

A Corpus-Independent Feature Set for Style-Based Text Categorization

Moshe Koppel
Navot Akiva
Ido Dagan

{koppel, navot, dagan}@cs.biu.ac.il
Computer Science Department, Bar Ilan University, Ramat Gan 52900, Israel

Abstract

We suggest a corpus-independent feature set appropriate for style-based text categorization problems. To achieve this, we introduce a new measure on linguistic features, called *stability*, which captures the extent to which a language element, such as a word or syntactic construct, is replaceable by semantically equivalent elements. This measure may be perceived as quantifying the degree of available “synonymy” for a language item. We show that frequent but unstable features are especially useful for style-based text categorization.

1. Introduction and Background.

Style-based text categorization tasks, such as authorship attribution [Mosteller and Wallace, 1964; Holmes, 1995], are in a sense orthogonal to the more common problem of categorization by topic [Lewis and Ringuette, 1994; Schutze et al., 1995]. For style-based categorization, we seek features that are roughly invariant within the documents of a given author (or, more broadly, style class) but variant from author to author. Typically, a promising set of discriminating features is chosen and then training examples from each category and some machine-learning algorithm are employed to produce a model for categorizing.

Since the number of potential discriminating features is often uncomfortably large, feature selection methods are sometimes used to select out a particularly promising set of features. A number of these methods, conveniently summarized in [Sebastiani, 2002], have proved to be reasonably successful for topical categorization. Typically, the features selected by these methods are those that, individually, discriminate well on the training corpus. While such methods do tend to eliminate useless features, they sometimes do harm by pre-empting the learning algorithms they are meant to serve: the learning algorithms themselves, by taking into account dependencies among features, eliminate useless features

in more subtle ways than these direct feature selection methods. Furthermore, in the case of style-based categorization, these methods can lead to over-fitting by focusing attention on features that are highly correlated with one of the authors for reasons that might be specific for the training corpus but unrelated to the author’s generic style.

A different approach to feature selection for style-based categorization is that used by Mosteller and Wallace [1964] in their seminal work on authorship discrimination. Mosteller and Wallace [1964] simply chose a set of features that are not dependent on the training corpus but rather have certain appropriate universal properties. In their case, the features chosen were function words, which were deemed topic-independent. This is a perfectly sensible approach but it is somewhat limiting: it does not offer the flexibility of ranking features so that more or fewer can be chosen and it limits consideration in advance to lexical features.

In this paper, we propose a more general method for choosing a corpus-independent feature set for style-based text-categorization. We introduce a measure of the “stability” of a linguistic feature: roughly speaking, a feature is stable if it can’t be replaced in a text without changing the meaning of the text. Conversely, the feature is “unstable” precisely to the extent that alternatives (e.g. synonyms) to it are available which do not alter the meaning of the text. We will show that unstable features (of which many function words are prime examples) constitute an excellent universal feature set for style-based text categorization. This makes intuitive sense: stable features are ones which don’t offer viable meaning-preserving alternatives so that differences in usage of stable features between authors more likely reflect irrelevant differences in content than differences in style; differences in usage of unstable features are likely to reflect different stylistic choices. Since instability can be ranked, we can choose more or fewer features as is needed, possibly using cross-validation to optimize. Moreover, we shall see that, unlike function words, the stability criterion can be applied to any type of linguistic feature, whether lexical, syntactic, complexity-based, etc.

2. Linguistic Feature Stability

Data-driven approaches to text and language processing apply various quantitative measures to linguistic elements, such as words and terms. These measures capture important properties of language items and are often utilized in various ways for different computational tasks, such as information retrieval, text classification, terminology extraction and statistical thesaurus construction.

In this paper, we introduce a new type of measure for linguistic elements that we call *meaning-preserving stability*, or *stability* for short. It captures the degree to which a linguistic element or construct can be substituted for in a piece of text without affecting the text meaning. This measure is applied most intuitively to words and terms, but is applicable also to other types of linguistic constructs, such as part-of-speech sequences. Stability captures interesting properties in language, and thus seems interesting from a purely scientific point of view, and also has potential use within applied tasks. Stability can be perceived as quantifying the typical degree of available “synonymy” for a language element, while generalizing the notion of synonymy to any type of linguistic feature rather than being restricted to its common use for words only.

To make this a bit more concrete, let us begin with a simple example. Consider the three sentences:

1. *John was lying on the couch next to the window.*
2. *John was reclining on the sofa by the window.*
3. *John had been lying on the couch near the window.*

The three sentences convey (approximately) the same message. Some of the words remain invariant in all three versions (*John, window*), while others are replaced by other words (*was : had been, couch : sofa, next to : by : near, lying : reclining*) which don't significantly change the meaning of the sentence. More generally, consider features of a text such as word or phrase frequencies, frequencies of syntactic structures or any other feature which can be measured in a given text. Roughly speaking, the *stability* of such a feature is the extent to which that feature tends to remain invariant across different texts that convey the same meaning. For example, proper nouns are very stable, while words with many synonyms are unstable.

Obtaining training material for measuring stability is a difficult challenge, requiring alternative versions of texts that carry, ideally exactly, the same meaning. “Parallel” (monolingual) texts of this nature have been used for automatic paraphrase extraction, using either multiple translations of the same text or news stories from different sources that describe the same event [Barzilay and McKeown 2001; Shinyama et al. 2002]. The disadvantage of these types of training material is that they rely on the

fact that different people have manually created different versions of the same text. Such manually created versions are not easily available for most texts (indeed, we did not have such corpora available). In order to lighten the supervision requirements, we have used a machine translation (MT) system to translate our training corpus back-and-forth via multiple languages, thus obtaining several English versions of the same text. The use of automatic translation to generate training data is problematic for several reasons, as discussed below. Yet, the combination of MT with other technologies that can leverage the stability measure, such as style-based text categorization, provides appealing advantages that reduce the overall dependency on particular training materials.

In the remainder of the paper, we first present the general notion of stability and define a concrete stability measure. We then present empirical measurements of stability and illustrate the properties that it captures. Finally, we demonstrate how stability can be utilized effectively for feature selection in style-based text classification.

3. A Stability Measure

A *stability measure* would be a quantitative measure that correlates with a feature's tendency to be preserved across different meaning-preserving variants of a text. We formally define a specific stability measure and consider empirical results of stability experiments on the Reuters 21578 corpus [Lewis, 1997].

Let $\{d_1, d_2, \dots, d_n\}$ be a set of documents (or text segments) and let $\{d_i^1, d_i^2, \dots, d_i^m\}$ be $m > 1$ different versions of d_i where the meanings of the m versions are all roughly identical.¹ For any measurable feature c , let c_i^j be the value of c in document d_i^j . Let us develop the final formula step by step.

Step 1. Define the stability of a feature in multiple versions of a single document.

Let $k_i = \sum_j c_i^j$. Then the *stability* of a feature c in document d_i is defined to simply be the usual entropy measure:

$$H(c_i) = -\sum_j [(c_i^j / k_i) \log (c_i^j / k_i)] / \log m.$$

If c_i^j is a frequency measure, we can think of c_i^j / k_i as the probability that a random appearance of c in d_i is in version d_i^j . Then $H(c_i)$ is just the usual entropy measure. (We normalize by $\log m$ to keep the range of stability

¹ Since stability is an empirical measure based on statistical data, its definition does not depend on the exact notion of “meaning preservation” that holds across the different text versions. Applying the stability measure would make sense whenever the different text versions represent alternative interchangeable linguistic choices.

values to [0,1].) Thus, for example, if a feature assumed the identical value in every version of a document, its stability would be 1. If a feature assumed a positive value in a single version of the document but 0 in all others, its stability would be 0.

Step 2. Extend the definition to multiple documents $\{d_1, d_2, \dots, d_n\}$.

The impact that a given document has on the overall measure is defined to be proportional to the average value of the feature in the various versions of that document. (Thus, for example, when c is a frequency measure of some attribute, documents in which the attribute is more frequent will contribute proportionately more to the overall stability of the feature in the corpus.)

Let $K = \sum_i k_i$. Then

$$S(c) = \sum_i [(k_i/K) * H(c_i)]$$

This formula can be transformed to the equivalent formula:

$$S(c) = \{\sum_i [k_i \log k_i - \sum_j c_j^i \log c_j^i]\} / K * \log m \quad (1)$$

4. Measuring Feature Stability Empirically Using Machine Translation

In order to empirically measure stability, one needs several text variants that convey the same meaning, or at least have substantial overlap. For example, we might consider translations of the same text by different translators, and to a much lesser extent, reports on the same event by different reporters or journalists. All these are hard to obtain for experimental purposes. One interesting way to generate text variants artificially is to use a machine translation (MT) system to translate the text to another language and then translate it back to the original language. For our experiments, we used SystranPro to translate each document in the Reuters 21578 corpus into each of five different languages (French, German, Spanish, Italian, Portuguese) and back into English. In order to check that the measure is not overly dependent on the base corpus, we did the same to a set of several hundred book length documents from the British National Corpus (BNC). We applied formula (1) to frequencies of words, parts-of-speech n-grams and other linguistic features.

An obvious weakness of this experiment is that it is subject to the idiosyncrasies of Systran and, to a lesser extent, of the particular corpus. We have attempted to mitigate this affect by using five translation packages. This is only a partial solution since it is likely that all of them share certain underlying methods and programming code. Yet, as the experiments below indicate, this setting does provide an interesting “grasp” of feature stability, probably because the vast knowledge encoded in the five translation systems does capture much of the inherent phenomena that determine linguistic stability.

We will consider the stability distributions of various classes of features. We will also consider some specific features and discuss why their respective stabilities are particularly high or low.

5. Stability Distributions and Examples




Figure 1. Stability histogram of all words. The x-axis denotes stability value ranges and the y-axis denotes the proportion of features receiving the corresponding stability values.

In Figure 1, we show a histogram of stabilities of all single words in the Reuters corpus. As is evident the number of words in a given stability range descends as the mean of the range descends. When we look at specific word classes a clearer picture emerges. As might be expected, certain features, such as proper names, are highly stable. All proper names have stability above 0.9. Other word classes are considerably less stable.

FW	Reuters	BNC	FW	Reuters	BNC
and	0.99	0.99	on	0.85	0.93
to	0.98	0.98	from	0.77	0.79
the	0.97	0.99	by	0.76	0.75
in	0.97	0.98	been	0.66	0.72
a	0.97	0.99	at	0.61	0.72
not	0.95	0.97	which	0.56	0.58
is	0.94	0.95	like	0.54	0.74
of	0.94	0.95	has	0.52	0.27
with	0.90	0.97	over	0.47	0.54
it's	0.87	0.88	out	0.43	0.61

Table 1. Examples of function words ranked by their respective stabilities using Reuters and BNC, respectively.

In Figure 2(a-c), we show histograms of stability of nouns, verbs and function words, respectively. While nouns follow the general pattern with a plurality in the highest stability range, verbs and function words distribute more normally. This is because, in English, additional semantic information – such as tense – is encoded in verbs, so that verbs are often replaceable with related forms. Likewise, function words are often replaceable with equivalent syntactic constructions.

In Table 1, we consider a selection of function words and their respective stabilities. A number of these examples are very instructive. Words like *and* and *the* don't offer more natural alternatives and are thus stable.




Figure 2(a). Stability histogram for nouns




Figure 2(b). Stability histogram for verbs




Figure 2(c): Stability histogram for function words




Figure 3: Stability histogram for part-of-speech trigrams

However, words like *has* and *been* are unstable because, for example, the present perfect *has been* is easily replaced by the past tense *was*. Similarly, a word like *over* is easily replaced either by synonyms (e.g. *above*) or alternative constructions (e.g. *go over there : go to there*). Note that with some notable exceptions (*has*, *like*), the stabilities yielded by Reuters and by BNC are very close to each other.

Let's now consider features other than words. In Figure 3, we show a histogram of stability on parts-of-speech trigrams. In Table 2, we consider a selection of trigrams of parts-of-speech and their respective stabilities. Note, for example, that the trigram *noun_noun_noun* is very unstable. A typical occurrence is "U.S. construction spending", a tightly-wound phrase invariably unwound into something looser like "spending on construction by the U.S."

Altogether, it can be seen that the stability measure captures in a unified way different types of semantic "uniqueness" that are related to both syntactic and lexical phenomena, including content words, function words and part-of-speech sequences.

6. Applying Stability to Style-based Text Categorization

We wish to test the hypothesis that the class of frequent but unstable features constitutes a useful universal feature set for authorship attribution experiments. In the first experiment, we will attempt to learn to distinguish the writing style of Anne Bronte from that of her sister, Charlotte Bronte. This is a particularly difficult attribution problem because the authors came from identical social and linguistic backgrounds and wrote in what appear to be very similar styles. We consider two books by Anne Bronte (*Agnes Gray*, *The Tenant of Wildfell Hall*) and two by Charlotte Bronte (*The Professor*, *Jane Eyre*). Each book is divided into between 100 and 150 equal-sized passages. We train on passages of one book by each author and test on passages of the remaining two books. We then run the experiment again with the training sets female authors. We will attempt to learn the gender of a document's author [Koppel *et al.* 2003]. We test the effectiveness of our learning methods using five-fold cross-validation.

POS	Reuters Stability	BNC Stability
NN_IN_DT	0.94	0.97
IN_DT_NN	0.94	0.98
DT_NN_IN	0.93	0.97
NN_IN_NN	0.92	0.89
IN_NN_IN	0.86	0.86
NN_IN_NNP	0.81	0.89
IN_CD_NN	0.76	0.72
CD_NN_IN	0.72	0.58
JJ_NN_NN	0.67	0.80
NN_NN_NN	0.43	0.59

Table 2. Examples of noun-related part-of-speech triples ranked by their respective stabilities using Reuters and BNC, respectively.

Word	Reuters Instability	ReutersFreq. $[F_R]$	$IN*F_R(*10^3)$
of	0.063	0.079	4.97
the	0.027	0.102	2.75
has	0.499	0.004	1.99
which	0.455	0.003	1.36
from	0.276	0.004	1.10
said	0.072	0.013	0.93
at	0.455	0.002	0.91
by	0.273	0.003	0.81
you	0.777	0.001	0.77
had	0.308	0.002	0.61

Table 3. Highest ranked words according to $IN*F_r$.

We begin with a feature set consisting of all features and then eliminate more and more features according to various criteria. For each reduced feature set we will use a learning algorithm to build a categorization model and then test the model on the chapters in the test set. We use Balanced Winnow [Littlestone, 1988; Dagan et al., 1997] as our learning method.




Figure 4(a): Categorization accuracy of Balanced Winnow on the Bronte corpus using six feature reduction methods to select single words. The x -axis represents the number of features used and the y -axis records accuracy.




Figure 4(b): Categorization accuracy of Balanced Winnow on the Bronte corpus using six feature reduction methods to select single words and POS triples. The x -axis represents the number of features used and the y -axis records accuracy.

In Figure 4(a), we show results on the Bronte experiment using Balanced Winnow on an initial feature set consisting of all 3500+ words that appear at least 4 times in the Bronte corpus. Feature reduction is performed by ranking features according to various measures (listed below). Since only 1300 words appear both in this list and in the Reuters corpus (which was used for measuring feature stability) the first data point we show for a reduced feature set is 1300. In Figure 4(b), we repeat the experiment using an initial feature set consisting of the full word list as well as all parts-of-speech triples that appear at least 5 times in the Bronte corpus. Each of the reduced sets consists of an equal number of words and parts-of-speech triples.

As our benchmark, we use the *odds ratio* (OR) measure [Mladenic, 1998; Ruiz and Srinivasan, 2002; Caropreso et al., 2002] which is a typical and particularly successful representative of discrimination-based feature reduction [Sebastiani 2002]. For a given feature c , let c_j be the frequency of c in category j and let c_j be the frequency of c in all other categories². Then OR ranks features according to the score $\sum_j c_j(1-c_j)/c_j(1-c_j)$. Other methods rearrange the same basic ingredients in different ways.

Altogether, six measures were used for ranking features:

OR – odds ratio in training corpus

F_t – average frequency in training corpus

$OR * F_t$ – odds ratio in training corpus * average frequency in training corpus

IN – instability ($= 1 - S(c)$)

$IN * F_t$ – instability * average frequency in training corpus

$IN * F_r$ – instability * average frequency in Reuters

In Figures 5(a) and (b), we show the analogous results for five-fold cross-validation experiments on the gender experiment.

Note that, in both experiments, without taking feature frequency into account both OR and IN fail miserably: too many rare features are ranked highly by each method. More importantly, these experiments show that, although $IN * F_r$ is completely independent of the training corpus, it is actually the best measure by which to choose features. The highest ranked features in the corpus according to the $IN * F_r$ ranking are shown in Table 3. It is also evident that

² Technically, the frequency used in OR as defined by Caropreso et al (2002) is the proportion of documents in the category in which the feature appears at least once. Among common features, which appear in every document, OR as so defined offers no discrimination. Thus it is inappropriate for style-based categorization. Instead, we use the actual relative frequency of the feature in the category.

(a) feature selection does lead to improvement over using the complete feature set (i.e. letting Winnow select features implicitly); (b) optimal performance is maintained with a quite small feature set. In fact, in both




Figure 5(a): Categorization accuracy of Balanced Winnow on the gender problem in BNC using six feature reduction methods to select single words. The x-axis represents the number of features used and the y-axis records accuracy.




Figure 5(b): Categorization accuracy of Balanced Winnow on the gender problem in BNC using six feature reduction methods to select single words and POS triples. The x-axis represents the number of features used and the y-axis records accuracy.

experiments, the best 400 features selected according to this criterion are better than using a standard list of 400 function words, which achieves 81% accuracy on Bronte and 72% on the gender problem.

Note also that when the number of features approaches the bottom of our range, it is better to use the frequency of a feature in the training corpus than in Reuters. This is because many of the highly-ranked features based on Reuters frequency do not appear sufficiently frequently in the training/testing corpora to yield optimal categorization. This suggests that a broader-based corpus than Reuters is advisable.

Overall, these results suggest that the set of features identified as useful by IN^*F_r may constitute, to a substantial degree, a compact *universal* feature set for style-based categorization. This set might be thought of as a generalization of a universal feature set like a list of function words, with the added advantages that the features can be of any type (not only lexical) and that they can be ranked.

7. Conclusions and Future Work

We have shown how the extent to which a linguistic feature may be replaced with a semantically equivalent feature can be effectively quantified. The resulting stability measure is useful for identifying promising candidates for style-based text categorization. The specific technique that we used for estimating stability – back and forth translation – is, admittedly, flawed due to its dependence on a particular software package but is, nevertheless, novel and effective. The results of our experiments indicate that we might regard frequent and unstable features as a generalization of function word lists that can serve as a universal feature set for style-based text categorization. Many more experiments – on other corpora, on more feature types, on other learning algorithms and using other existing feature reduction methods as benchmarks – must be performed to strengthen and extend these conclusions.

References

[Barzilay and McKeown, 2001] Barzilay, R. and McKeown, K. 2001. Extracting Paraphrases from a Parallel Corpus. Proceedings of the Annual Meeting of the Association for Computational Linguistics (pp. 50–57).

[Caropreso *et al.*, 2001] Caropreso, M. F., Matwin, S. and Sebastiani, F. 2001. A learner-independent evaluation of the usefulness of statistical phrases for automated text categorization. In A. G. Chin Ed., *Text Databases and Document Management: Theory and Practice*, pp. 78–102. Hershey, US: Idea Group Publishing.

[Dagan *et al.*, 1997] Dagan, I., Karov, Y. and Roth, D. 1997. Mistake-Driven Learning in Text Categorization.

In *Proceeding of Second Conference on Empirical Methods in Natural Language Processing (EMNLP-2)*, Providence, Rhode Island.

[Holmes, 1995] Holmes, D. 1995. Authorship attribution. *Computers and the Humanities*, 28.

[Koppel *et al.*, 2003] Koppel, M., Argamon S. and Shimoni, A. R. 2003. Automatically Categorizing Written Texts by Author Gender. *Literary and Linguistic Computing* 17(3), in press.

[Lewis, 1997] Lewis, D. 1997. *Reuters-21578 text categorization test collection*.

[Lewis and Ringuette, 1994] Lewis, D. and Ringuette, M. 1994. A Comparison of Two Learning Algorithms for Text Categorization. *Proceedings of the Third Annual Symposium on Document Analysis and Information Retrieval* (pp. 81–93), Las Vegas, USA.

[Littlestone, 1988] Littlestone, N. 1988. Learning quickly when irrelevant features abound: A new linear-threshold algorithm. *Machine Learning*, 2.

[Mladenic, 1998] Mladenic, D. 1998. Feature subset selection in text learning. In *Proceedings of ECML-98, 10th European Conference on Machine Learning* (pp. 95–100).

[Mosteller and Wallace, 1964] Mosteller, F. and Wallace, D. L. 1964. *Applied Bayesian and Classical Inference: The Case of the Federalist Papers*. Springer.

[Ruiz and Srinivasan, 2002] Ruiz, M. and Srinivasan, P. 2002. Hierarchical text classification using neural networks. *Information Retrieval* 5, 1, 87–118.

[Schutze *et al.*, 1995] Schutze, H., Hull, D. A. and Pedersen, J. O. 1995. A comparison of classifiers and document representations for the routing problem. In *Proceedings of SIGIR-95, 18th ACM International Conference on Research and Development in Information Retrieval* (Seattle, US, 1995), pp. 229–237.

[Sebastiani, 2002] Sebastiani, F. 2002. Machine learning in automated text categorization. *ACM Computing Surveys*, 34(1), 2002.

[Shinyama *et al.*, 2002] Shinyama, Y., Sekine, S., Sudo, K. and Grishman, R. 2002. Automatic Paraphrase Acquisition from News Articles. Proceedings of Human Language Technology Conference, San Diego, USA.

[Yang and Pedersen, 1997] Yang, Y. and Pedersen, J. O. 1997. A comparative study on feature selection in text categorization. In *Proceedings of ICML-97, 14th International Conference on Machine Learning* (pp. 412–420), Nashville, USA.