Trnka, Keith and Kathleen F. McCoy. *Adaptive Word Prediction for AAC.*ISAAC 2008.

Adaptive Word Prediction for AAC

ABSTRACT

Word prediction in AAC devices can increase communication rate and lower physical fatigue. However, predictions are often inappropriate for the user, so many users lose trust in the predictions. We propose a statistical word prediction system which, in addition to being trained on generic text like most systems, will be automatically customized to the user's style and the conversational context. We can also refine the predictions using actual user text, allowing the user to easily reuse words. The combination of these adaptations should help personalize the predicted words and regain user trust in the system as well as save additional keystrokes.

Keywords: Technology, Theoretical models and issues

INTRODUCTION

High-tech AAC devices assist users with communication difficulties, especially in conversational settings. The user often enters letters and words, which are vocalized by an integrated speech synthesizer. However, developers of AAC devices face numerous challenges. One of the chief difficulties is the communication rate divide between AAC speech and non-AAC speech. Also, hours of typing throughout the day contributes to physical fatigue.

Typical devices seek to address these issues by reducing the number of keys to type a sentence. One way to address the problem is to provide one-hit access to very common words. Another way to reduce the number of keystrokes is to predict the word being typed and provide one-hit access to words the user is likely to be typing. Word prediction generates the predicted words primarily by statistical methods (e.g., word frequency), which processes a large amount of text (training text) to collect the necessary statistics.

However, the data used to create the predictions is oftentimes very different than the user's text. The mismatch between the predictions and user's language can cause the predicted words to be unrelated or inappropriate for the user. Poor quality may be cognitively jarring to the user which may slow him/her down scanning the prediction list. Over time, the user may decide to ignore the prediction list and thus not benefit from the potential keystroke savings.

ADAPTING TO THE USER

We seek to address the mismatch between training text and actual user text by adapting the predictions to be appropriate for the specific user of the device. One way in which the system can adapt to the user is by focusing the predictions on the parts of the training text that are most similar to actual user text. Another way in which the system can improve predictions for the user is by dynamically integrating user text into the statistics.

In this paper, we address the problem of adaptive predictions in three ways: adapting to the topic of conversation (e.g., work, sports), adapting to the style of conversation (e.g., formal/informal, academic/ordinary), and adapting by re-training on actual user text.

Topic Adaptation

The topic of conversation is an important aspect that most prediction systems overlook. In a conversation about sports, we would expect the speaker to use words from a relatively small set relevant to the current topic, such as *ball* and *bat,* but not use irrelevant words such as *pot* or *stove*. We have developed a method for adapting the predictions to the topic by focusing on texts in training that use the same words as the speaker (Trnka et al., 2006). This approach seamlessly adapts the predictions without any explicit user interaction, providing full access to the user's vocabulary at any time, but giving preference to on-topic words.

Our approach to topic adaptation records all words used in the conversation and determines the relevance of each topic by comparing the words used in each topic to the words used in the actual conversation. Then each topic influences the prediction list according to the similarity between the words used in each topic and the conversation. For example, suppose that the training data contains a cooking topic and the user has said *Put the bundt cake in the oven.* The method would give strong preference to topics containing content words such as *bundt*, *cake*, and *oven*, which are good indicators of the cooking topic. This approach allows highly related topics to highly influence the predicted words. However, even unrelated topics in training will have some small influence. Therefore, topic modeling can focus on relevant words when possible and safely fall back to general predictions when an appropriate topic is unavailable. We have tested our approach in topic adaptation and found that it increases keystroke savings and provides more natural predictions (Trnka et al., 2006), even when applied to very different user text (Trnka and McCoy, 2007).

Style Adaptation

The style of conversation is another aspect of language that existing word prediction methods overlook. Unlike topic adaptation, where the conversation itself is being adapted to, style modeling adapts to both the conversation and the user. Also, the style of conversation seems to affect the grammar of language more than the vocabulary. We hope to capture this by investigating the variations in part of speech tags (e.g., singular noun, adjective) across styles. In a small comparison of email and technical writing, we found that pronouns are much more common in emails, while past participle verbs (e.g., *used*, *shown*, *named*) are much more common in technical writing. Pairs of part of speech tags may offer additional insight into stylistic variations. For example, we found that noun-noun compounds (e.g., *word prediction, keystroke savings, topic modeling*) were more common in technical writing, while phrasal verbs (e.g., *write up, figure out, pick up*) were much more common in emails.

Our previous approach in adaptive language modeling focused mainly on sequences of words, but modeling style in text requires some grammatical processing. Fazly and Hirst (2003) demonstrated that a part of speech model can be very effective for word prediction. Therefore, we plan to model language using a similar two-step process: In the first step, the system selects the part of speech of the word based on the previous parts of speech. The second step predicts the word being typed on the basis of the selected part of speech. We plan to apply our experience in topic adaptation to a two-step process for style adaptation: first, by identifying the training styles that are more similar to the current text. Then, once the relevance of each style is determined, adjusting the selected parts of speech to be more stylistically appropriate. The system's prediction about which words are best for each part of speech would remain unaffected by the style of conversation, but because the likelihood of each part of speech is more stylistically appropriate, the system will focus on words that are more appropriate for the grammatical style. We envision that stylistically-adapted predictions will be much more appropriate for the user and conversation, providing a more natural interface to the user and increasing keystroke savings.

User Text Adaptation

The third way in which we envision adapting to the user is to directly apply user text to generating the predictions. We feel that this will help to compensate for the lack of appropriate training text. The primary difficulty is that the amount of user text is likely to be very small when compared to the amount of text used for training the system. However, we have shown in previous work that a small amount of highly appropriate training text in combination with a large amount of general training text can substantially improve keystroke savings over either text collection individually (Trnka and McCoy, 2007). This gives us hope that an adaptive model that combines a small amount of user text with the general-purpose training text will offer a significantly better experience. Also, Wandmacher and Antoine (2006) demonstrated that user text adaptation for word prediction has the potential to increase keystroke savings substantially, especially when the training text is very different from actual usage.

We plan to implement an adaptive approach that is based on a part of speech (POS) model, following Kuhn and de Mori (1990). The POS model utilizes fewer parameters compared to a more traditional model, which makes the POS model more appropriate for dealing with small amounts of text. As with the style model, the system selects the part of speech first, then predicts the word given that it occurs in the hypothesized part of speech. A part of speech system using the user's text would assign part of speech tags to each word and collect statistics as the users types. In predicting a word, the system chooses the part of speech normally, but when predicting the word in the part of speech, the system would consider not only words that occurred in training, but also words that have been typed by the user. The result of the user text adaptation is that the user should be able to reuse words with very little effort.

COMBINING ADAPTATIONS

We have presented three ways in which user text can be integrated into a word prediction system and we feel that these techniques will provide a more natural experience for AAC users while also saving keystrokes. Each adaptation should individually benefit the user, but we are very interested in the combination of all three adaptations. We envision a model in which the text of the current conversation is used to adapt the predictions both topically and stylistically, affecting both the vocabulary and grammar of the predictions. In addition to the topic and style adaptations, we feel that a part of speech model which integrates user text into the model will allow the user to easily access words that he or she has typed already as well as words that are topically and stylistically appropriate.

REFERENCES

Afsaneh Fazly and Graeme Hirst. 2003. Testing the efficacy of part-of-speech information in word completion. In *EACL-03 Workshop on Language Modeling for Text Entry*, pages 9–16.

R. Kuhn and R. De Mori. 1990. A cache-based natural language model for speech recognition. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(6):570–583.

Keith Trnka and Kathleen F. McCoy. 2007. Corpus Studies in Word Prediction. To appear in*ASSETS*.

Trnka, Keith and Kathleen F. McCoy. *Adaptive Word Prediction for AAC.* ISAAC 2008.

Keith Trnka, Debra Yarrington, Kathleen McCoy, and Christopher Pennington. 2006. Topic Modeling in Fringe Word Prediction for AAC. In *ISAAC*.

Tonio Wandmacher and Jean-Yves Antoine. 2006. Training Language Models without Appropriate Language Resources: Experiments with an AAC System for Disabled People. In *ASSETS*.