

# Building Semantic Networks to Improve Word Finding in Assistive Communication Tools

**Sonya Nikolova**

Princeton University, Princeton, NJ 08540 U.S.A.  
nikolova@cs.princeton.edu

**Perry Cook**

Princeton University, Princeton, NJ 08540 U.S.A.  
prc@cs.princeton.edu

## ABSTRACT

Finding words in an assistive communication device can be challenging and time-consuming for individuals with lexical access disorders like those caused by aphasia. These users have persistent difficulties accessing and retrieving words due to impaired semantic links in their mental lexicon. As a result, they can easily get lost in a vocabulary hierarchy or become confused and discouraged if extensive browsing of large word collections is required. We describe the design of the Visual Vocabulary for Aphasia (ViVA) which attempts to provide effective word finding by organizing words in a dynamic semantic network where links between words reflect word association measures, human judgments of semantic similarity, and past vocabulary usage. We present results from preliminary evaluation of ViVA and discuss the challenges inherent to evaluating adaptive assistive communication tools.

## INTRODUCTION

Aphasia is an acquired disorder that impacts an individual's language abilities. It is often acquired as a result of stroke, brain tumor, or other brain injuries. The resulting impairments to the ability to understand and produce language vary in severity and can affect an individual in any combination. Even though rehabilitation can alleviate the level of impairment, a significant number of people with aphasia are left with a life-long chronic disability that impacts a wide range of activities and prevents full re-engagement in life. Research and commercial efforts have shown that technology has the potential to help individuals with aphasia communicate and thus regain some of their independence and social life. Our work is also guided by the idea that high-tech tools can assist people find the words they wish to express. Furthermore, we believe that the effectiveness of such aids can be enhanced by our knowledge of human semantic memory.

Assistive communication tools often fail to meet the functional needs of users because they do not provide an intuitive and quick way for selecting words when composing phrases for communication. This is due to the common choice of organizing the words in such tools in either hierarchies, which tend to be deep and confusing, or simple word collections, which tend to result in extensive scrolling. We have developed a prototype for a visual vocabulary that implements a novel approach to word organization and enables the user to browse for words

efficiently. The visual vocabulary for aphasia (ViVA) organizes words in a semantic network where links between words reflect word association measures and vocabulary usage statistics. In this paper, we describe the design of the vocabulary and share some results from ViVA's initial evaluation. We also discuss the different approaches we have taken to address some of the challenges inherent to evaluating an adaptive tool, especially one that is designed for people who have communication difficulties.

## BACKGROUND

In this section we highlight one common disadvantage of existing assistive communication tools: extensive vocabularies that are challenging to navigate for people with lexical access disorders. To address this problem we propose an assistive vocabulary based on theories that explain how our mind stores and retrieves words. These theories are also introduced in this section.

### Assisting People with Language Impairments

Existing assistive communication tools for people with language impairments share a number of attributes. They make use of picture-based representation of concepts and often provide a multi-modal interface that combines images, text, and speech audio. Most assistive devices also enable phrase composition by assembling words in a linear fashion. For example, if a user wants to communicate the phrase *I am hungry*, she needs to find the icons for the pronoun *I*, the verb *am*, and the adjective *hungry*, and arrange them in the correct order. In order to be expressive, such tools provide vocabularies consisting of thousands of entries. However, previous research has shown that such large word collections are difficult to navigate for most users [4, 6]. It is a true challenge to provide functional communication assistance by enabling the user to intuitively and quickly select words. Many people with aphasia cannot always type the desired word in a search box, but need to browse through the vocabulary until they find a concept that expresses their need. Thus, minimizing the complexity of navigating the vocabulary and supporting efficient word finding is essential to the usability of assistive communication tools.

Naturally, the ease of word finding depends to a large extent on how the words are organized and what the search and retrieval options are. Most existing assistive visual

vocabularies have a lexical organization scheme based on a simple list of words. Some word collections are organized in hierarchies, which often leads to deep and non-intuitive searches; others are simply a list of arbitrary categories which cause excessive scrolling and a sense of disorganization. VocaSpace, featured in a new assistive software product Proloquo2Go [3], for example, organizes the words in functional categories such as greetings and questions, and common word categories such as colors, places, and clothes. The vocabulary can also be extended and personalized, and provides phrase starters such as *I want to* and *I need* [3]. No results from a formal evaluation have been reported yet.

Both TouchSpeak [16,18] and Lingraphica [9], two other commercial assistive communication tools, offer a hierarchical vocabulary that can be enhanced with images from the user's personal collection or ones that have been taken with a digital camera. Lingraphica is designed specifically for people with aphasia and its vocabulary attempts to mimic real-life situations by grouping words according to shared contexts. If you need to find *milk*, for example, you select *kitchen*, then the *fridge* category, then you find *dairy*. The icon for *milk* is in *dairy*. This organization is not necessarily intuitive for all users since, for example, some people associate the word *milk* with the *drinks* category while others may prefer to have it in the *breakfast* category along with *cereal*. This problem is somewhat alleviated by allowing the user to customize the vocabulary categories, but this requires additional effort on the user's behalf. If browsing for words tends to be time-consuming and confusing, it can cause frustration and discourage people from exploring the vocabulary in the future. A few informal trials we performed, asking elderly people to find words using Lingraphica, demonstrated that searching for words is hard even for able people. Having to perform multiple clicks to reach one single words is an additional disincentive for the user. These informal studies and feedback from speech-language pathologists [6], who evaluate and recommend assistive tools to people with language impairments, motivated us to investigate effective approaches to address these problems.

### Speaker's Mental Lexicon

To address the problem of cumbersome vocabulary navigation, we propose a vocabulary organization based on theories that explain how the human mind organizes words. We first appeal to the psychological literature on speakers' "mental lexicon," where words are stored and organized in ways that allow efficient access and retrieval. Every speaker has experienced the inconvenience of temporarily impaired semantic connections (the so-called tip-of-the-tongue (TOT) phenomenon). This inability to retrieve a specific word needed to express a given concept can be due to a variety of causes such as fatigue or interference from a word that is morphologically or phonologically similar to the target word.

Experimental evidence — including evidence from TOT states induced in the laboratory — suggests that words are organized in a speaker's mental lexicon by various similarity relations, in particular phonological and semantic similarity. For example, subjects in word association experiments overwhelmingly respond with *husband* to the stimulus *wife* [11]. Semantic priming [14], a robust and powerful tool for the experimental investigation of cognitive processes, relies on the semantic relatedness of the prime and an experimental target: responses to the target are faster when it is related to the prime as in the classic case *doctor-nurse*. Spreading network activation models [7] assume that presenting a prime stimulus word activates the corresponding representation in lexical memory and that this activation spreads to other related nodes, thus facilitating the processing of related target words. The semantic network WordNet [8, 10] is a large-scale lexical database inspired by network theories of semantic memory that accommodate the spreading activation paradigm among related words and concepts.

To build our vocabulary, we apply these theories and take advantage of WordNet's structure by augmenting a basic hierarchical vocabulary with a collection of semantic associations. Evocation, i.e. how much one word brings to mind another word, is one such association. Evocation is particularly useful for adding cross-part-of-speech links that allow for connections among entities (expressed by nouns) and their attributes (encoded by adjectives); similarly, events (referred to by verbs) can be linked to the entities with which they are characteristically associated. For example, the intuitive connections among *traffic*, *congested*, and *stop* can be clearly conveyed using evocation. We hypothesize that creating such meaningful semantic links between words in the vocabulary will enable effective vocabulary navigation, and improve word search and user satisfaction.

### THE VISUAL VOCABULARY FOR APHASIA

The visual vocabulary for aphasia (ViVA) addresses the problem of cumbersome vocabulary organization and navigation by modeling a speaker's mental lexicon. ViVA is designed to reorganize and update the vocabulary structure according to user preferences and system usage statistics. It is both *adaptable*, able to be customized by the user, and *adaptive*, able to dynamically change to better suit the user's past actions and future needs. This mixed-initiative approach enables the user to feel in control by making changes and anticipating ones that have been initiated by the tool while still allowing adaptive methods to help determine where and when changes are required. The vocabulary's adaptable component allows the user to add and remove vocabulary items, group them in personalized categories, enhance words with images and sounds, and associate existing phrases with a concept. The adaptive component updates the vocabulary organization based on the usage of the system, user preferences and a number of semantic association measures. The result is a semantic

network that reflects user-specific information and general knowledge of human semantic memory. For example, if the user wishes to compose the phrase *I need an appointment with my doctor* and she searches for *doctor* first, the words *medication* and *appointment* may surface (see Figure 1), because they have been linked to *doctor* due to past usage, while *hospital* and *doctor* could be linked due to prediction based on known semantic measures. In addition, the user may be able to find the phrase *Need appointment with my doctor* right away if it has been composed in the past.

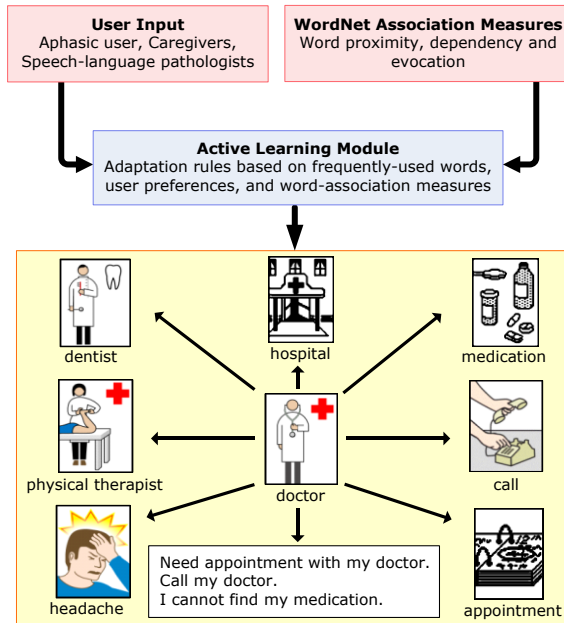


Figure 1. Schematic of components of proposed vocabulary showing an example semantic network centered on *doctor*.

### Core Vocabulary with Evocation Associations

Customization and adaptation of the vocabulary based on its usage can be a powerful feature, but we still need an initial organization to allow the user to successfully use ViVA from day one. We selected ViVA’s core vocabulary such that it is a collection of commonly used words as well as ones relevant to our target population, people who have aphasia. The vocabulary originates from two sources: the “core” WordNet consisting of frequent and salient words and the visual vocabulary provided by Lingraphica [9], which we described briefly earlier in the paper. We used all synsets (synonym sets) from the core 1000 synsets [5], all verbs in Lingraphica’s vocabulary, and all nouns and adjectives in both Lingraphica’s vocabulary and the core 5000 synsets. This initial collection of words was organized to model Lingraphica’s hierarchy.

To compensate for some of the impaired semantic connections in a user’s mental lexicon, we first incorporated in ViVA’s vocabulary network links between words reflecting strength of evocation. While the evocation set created by Boyd-Graber et al. [5] provided us with an initial collection of human semantic association ratings,

many of those ratings were zeros since the synsets were selected and paired randomly. Using machine learning techniques, the structure of WordNet, and this initial collection of evocation ratings, we generated a list of word pairs with predicted high evocation. Over the period of three months, we collected ratings for this new set of 107,550 synset pairs through an online experiment posted on Amazon Mechanical Turk [2]. To ensure the quality of the ratings and a consistency with previous results, we used embedded checks to decide which submitted ratings were valid. The ratings for four of those checks were collected from the dataset available from Boyd-Graber et al. [5]. The fifth check asked subjects to rate the strength of evocation between the same words, e.g. *hungry* and *hungry*. The subset of the data for which the responses to the checks correlated well with the original ratings was then used to augment the basic vocabulary hierarchy with links reflecting how strongly people associate two words [12].

### INITIAL EVALUATION

It would have been challenging to evaluate our adaptive vocabulary, which is designed to assist people with everyday communication, in short laboratory studies. On the other hand, asking users to incorporate an early prototype in their everyday life can cause them much frustration, especially when they experience cognitive difficulties. Thus, we chose to first evaluate ViVA using two alternative approaches.

### Simulating Usage of the Vocabulary

As a first step, we concentrated on evaluating the backend adaptive functionality of ViVA. We simulated sentence composition in ViVA and compared searching for words in ViVA to searching for words in Lingraphica’s hierarchy. Our hypothesis was that reflecting usage statistics and providing word associations will form meaningful and user-specific connections which will make linking words in the process of composing a sentence easier.

We started off with two basic lexical inventories. One was Lingraphica’s current hierarchy of words and the other one, ViVA’s core vocabulary augmented with links between words based on the evocation data that we collected. If the average evocation rating between two words was moderate or high, a link between them was introduced. To simulate usage of the vocabulary, we collected text from five elderly individuals whose blogs were readily available online [1]. They covered topics such as cooking, gardening, health and family. The text from each blogger was broken into 1000 sentences. We then extracted the nouns, verbs, and adjectives from each sentence and created word pairs connecting every two neighbouring words. We used 80% of the data to train ViVA and the rest of the data to test how the system performs when linking words to compose a new sentence from the same blogger.

The links in ViVA’s vocabulary network which reflected evocation and simulated usage shortened the browsing

paths between 51% of the words that appeared next to each other in a sentence from the usage sets (the rest of the paths were the same length as in Lingraphica). On average, 22% of the paths became shorter by two or more steps due to ViVA's vocabulary organization. A naïve baseline test showed that such improvement in shortening the distances between related words cannot be achieved with a random increase in the density of the vocabulary network [12].

While these results were encouraging, they illustrated only the theoretical potential of having an adaptive vocabulary organization. It is difficult to model how a user navigates the vocabulary in practice and how she forms decisions about what to click on next. Having predicted, related words popup during a selection can help the user find the path to the target word more quickly, but it can also be confusing depending on the choice of the related words.

### **Evaluation with Able Users**

Next we evaluated our approach to vocabulary organization with users who do not have language impairments. The goal of the study was to determine whether providing users with related words adapted according to evocation enables them to find words faster than searching through hierarchically organized words. If users have shorter search paths when using ViVA, we argue that this search path efficiency will require less cognitive load and enable users to compose phrases more effectively.

We recruited sixteen people (seven female) from the Princeton University campus. Participants were asked to find a total of 164 missing words in 78 phrases using one of two vocabularies. Each participant was randomly assigned one of two vocabulary organization conditions first and the remaining one second. The first vocabulary which we call LG is a subset of Lingraphica's hierarchy. It consists of 270 words chosen such that it is possible to navigate to all missing words. The second vocabulary, ViVA, inherits the LG hierarchy of 270 words but also provides associations between words based on the evocation data collected by [12]. The associated words are displayed in a related-words panel after a user has clicked on a word.

Within-subjects analysis of variance (ANOVA) showed that significantly fewer clicks were required to find a word in ViVA than in LG ( $F(1, 83) = 33.70, p < .01$ ). In LG, the actual number of clicks used in the study increased consistently as the optimal path length (shortest path in the hierarchy) increased ( $F(1, 83) = 10.82, p < .05$ ). The actual path length in ViVA was not affected by the hierarchy's optimal path length in the same way, indicating that people took shorter paths enabled by the provided related words.

ANOVA results also showed that it took significantly less time to complete all phrases using ViVA compared to using LG ( $F(1, 31) = 35.46, p < .01$ ). For both conditions, it took significantly longer if being tested first indicating an order effect ( $F(1, 31) = 33.70, p < .01$ ). This was anticipated because both vocabularies rely on the same hierarchy.

All participants agreed that having related words automatically suggested helped them find words faster and most thought that finding words in ViVA was less confusing than searching in LG. With ViVA people tended to search for words utilizing the related-words panel instead of trying to locate what category the word should belong to. Five people commented that they did that without analyzing the associations and merely took advantage of what was given. Further details about the experiment and the results could be found in [13].

### **DISCUSSION**

The experiments we described serve as a stepping stone to further evaluation with aphasic users. There are a number of challenges in evaluating tools with people with aphasia. The first obvious one is the difficulty of communicating with users when gathering feedback. In addition, it is often hard to recruit subjects because even though there are approximately 1 million people with aphasia in the US, they are dispersed and most tend to be isolated from society due to their disability. Since most people with aphasia tend to be older, they often shy away from using technology and are intimidated by its introduction in their daily lives. Often the motivation of using technology and its adoption depend on family members and caregivers who are already overwhelmed with other responsibilities. Finally, these users tend to get tired faster than regular subjects which limits the duration of study sessions.

To abstract from some of these challenges, we first evaluated ViVA with able users. Showing significant difference in user performance and satisfaction with fluent English speakers who have intact mental lexicons suggests aphasic people can benefit even more from a tool that aims to compensate for impaired semantic connections. There are a number of interesting issues to consider in evaluating ViVA further. It is unclear how much cognitive load an adaptive vocabulary structure would add to the interaction. On one hand, stability of the user interface is important for people who are already overwhelmed due to age-related difficulties of using technology and the inability to find or associate words. On the other hand, a dynamic vocabulary has the potential to help users find words faster, compose more complex phrases, explore the vocabulary better, and possibly serve for rehabilitation.

The adaptive nature of ViVA introduces yet another challenge in evaluating it. Assistive communication tools are designed to aid users in conversation settings and when expressing daily needs and desires. To evaluate the potential of the adaptive vocabulary, you would need to train and test it on a large scale corpus of text generated by users of assistive communication tools or text used for assistive communication and covering a variety of topics. Unfortunately, as explored by other researchers, no such corpus exists [17]. As a starting point, we chose Lingraphica's vocabulary as our basic organization, because we already had a certain amount of experience working with the tool and because it provides us with a

well-structured vocabulary that has been carefully designed and proven useful (via application in a commercial assistive device widely used in the US) despite its disadvantages. It also consists of a set of common words used in daily communication which has evolved over the years to reflect the needs of the user population as well as feedback from professionals such as speech-language pathologists. We also created a corpus from text found on blogs of elderly people, describing everyday life experiences. All bloggers were above the age of 60 which places them in an age range where there is a higher risk of suffering from stroke and thus acquiring aphasia. Even though they probably lead more active lives than people with aphasia, they have similar social interactions and needs, which makes their communications relevant to the type of data ViVA will be handling. Since we envision a tool that will help compose phrases for communication, we also thought about using data from repositories of switchboard exchanges [15]. However, there would not have been an easy way to filter the data such that they are appropriately representative of our target user population.

It is often less challenging and thus more appealing to test the potential of an adaptive interface theoretically (as we did with our first experiment). The crucial test though is the one involving a real user performing a real task in real context. Thus, we will evaluate ViVA further with people who have aphasia and would benefit from it. We plan to modify the experiment we ran with able subjects in order to run it with people with aphasia. We will also run a longitudinal case study, involving a small number of aphasic users, during which we will measure communication improvement over time and observe how the vocabulary evolves and tailors itself to a user. The results from the case study will give us a more realistic view of the tradeoffs and benefits of long term adaptation applied to real-life situations. We are looking forward to receiving feedback on the described research from the SEMAIS 2010 Workshop participants and are particularly interested in discussing and finding out about other alternative methods for evaluating adaptive tools and user interfaces where training corpus data is not easily available.

## CONCLUSION

We described the design of the Visual Vocabulary for Aphasia which is an assistive communication tool for people with lexical access impairments. ViVA attempts to model a speaker's mental lexicon by incorporating semantic links between words based on word association measures and vocabulary usage statistics. Results from the evaluation of the current prototype of ViVA demonstrate the tool's potential to improve vocabulary navigation and make word finding more efficient. To address some of the challenges inherent to evaluating adaptive technology with people with communication disabilities, we first simulated vocabulary usage by using sentences collected from blogs of elderly

people. In addition, we tested ViVA with able subjects in a study that compared finding words in ViVA to finding words in a commercial assistive vocabulary. While encouraging, these preliminary results are inconclusive so we plan to evaluate our tool involving users with aphasia.

## REFERENCES

- [1] Ageless Project. Accessed November 2009 from <http://jenett.org/ageless/>.
- [2] Amazon.com, Inc. Amazon Mechanical Turk. Accessed November 2009 from <https://www.mturk.com>.
- [3] AssistiveWare. Proloquo2go. Accessed January 2010 from <http://www.assistiveware.com/>.
- [4] Beukelman, D. R. & Mirenda, P. Augmentative and alternative communication: Management of severe communication disorders in children and adults. Brooks Publishing Company, 1998.
- [5] Boyd-Graber, J., Fellbaum, C., Osherson, D. & Schapire, R. 2006. Adding dense, weighted connections to WordNet. *Proc. 3rd Global WordNet Meeting*.
- [6] Boyd-Graber, J., Nikolova, S., Moffatt, K., Kin, K., Lee, J., Mackey, L., Tremaine, M. & Klawe, M. 2006. Participatory design with proxies: developing a desktop-PDA system to support people with aphasia. *Proc. CHI '06*, 151–160.
- [7] Collins, A. M. & Loftus, E. F. A spreading-activation theory of semantic processing. *Psychological Review*, 82(6):407–428, 1975.
- [8] Fellbaum, C. WordNet: An Electronic Lexical Database, chapter: A semantic network of English verbs. MIT Press, Cambridge, MA, 1998.
- [9] Linggraphicare. Accessed January 2010 from <http://www.linggraphicare.com>
- [10] Miller, G. A. 1990. Nouns in WordNet: A lexical inheritance system. *Int. Journal of Lexicography*, 3(4), 245–264.
- [11] Moss, H. & Older, L. Birkbeck Word Association Norms. Psychology Press, 1996.
- [12] Nikolova, S., Boyd-Graber, J., Fellbaum, C. & Cook, P. Better Vocabularies for Assistive Communication Aids: Connecting Terms using Semantic Networks and Untrained Annotators. *Proc. ASSETS'09*. 171–178.
- [13] Nikolova, S., Ma, X., Tremaine, M. & Cook, P. Vocabulary Navigation Made Easier. *Proc. IUI'10 (to appear)*.
- [14] Swinney, D. Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, (18):645–659, 1979.
- [15] SWITCHBOARD: A User's Manual, 2007. Accessed January 2010 from [http://www ldc.upenn.edu/Catalog/readme\\_files/switchboard.readme.html/](http://www ldc.upenn.edu/Catalog/readme_files/switchboard.readme.html/)
- [16] TouchSpeak & TypeSpeak. Accessed January 2010 from <http://www.touchspeak.co.uk/>
- [17] Trnka, K. and McCoy, K. F. 2007. Corpus studies in word prediction. *Proc. ASSETS'07*. 195–202
- [18] van de Sandt-Koenderman, M., Wiegers, J., & Hardy, P. 2005. A computerized communication aid for people with aphasia. *Disability and Rehabilitation*, 27, 529–533