

# Topic Modeling in Word Prediction for AAC

**Keith Trnka, Debra Yarrington, Kathleen McCoy**  
Computer Science Department  
University of Delaware  
Newark, DE 19716  
{trnka,yarringt,mccoy}@cis.udel.edu

**Christopher Pennington**  
AgoraNet, Inc.  
314 E. Main St., Suite 1  
Newark, DE 19711  
penningt@agora-net.com

## ABSTRACT

Word prediction is a method for enhancing the communication ability of persons with speech and language impairments. In this work, we explore one method of adjusting the language model based on the content of a conversation.

## Keywords

Word prediction, keystroke savings, alternative and augmentative communication (AAC), topic modeling, language modeling

## INTRODUCTION

Alternative and Augmentative Communication (AAC) is the field of research concerned with finding ways to help those with speech difficulties communicate more easily and completely. Today there are approximately 2 million people in the United States with some form of communication difficulty. One means to help ease communication is the use of an electronic communication device, many of which have synthetic speech as output. However, one issue in using a communication device is that communication rate is generally slower than the common speaking rate. Whereas speaking rate is estimated at 180 words per minute (wpm) and experienced typists can manage 100 wpm, a disabled user's input rate is estimated at roughly 15 wpm (Beukelman and Mirenda, 1998; Copestake, 1997; Newell et. al., 1998). Thus one goal of developers of AAC devices is to find ways to increase the rate of communication output. Developers cannot significantly increase the rate of input of users, but instead we seek to increase the amount of communication output by requiring less input for common tasks.

This paper investigates the use of a word prediction system. In word prediction, we assume that the user enters letters using a standard keyboard. The system predicts full words that are likely to be desired, and provides them to the user for selection with one additional keystroke.

A word prediction system predicts the word currently being typed on the basis of what has already been typed. Suppose that the user wants to enter "I want a home in the country." After typing, "I want a h", they might see something like shown below. The system has created a

*prediction window* containing the five words which it thinks the current word is most likely to be. In this example, the user can press F2 to complete the word "home" and the system will automatically enter a space afterwards. So in this example, the user needed 3 keystrokes to enter what would normally take 5 keystrokes, when the space is considered.

```
I want a h
-----
           house   F1
           home    F2
           hot     F3
           horse   F4
           hoogie  F5
```

*Figure 1: An example of what a word prediction interface might look like*

The prediction list can vary in length, but most systems tend to use lists of length between five and seven. The prediction list can occur in-line (it can appear within the line being entered as it is entered), or it can occur somewhere separately on the interface screen. For row-column scanning devices, the word list often appears in an extra row or column on the scanning grid.

It is difficult to judge how much word prediction can speed communication rate. Much of this determination is dependent on both the characteristics of the user, such as their physical and cognitive abilities, and characteristics of the user interface, such as where the prediction list is displayed and how a word in the list is selected.

It quickly becomes apparent that many factors affect the efficacy of word prediction, and more studies are needed to determine the effect different factors have on the success of word prediction. It is equally apparent, however, that unless the word prediction system is able to successfully predict words and thus decrease the number of keystrokes necessary, other factors are irrelevant. Therefore, before tackling the added issues involved in user interfaces, it is instructive to look at the percentage of keystrokes saved, since this measure provides an upper bound on any communication rate increases from word

prediction. Thus our work here concentrates on investigating keystroke savings in word prediction.

Our long-term goal is to investigate methods for increasing keystroke savings in word prediction by taking various amounts of contextual information into account during the prediction process. Of course, in doing this we must have a way of evaluating whether or not our various attempts at capturing contextual information are fruitful - so first we must establish a baseline prediction system and a method for calculating keystroke savings against which our future systems can be tested.

Clearly this paper is not the first to discuss the use of word prediction in AAC (Bogges, 1988; Carlberger et. al., 1997; Copestake, 1997; Fazly and Hirst, 2003; Garay-Vitoria and González-Abascal, 1997; Leshner and Rinkus, 2002), and each of these systems has been presented with evaluations. However, as explained in (Trnka et. al., 2005), we find results difficult to interpret due to the lack of a standardized approach in evaluation. Similarly, this paper is not the first to discuss topic modeling (Bellegarda, 2000; Florian and Yarowsky, 1998/1999, Seymore and Rosenfeld, 1997; Chen et. al., 1998) However, this is only the second paper to discuss topic modeling in the context of word prediction. (Leshner and Rinkus, 2002) explored this area of research, but used an impractical evaluation.

In this paper we first give some background in statistical approaches to word prediction. We also give the details of our baseline word prediction system. We present the full details of our approach to integrating topic modeling into word prediction, then present comparative evaluations between the two systems. Finally, we discuss future improvements to the language model used for word prediction and conclude.

## METHODS

Like other approaches, we apply statistical language modeling techniques to word prediction (Bogges, 1988; Carlberger et. al., 1997; Copestake, 1997; Fazly and Hirst, 2003; Garay-Vitoria and González-Abascal, 1997; Leshner and Rinkus, 2002). Our baseline is a pure n-gram based method. (Copestake, 1997; Fazly and Hirst, 2003; Garay-Vitoria and González-Abascal, 1997) additionally integrated syntactic knowledge into their language models, which was found to improve prediction somewhat.

Basic word prediction treats each sentence of the user's conversation as independent. At any given point in the sentence, the user has some word that they are typing. At that point, the word prediction system presents a list of words, called a *prediction window*, that the user may be typing. If the desired word appears in the list, the user selects it using a key reserved for that position in the list. If the word doesn't appear in the list, the user must enter another character. Then the word prediction system updates the list and the process repeats.

To present the user with a list of possible words, the word prediction system needs to know all of the words in the language. The vocabulary is constructed by considering all words that occur in some training corpus. If a word being typed isn't in the vocabulary, it can't be predicted.

The second requirement is a statistical language model. The purpose of the language model is to compute the probability  $P(\text{word} \mid \text{history})$ , where the history is the words that have already been entered. Given a vocabulary and a language model, the list of predictions is generated as follows: The vocabulary is first filtered to remove words that don't match the partially entered word. For example, if 'a' has been entered, the system will only consider words beginning with 'a'. Then this list of candidates is sorted by  $P(\text{word} \mid \text{history})$ . The top  $W$  words are presented to the user, where  $W$  is the prediction window size.

The remainder of this section is devoted to the construction of a statistical language model.

## Corpus

Statistical approaches require a collection of text to construct a language model. Ideally, our corpus would be a collection of conversations involving one or more people using an AAC system. Such a corpus is unavailable, so we follow (Leshner and Rinkus, 2002) in using the Switchboard corpus, which is a collection of telephone conversations and their transcriptions.<sup>1</sup> The corpus is divided into two sections: one section to train the statistical language model and one section to evaluate the word prediction method. The training section contains a randomly selected 2217 conversations and the testing section contains the remaining 221 conversations.

Because the corpus is a collection of transcribed conversations, it contains many speech repairs. Consider the example

*is there um an- is there a like a code of dress ...*

However, a person in a textual conversation would write

*is there a code of dress ...*

(Hindle, 1983) categorized these sorts of self-corrections and integrated his solution into a parser. In short, his approach identified the words that are being corrected and the correction. The words being corrected are then removed for syntactic processing. Hindle's full set of editing rules was not practical for our purposes, so we implemented a subset of the rules: remove uh/um, exact repetitions, and repetitions in which the last word of the corrected part is abandoned, as in "an-". These editing rules bring the Switchboard conversations closer to what we envision an AAC user would type. This agrees with the

<sup>1</sup> The Switchboard transcriptions were available from <http://www.isip.msstate.edu/projects/switchboard/>

preprocessing that (Leshner and Rinkus, 2002) did according to personal communication with Dr. Leshner.

### Baseline Language Model

N-grams have been shown to give better performance with larger amounts of training text and roughly better performance with larger  $n$ . (Leshner et. al., 1999; Manning and Shütze, 2000) For an introduction to n-grams, refer to (Manning and Shütze, 2000) or (Jurafksy and Martin, 2000). The main weakness of n-grams is that they require substantial amounts of data. Using higher-order n-grams such as trigrams increases the problem of data sparseness. To combat this, we implement backoff with Good-Turing smoothing, the current best practice in statistical language modeling according to Manning and Shütze.

### Backoff

Backoff is applied here much like in (Katz, 1981). The idea is that the probability will be determined by trigram probability if the trigram probability is defined, otherwise we'll try the bigram probability. If that's undefined, we'll try the unigram probability. The crux of backoff is adjusting these probabilities so that all probabilities sum to one.

In order to apply backoff, some probability must be removed from trigrams which were seen in the corpus. The process of removing probability from seen trigrams and giving it to unseen trigrams is called smoothing and is described in the following section. For now, appreciate that each conditional trigram distribution will sum to less than one.

$$1 - \delta = \sum_{\text{word} \in V} P(\text{word} | \text{word}_{-2} \text{word}_{-1})$$

Then the probability obtained by backoff is defined as

$$P'(w | w_{-2} w_{-1}) = \begin{cases} P(w | w_{-2} w_{-1}) & \text{if } P(w | w_{-2} w_{-1}) > 0 \\ \delta \times P'(w | w_{-1}) & \text{otherwise} \end{cases}$$

Similarly, we define the bigram backoff probability as:

$$P'(w | w_{-1}) = \begin{cases} P(w | w_{-1}) & \text{if } P(w | w_{-1}) > 0 \\ \delta \times P'(w) & \text{otherwise} \end{cases}$$

Note that  $\delta$  is different for bigrams and trigrams.

### Smoothing

Good-Turing smoothing can be summarized as taking probability away from low-frequency words and giving it to words that never occurred in the training corpus. For a technical account of Good-Turing smoothing, see (Manning and Shütze, 2000) or (Jurafksy and Martin, 2000). For the intuitive account, refer to the graph below, which shows the probability of a trigram computed

normally and computed after applying Good-Turing smoothing.




Figure 2: Smoothing reduces the probability of low frequency events and gives it to unseen events.

### Sentence-initial Model

The trigram backoff model just described is applied to all words in a sentence but the first word. The first word in the sentence is modeled using a special unigram model to capture the notion that sentences tend to begin with a small set of words such as determiners and discourse markers. If the probability of a given word is undefined, then we backoff to the normal unigram model, using the same type of backoff as for bigrams and trigrams.

Our approach for treating the beginning of a sentence differs slightly from other researchers: The common method employed is to add a fake word at the beginning of a sentence and skip this symbol in training and testing; relying on it as a symbol to signal the particular context of a word. In fact, our approach and the traditional approach will perform identically on the very first word of a sentence. However, if both are trigram approaches, they will differ slightly on the second word of a sentence. As tetragrams (4-grams) are generally considered intractable, the two approaches will be identical for the remainder of the sentence. This difference reflects our intuition that only the second word in a sentence is influenced almost entirely by the preceding word; adding an additional conditioning event (the special symbol) would only serve to make the distribution of possible second words more sparse.

### Topic Modeling

The goal of topic modeling is to identify the current topic of conversation, then increase the probability of related words and decrease the probability of unrelated words. The topic identification and topic application steps can be treated somewhat independent of one another, but they both depend on the representation of a topic. Like

other researchers, we represent a topic using a collection of text. In this work, we use a collection of conversations to represent a topic, while others have primarily used collections of documents. However, researchers differ on the best way to represent a collection of topics: Some researchers have created a hierarchical collection of documents (Florian and Yarowsky, 1999), while others have created a set of topics (Mahajan et. al., 1999; Bellagarda, 2000; Seymore and Rosenfeld, 1997). The primary advantage of a hierarchical approach is that a more general topic can be selected when the topic of conversation is difficult to identify, and can be very specific when possible. However, as Mahajan et. al. point out, a conversation may simultaneously discuss multiple topics. For instance, a discussion of the economic impacts of baseball involves multiple topics. The set approach is more amenable to allowing the selection of multiple topics. We feel that the primary lure of a hierarchical approach, the ability to generalize, can be captured in the set approach as well. Rather than wait to construct a language model until the topic of conversation is known, we precompute a language model for each of the topics in our corpus.

#### Topic Identification

The current topic of conversation must be identified from the part of the conversation that has taken place so far, and updated periodically in the conversation. Further, we must devise a representation for a partial conversation that is conducive to identifying the topic. We follow other researchers in maintaining something like a unigram model of the current conversation and applying document similarity measures to identify appropriate topics.

In representing the conversation so far, we choose to implement an exponentially decayed cache, like (Bellegarda, 2000), using TF-IDF values rather than raw frequencies. The algorithm proceeds as follows: When a word is added to the cache, first all words that have occurred previously are decayed. This is done by multiplying by  $\lambda$ , a constant between 0 and 1. For our experiment, we followed Bellegarda in using .975. Secondly, the new word is added with its IDF as the weight instead of 1. As our approach is for topic identification, we ignore words that occur in 85% or more of the topics, with the intuition that such words are function words.

As the step to bridge our model of the current conversation to the application of our topic model, we compute the document similarity between the cache and the unigram model for each topic. We do not feel that the words in the language model for each topic needed to use the TF-IDF method, as incorporating IDF again would be redundant. More importantly, however, is the capacity for variation in computing document similarity. Many different measures of similarity have been used previously in literature, see (Lee, 1999) for a comparison in the context

of smoothing using terms that appear in similar contexts. However, we chose to use the cosine metric, following (Florian and Yarowsky, 1999).

#### Topic Application

Given that we have computed a similarity score for each topic to the current conversation, there are two main variations on how to construct a new language model. (Mahajan et. al., 1999) chose to implement a  $k$ -nearest solution, constructing the topic model from the most similar  $k$  topics. Each topic's language model was weighted equally for their experiments. We chose to follow Florian and Yarowsky's approach. (Florian and Yarowsky, 1999) They expand the probability for a word ( $w$ ) given a history ( $h$ ) as follows:

$$P(w|h) = \sum_{t \in \text{topics}} P(t|h) \times P(w|t, h)$$

$P(w|t, h)$  is simply the probability of  $w$  taken from the language model constructed for topic  $t$ . The probability of the topic involves a little work to estimate:

$$P(t, h) \approx \frac{S(t, h)}{\sum_{i \in \text{topics}} S(i, h)}$$

where  $S(t, h)$  is the similarity of the topic to the current part of the conversation.

#### Discussion

In this approach, all topics will have some influence, albeit minor, over the interpolated language model. This avoids the problem noted by (Leshner and Rinkus, 2002) that selecting the single best topic model leads to data sparseness problems, which nullify the benefit of having a better expectation of the language use. Now that the model is explained in detail, reconsider the activation of multiple topics: in this model, it simply means that multiple topics will have high probabilities. Additionally, we feel that this model has the ability to generalize topics. In the event that the context isn't very specific, we propose that the keywords in the conversation will reflect a more general category of topics. Given that this is the case, the more specific versions of this general topic should all be weighted somewhat higher than unrelated topics. Also, we feel that this model encapsulates the boosting and depressing of probabilities used by (Chen et. al., 1998). Apart from the use of Latent Semantic Analysis, there is one major deviation from Bellegarda's approach. (Bellegarda, 2000) In their work, the topic modeling is used to determine the vocabulary, and then integrated with a standard language model. Our approach, however, allows for stylistic and collocational peculiarities of a topic to be interpolated into the overall language model.

## PRACTICAL CONSIDERATIONS

Although we are primarily interested in investigating whether or not topic modeling improves keystroke savings in word prediction, a number of practical considerations have both influenced our approach as well as our evaluation.

### Smoothing and Topic Modeling

Smoothing to use for backoff can be applied at two different points in processing with topic modeling. Firstly, each individual topic language model may be smoothed. The alternative is to smooth the interpolated topic model. As there may be many topics, it's likely that each topic model will be rather sparse, whereas the interpolated model will be less sparse. If we apply smoothing before interpolation, bigrams and trigrams may be given less weight than is appropriate for the interpolated model. It is for this reason that we choose to apply smoothing after interpolation. Normally, this would lead to increased time of execution, but our implementation only smooths distributions that are reached in testing, as they are reached. However, smoothing requires frequencies rather than probabilities. To solve this, we approximate the interpolated topic model by using frequencies in the interpolation equation instead of probabilities.<sup>2</sup> However, the interpolated frequencies are not whole numbers do to the interpolation process. We address this problem by creating bins of size one to use for smoothing. This is equivalent to taking the ceiling of each frequency before smoothing.

### Number of Topics and k-nearest Topics

We had originally intended to follow (Mahajan et al., 1999) in abandoning the notion of an explicit topic. Their approach treats each document in the training set as it's own topic, and relies upon the interpolation model to select the relevant topics and make appropriate generalizations. This approach is attractive for two main reasons. Firstly, they achieve an astounding 37.6% reduction in test set perplexity over a baseline trigram model. More importantly, however, their approach doesn't require a document or conversation to be labeled with a topic. This avoids the problem of either finding a corpus annotated for topic or the problem of performing appropriate automatic topic clustering.

Unfortunately, our preliminary tests revealed that interpolation of a topic model from about 2,200 smaller topic language models was too computationally demanding for the resources available to us. We realize now that this is

<sup>2</sup> The problem which this causes is that topics are now not only weighted by their topical similarity, but by their size, as topics with higher frequencies will tend to dominate the model. We are looking into the severity of this problem at the moment.

one of the main advantages of using a *k*-best approach in interpolation. To help alleviate the computational demands, we instead used the topic assignments given in the Switchboard corpus. This reduces the number of topics to roughly 40.

### Realistic Dynamic Topic Modeling

As we intend our research to be used for practical word prediction systems, we designed our topic model as practically as possible. That means, unlike Mahajan et. al., we only have access to the parts of the conversation that have been already said. Their approach assumes that the first 100 words of a conversation are available a priori. Also, unlike (Leshner and Rinkus, 2002), we do not assume to know the topic of a conversation ahead of time in testing. Their investigation used Switchboard, like ours, but read the topic number of each conversation to select the appropriate topic in testing.

However, dynamically re-interpolating the topic model is computationally expensive. Although we re-interpolate the topic modeling as the conversation proceeds, we only recompute the model after every second turn to mitigate the computational cost of interpolation.

### Memory Requirements

We also discovered that topic modeling requires a substantial amount of memory. This could be alleviated somewhat by using a complex data structure such as a character-based trie, but that would further decrease runtime speed. Instead, we decided to only use bigrams and unigrams for topic modeling. We evaluate the impact of this decision in the following section.

## EVALUATION

For a more thorough account of evaluation methods and the impact of user interface decisions on evaluation, see (Trnka et. al., 2005). In this work, the speak key is included in all simulations. Although immediate prediction gives much better keystroke savings, we use delayed prediction unless otherwise specified because it runs faster.

Like our other evaluations, the user interface simulates prediction using window sizes of 1-10 and adds a space after the word when prediction is used. All evaluation is done using keystroke savings (KS):

$$KS = \frac{\text{keys}_{\text{normal}} - \text{keys}_{\text{with prediction}}}{\text{keys}_{\text{normal}}}$$

Due to the computational cost of topic modeling, the evaluations presented here are preliminary; they are the results of evaluation on 2 out of the 221 conversations available for testing. The baseline for comparison is a bigram model unless otherwise specified. Rather than

present the results of all ten window sizes simulated, we present results of what we feel are the common sizes.

### Which sort of topic modeling is better?

As our first simulation, we chose to compare the  $k$ -best conversations approach that Mahajan et. al. used with the approach of using the topics as annotated by Switchboard, which we call *true topic modeling*. For these tests, we use delayed prediction.

<i>Window size</i>	<i>Baseline</i>	<i>100-best</i>	<i>700-best</i>	<i>True topic</i>
3	48.6%	46.3%	48.8%	49.2%
5	51.7%	49.5%	51.5%	52.2%
6	52.7%	50.5%	52.5%	53.0%

Table 1: The keystroke savings of different forms of topic modeling at three window sizes.

Here, we find that true topic modeling outperforms the baseline and both  $k$ -best approaches, although the improvement over the baseline is somewhat small. However, given that the maximum keystroke savings for this test set is 59.5%, the average error reduction due to true topic modeling is 5.4%

### Immediate Prediction

Although the benefit from topic modeling with delayed prediction is somewhat small, we have demonstrated in (Trnka et. al., 2005) that delayed prediction is partially limited by the user interface. Below is a comparison between the baseline and topic modeling using immediate prediction.

<i>Window Size</i>	<i>Baseline</i>	<i>True topic</i>	<i>Error reduction</i>
3	54.0%	54.9%	3.6%
5	58.6%	59.2%	3.0%
6	60.0%	60.4%	2.1%

Table 2: Comparison of the keystroke savings of true topic modeling to the baseline using immediate prediction. Error reduction is also shown for reference.

With more room for improvement, topic modeling produces more improvement at each window size than under delayed prediction. Although the absolute improvement in keystroke savings is greater, the error reduction is lower, as the limit under immediate prediction is 78.9% savings for this test set. However, we can offer no explanation for why error reduction is reduced as the window size increases.

### Comparison to a Trigram Baseline

It is clear that topic modeling improves word prediction in this study, although at great computational cost. Here, we compare the bigram-based true topic model to a trigram model under immediate prediction.

<i>Window Size</i>	<i>Bigrams</i>	<i>Trigrams</i>	<i>True topic</i>
3	54.0%	55.0%	54.9%
5	58.6%	59.3%	59.2%
6	60.0%	60.6%	60.4%

Table 3: Comparison of keystroke savings between bigram model, trigram model, and true topic model using immediate prediction.

Topic modeling is marginally outperformed by a trigram model in terms of keystroke savings. However, the trigram model processed the testing set at 1,230 word per minute (wpm), whereas topic modeling was much slower at 132 wpm.

### Subjective Evaluation

Although topic modeling doesn't improve the savings due to word prediction much, manual inspection of the simulation shows that the prediction lists were generally more appropriate to the current topic of conversation. The follow examples are taken from Switchboard conversation 2001, in which the participants were asked to discuss corporate dress codes. The excerpts are taken after the topic model had been interpolated once. The desired word is shown bold.

well i w
was
would
went
<b>work</b>
want
wouldn't
<b>Example 1, with the bigram baseline</b>

```
well i w
      was
      would
      work
      went
      want
      wouldn't
```

*Example 1, with true topic modeling*

```
so i usually w
              when
              we
              what
              watch
              work
              wear
```

*Example 3, with the bigram baseline*

The results are clear: topic modeling boosted only the list placement of the on-topic word and left all other words in the same order. In fact, the probability of “work” is 1.71 times higher with topic modeling.

```
we have to d
           do
           drive
           death
           deal
           decide
           dress
```

*Example 2, with the bigram baseline*

```
so i usually w
              wear
              when
              what
              we
              watch
              work
```

*Example 3, with true topic modeling*

Example 3 shows the desired, on-topic word being boosted up the prediction list, while the rest of the list remains mostly the same. The only other difference is that “when” and “we” switched positions. In this case, “wear” is 5.23 times more likely with the topic model than without it.

```
we have to d
           do
           dress
           drive
           deal
           death
           decide
```

*Example 2, with true topic modeling*

In example 2, “dress” is brought far up the list by topic modeling. The other change is that “death” is ranked below “deal”, which seems reasonable given the topic. In this case, the probability of dress with topic modeling is 4.73 times higher.

## FUTURE WORK

The approach we took to topic modeling was neither a great success nor a great failure. To be certain of its performance, we will run tests on a larger test set. We also plan to vary the algorithm slightly. In particular, we will experiment with other ways of fixing the smoothing-topic interaction that don't give certain topics too much weight. Additionally, we would like to evaluate Bellegarda's (Bellegarda, 2000) approach by only adjusting a unigram model by topic and integrating it into a trigram model. Additionally, we plan to evaluate the end-user effects of topic modeling. We hypothesize that the more appropriate predictions will reduce cognitive load, and thus speed up typing rate. User studies will be necessary to validate or invalidate our intuition.

The idea of language modeling is to understand what words to use in what situations. Clearly, n-gram methods do not capture the full extent of why a person chooses to say particular words. Topic modeling helps to model more of what isn't captured by n-grams, namely, long-distance and variable-distance dependencies between words. However, much work remains in adequately accounting for the particular use of words in conversation. One of the several alternative approaches we intend to investigate is style modeling. While topic modeling will

primarily influence the vocabulary usage, we envision that a style model would influence higher-order n-grams in a similar fashion to topic modeling. Another avenue of research we intend to investigate is the integration of part-of-speech modeling into our topic language model. Other researchers have investigated part-of-speech modeling in the context of word prediction (Copestake, 1997; Fazly and Hirst, 2003; Garay-Vitoria and González-Abascal, 1997), and found that it slightly improves word prediction. In particular, we noticed that the topic model's prediction lists were largely good predictions with the exception that several predictions were of an incorrect part-of-speech. We postulate that this is due to the use of backoff to the unigram level. If part-of-speech were either integrated into the backoff model or used separately, we think topic modeling would be more useful.

Finally, our long term goal is to incorporate several different and complimentary language models into a single word prediction engine. It is our intuition that word choice in natural language is determined by several different factors. For example, terminology is clearly affected by topic. However, choice amongst synonyms is largely determined by user preference and style.

## CONCLUSIONS

Topic modeling is a clear, but small, improvement over an n-gram model of the same order. We think that our approach to topic modeling can be improved by adequately addressing the smoothing-topic problem, as described in the previous section. However, our evaluation must be performed on a larger test set to draw reliable conclusions.

## ACKNOWLEDGMENTS

We would like to thank the US Department of Education for funding this research under grant H113G040051, under the National Institute on Disability and Rehabilitation Research program. Also, thanks to Gregory Lesher for correspondence regarding his work.

## REFERENCES

1. Bellegarda, Jerome. Large Vocabulary Speech Recognition with Multispan Language Models. *IEEE Trans. On Speech and Audio Processing*, 8(1): 2000.
2. Beukelman, D. R. and Mirinda, P. *Augmentative and alternative communication: Management of severe communication disorders in children and adults*. Baltimore: Paul H. Brookes Publishing Co., 1998.
3. Boggess, Lois. Two simple prediction algorithms to facilitate text production. *Proceedings of Applied Natural Language Processing*, 1988.
4. Carlberger, Alice, John Carlberger, Tina Magnuson, M. Sharon Hunnicutt, Sira Palazuelos-Cagigas, and Santiago Aguilera Navarro. Profet, A New Generation

- of Word Prediction: An Evaluation Study. *Proceedings of Natural Language Processing for Communication Aids*, 1997.
5. Copestake, Ann. Augmented and alternative NLP techniques for augmentative and alternative communication. *Proceedings of Natural Language Processing for Communication Aids*, 1997.
6. Fazly, Afsaneh and Graeme Hirst. Testing the Efficacy of Part-of-Speech Information in Word Completion. *Proceedings of the 10<sup>th</sup> Conference of the European Chapter of the Association for Computational Linguistics*, 2003.
7. Florian, Radu and David Yarowsky. Dynamic Nonlocal Language Modeling via Hierarchical Topic-Based Adaptation. *Proceedings of the 37<sup>th</sup> Annual Meeting of the Association for Computational Linguistics*, 1999.
8. Garay-Vitoria, Nestor and Julio González-Abascal. Intelligent Word-Prediction to Enhance Text Input Rate. *Proceedings of the second international conference on Intelligent User Interfaces*, 1997.
9. Hindle, Donald. Deterministic Parsing of Syntactic Non-fluencies. *Proceedings of the 21<sup>st</sup> Annual Meeting of the Association for Computational Linguistics*, 1983.
10. Jurafsky, Daniel and James Martin. *Speech and Language Processing*. Prentice Hall, Upper Saddle River NJ, 2000.
11. Katz, Slava. Estimation of Probabilities from Sparse Data for the Language Model Component of a Speech Recognizer. *IEEE Trans. On Acoustics, Speech, and Signal Processing*, 35(3): 1981.
12. Lesher, Gregory and Gerard Rinkus. Domain-specific word prediction for augmentative communication. *Proceedings of the RESNA '02 Annual Conference*.
13. Lesher, Gregory, Bryan Moulton, and Jefferey Higgonbotham. Effects of ngram order and training text size on word prediction. *Proceedings of the RESNA '99 Annual Conference*.
14. Mahajan, Milind, Doug Beeferman, and X. D. Huang. Improved topic-dependent language modeling using information retrieval techniques. *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, 1999.
15. Manning, Christopher and Hinrich Shütze. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge MA, 2000.
16. Newell, Alan, Stefan Langer and Marianne Hickey. The rôle of natural language processing in alternative and augmentative communication. *Natural Language Engineering*, 4(1): 1998, 1-16.
17. Seymore, Kristie and Ronald Rosenfeld. Using story topics for language model adaptation. *Proceedings of*

*5th European Conference on Speech Communication and Technology (Eurospeech), 1997.*

18. Lee, Lillian. Measures of Distributional Similarity. *Proceedings of the 37<sup>th</sup> Annual Meeting of the Association for Computational Linguistics, 1999.*

19. Trnka, Keith, Debra Yarrington, Kathleen McCoy, and Christopher Pennington. The Keystroke Savings Limit in Word Prediction for AAC. *Technical report, 2005.*