave-one-out cross validation over entire data set with sample size, n = 40.

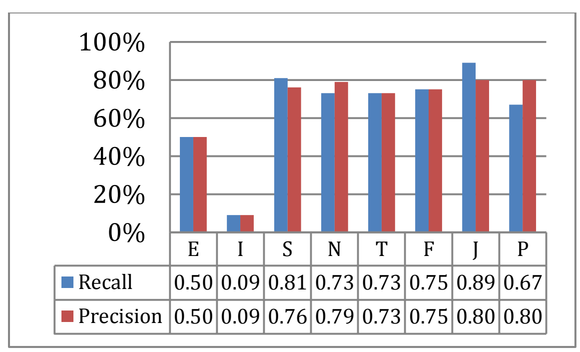

Figure 2: Word-based classification results using naïve Bayes with leave-one-out cross validation over reduced data set based on clarity scores with sample size, n = 30.

LIWC-based naïve Bayes Classification Results

Our next hypothesis was to test whether the semantic categories of the words, as specified by the LIWC database, would also provide the ability to predict MBTI personality type. Using a similar methodology applied to individual words, we used the word categories (e.g., positive emotion, work, money, family). Again, we used the full dataset with leave-one-out testing, and then used a modified dataset with 25% of the data removed with the lowest clarity scores. The results of this experiment are presented in Figure 3 and Figure 4. The precision and recall rates are not as strong compared with the use of stemmed word tokens in a naïve Bayes classifier.

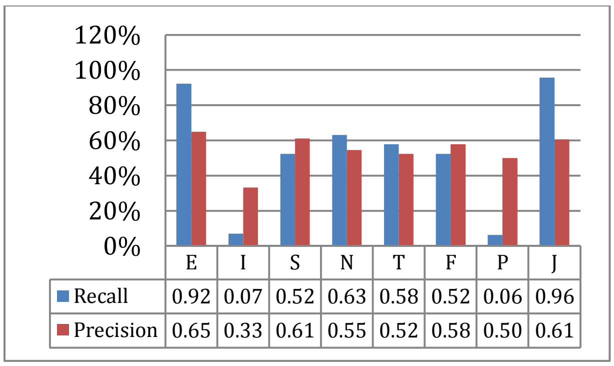

Figure 3: LIWC-based classification results using naïve Bayes with leave-one-out cross validation over entire data set with sample size, n = 40.

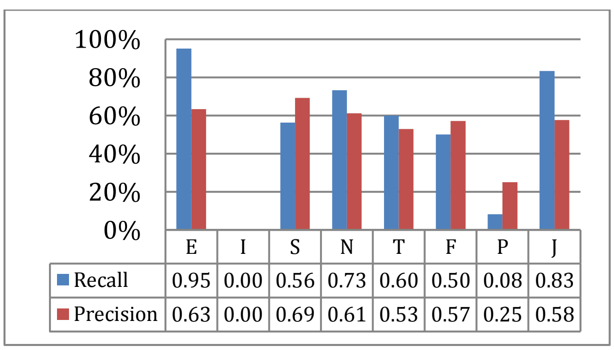

Figure 4: LIWC-based classification results using naïve Bayes with leave-one-out cross validation over reduced data set based on clarity scores with sample size, n = 30.

SVM Classification Results

Our next experiment was to determine whether a non-linear classifier might perform better than the stochastic Bayesian classifier. Two support vector machine classifiers were used: one using the individual words as features and the other classifier using LIWC word categories as features. As in the previous experiments, we looked at the full dataset and the reduced dataset with the 75% highest clarity scores. We only present the results from the reduced dataset (Figure 5 and Figure 6) because the performance was significantly worse than using the Naïve Bayes approach. The combination of the large feature space of the words appearing in each dichotomy and the small training set resulted in very poor performance of the word-based classification problem. Also, the SVM classifiers were vulnerable to the unbalanced training sets (i.e., the Introversion/Extroversion set and the Judging/Perceiving set) – a known problem with SVMs. The SVM performance using the LIWC features (a smaller feature space than the number of unique words) was somewhat better. In future experiments, we plan to use alternative classifiers and a larger corpus.

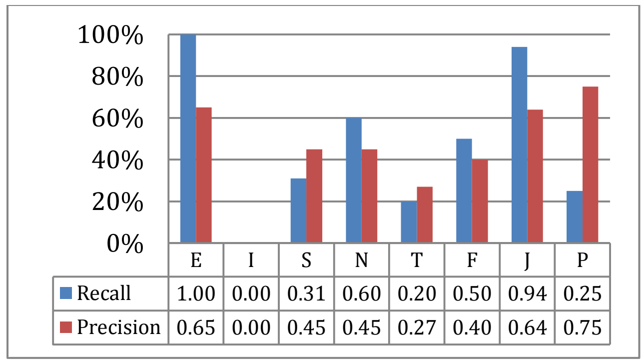

Figure 5: LIWC-based classification results using SVM with leave-one-out cross validation over reduced data set based on clarity scores with sample size, n = 30.

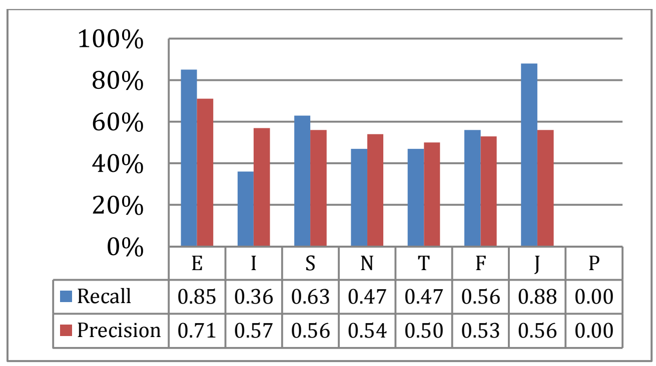

Figure 6: LIWC-based classification results using naïve Bayes with leave-one-out cross validation over reduced data set based on clarity scores with sample size, n = 30.

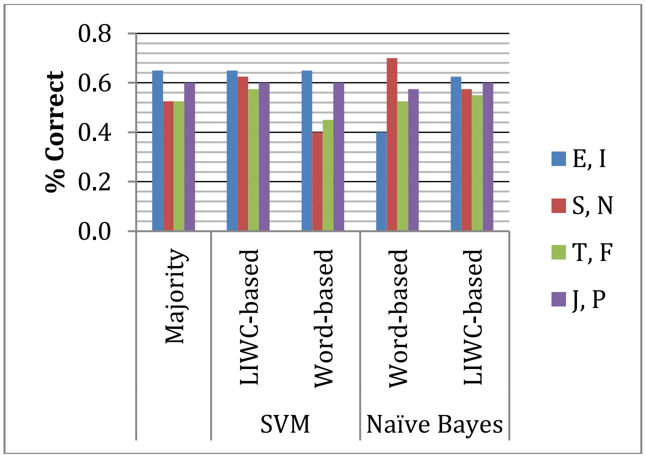

Figure 7: Summary of results with leave-one-out cross validation and sample size, n = 40

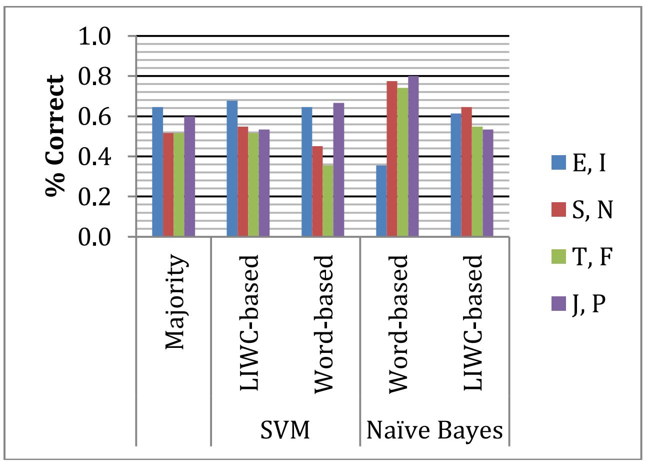

Figure 8: Summary of results with leave-one-out cross validation and reduced sample size, n = 30; lowest clarity scores removed

Conclusions

Figure 7 and Figure 8 summarize the results of our trials. Using a relatively small dataset and leave-one-out cross-validation, a Naïve Bayes classifier using a bag-of-words approach was able to predict three of the Myers-Briggs personality dichotomies with substantially higher precision and recall than those obtained using a simple majority discussion. The dichotomies of Sensing/Intuitive and Thinking/Feeling were predicted with precision and recall above 75% when using the dataset with the top 75% of clarity scores. The Judging/Perceiving category was predicted with higher precision but slightly lower recall rates. Using LIWC features resulted in less successful predictions although the Sensing/Intuitive dichotomy was moderately distinguishable. SVMs did not prove to be a useful classifier with our dataset. Both the Extrovert/Introvert set and the Judging/Perceiving set were unbalanced. SVMs are known to be vulnerable to training errors when presented with unbalanced datasets. Other researchers suggest that the poor performance of the SVMs could be attributed to the sparseness of a data set relative to a large number of features (Ng & Jordan, 2002). Thus, if SVMs generally perform better on densely populated data sets, a possible indication of such an attribute is the more balanced performance during LIWC-based SVM classification compared with word-based SVM trials. In our future work, we will examine Random Forest classifiers, which may perform better given the constraints of our dataset.

References

Bird, S., Klein, E. & Loper, E. (2009). *Natural Language Processing with Python – Analyzing Text with the Natural Language Toolkit*. O'Reilly Media.

Center for Applications of Psychological Type [CAPT]. (2010). *Jung's theory of psychological types and the MBTI® instrument*. Retrieved from http://www.capt.org/take-mbti-assessment/mbti-overview.htm

Chen, S. F. & Goodman, J. (1999). An empirical study of smoothing techniques for language modeling. *Computer Speech & Language*, 13.4, 359-93.Chung, C.K. & Pennebaker, J.W. (2008). Revealing dimensions of thinking in open-ended self-descriptions: An automated meaning extraction method for natural language. *Journal of Research in Personality*, 42.1, 96-132.

Cortes, C. & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20.3, 273-97.

Elisseeff, A. & Pontil, M. (2003). Leave-one-out error and stability of learning algorithms with applications. *Learning Theory and Practice*. IOS Press.

Furnham, A. 1996. The big five versus the big four : The relationship between the Myers-Briggs Type Indicator (MBTI) and NEO-PI five factor model of personality. Personality and Individual Differences 21 :303 - 307.

King, L. A. (2001). The health benefits of writing about life goals. *Personality and Social Psychology Bulletin*, 27.7, 798-807.

Lee, C., Kim, K., Seo, Y., & Chung, C. (2007). The relations between personality and language use. *The Journal of General Psychology*, 134.4, 405-13

Mairesse, F., Walker, M. (2006) *Automatic Recognition of Personality in Conversation*. In *Proceedings of HLT-NAACL* 2006, New York City, June 2006.

Myers, I. (1998). *MBTI Manual: A guide to the development and use of the Myers-Briggs Type Indicator*. Palo Alto, CA: Consulting Psychologists.

Pennebaker, J. W. & King, L. A. (1999). Linguistic styles: Language use as an individual difference. *Journal of Personality and Social Psychology*, 77.6, 1296-312.

Porter, M.F. (1980). An algorithm for suffix stripping, *Program*, 14.3, 130–137.